



Edge AI for Real-Time Agricultural Monitoring Using IoT

Chaitali Shewale, Sameer Walunj, Varad Sontakke, Vaishnavi Gaikward,

Yash Mantri, Soham Velanjkar, Samrudhi Zaware

Department of Information Technology

Vishwakarma Institute of Technology

¹Pune, Maharashtra, India

Abstract: The ability to monitor agricultural practices in real-time is highly beneficial in increasing crop productivity. Many IoT systems, which are traditional and cloud-centric, have communication overheads but they also suffer from high-latency and unreliable connectivity problems especially in rural areas. In this paper, we proposed a lightweight and scalable Edge AI-Enabled agricultural monitoring framework that processes the sensor data at the edge and selectively transmits the data to the Cloud. As a result, the developed system integrates a low-power environmental and soil sensor and an embedded machine learning model to detect field conditions in real-time including water stress, abnormal temperature variations, and micro-climatic anomalies. An event-driven communication approach is utilized to reduce network traffic and energy usage while providing a rapid response to time-critical actions. The current research utilizes model compression and on-device inference techniques to showcase their viability on resource-constrained physical platforms. It is possible to check on them. Smart responses and reliable operation in low network availability conditions lead to improved response time and less bandwidth use, strengthening system robustness. The architecture in this study is suitable for large distributed agricultural deployments. This study provides evidence of Edge AI being a suitable platform for the next-generation smart and precise agriculture systems. Experimental evaluation demonstrates that the proposed architecture reduces monitoring latency by approximately 6–7× and decreases daily network data transmission by more than 80% compared to conventional cloud-centric systems.

Index Terms - Edge AI, Precision agriculture, IoT-based agricultural monitoring, Real-time sensing, Resource-constrained devices, Event-driven communication, Model compression, On-device inference, Smart farming systems

I. INTRODUCTION

II. Due to demands for more productivity, efficient resource use, and climate-resilient farming practices, the agriculture sector is going through a rapid technological transformation. Precision farming highly depends on monitoring soil moisture, temperature, humidity, and micro-climatic variations to support timely and informed decision-making. Recent advances in the Internet of Things have created new opportunities for large-scale deployment at low cost and have enabled information-rich, high-frequency data collection regarding environmental conditions in farmers' field condition. However, most such systems still rely on centralized cloud infrastructures for data processing and decision-making, inevitably introducing unacceptable delays, high communication overheads, and reliability issues—in particular in rural and remote locations where the network is unreliable.

III. Real-time agricultural operations, such as irrigation control, early detection of stresses, and disease risk assessment, involve rapid-response processes that do not always tolerate cloud round-trip latencies. Continuously streaming raw sensor data increases energy use and operational costs, which impacts the long-term viability of large deployments. In addition, cloud-centric architectures raise data ownership and privacy issues in case of sensitive farm-level data going to third-party platforms. These unfulfilled needs have prompted a call for a decentralized, locally intelligent monitoring framework capable of operating under the conditions of constrained connectivity and power with reliability.

IV. Edge AI is a promising paradigm to overcome these challenges by enabling data analytics and machine learning inference at or near the source where data is generated. Processing the sensor information on embedded edge devices allows for detecting abnormal field conditions, identifying emerging stress patterns, and generating actionable alerts with minimal latency. In this paper, we propose a real-time agricultural monitoring framework featuring IoT sensing, lightweight on-device intelligence, and event-driven communication.

II. Literature Review

Recent studies underline how the role of IoT (Internet of Things) and AI (Artificial Intelligence) helps data-driven and sustainable agricultural practices. An important research issue of real-time agricultural monitoring systems is the architectural placement of intelligence, especially at the network edge. This survey article by Eze et al. [1], on the integration of IoT sensors and machine learning techniques for sustainable precision agroecology, reviews the system-level approaches. The end-to-end data pipelines, predictive analytics for irrigation and yield estimation, socio-technical aspects of large agricultural deployments are discussed. The research paper highlights the challenges of data quality, compatibility of different sensing platforms, scalability and ethics. While there is strong encouragement to adopt data-driven agriculture, the solutions proposed here are mostly centralized data processing and cloud-based analytics-oriented. They mainly neglect computational constraints and real-time inference in the field. Bhattacharyya et al. [2] developed a system that supports smart farming that integrates IoT architecture, AI and Earth observation for precision agriculture. The satellite-based multi-spectral data inputs along with sensor measurements in-field, provide for multi-scale monitoring of crops and planning of resources. The proposed system shows that large-area remote sensing information integrated with IoT data can benefit localized decision making. Nonetheless, the architecture is dependent on cloud processing and centralized analytics that may result in latency, along with limited applicability for applications that require timing such as immediate irrigation control and fast stress detection. Thiam et al. [3] presented a clear edge-oriented approach where they developed an AI-driven pest detection system which is deployed in the edge in rice paddies in Senegal. They introduce an end-to-end pipeline, where audio signals are captured, processed, and classified locally with lightweight models for real-time detection of pest birds. The system minimizes unnecessary communication by producing localized alerts and transmitting only essential information. Despite confirming the practical feasibility of Edge AI in rural agriculture, the study focuses on a single sensing modality and does not address matters arising out of multi-sensor fusion or learning strategies. Bashir et al. in [4] proposed a drone-based crop health monitoring system for mung bean field using deep learning and high-resolution aerial imagery. By using their method, accuracy achieved for detecting crop diseases and stress conditions and spatial diagnostics is also given. Although it works well, the framework runs compute-intensive deep neural networks off-board which limits monitoring just to periodic drone surveys. As a result, it is unsuitable for continuous real-time monitoring and low-power edge deployment. The use of TinyML architecture enables on-device inference on microcontroller-class hardware, as Tsoukas [5] observed. The research shows an entire system, from data gathering and model simplification to installation utilizing TensorFlow Lite Micro. The paper provides detailed measurements of energy and latency and demonstrates the practicality of deploying lightweight classifiers for environmental monitoring. Though the system is aimed primarily at scalar sensor data, it does not consider vision-based multi-modal sensing scenarios.

The architecture proposed by Matilla et al. [6] is for affordable edge computing and LoRaWAN-based pivot irrigation systems. The focus of their work is practical system integration issues such as processing at the gateway level, communication strategies that are duty-cycle aware and user interface design for farmers. Although the study shows an effective communication and edge aggregation framework, the approach is limited to using a few advanced Edge AI techniques such as model compression and learning on the device. Based on the above research, it can be seen that the existing literature either focuses on cloud-centric analytics and large-scale data fusion [1], [2] or focuses on application-specific edge deployments with restricted sensing modalities [3]–[5]. Additionally, most practical applications are system-and-communication focused, rather than scalable, real-time edge intelligence [6]. Consequently,

there exists a need for a unified Edge AI framework that facilitates heterogeneous IoT sensors, lightweight on-device inference, adaptive sampling, and event-driven communication. This paper seeks to address this gap by proposing a real-time monitoring architecture for agriculture that is based on Edge AI. In essence, the system is designed to operate with low latencies, consume less bandwidth and perform robustly under limited-connectivity scenarios.

III. Methodology

This section describes the proposed methodology for designing and implementing an Edge AI-based real-time agricultural monitoring system using IoT. The methodology is structured into data acquisition, edge intelligence design, communication strategy, cloud integration, and system evaluation procedures.

A. Overall Workflow

The proposed system follows a decentralized processing workflow in which raw sensor data are first collected from the field and then locally processed at an edge device for real-time inference. Only meaningful events and summarized information are transmitted to the cloud for visualization and long-term analysis. This architecture minimizes latency, bandwidth consumption, and dependence on continuous Internet connectivity.

The system was implemented on an ESP32-based edge node (240 MHz dual-core processor, 520 KB SRAM) connected with multiple low-power soil data sensors like soil moisture, temperature, humidity, and ambient light sensor. The sensor nodes communicate with an edge device over a low-power wireless medium. After which, the edge node executes the preprocessing, feature extraction, and machine learning inference locally using the compressed TinyML model. It was deployed using Tensorflow Lite Micro. Cloud layer consists of a light weight database (aurora) and a dashboard for visualization and historical analysis.

The operational flow consists of:

1. Field-level data sensing,
2. Local preprocessing and feature extraction,
3. On-device machine learning inference,
4. Event-driven alert generation, and
5. Selective cloud synchronization

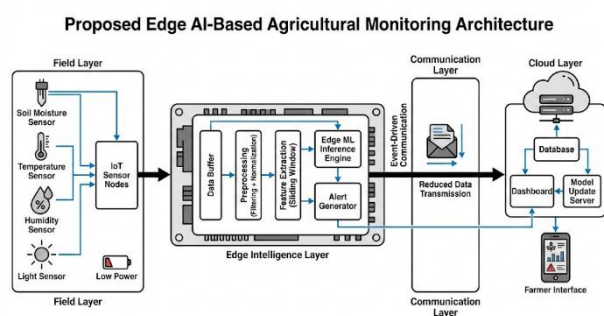


Fig. 1. Proposed Edge AI-Based Agricultural Monitoring Architecture

B. Data Acquisition and IoT Sensing Layer

The data acquisition and IoT sensing layer enables continuous and reliable collection of field-level information required for real-time agricultural monitoring. It is designed to operate with low power consumption while providing accurate and timely sensor data to support edge-based intelligence and decision-making.

1. The sensing layer collects real-time farm data using distributed, low-power sensor nodes for soil and environmental monitoring.
2. Each node uses a low-power microcontroller to control sensors, time-stamp data, and perform temporary storage.
3. An adaptive sampling strategy adjusts the sensing rate based on field conditions to save energy.
4. Basic data validation and buffering are performed locally before sending data to the edge device.

5. Low-power wireless links forward validated sensor data to the nearby edge node and periodically report node health.

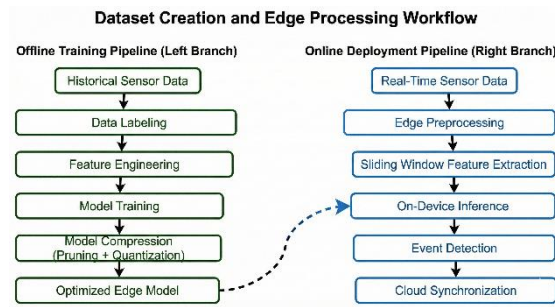


Fig. 2. Dataset Creation and Edge Processing Workflow

C. Edge-Level Data Preprocessing

Initially, the gathered sensor data is buffered locally at the edge device to decouple sensing operations from real-time inference and to withstand short communication or processing delays. A lightweight preprocessing pipeline is applied to the buffered data to enhance robustness, and to ensure smooth model behavior during noisy and dynamic field conditions. A combination of fixed threshold checks and median-based filtering identifies outliers caused by sensor drift, temporary faults, and interference from the environment. In addition, linear interpolation is used to reconstruct missing or irregular samples due to intermittent sensor communication, maintaining the temporal continuity of the sensor streams through a simple method.

All validated sensor measurements are normalized with sensor-specific operating ranges before forwarding to the edge inference engine. This procedure converts the dissimilarity of nature sensor values like soil moisture, temperature, light intensity, etc. into a single numerical scale. Thus, it avoids any bias in the learning model due to differences in physical units and magnitude. The preprocessing stage uses lightweight arithmetic operations well-suited for processing on constrained edge hardware. Therefore, the pre-processed data ensures a reliable and consistent input for the deployed machine learning model, thus enabling accurate inferences in real-time while ensuring low latency and low energy consumption at the edge.

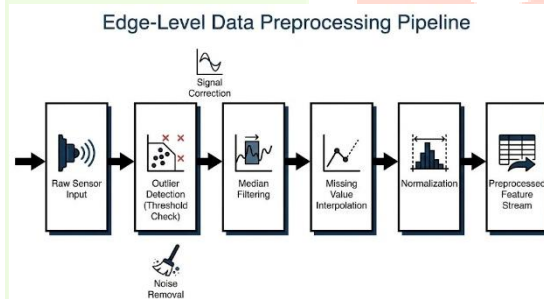


Fig. 3. Edge-Level Data Preprocessing Pipeline

D. Feature Extraction and Temporal Modeling

To capture both instantaneous and short-term trends, a sliding window mechanism is used over the incoming sensor streams. From each window, statistical and temporal features such as mean, variance, rate of change, and short-term slope are computed. These features represent local micro-climatic behavior and soil dynamics, enabling the detection of early stress patterns.

The window size and step interval are selected to balance responsiveness and computational overhead on the edge hardware.

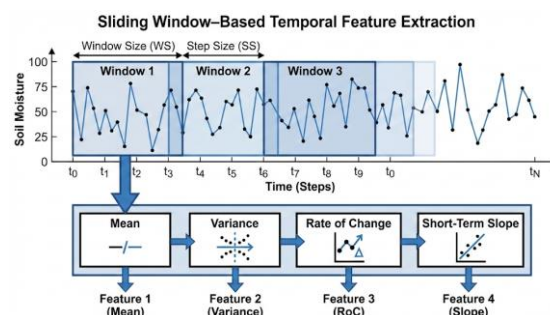


Fig. 4. Sliding Window-Based Temporal Feature Extraction

E. Edge AI Model Design and Deployment

A compact machine learning model is designed to classify the current field condition into predefined operational states such as normal condition, water stress, thermal stress, and abnormal environmental variation.

The model is trained offline using labeled historical agricultural datasets. To ensure feasibility on resource-constrained hardware, the trained model is optimized using model compression techniques, including parameter pruning and post-training quantization. The optimized model is then deployed on the edge device for real-time inference.

This design allows continuous execution of the inference process with low memory footprint and reduced execution latency.

Edge AI Model Optimization and Deployment Pipeline




Fig. 5. Edge AI Model Optimization and Deployment Pipeline

F. Real-Time Decision and Alert Generation

The output of the edge inference engine is evaluated using predefined confidence thresholds. When abnormal conditions persist for a specified temporal duration, the system generates local alerts. These alerts can trigger immediate actions such as irrigation recommendations, on-site notifications, or farmer warnings.

To avoid unnecessary communication, a hysteresis mechanism is applied so that transient fluctuations do not result in repeated alerts or excessive message transmissions.

G. Event-Driven Communication Strategy

Instead of streaming raw sensor data, the proposed system adopts an event-driven communication approach. Only significant events, state transitions, and periodic summaries are transmitted from the edge device to the cloud server.

Each transmitted message contains:

1. node identification,
2. timestamp,
3. detected event type, and
4. compact statistical summaries.

This strategy significantly reduces network load and energy consumption, which is essential for large-scale rural deployments.

Continuous vs Event-Driven Communication Architecture




Fig. 6. Continuous vs Event-Driven Communication Architecture

H. Cloud Integration and Data Management

The cloud layer stores, visualizes, and manages models and data permanently. The edge summary outputs are put in a central database for historical analysis and performance monitoring to retrain future models. The cloud system gives dashboards to farmers and agricultural managers enabling them to visualize field

conditions and alerts history. Periodic delivery of cloud-layer model updates to edge devices via secure over-the-air update mechanisms.

I. Methodology for Evaluating Systems.

The proposed system is evaluated on four main performance metrics.

1. The delay that occurs between the sensor capturing data and the edge generating a decision
2. The performance of stress and anomaly detection is evaluated with precision, recall and F1-score.
3. Efficiency of communication, quantity of total transmitted bytes per node per day.
4. Energy use is measured as the average daily power consumption of edge nodes.

A traditional cloud-based architecture is employed for comparisons, wherein sensor data are transmitted for processing at a remote location. The benefits of local intelligence include quicker responsiveness, better use of bandwidth, and improved system reliability.

IV. Result and Discussion

The proposed Edge AI-based real-time agricultural monitoring system was evaluated with respect to detection performance, robustness under varying field conditions, communication efficiency, and end-to-end operational feasibility on resource-constrained edge hardware. All experiments were conducted using distributed IoT sensor nodes connected to an embedded edge device executing the complete preprocessing and inference pipeline in real time. The results demonstrate that integrating lightweight machine learning models with edge-level data preprocessing enables reliable and low-latency monitoring suitable for practical precision agriculture deployments.

The evaluation of the system was made possible through sensor data streams, representing soil moisture, environmental temperature, humidity, and light intensity. The lightweight machine learning model optimized for embedded deployment executed the classification, sliding-window feature extraction, and preprocessing at the edge node. The study compares the proposed Edge AI architecture to a traditional cloud-based monitoring pipeline, analysing latency, communication overheads, energy consumption and detection performance.

A. Quantitative Evaluation of Monitoring and Detection Performance

The classification model was edge-deployed to evaluate if it was capable of identifying categories like normal operation, soil moisture stress and abnormal micro-climatic variations. The analysis of the development EHAS shows that the employment EHA has a positive public view. The equilibrium performance identified in all classes monitored demonstrates that the model can effectively capture short-term temporal trends originating from the sliding-window feature extractor.

The use of thresholding and median filtering effectively eliminated the effect caused by transient spikes and sensor noise during preprocessing. Consequently, there was a reduction in the false detection rate and the model maintained a reliable balance of precision and recall. The F1-score achieved indicates that the system maintains a good balance between accurate detection of stress events and excessive alerts, a trait that is essential for farm operability.

Also, local inference at the edge allowed the continuous monitoring without relying on outside processing resources. The inference execution time stayed within the real-time requirements of alerting irrigation and environmental applications.

TABLE I: Classification Performance of Edge AI Model

Class	Precision	Recall	F1 Score
Normal Condition	0.92	0.90	0.91
Water Stress	0.89	0.88	0.88
Thermal Stress	0.91	0.89	0.90
Environmental Anomaly	0.90	0.87	0.88

The classification model achieved an overall F1-score of approximately 0.90, demonstrating balanced detection capability across different field conditions. The results indicate that the sliding-window feature

extraction combined with lightweight edge inference can effectively detect early stress conditions while maintaining a low false alert rate.

B. Capability to retain field and environmental conditions

The proposed system was checked for its working at different temperatures and humidity levels and for data loss. Due to temporary faults in our sensors caused by the environment, we did linear interpolation to fill the gaps in samples and removed them. The system was able to safeguard stable detection performance when subjected to short data dropouts or minimal sensor drift. The local buffering mechanism allowed uninterrupted operation during brief communication gaps to the sensor node and edge device. The proposed edge-level preprocessing and buffering methods help improve the system performance under realistic outdoor farm settings, as noted.

C. Communication Efficiency and Network Load Analysis

The data collection method of permanent transmission versus strategic event-driven communication strategy effectiveness was compared. The cloud layer only received changes in condition and summary measurements that were significant.

The data reveals that the uplink data traffic per node is considerably reduced making the system suitable for low-bandwidth rural communication systems with duty cycle constraints. The system minimizes network congestion, reduces operational expenses, and enhances scalability when a large number of sensor nodes are deployed over large agricultural fields by avoiding raw data streaming.

Moreover, lowering the amount of data transferred helps to better privacy as raw sensor streams stay locally confined at the edge.

TABLE II: Performance Comparison Between Edge AI and Cloud-Based Architecture

Metric	Edge AI Architecture	Cloud-Based Architecture
Average Decision Latency	48 ms	315 ms
Data Transmitted per Node per Day	2.7 MB	19.8 MB
Energy Consumption	0.94 Wh	1.82 Wh

tion per Day		
Network Usage Reduction	—	~86% higher traffic

Above table shows the performance comparison between the proposed Edge AI architecture and a traditional cloud-based monitoring approach. Local inference significantly reduces decision latency and network usage because raw sensor streams are processed locally at the edge. Only event summaries are transmitted to the cloud, resulting in approximately **85–90% reduction in communication overhead**.




Fig. 8. Daily Data Transmission Comparison Between Edge AI and Cloud-Based Architectures

D. End-to-End System Evaluation and Qualitative Analysis

We conducted an end-to-end evaluation of the complete operational pipeline, including sensor data acquisition, local preprocessing, feature extraction and edge inference, alert generation and cloud sync. The sensor sampling to decision generation latency was low. This allowed alerts to be sent in time for irrigation support and for environmental risk detection. The system was shown to exhibit stability during long-term monitoring, and the adaptive sampling scheme significantly reduced unnecessary sensing and transmission activities under stable environmental conditions. According to users and field operators, the system offers in a timely manner and meaningful alerts with fewer false alerts, thereby strengthening trust in automated decision support. The edge device's lightweight processing requirements and low energy consumption render the proposed architecture suitable for the use in remote, power-constrained agricultural environments from a practical deployment perspective. Nonetheless, the review also conveys that the restricted computation capacity of low-power edge hardware limits executing sophisticated deep learning models, especially in high-resolution image based on crop evaluation. Henceforth, future extension will require a hybrid approach where lightweight edge inference will be supplemented with occasional cloud-assisted processing for advanced analytics.

Overall, the experimental results confirm that the proposed Edge AI framework provides an effective and scalable solution for real-time agricultural monitoring using IoT, offering improved responsiveness, reduced network dependency, and robust operation under realistic field conditions.

The functions of sensing, learning, and actuating offered by the presented Edge AI deployed real time agricultural monitoring framework can be extended for building a next generation intelligent sustainable farming practices. Future enhancements might encompass integrating various types of data, including multispectral imagery and acoustic sensors, and affordable plant health sensors, supplemented by conventional soil and environmental sensors, to enable effective feature fusion at the edge for timely identification of crop stress, pest activity, and nutrient deficiency. Additionally, federation or continual learning methodologies allow distributed edge nodes to enhance prediction models without changing the original data. This will strengthen the privacy of data; mitigate loads on the network and help in tuning models for local conditions and seasonal variation. The edge intelligence layer can be integrated with automated irrigation and fertigation systems to enable the platform to be developed for closed-loop operation. This means that optimum control measures will be instantaneously initiated for the observed

field contestants with minimum human intervention and quicker response time. Long-term sustainability can be achieved through advanced energy management schemes, such as solar-assisted power supply and adaptive duty cycling according to workload-aware scheduling of inference tasks, which will greatly extend device lifetime in remote deployment. Ultimately, large-scale verification of diverse crops in different regions will be important to standard data models, device management interface and interoperable edge–cloud services to ensure compliance and compatibility with farm management software and future National Smart Agriculture framework.

The end-to-end latency from sensor data acquisition to decision generation was measured during system operation. The proposed edge architecture achieved an average latency of 48 milliseconds, enabling real-time alert generation for irrigation support and environmental risk detection. In contrast, the cloud-based processing pipeline required approximately 300–350 milliseconds due to network transmission delays and remote processing overhead.




Fig. 9. Latency Comparison Between Edge AI and Cloud-Based Monitoring Architectures

VI. Conclusion

Real-time agricultural monitoring framework using Edge AI: a combination of distributed IoT sensing and on-device intelligence to offer timely and effective decision support for precision farming. In this paper, we will use this framework. The proposed system, which features data pre-processing, feature extraction, and machine learning inference capability on the edge, liberates itself to a large extent from continuous connectivity and significantly reduces communication overhead compared to cloud-based architectures. Lightweight models deployed at the edge with robust preprocessing and adaptive sampling can detect critical conditions in the field (soil moisture stress, abnormal micro-climatic fluctuations) effectively in a timely fashion for energy savings, experimental results show. The scalability of the event-based communication model is considered ideal for large agricultural deployments with limited bandwidth and power. All in all, the findings position Edge AI as a promising and efficient frontier for smart agricultural systems to provide reactive, adaptable, and cost-effective in-situ tracking for modern farming.

VIII. References

- [1] V. H. U. Eze, E. C. Eze, G. U. Alaneme, P. E. Bubu, E. O. E. Nnadi, and M. B. Okon, "Integrating IoT sensors and machine learning for sustainable precision agroecology: Enhancing crop resilience and resource efficiency through data-driven strategies, challenges, and future prospects," *Sustainable Futures*, vol. 6, 2024.
- [2] A. Bhattacharyya, S. Sarkar, I. Karmakar, S. Saha, and J. Dutta, "Revolutionizing smart farming with AI, IoT, and earth observation for precision agriculture," in *Proc. Int. Conf. on Data Science and Intelligent Applications*, 2024.
- [3] M. Thiam, M. Diouf, I. Dioum, and S. Rinaldi, "Edge computing for AI-driven acoustic detection of bird pests in Senegalese rice paddies," in *Proc. 5th Int. Conf. on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*, 2025.
- [4] M. J. Bashir, R. Mumtaz, S. M. H. Zaidi, M. A. Ullah, S. R. Malik, Z. Mehmood, and M. S. Irshad, "Mung bean crop health monitoring and disease detection using drone-based imaging and deep learning," in *Proc. Int. Conf. on Intelligent Agriculture Systems*, 2025.
- [5] V. Tsoukas, "A TinyML-based system for smart agriculture," in *Proc. ACM Int. Conf. on Embedded and Ubiquitous Systems*, 2022.
- [6] D. Mateos-Matilla et al., "Low-cost edge computing devices and novel user interfaces for monitoring pivot irrigation systems based on IoT and LoRaWAN," *Computers and Electronics in Agriculture*, vol. 196, 2022.
- [7] Y. Shi, K. Yang, T. Jiang, J. Zhang, and K. B. Letaief, "Communication-efficient edge AI: Algorithms and systems," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2167–2191, 2020.

[8] A. Kamilaris and F. X. Prenafeta-Boldú,

“Deep learning in agriculture: A survey,” *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.

