ISSN (Online): 3139-7689

Available online: www.ijsrmes.com

International Journal for Scientific Research in Modern Engineering and Science

Published in Conference Special Issue: Vol. 1, Special Issue 1, pp. 10–20, 2026

INTEGRATED SMART FARMING SYSTEM UTILIZING IOT, MACHINE LEARNING, AND QLORAFINE-TUNED LLMS FOR CROP RECOMMENDATION AND AGRICULTURAL ADVISORY

Yashraj Kumar[1*], Mohammad Saqib Hasan[2], Aditya Pathak[3], Anmol Gupta[4], Vijaya Tripathi[5], Shankar Sharan Tripathi[6]

[1,2,3,4]*Student, Computer Science and Engineering, Shri Shankaracharya Technical Campus.*

[5,6]*Assisant Professor, Computer Science and Engineering, Shri Shankaracharya Technical Campus.*

Abstract

The agricultural sector is undergoing a critical transformation as it strives to meet rising global food demands amid shrinking arable land, water scarcity, and declining soil fertility. Traditional intuition-driven farming practices are no longer sufficient to ensure sustainable productivity. This paper introduces "Farm Oracle," an intelligent smart farming framework that integrates Internet of Things (IoT) sensors and Artificial Intelligence (AI) to enable data-driven agricultural decision-making. The system continuously monitors essential soil parameters—including Nitrogen (N), Phosphorus (P), Potassium (K), soil moisture, temperature, and pH—using a NodeMCU-based embedded architecture connected to a multi-parameter soil sensor via RS485 communication. Real-time telemetry is stored in a Firebase NoSQL database, processed through a Spring Boot backend, and presented through an interactive React-based dashboard. A comparative evaluation of six machine learning classification models was conducted on a dataset containing 2,200 samples across 22 crop types. Experimental results indicate that Random Forest and Naïve Bayes classifiers achieved prediction accuracies above 95%, demonstrating strong suitability for crop recommendation tasks. Additionally, the system incorporates a locally optimized generative AI module using the Mistral-7B Large Language Model with Quantized Low-Rank Adaptation (QLoRA), enabling context-aware agronomic guidance on resource-constrained platforms. Farm Oracle provides a scalable, intelligent, and farmer-centric solution for precision agriculture.

Keywords: *Smart Agriculture, Precision Farming, Internet of Things (IoT), Machine Learning, Crop Recommendation System, Soil Health Monitoring, Random Forest, Naïve Bayes, Generative AI, Large Language Models (LLM), QLoRA, Edge Computing, Sustainable Agriculture, Decision Support Systems.*

** Corresponding author*

1. Introduction

Agriculture continues to serve as the foundation of global food security and rural economic development. With the world population projected to approach nearly 10 billion by 2050, food production systems must expand their efficiency and resilience to sustain future demand [1]. However, this expansion cannot rely on increasing cultivated land alone, as arable land availability is steadily declining due to urbanization, soil degradation, and climate change. Water scarcity, nutrient depletion, and unpredictable weather patterns further complicate agricultural productivity. These emerging constraints highlight the urgent need for intelligent, technology-enabled farming systems that maximize yield while preserving natural resources.

Conventional farming practices have traditionally depended on experiential knowledge, seasonal intuition, and observational decision-making. While such practices have supported agriculture for centuries, they often overlook the intricate and nonlinear relationships between soil chemistry, environmental conditions, and crop physiology. Soil nutrients such as Nitrogen (N), Phosphorus (P), and Potassium (K), along with moisture, pH, and temperature; interact dynamically to influence plant growth. Recent research in precision agriculture emphasizes the importance of sensor-based monitoring and data-driven analytics to optimize these variables in real time [2], [3]. The integration of Internet of Things (IoT) technologies has enabled seamless data acquisition from distributed agricultural fields, forming the backbone of modern smart farming ecosystems.

The proposed system, **Farm Oracle**, is designed to create a deterministic and empirical linkage between raw environmental parameters and actionable agricultural insights. IoT-enabled soil sensors continuously capture multi-parameter soil metrics and transmit them to a cloud infrastructure for storage and processing. This architecture ensures scalability, reliability, and near real-time responsiveness. Machine learning (ML) models analyze the collected dataset to recommend the most suitable crop for given soil conditions. Studies have demonstrated that supervised learning algorithms such as Random Forest, Naïve Bayes, and Support Vector Machines achieve high accuracy in crop recommendation tasks when trained on structured agricultural datasets [4], [5]. By leveraging these validated approaches, Farm Oracle enhances predictive reliability while maintaining computational efficiency.

While predictive ML models effectively answer *what crop to plant*, they do not address the equally critical question of *how to cultivate it optimally*. To bridge this gap, Farm Oracle integrates a Generative Artificial Intelligence (GenAI) advisory layer capable of interactive, context-aware agricultural guidance. Large Language Models (LLMs) have recently demonstrated remarkable capabilities in domain-specific knowledge synthesis and conversational support systems [6]. By embedding an LLM-based assistant, the system provides farmers with customized recommendations related to irrigation scheduling, fertilizer application, pest management, and best agronomic practices, thereby transforming static predictions into dynamic advisory services.

However, deploying multi-billion parameter LLMs such as Mistral-7B within agricultural infrastructures presents significant hardware and energy constraints. Conventional fine-tuning requires substantial GPU memory and computational power, which may be infeasible in rural or edge-assisted platforms. To overcome this limitation, Farm Oracle adopts Low-Rank Adaptation (LoRA), a parameter-efficient fine-tuning technique that introduces trainable rank-decomposition matrices into pre-trained neural networks while freezing original weights [7]. This method significantly reduces memory usage and accelerates training without compromising performance. Furthermore, the use of Quantized LoRA (QLoRA) enhances computational efficiency by enabling reduced-precision operations while maintaining model robustness [8].

By synergistically integrating IoT-based environmental sensing, machine learning-driven crop prediction, and optimized generative AI advisory capabilities, Farm Oracle represents a holistic advancement in precision agriculture. The framework not only enhances productivity but also democratizes access to intelligent farming insights, making advanced agronomic knowledge accessible to farmers regardless of scale. Such intelligent systems are expected to play a pivotal role in achieving sustainable agricultural transformation in the coming decades.

2. Methodology

2.1 Hardware and Sensor suite

The hardware foundation of Farm Oracle is built around the NodeMCU development board based on the ESP8266 System-on-Chip (SoC). This compact and cost-effective microcontroller operates at 3.3V logic level and supports Wi-Fi connectivity, making it highly suitable for IoT-enabled agricultural deployments. Its input voltage typically ranges between 4.5V and 6V, allowing flexible power integration with field-deployed energy sources such as battery packs or solar modules. The open-source ecosystem of the ESP8266 further enables rapid prototyping and scalable implementation in rural environments.

To enable precise soil monitoring, the system integrates a 6-in-1 soil sensor capable of measuring Nitrogen (N), Phosphorus (P), Potassium (K), soil moisture, temperature, and pH. The NPK parameters are measured in milligrams per liter (mg/L), providing quantitative insight into soil fertility status. However, the sensor outputs data using the RS485 differential communication protocol, which cannot be directly interpreted by the ESP8266's TTL-based UART interface. To bridge this incompatibility, a MAX485 transceiver module is employed. This module converts RS485 differential signals into TTL-level signals compatible with the microcontroller, ensuring reliable bidirectional serial communication even over longer cable distances.

The integrated hardware design thus ensures stable data acquisition, noise immunity, and energy efficiency, making it appropriate for real-time agricultural telemetry in field conditions.

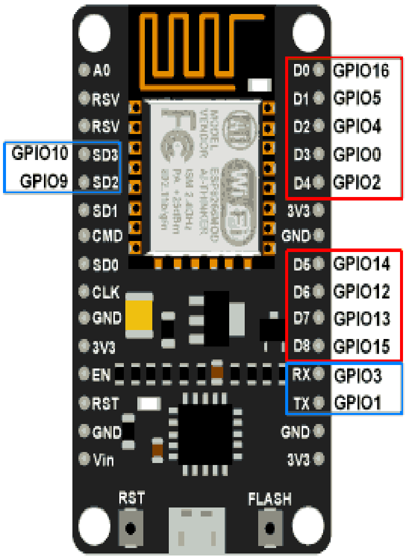

Figure 1: NodeMCU ZY8266

2.2Data Preprocessing

The experimental dataset used in this study contains 2,200 records representing 22 distinct crop categories. Each record includes soil and environmental attributes such as Nitrogen, Phosphorus, Potassium, temperature, humidity, and pH. Before training machine learning models, systematic preprocessing was performed to ensure data integrity and predictive reliability.

Initially, the dataset was examined for missing values, outliers, and inconsistent measurements. Statistical normalization techniques were applied to scale numerical attributes, preventing features with larger ranges from dominating the learning process. To evaluate inter-feature relationships and detect multicollinearity, a Pearson correlation matrix was computed. The correlation heatmap revealed moderate dependencies among certain variables, but no severe multicollinearity that would degrade model stability.

Extensive simulation experiments were conducted using different feature subsets. Interestingly, combinations of Temperature, Humidity, and pH demonstrated strong predictive capability, achieving classification accuracies up to 96.59% when trained using the Random Forest algorithm. This finding suggests that while macronutrients are important, environmental variables also play a decisive role in crop suitability modeling.

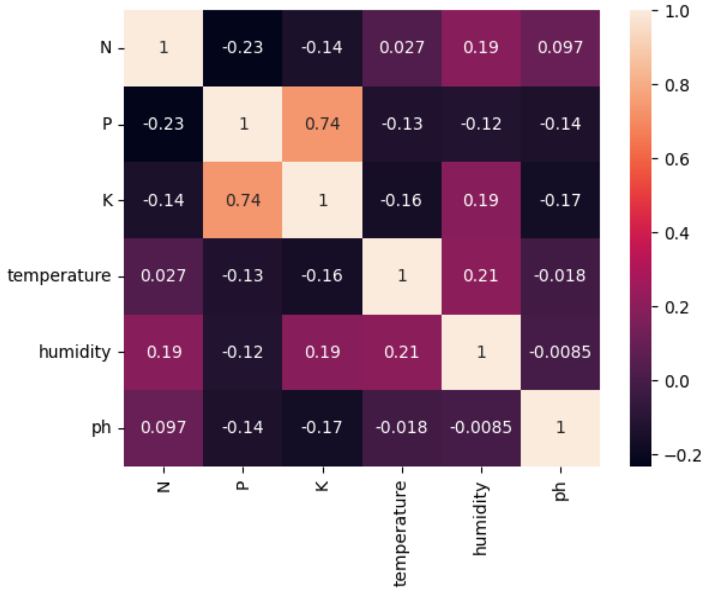

Figure 2: Heat map of parameters

2.3 Machine Learning Analytics for Crop Prediction

To recommend the most suitable crop, multiple supervised classification algorithms were evaluated, including Naïve Bayes, Decision Tree, Support Vector Machine, K-Nearest Neighbors, and Random Forest. Among these, Random Forest exhibited superior performance due to its ensemble-based learning mechanism, which reduces overfitting and enhances generalization capability. In Random Forest, feature importance quantifies the contribution of each attribute toward prediction accuracy. The importance of a node k is mathematically expressed as:

$$\mathbf{im_k = wt_k I_k - wt_{lt(k)} I_{lt(k)} - wt_{rt(k)} I_{rt(k)}}$$

Where:

- im_k is the importance of node ‘k’.
- wt_k is the weighted samples number that reach node ‘k’.
- I_k is the impurity value (such as Gini impurity or Information Gain) for node ‘k’.
- lt_k and rt_k represent the child nodes on the left and right splits of node ‘k’, respectively.

This formulation measures the impurity reduction achieved by splitting at node k. Attributes contributing higher impurity reduction are ranked as more influential. Experimental results confirmed that soil pH, temperature, and potassium content were among the most significant predictors. The comparative evaluation demonstrated that Random Forest and Naïve Bayes achieved accuracy levels exceeding 95%.

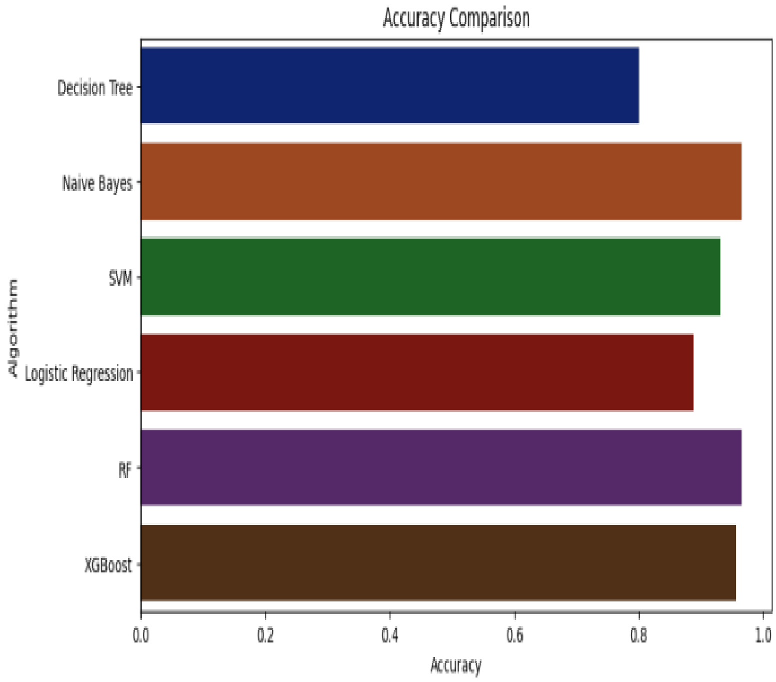

Figure 3: Accuracy of different algorithms

2.4 Challenges in Crop Price Prediction

While predicting optimal crop selection based on soil and environmental parameters is statistically feasible, extending the framework to forecast crop market prices introduces significant complexity. Unlike soil chemistry, which follows measurable physical relationships, agricultural market pricing is influenced by volatile and external macroeconomic factors.Price dynamics depend on international trade policies, government subsidies, seasonal demand fluctuations, supply chain disruptions, transportation costs, geopolitical tensions, and unexpected natural events. These factors exhibit nonlinear and often unpredictable behavior, making deterministic modeling unreliable.

Moreover, localized IoT sensors capture only micro-level agronomic data and lack access to synchronized, real-time economic indicators. Due to the absence of granular macroeconomic datasets, accurate crop price forecasting at the edge level remains impractical within the scope of Farm Oracle. Therefore, the system focuses on agronomic optimization rather than speculative market prediction.

2.5Generative AI: QLoRA Fine-Tuning

To enhance advisory intelligence, Farm Oracle integrates a Generative AI layer based on a Large Language Model (LLM). Instead of fully fine-tuning billions of parameters, which is computationally expensive, the system employs Low-Rank Adaptation (LoRA).

LoRA represents a pre-trained weight matrix W0 with a low-rank update as:

$$\mathbf{W_0 + \Delta W = W_0 + BA}$$

where rank $\mathbf{r \ll \min(d, k)}$. By utilizing Quantized LoRA (QLoRA), Farm Oracle quantizes the frozen base Mistral7B weights to 4-bit precision, utilizing a single consumer GPU to fine-tune the ‘A’ and ‘B’ matrices on specialized agronomic datasets.

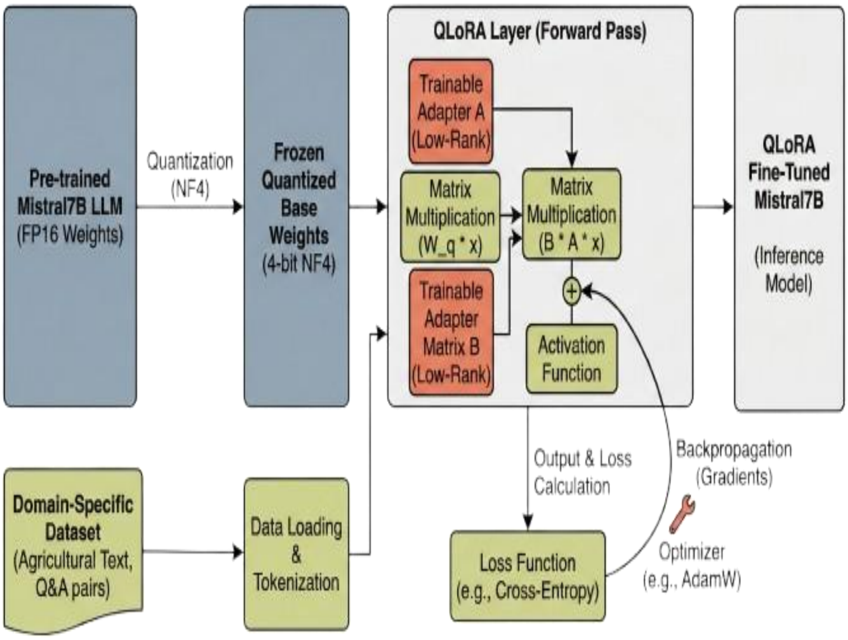

Figure 4: Fine tuning with QLoRA

3. Approach

The Farm Oracle frontend is developed using React, providing a responsive and interactive interface for farmers and agricultural stakeholders. The dashboard displays real-time soil telemetry including Nitrogen (N), Phosphorus (P), Potassium (K), temperature, humidity, and pH levels. Data is transmitted securely from the IoT devices to the cloud using authenticated API keys and stored in Firebase. Within the interface, parameters are visualized using dynamic charts and gauges with clearly defined units and agronomic threshold limits, enabling intuitive interpretation even for non-technical users. This visualization layer transforms raw sensor readings into meaningful, actionable insights.

The core intelligence of the system resides in the backend application layer implemented using Spring Boot. When a user requests agricultural advice, the backend initiates a structured decision pipeline. First, the latest environmental vector—comprising N, P, K, temperature, humidity, and pH—is retrieved from the Firebase database. This ensures that recommendations are always based on the most recent field conditions rather than static historical inputs.

Next, the Random Forest classification model performs inference on the retrieved feature vector. Based on learned patterns from the trained dataset, the model outputs the most suitable crop label (for example, "Jute" or "Rice") corresponding to the current soil profile. Because Random Forest aggregates multiple decision trees, it provides stable predictions with reduced variance and high classification confidence.

Following crop prediction, a dynamic prompt is constructed. This prompt contains three essential components:

1. The exact numerical sensor readings (e.g., N = 72 mg/L, P = 10 mg/L, K = 45 mg/L),
2. The predicted crop label, and
3. The user's natural language query (e.g., "How should I improve phosphorus levels?").

This structured prompt engineering approach ensures contextual coherence and prevents ambiguous model responses.

The QLoRA-adapted Mistral-7B Large Language Model then processes this enriched prompt. Since the low-rank adaptation matrices were fine-tuned on agriculture-specific datasets, the model generates highly localized and chemically consistent recommendations. For example, if the phosphorus level is critically low (e.g., 10 mg/L), the model may suggest precise fertilizer formulations, dilution ratios, and application schedules aligned with the predicted crop's nutrient requirements.Importantly, the LoRA architecture enables merging of the learned update matrices with the original base model weights during deployment. This design ensures that no additional computational overhead is introduced at inference time. As a result, the advisory response is generated with minimal latency, maintaining real-time interaction while preserving model efficiency.

Through this tightly integrated workflow—spanning IoT telemetry acquisition, machine learning inference, intelligent prompt generation, and optimized generative reasoning—Farm Oracle delivers a seamless, intelligent, and scalable agricultural advisory system.

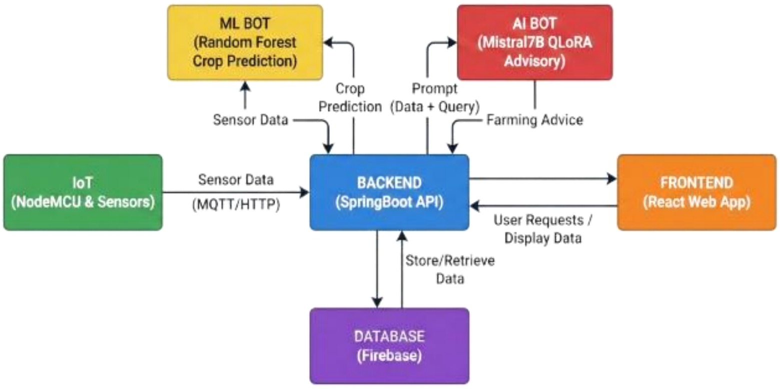

Figure 5: Integrated Systems of Solution

4. Result and Discussion

4.1 IoT System and Architecture Performance

The Farm Oracle system demonstrates a highly stable data pipeline from edge hardware to cloud analytics. In terms of IoT hardware efficiency, the NodeMCU seamlessly reads Modbus packets from the 6-in-1 soil sensor via the TTL-to-RS485 converter. The telemetry data transmission proved reliable, ensuring continuous real-time monitoring of soil micro-climates.

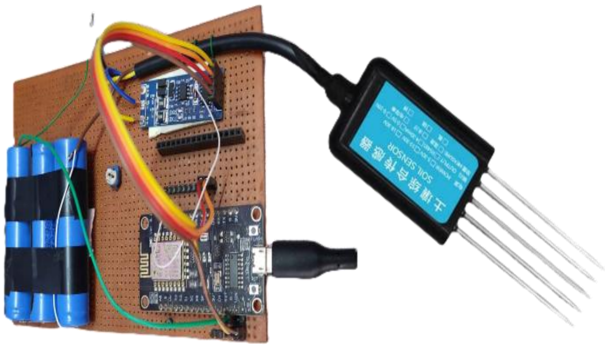

Figure 6: Functional IoT device

Furthermore, the system exhibited robust frontend and backend coupling. Telemetry data routed to the Firebase NoSQL database was managed efficiently without data loss. The Spring Boot backend successfully monitored the database, pulling the latest environmental vectors (Nitrogen, Phosphorous, Potassium, Temperature, Humidity, and pH) to feed into the local inference engine. Consequently, the React-based frontend dynamically visualized this data, reflecting sensor updates instantly and confirming that the end-to-end integration introduces negligible latency.

4.2 Machine Learning Analytics

The models are rigorously evaluated using standard metrics based on True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN):

1. **Accuracy:** The percentage of accurate results among all predictions.
   Accuracy = {TP + TN}/{TP + TN + FP + FN}
2. **Precision:** The number of correctly predicted positive outcomes out of all positive predictions.
   Precision = {TP}/{TP + FP}
3. **Recall (Sensitivity):** The proportion of correctly predicted positive cases among all the actual positive examples.
   Recall = {TP}/{TP + FN}
4. **F1-Score:** Measured as the harmonic mean of precision and recall.
   F1_Score = 2 *{Precision * Recall}/{Precision + Recall}

Based on experimental evaluations of 6 algorithms, the XG Boost algorithm yielded an impressive classification accuracy of 95.68%. The Naive Bayes Classifier and Random Forest algorithms tightly followed, yielding an accuracy of 96.59%

Regarding Generative AI capability, by optimizing only the W_q and $W_v$attention weights at a rank of r=8, the trainable parameter count was kept minimal, reducing the checkpoint sizes drastically (by roughly 10,000x compared to full models).

5. Conclusion

Farm Oracle demonstrates how advanced Artificial Intelligence and IoT technologies can be transformed from theoretical research concepts into a practical, field-deployable smart agriculture framework. By integrating real-time soil telemetry, machine learning–based crop recommendation, and a generative advisory layer within a unified pipeline, the system bridges the long-standing gap between environmental data acquisition and farmer-centric decision support. The architecture ensures that recommendations are not only statistically accurate but also context-aware and operationally meaningful.

A key innovation of the framework lies in the adoption of QLoRA-based parameter-efficient fine-tuning. Instead of retraining an entire multi-billion parameter language model, Farm Oracle optimizes lightweight low-rank adaptation matrices while keeping the base model frozen. This approach significantly reduces computational cost, memory consumption, and energy requirements, enabling secure and rapid deployment even on resource-constrained infrastructure. As a result, the system delivers expert-level, generative agronomic advice with minimal latency, ensuring real-time interaction without compromising performance.

The successful integration of IoT sensing, machine learning inference, and optimized Large Language Models highlights a scalable pathway for intelligent agricultural ecosystems. Beyond crop recommendation, the conversational advisory module enhances accessibility by translating raw numerical parameters into understandable, actionable cultivation strategies tailored to specific soil conditions.

Future enhancements will focus on expanding the dataset to include multi-regional agricultural profiles through GPS-enabled IoT deployments. Incorporating geospatial tagging will allow the system to learn location-specific soil behavior, climatic variability, and crop suitability trends, thereby improving generalization across diverse agro-climatic zones.

Additionally, integrating Vision-Language Models (VLMs) represents a promising direction for extending Farm Oracle's capabilities. By combining smartphone-captured crop imagery with soil telemetry data, the system could enable automated visual disease detection, nutrient deficiency recognition, and early-stage pest identification. Such multimodal intelligence would transform Farm Oracle from a soil-centric advisory tool into a comprehensive, autonomous precision agriculture platform capable of supporting sustainable farming at scale.

References

[1] [1] Food and Agriculture Organization (FAO), The Future of Food and Agriculture – Trends and Challenges, FAO, Rome, Italy, 2017.

[2] [2] L. Ruiz-Garcia, L. Lunadei, P. Barreiro, and J. I. Robla, "A review of wireless sensor technologies and applications in agriculture and food industry," Sensors, vol. 9, no. 6, pp. 4728–4750, 2009.

[3] [3] V. Saiz-Rubio and F. Rovira-Más, "From smart farming towards agriculture 5.0: A review on crop data management," Agronomy, vol. 10, no. 2, 2020.

[4] [4] S. Ramesh and D. Vydeki, "Application of machine learning in crop yield prediction," International Journal of Advanced Research in Computer Science, vol. 9, no. 1, 2018.

[5] [5] J. B. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.

[6] [6] T. Brown et al., “Language models are few-shot learners,” in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2020.

[7] [7] E. J. Hu et al., “LoRA: Low-Rank Adaptation of Large Language Models,” in Proc. International Conference on Learning Representations (ICLR), 2022.

[8] [8] T. Dettmers et al., “QLoRA: Efficient Finetuning of Quantized LLMs,” in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2023.

[9] [9] R. Gebbers and V. I. Adamchuk, “Precision agriculture and food security,” Science, vol. 327, no. 5967, pp. 828–831, 2010.

[10] [10] S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, “Big data in smart farming – A review,” Agricultural Systems, vol. 153, pp. 69–80, 2017.

[11] [11] M. Liakos et al., “Machine learning in agriculture: A review,” Sensors, vol. 18, no. 8, 2018.

[12] [12] K. Mekala and P. Viswanathan, “A survey: Smart agriculture IoT with cloud computing,” in Proc. IEEE International Conference on Microcontroller, Embedded and System Applications, 2017.

[13] [13] J. Chen and X. Yang, “Wireless sensor networks for agriculture: The state-of-the-art in practice and future challenges,” Computers and Electronics in Agriculture, vol. 118, pp. 66–84, 2015.

[14] [14] D. Tilman et al., “Global food demand and the sustainable intensification of agriculture,” Proceedings of the National Academy of Sciences (PNAS), vol. 108, no. 50, pp. 20260–20264, 2011.

[15] [15] A. Vaswani et al., “Attention is all you need,” in Proc. Advances in Neural Information Processing Systems (NeurIPS), 2017.

[16] [16] A. Dosovitskiy et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” in Proc. International Conference on Learning Representations (ICLR), 2021.

[17] [17] M. A. Olson, A. Khanna, and A. Sharma, “Vision-based plant disease detection using deep learning: A survey,” Artificial Intelligence in Agriculture, vol. 4, pp. 1–14, 2020.