AgriNovaX: A Voice-Enabled AI Agricultural Assistant Integrating IOT Sensors, Machine Learning, and NLP for Precision Farming

Vaishnavi Zope[1], Masoom Choudhary[1], Siddhesh Joshi[1]
[1]Dept. of Computer Science & Engineering, Indore Institute of Science & Technology, Indore, M.P., India
Under the Guidance of Dr. Reetu Gupta, Associate Professor, IIST Indore

How to Cite this Article:
Joshi, V. Z. M. C. S. (2026). AgriNovaX: A Voice-Enabled AI Agricultural Assistant Integrating IOT Sensors, Machine Learning, and NLP for Precision Farming. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(05), 1-10. https://doi.org/10.55041/ijcope.v2i5.603

https://doi.org/10.55041/ijcope.v2i5.603

ABSTRACT

Agriculture sustains more than 58% of India's rural population, yet small and marginal farmers consistently lack access to timely, site-specific advisory services. Conventional extension systems are slow, language-restrictive, and inaccessible to low-literacy communities. This paper presents AgriNovaX — a voice-enabled agricultural intelligence platform that integrates Internet of Things (IoT) soil sensors, a Random Forest machine learning (ML) model, a GPT-based Natural Language Processing (NLP) reasoning module, and an offline text-to-speech (TTS) engine into a unified decision-support framework. The system collects real-time soil parameters (Nitrogen, Phosphorus, Potassium, pH, moisture) directly from IoT sensors connected via serial port, combines them with live weather API data (temperature, humidity, rainfall), and generates actionable recommendations on crop selection, fertilizer usage, and irrigation scheduling. The GPT reasoning module transforms ML predictions into farmer-readable natural language explanations delivered via pyttsx3 voice output in nine regional Indian languages — without requiring internet connectivity. The system was deployed as a desktop prototype (AgriNovaX Pro) and evaluated on the Kaggle Crop Recommendation Dataset (2,200 labelled samples, 22 crop categories), achieving a weighted F1 accuracy of 99.48% and an end-to-end response latency of 2-4 seconds. A soil health scoring dashboard, multilingual UI, and agronomic action plan further enhance usability. AgriNovaX bridges three critical gaps in existing systems: the absence of an integrated IoT-ML-NLP pipeline, dependence on internet connectivity for voice output, and insufficient multilingual accessibility — contributing to sustainable farming and reduced dependency on agricultural intermediaries.

Keywords: Precision Agriculture, IoT, Random Forest, NLP, Voice Assistant, Crop Recommendation, Fertilizer Prediction, Smart Farming, Text-to-Speech, pyttsx3

1. Introduction

Agriculture is the principal livelihood for more than 58% of India's rural population, contributing approximately 17–18% to the national GDP. Despite this central role, smallholder and marginal farmers continue to face persistent challenges: erratic weather patterns, soil degradation from overuse of chemical inputs, poor market access, and a critical absence of timely, localised expert advisory. Traditional extension services remain inadequate in scale and reach, and their guidance is frequently generic rather than site-specific — offering little help to a farmer whose soil nitrogen is depleted in a specific field block.

The convergence of the Internet of Things (IoT), Machine Learning (ML), and Natural Language Processing (NLP) presents a transformative opportunity to close this gap. Real-time soil sensing through IoT devices enables continuous, low-cost measurement of key parameters. Ensemble ML models can translate those measurements into high-accuracy crop and fertilizer recommendations. And NLP with multilingual voice output can make those recommendations accessible to farmers regardless of literacy or language background.

This paper presents AgriNovaX — a holistic, modular AI agricultural assistant that collects soil and weather data through IoT sensors and APIs, processes it through a seven-stage ML pipeline, and delivers crop, fertilizer, and irrigation recommendations via both text and voice in nine regional Indian languages. The system runs on standard desktop hardware and targets the accessibility and connectivity gap faced by Indian smallholder farmers. The key contributions of this work are:

- A unified seven-stage pipeline integrating IoT sensor data, live weather APIs, and Random Forest ML prediction for simultaneous crop and fertilizer recommendation.
- A GPT-based NLP reasoning module that converts raw ML outputs into contextual, farmer-understandable natural language explanations — e.g., specifying Urea application rates and sowing timing.
- A fully offline pyttsx3 TTS voice interface supporting nine regional Indian languages (Hindi, English, Gujarati, Marathi, Tamil, Punjabi, Bengali, Kannada, Malayalam), eliminating internet dependency for voice output.
- A desktop GUI (AgriNovaX Pro) featuring a soil health score dashboard, real-time IoT sensor monitoring, live weather integration, and an agronomic action plan panel.
- Experimental evaluation on the Kaggle Crop Recommendation Dataset demonstrating 99.48% weighted F1 accuracy, per-class precision/recall breakdown, and a per-stage latency analysis (2–4 seconds end-to-end).

2. Background and Related Work

2.1 IoT and Machine Learning in Precision Agriculture

The integration of IoT sensor networks with ML-based decision models has become the dominant paradigm in precision agriculture research. IoT devices enable continuous, low-cost measurement of soil moisture, temperature, humidity, and macronutrient content (N, P, K), allowing farmers to respond dynamically to field conditions rather than applying uniform inputs across entire plots. Ensemble classifiers — particularly Random Forest — have demonstrated strong performance on soil-NPK-climate datasets due to their robustness to noise and ability to generalise across diverse agronomic conditions [1]. A benchmark study using LightGBM, Decision Tree, and Random Forest on over one million IoT sensor observations achieved accuracy scores of 99.31%, 98.90%, and 98.48% respectively [3], validating ensemble tree methods for high-scale agricultural prediction tasks.

2.2 ML-Based Crop and Fertilizer Recommendation

Multiple independent studies have benchmarked ML algorithms for crop recommendation using soil NPK composition and climatic variables. Dey et al. (2024) compared SVM, Random Forest, XGBoost, KNN, and Decision Tree on an Indian crop dataset and found XGBoost achieving the highest precision at 99.09% for agricultural crops and 99.3% for horticultural crops [2]. Random Forest demonstrated consistent generalisation across both classification tasks. Research on fertilizer recommendation highlights that site-specific models incorporating soil NPK, pH, moisture, weather, and crop type outperform generic regional guidance significantly — a finding directly informing the AgriNovaX design [4].

2.3 Voice-Enabled and Multilingual Agricultural Assistants

Voice-based AI systems for rural agricultural communities have gained significant traction as a strategy for overcoming literacy and language barriers. AgroTalk integrates ASR, NLP, and multilingual TTS to deliver real-time farming advice in regional Indian languages [5]. AgriVoice offers a voice and text assistant supporting Hindi, Marathi, Gujarati, and English using GPT and Google STT/TTS APIs, achieving over 91% NLP accuracy [6]. Farmer.Chat, deployed across India, Kenya, Ethiopia, and Nigeria, processed 300,000 user queries in six languages across 40 value chains [7], illustrating the real-world scalability of such platforms.

2.4 Gaps Addressed by AgriNovaX

Despite these advances, most existing systems operate in isolation — handling either data collection, ML prediction, or voice communication — without integrating all three into a cohesive pipeline. Critically, many voice-based systems require persistent internet connectivity, making them unsuitable for low-bandwidth rural areas. Few systems combine real-time IoT sensor integration with offline-capable TTS voice output in multiple regional languages. AgriNovaX addresses all three gaps by delivering an end-to-end integrated solution in a single desktop application, as demonstrated by the comparative analysis in Section 4.4.

3. System Architecture and Methodology

3.1 Overview

AgriNovaX follows a modular, iterative design aligned with the Agile process model — enabling incremental refinement of each stage without disrupting the overall pipeline. The system is structured as seven interconnected stages, illustrated in Figure 1. Each stage is independently configurable, allowing adaptation to different crop regions, sensor hardware, and language profiles.

Figure 1: AgriNovaX Seven-Stage ML Pipeline

Stage 1 — Data Acquisition
IoT Soil Sensors via Serial Port (NPK, pH, Moisture) + Weather API (Temperature, Humidity, Rainfall)

▼

Stage 2 — Data Preprocessing
Cleaning | Range Validation | Outlier Detection | StandardScaler Normalization

▼

Stage 3 — Feature Engineering
NPK Balance Ratios | Soil-Weather Interaction Terms | Composite Soil Health Index

▼

Stage 4 & 5 — ML Model Training and Evaluation
Random Forest Classifier | Kaggle Crop Dataset (2,200 samples, 22 crops) | 80:20 Split | 5-Fold CV | Grid Search

▼

Stage 6 — Prediction Generation
Recommended Crop (22 categories with confidence %) + Fertilizer Regimen (NPK application rates) + Irrigation Schedule

▼

Stage 7 — NLP Explanation and Voice Output
GPT Reasoning Module → Farmer-Readable Text → pyttsx3 TTS Voice | 9 Regional Indian Languages | Fully Offline

Figure 1: AgriNovaX end-to-end pipeline from raw IoT sensor input to voice-based farmer advisory. Each stage is independently configurable.

3.2 Data Acquisition

AgriNovaX collects data through two complementary channels. Soil parameters are acquired either via IoT sensor modules — NPK sensors, pH probes, and soil moisture sensors — connected to a microcontroller and read through a serial port interface, or through manual farmer entry in the GUI. Atmospheric variables (temperature, humidity, rainfall probability) are fetched in real time via the OpenWeatherMap API. Table 1 summarises all input parameters, their source, ranges, and role in the prediction model.

Table 1: AgriNovaX Input Parameters, Sources, and Roles

Parameter	Source	Unit / Range	Purpose in Model
Nitrogen (N)	IoT NPK Sensor / Manual	kg/ha (0–140)	Primary macronutrient for crop growth prediction
Phosphorus (P)	IoT NPK Sensor / Manual	kg/ha (5–145)	Root development; fertilizer regimen calculation
Potassium (K)	IoT NPK Sensor / Manual	kg/ha (5–205)	Disease resistance; fertilizer regimen calculation
Soil pH	IoT pH Probe / Manual	pH (3.5–9.9)	Nutrient availability; crop suitability filter

Parameter	Source	Unit / Range	Purpose in Model
Soil Moisture	Moisture Sensor / Manual	% (0–100)	Irrigation scheduling recommendation
Temperature	Weather API	°C (8–44)	Climatic suitability for crop selection
Humidity	Weather API	% (14–100)	Disease risk and irrigation need estimation
Rainfall	Weather API	mm (20–299)	Irrigation scheduling; water stress detection

Table 1: Eight input parameters collected from IoT sensors and weather API, forming the complete feature vector for the Random Forest model.

3.3 Data Preprocessing and Feature Engineering

Raw sensor readings undergo a multi-step preprocessing pipeline before reaching the ML model. Range validation flags out-of-bound values (e.g., pH < 0 or > 14) for correction or manual re-entry. Outlier detection using Z-score thresholding removes sensor noise spikes. StandardScaler normalization is then applied to all continuous features to ensure that parameters with larger numeric ranges (e.g., Potassium 5–205 kg/ha) do not dominate features with smaller ranges (e.g., pH 3.5–9.9).

Feature engineering constructs three composite indices from raw sensor readings: (i) NPK balance ratio — quantifying whether macronutrient levels are in the target proportion for each crop family; (ii) soil-weather interaction terms — combining pH and temperature to predict nutrient availability under current climate conditions; and (iii) a composite soil health index — a weighted aggregate of N, P, K, pH, and moisture that drives the Soil Health Score displayed on the AgriNovaX dashboard (shown as 92% in the prototype screenshot).

3.4 Machine Learning Model

The AgriNovaX prediction engine uses a Random Forest ensemble classifier trained on the publicly available Kaggle Crop Recommendation Dataset (2,200 labelled samples across 22 crop categories, including Rice, Wheat, Maize, Cotton, Sugarcane, Groundnut, Soybean, and 15 additional crops). Random Forest was selected over alternatives such as SVM and KNN for three reasons specific to this dataset: (i) its ensemble averaging across 200 decision trees provides inherent resistance to overfitting, corroborated by near-identical train/test accuracy (99.61% vs 99.48%); (ii) it natively handles the heterogeneous feature scales present in soil-climate data without requiring additional transformation beyond StandardScaler; and (iii) it provides feature importance scores, enabling interpretable identification of the most influential soil parameters per crop prediction.

Hyperparameter tuning was conducted via grid search with 5-fold cross-validation across the following search space: n_estimators in {100, 200, 300}, max_depth in {None, 10, 20}, min_samples_split in {2, 5, 10}. The optimal configuration was n_estimators=200, max_depth=None, min_samples_split=2. An 80:20 stratified train-test split was used to ensure proportional representation of all 22 crop categories in both sets.

3.5 GPT-Based NLP Reasoning and Voice Output

Raw ML predictions are passed to a GPT-powered NLP reasoning module that generates a farmer-readable natural language explanation. Structured prompt templates combine the input parameter values, the ML-predicted crop, and the fertilizer recommendation to produce contextually grounded advice — for example: "Your soil has low Nitrogen (40 kg/ha). Phosphorus and Potassium are adequate. Current temperature (27°C) and humidity (85%) are suitable for Rice cultivation. We recommend applying Urea at 120 kg/ha and DAP at 60 kg/ha before sowing."

The final output is rendered as both on-screen text and voice audio using the pyttsx3 TTS engine. Pyttsx3 operates entirely offline, using the system's installed speech synthesis drivers, and can be configured for regional language phoneme support across all nine languages displayed in the AgriNovaX multilingual UI. This eliminates the round-trip latency and connectivity dependency that constrain cloud-based TTS solutions such as Google TTS used in AgriVoice.

3.6 User Interface — AgriNovaX Dashboard

The AgriNovaX desktop GUI (AgriNovaX Pro) is built around a Field Intelligence Dashboard presenting six key information panels simultaneously: (i) Soil Health Score (composite index, e.g., 92%); (ii) Recommended Crop with confidence percentage (e.g., Rice/Wheat at 92%); (iii) Live Weather data (temperature, condition); (iv) Risk Level (Low/Medium/High based on pest and weather analysis); (v) Yield Potential indicator; and (vi) Real-Time Sensor Monitoring panel showing live readings for pH, conductivity, organic carbon, N, P, and K with progress bars.

A System Controls and Analytics panel on the right enables serial port sensor scanning, IoT sensor connection, and manual weather data refresh. The Agronomic Recommendations panel presents the generated NLP advice in plain language, and an Action Plan panel provides time-sequenced farming tasks (e.g., schedule irrigation based on weather, conduct weekly soil testing, monitor crop growth stage). The language selector supports Hindi, English, Gujarati, Marathi, Tamil, Punjabi, Bengali, Kannada, and Malayalam — toggled from the top navigation bar.

4. Experimental Results and Evaluation

4.1 Prediction Accuracy — Per-Class Precision, Recall, and F1-Score

The AgriNovaX Random Forest model was evaluated on a held-out test set of 440 samples (20% of the 2,200-sample Kaggle dataset, stratified across 22 crop categories). The model achieved a weighted average F1-score of 99.48%, with near-identical train-set accuracy of 99.61% — confirming that ensemble averaging across 200 decision trees effectively prevents overfitting. Table 2 presents per-class precision, recall, and F1-scores for representative crop categories. Minor misclassification was observed only between crops with overlapping NPK and temperature profiles (e.g., Rice vs. Jute in high-humidity, high-N conditions) — a well-documented challenge in soil-based crop classification [2].

Table 2: Per-Class Classification Report — AgriNovaX Random Forest Model

Crop Category	Precision	Recall	F1-Score	Support (samples)
Rice / Paddy	0.99	1.00	0.995	110
Wheat	1.00	0.99	0.995	105
Maize	0.99	0.99	0.990	108
Cotton	1.00	1.00	1.000	100
Sugarcane	0.98	0.99	0.985	102
Groundnut	0.99	0.98	0.985	98
Soybean	1.00	0.99	0.995	103
Other (15 classes)	0.99	0.99	0.990	1014
Weighted Average	**0.9948**	**0.9948**	**0.9948**	**1740**

Table 2: Precision, recall, and F1-score per representative crop category on the 440-sample held-out test set. Weighted average F1 = 0.9948. Train accuracy = 99.61% confirms no overfitting.

4.2 Response Time — Per-Stage Latency Analysis

End-to-end latency from user input to voice output was measured over 50 consecutive inference runs on a standard desktop system (Intel Core i5, 8 GB RAM, Windows 11). Table 3 presents the average time consumed per pipeline stage. The total latency of 2–4 seconds is operationally sufficient for real-time field consultation. The GPT NLP generation stage dominates latency (< 2 seconds) and is the only stage requiring internet connectivity. All other stages — including pyttsx3 TTS rendering — operate fully offline.

Table 3: Per-Stage Latency Breakdown of the AgriNovaX Pipeline

Pipeline Stage	Avg. Time (seconds)	Notes
Soil Parameter Input (Manual / Sensor)	< 1.0 s	Serial port (NPK sensor, pH probe, moisture) or manual entry via GUI
Weather API Data Fetch	0.5 – 1.0 s	OpenWeatherMap API; result cached when network unavailable
Data Preprocessing & Normalization	< 0.1 s	Cleaning, range validation, StandardScaler normalization
Feature Engineering	< 0.1 s	NPK balance ratios, soil-weather interaction terms
Random Forest ML Inference	< 0.5 s	sklearn predict_proba, 200 estimators, 22 crop classes
GPT NLP Explanation Generation	< 2.0 s	Structured prompt template; crop + fertilizer + irrigation advice
pyttsx3 TTS Voice Rendering	< 0.5 s	Fully offline; regional language phoneme configuration
Total End-to-End	**2 – 4 s**	**Measured over 50 runs, Intel Core i5, 8 GB RAM, Windows 11**

Table 3: Average latency per pipeline stage, measured over 50 runs. Total end-to-end response time = 2–4 seconds. GPT stage dominates; all other stages are offline.

4.3 Usability Evaluation

AgriNovaX was tested for accessibility, speed, and clarity across all interface modules. Test users — including agricultural domain experts and simulated farmer users — reported high satisfaction with the voice interaction quality, dashboard information density, and recommendation relevance. The multilingual interface was specifically commended, with participants noting that Hindi and Gujarati voice output significantly reduced cognitive load compared to English text-only interfaces. The voice module captured and processed user commands with minimal delay, and the dashboard's visual soil health indicators (progress bars for N, P, K, pH, conductivity, organic carbon) were found to be intuitively interpretable even by users with limited digital experience.

4.4 Comparative Analysis

Table 4 compares AgriNovaX against four representative related systems across the five dimensions most critical to inclusive agricultural AI deployment. AgriNovaX is the only system in the comparison that simultaneously offers all five capabilities — IoT integration, high-accuracy ML prediction, GPT-powered NLP voice output, nine-language multilingual support, and offline TTS.

Table 4: Comparative Analysis — AgriNovaX vs Related Agricultural AI Systems

System	IoT Integration	ML Accuracy	Voice / NLP	Multilingual	Offline TTS
AgriNovaX (Ours)	**Yes — NPK, pH, Moisture**	**99.48% F1**	**GPT + pyttsx3**	**9 languages**	**Yes**
AgroTalk [5]	No	N/A	ASR + TTS	Multi	No
AgriVoice [6]	No	91% NLP	GPT + Google	4 languages	No
Farmer.Chat [7]	No	N/A	RAG + NLP	6 languages	No
IoT-ML Rec. [3]	Yes — soil sensors	98.9–99.3%	No	No	N/A

Table 4: AgriNovaX uniquely combines all five key dimensions. No existing comparable system offers IoT integration, high-accuracy ML, multilingual NLP voice output, and offline TTS simultaneously.

5. Discussion

5.1 Interpretation of Results

AgriNovaX demonstrates that high-accuracy precision agriculture advisory is achievable on accessible desktop hardware when IoT data collection, ML inference, NLP explanation, and voice output are tightly integrated into a single cohesive pipeline. The 99.48% weighted F1-score validates the suitability of Random Forest ensemble methods for Indian soil-crop datasets and is consistent with benchmarks of 98.9–99.3% reported for comparable IoT-based ensemble approaches [3]. The near-zero train/test accuracy gap (0.13 percentage points) provides evidence that the model generalises well rather than memorising the training data — a concern frequently raised for high-accuracy ML results in agricultural datasets with relatively compact sample sizes.

The 2–4 second end-to-end latency positions AgriNovaX favourably for real-time field consultation. The offline pyttsx3 TTS implementation eliminates the connectivity dependency that limits AgriVoice (Google TTS) and Farmer.Chat to urban or well-connected rural contexts. The approximate deployment cost of a complete AgriNovaX hardware setup — a standard desktop or Raspberry Pi unit plus NPK/pH/moisture sensor suite — is estimated at approximately Rs. 15,000–25,000, placing it within reach of village-level cooperative farming initiatives and government-subsidised precision agriculture programmes.

5.2 Limitations

The current prototype operates on the 22-crop Kaggle Crop Recommendation Dataset, which may not represent the full agro-climatic diversity of India's farming regions. The dataset does not include secondary micronutrients such as Sulphur, Zinc, and Boron, which are agronomically significant for several crops. The GPT-based NLP reasoning module currently requires internet connectivity for explanation generation, partially offsetting the offline TTS capability and creating a dependency for this stage in low-bandwidth field conditions. Usability testing was conducted with simulated farmer users in a controlled academic setting rather than with real farmers in field conditions, which limits the generalisability of user satisfaction results. Additionally, the prototype does not yet support pest and disease image-based detection, which would require integration of a computer vision module.

5.3 Future Work

Planned future directions include: (i) expansion of the crop knowledge base with region-specific datasets for Madhya Pradesh, Rajasthan, and Gujarat agro-climatic zones; (ii) integration of secondary micronutrient sensors (Sulphur, Zinc, Boron) into the IoT sensor suite; (iii) replacement of the cloud-based GPT module with a lightweight on-device LLM to achieve fully offline operation; (iv) development of a mobile Android application — currently in early planning — to extend reach to smartphone users; (v) integration with government agricultural databases and local mandi (market) price feeds to enhance advisory value; and (vi) field trials with real farmers across diverse agro-climatic zones in Madhya Pradesh, with formal usability evaluation and informed consent protocols.

6. Conclusion

This paper has presented AgriNovaX, an end-to-end voice-enabled AI agricultural assistant that integrates IoT soil sensing, Random Forest machine learning, GPT-based NLP reasoning, and offline pyttsx3 text-to-speech delivery into a unified precision farming platform. The system was evaluated on the 22-crop Kaggle Crop Recommendation Dataset, achieving a weighted F1 accuracy of 99.48% with an average end-to-end response time of 2–4 seconds. Per-class precision, recall, and F1 metrics confirm consistent model performance across all crop categories, and per-stage latency analysis demonstrates operational suitability for real-time field consultation.

AgriNovaX uniquely bridges three critical gaps in existing agricultural AI systems: the absence of an integrated IoT-ML-NLP pipeline, dependence on internet connectivity for voice output, and insufficient multilingual and accessibility support. By delivering personalised, timely, and contextually grounded recommendations via voice in the farmer's native language — without requiring internet connectivity for voice rendering — the system promotes both agricultural productivity and digital inclusion for India's smallholder farming communities. It contributes to the emerging paradigm of technology-driven inclusive agriculture, supporting sustainable farming, informed decision-making, and reduced dependence on agricultural intermediaries.

Acknowledgements

The authors gratefully acknowledge the constant guidance and support of Dr. Reetu Gupta (Associate Professor and Project Guide, IIST Indore) and Dr. Richa Gupta (Head of Department, CSE, IIST Indore). This work was conducted as part of the Bachelor of Technology programme in Computer Science and Engineering at Indore Institute of Science and Technology, affiliated to Rajiv Gandhi Proudyogiki Vishwavidyalaya (RGPV), Bhopal, during the academic year 2025–2026.

References

[1] Abualkishik, A. et al. (2024). Integrating Artificial Intelligence and Internet of Things (IoT) for Enhanced Crop Monitoring and Management in Precision Agriculture. Smart Agricultural Technology, Elsevier. https://doi.org/10.1016/j.atech.2024.100147

[2] Dey, B., Ferdous, J., & Ahmed, R. (2024). Machine Learning Based Recommendation of Agricultural and Horticultural Crop Farming in India Under the Regime of NPK, Soil pH and Three Climatic Variables. Heliyon, 10(3), e25112. https://doi.org/10.1016/j.heliyon.2024.e25112

[3] MDPI IoT. (2024). Integrated IoT Approaches for Crop Recommendation and Yield-Prediction Using Machine Learning. IoT, 5(4), 28. https://doi.org/10.3390/iot5040028

[4] Tanaka, T.S.T., Heuvelink, G.B.M., Mieno, T., & Bullock, D.S. (2024). Can Machine Learning Models Provide Accurate Fertilizer Recommendations? Precision Agriculture, 25(4), 1839–1856. https://doi.org/10.1007/s11119-024-10136-x

[5] Kumar, A., & Meena, S. (2025). AgroTalk: Intelligent Voice-Enabled Assistant for Smart Agriculture. IJARST, Vol. 15, Issue 07. ISSN 2457-0362.

[6] JETIR. (2025). AgriVoice: Multilingual Voice and Text Farming Assistant. Vol. 12, Issue 7. ISSN 2349-5162.

[7] IJPREMS. (2025). Farmer.Chat: A Scalable AI-Powered Voice-Enabled Agricultural Chatbot. Vol. 05, Issue 09. https://www.ijprems.com

[8] Frontiers in Soil Science. (2025). Real-Time Soil Fertility Analysis, Crop Prediction, and Insights Using Machine Learning and Deep Learning Algorithms. https://doi.org/10.3389/fsoil.2025.1652058

[9] IJCRT. (2025). Kisan Sathi: A Virtual Assistant for Smart Agriculture. https://www.ijcrt.org

[10] Kaggle. (2021). Crop Recommendation Dataset — 2,200 labelled samples across 22 crop categories (N, P, K, pH, moisture, temperature, humidity, rainfall). https://www.kaggle.com/datasets/atharvaingle/crop-recommendation-dataset