iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.78404

AgroAI: An AI-Powered Intelligent Smart Agriculture Assistant Using Machine Learning and Multimodal Interaction

Shivam Sharma[1], Swayam Jaiswal[2], Ujjwal Ahlawat[3]
Computer Science and Engineering, Galgotias College of Engineering and Technology, Greater Noida, India

Abstract: The productivity of agriculture in developing economies like India is not yet advanced owing to the disjointed access to advice, soil erosion, unpredictable weather conditions, and continuous absence of professional advice in the rural areas. The case described in this paper is AgroAI that is a multimodal AI-based smart farming assistant comprising of a YOLOv8-CLS [11] convolutional neural network (CNN) to detect crop diseases, supervised machine learning models to recommend crops and fertilizers based on soil, and a large language model (LLM) chatbot to interact with the region using voice recognition. The module of disease detection was trained on the dataset of PlantVillage [12] that covered 15 classes of diseases and healthy plants. The model obtained a total classification accuracy of 98.8 using transfer learning, data augmentation, and AdamW optimiser with automatic mixed precision with a macro-averaged precision, recall, and F1-score of 0.99. The crop recommendation module uses parameter data of the soil nutrients, weather history, and market price feeds to produce explainable profit-conscious advisory decisions. Progressive Web App (PWA) architecture provides offline capabilities to rural areas with low-connectivity conditions, caching services workers and fallback logic based on rules. The confusion matrix analysis test proves that there is little non-inter-class misclassification mainly between the similar tomato pathologies visually. These results make AgroAI a comprehensive, all-inclusive and deployable platform in support of Indian smallholder farmers.
Keywords: Smart Agriculture, Crop Disease Detection, YOLOv8-CLS, Convolutional Neural Network, Transfer Learning, Multilingual Chatbot, Progressive Web App, Precision Farming, Explainable AI, Plant Village Dataset.

I. INTRODUCTION

Small agronomic research and the growing disconnect between innovative farming practices and the effective decision-making of small farmers in India represent a significant challenge in agriculture, since the field supports the livelihoods of approximately 600 million Indians. In rainfed districts of Maharashtra, Bihar and Andhra Pradesh, growers are still using experience-based heuristics in making sowing decisions, fertilizer management and pest management decisions driven more by tradition than by soil chemistry, weather predictions and market forces. The outcomes of such lack of knowledge are properly recorded: yield losses, wastage of inputs, after-harvest losses, and financial instability that leads to distress among rural citizens (Liakos et al., 2018) [5]. The lack of scientific guidance is not the main challenge that makes this issue especially hard to solve, but the institutional obstacles to its provision: the unequal spread of agricultural extension services, language heterogeneity of agricultural population, and the low reliability of rural telecommunication infrastructure that can be continuously observed.

Combination of machine learning (ML), deep learning, natural language processing (NLP) and mobile computing has provided a technically plausible channel to bridge this shortcoming. Convolutional neural networks currently detect plant diseases in photos of leaves with accuracy close to clinical accuracy (Mohanty et al., 2016) [6]. Gradient-boosted and random forest can be used to map the nutrient content of soil to crop suitability with a precision of over 85 percent (Sardeshmukh et al., 2025) [10]. With speech-to-text (STT) and text-to-speech (TTS) pipelines, large language models can be used to provide personalised advisory text in local languages with limited user friction (Rajanala et al., 2025) [9]. These are well known capabilities in literature individually. What is conspicuously missing is a coherent system that brings together all of them in one, offline enabled, explicable and India centric advisory system.

The utopian model envisioned by the proponents of precision agriculture all farmers receiving access to real time, contextualised, evidence based information is not effective in the real world due to three reasons, which are interconnected. First, the vast majority of the tools available deal with individual tasks: a disease detector in this case, a crop adviser in that case, a voice chatbot in this case (Kamilaris and Prenafeta-Boldu, 2018) [4].

Second, almost every implemented system makes the assumption of a stable internet connection and cloud computing inference, which restricts reach to regions with weak or missing mobile signals (Kamduri and Gupta, 2025) [3]. Third, most ML models are untrusted, giving predictions without explanation context, which is the key drawback to farmers whose confidence in digital applications is preconditioned with the ability to know the arguments underneath a suggestion (Bansal and Singla, 2025) [1].

The given paper introduces AgroAI: Intelligent Smart Agriculture Assistant, a system that is aimed at filling exactly this gap. AgroAI unites: (1) a YOLOv8-cls CNN trained on the PlantVillage [12] dataset to diagnose crop diseases in real-time; (2) supervised ML models to provide explainable crop, fertilizer, and irrigation advice to soil, weather, and history of previous yields; (3) a multilingual chatbot powered by an LLM, supporting text, voice, and image inputs in Indian regional languages; (4) a PWA architecture with fallback to rules.

II. RELATED WORK

The AI literature is wide-ranging and technically diverse in the agriculture domain, including crop recommendation systems, CNN-based disease detection systems, conversational advice systems, and precision irrigation schedulers. Placing AgroAI in the framework of this body of work involves looking at what each category has done as well as where the edges of individual work have created gaps that a system that is collectively focused needs to cover.

The smart crop recommendation framework by Sardeshmukh et al. (2025) [10] is one of the directly most relevant new entries, combining soil nutrient parameters, i.e. nitrogen, phosphorus, and potassium, with soil pH, rainfall, temperature, and live market prices available via the Agmarknet portal [13] in India. Their Random Forest classifier was found to be about 87 percent accurate and a distance based algorithm was used to extend recommendations to profitable alternative crops. Natural language assistant based on the Gemini-API provided customized local advice by using location-specific guidance and in local languages. Its multi-source integration and sensitivity to market are quite admirable, but it has formal dependency on cloud infrastructure that makes it functionally unreachable on places where the internet is not stable, which the authors fail to mention directly.

Another design concern Bansal and Singla (2025) [1] pursued is interpretability. Their AI-powered agent gets real-time weather information and uses decision-tree reasoning to suggest crops and irrigation times and the results are converted to Hindi, Tamil and Marathi. The method of preferring decision trees to ensemble methods ensures that every recommendation has a traceable conditional rule which can be understood by non-experts. This goes in line with the explainability goal of AgroAI. Nonetheless, the authors indicate no empirical user-tests and field performance indicators, and the use of LangGraph and weather API services presupposes again uninterrupted connectivity, which cannot be considered true in vast areas of the rural India.

Kamduri and Gupta (2025) [3] specifically address offline constraint by storing all the data locally in the device and using previous yield and soil history to provide crop, fertilizer and disease-related advisory without need of sensors or network connectivity. The authors claim convincingly that ML on the device can democratize precision agriculture among small farmers with limited resources. However, the system does not offer any usability testing, no class error rates, and minimal discussion of multilingual support features - weaknesses that have not been further validated.

The baseline as per disease detection is Mohanty et al. (2016) [6] who trained a CNN on 54,306 PlantVillage images by 14 crops and 26 disease classes with an accuracy of 99.35 on a controlled held-out set. The given outcome is frequently referred to as evidence of the technical feasibility of smartphone-based plant diagnostics. The same authors did point out, however, that when tested on uncontrolled, real-world imagery, accuracy dropped to about 31% when the model had been used, a dramatic illustration that there is no such thing as laboratory performance and field performance, which are interchangeable. The latter result is supported by the further research by Kamilaris and Prenafeta-Boldu (2018) [4] and Pantazi et al. (2019) [7], who state that domain-specific augmentation schemes, diversity of background, and multi-stage lighting change are the requirements of deployable models. AgroAI directly solves this issue by using YOLOv8-CLS fine-tuning together with RandAugment and a 70/15/15 train-validation-test split that strives to recreate realistic state of farm imaging.

A third dimension of literature is conversational interfaces. Rajanala et al. (2025) [9] state that an Agricultural Chatbot Voice Assistant (ACVA) is a voice assistance that uses a multilayer perceptron and NLP module to decode verbal queries of farmers. The hands free functionality of the system is a true development especially to semi literate users. AgriGO.AI, an advisory app that uses WhatsApp launched in Turkey shows that low-literacy populations can be successfully overcome with familiar messaging interfaces (AgriGO.AI, 2024). Both systems, nevertheless, do not touch upon offline use and integrate conversation with on-device diagnosis of images.

In general, reading the literature, it is possible to observe a trend: high-accuracy components are available separately, yet, not a single published system combines CNN disease detection, soil-ML recommendation, multilingual LLM conversation, market data, and offline PWA functionality into the unified platform that was oriented toward Indian agriculture. The literature gap is not just an aggregation gap rather it is an architectural coherence gap, offline resiliency gap and explainability are the three pillars of AgroAI that have been specifically designed.

III. METHODOLOGY

AgroAI is oriented to the structured, application-focused experimental design which integrates CNN modelling, supervised development of the ML, the construction of chatbot with NLP support, and the integration of all data in real-time in a deployable Progressive Web App. The overall objective is to create a system that can work in both connected and disconnected environments benefiting farmers via text, voice and image based systems in a language that they understand in their regions. The approach follows nine consecutive steps, all of them cumulative in terms of the results of the former.

Phase one entails an analysis of requirements and a collection of data. The Indian Council of Agricultural Research (ICAR), Kaggle [8] and FAOSTAT were used to obtain soil and crop data. The information on disease images was selected using the PlantVillage dataset, which includes 15 disease and healthy classes, about 54,000 labelled disease image images in tomatoes, potato, and bell peppers. Livestock weather parameters were acquired in real time through the OpenWeatherMap API whereas the price of commodities in the market was accessed on the Agmarknet [13] government portal. This multi-source data approach will make sure that the suggestions are based on both the biophysical and economic reality.

Phase two is concerned with preprocessing data and engineering features. The anomalous records of pH and nutrients in the soil were scrubbed, normalized to the standard range, and coded to categorical data, e.g. soil type. Crop images were resized to 224 x 224 pixels in RGB and ran through the RandAugment pipeline of YOLOv8 that uses random horizontal flipping, rotation, variation of brightness and contrast, random cropping, and scaling. These additions enhance the resistance of the model to the fluctuations of lighting and the complexity of the background typical of field photography, which are the conditions in which Mohanty et al. (2016) [6] recorded an accuracy loss.

The third stage is the functioning of the machine learning crop recommendation module. A monitored learning pipeline was trained on the soil nutrient characteristics (N, P, K, pH), the records of previous rainfall and temperature, and the records of crop yield, using the primary model of a Random Forest classifier, and the secondary models of K-Nearest Neighbours and Decision Tree. It was done with scikit-learn (Pedregosa et al., 2011) [8]. Agmarknet [13] market price feeds were also introduced as a second source of input to bias recommendations towards profit conscious crop selection and past the more biophysical advisory common to the literature.

The fourth phase will be the application of the disease detection module. Aspybota A pretrained YOLOv8 Nano Classification model (yolov8n-cls.pt) was fined tuned on the PlantVillage dataset with a fine-tuning transfer learning technique. The ImageNet-derived representations were first frozen at the backbone weights and the classification head was trained. At a second stage, the whole network was optimized at a smaller learning rate. The AdamW optimiser was used with Automatic Mixed Precision (AMP) training, early stopping using patience of 20 epochs, and automatic saving of the best-model using validation accuracy. Stratified sampling was used to maintain the proportion of classes in the dataset splits with 70% training, 15% validation, and 15% test.

The fifth phase is the building of the multimodal chatbot. Speech-to-text Whisper and Google TTS text-to-speech pipes were added to NLP pipelines based on OpenAI, Gemini, or IndicLLM API to allow voice interaction. An intent classification layer based on Dialogflow does routine query rule matching with low latency and complex contextual queries are sent to the LLM backend. The responses are created in the regional language that the farmer prefers, including Hindi, Tamil, Marathi, Telugu, and Bengali among other languages.

In phase six, the PWA offline fallback is developed. With IndexedDB,service workers store the latest soil, weather and advisory information on the device. In the absence of internet connectivity, the Decision Engine processes all inputs through a rule-based advisor which uses agronomic heuristics - e.g. saying drought-tolerant crops when cached rainfall drop is below some threshold - to produce simple advisory with no API calls.

The integration pipeline is carried out in phases seven through nine, which involves connecting all the modules with a backend of Node.js/Express.js with the use of REST API endpoints and deploying a TensorFlow CNN model on Python ML server and evaluation based on the accuracy, precision, recall, F1-score, and confusion matrix analysis of all 15 disease categories. The usability of the system was evaluated in the light of the System Usability Scale (SUS).

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

The study variables can be defined as the following: independent variables are the soil parameters, weather patterns, and the crop images; dependent variables are defined as the prediction accuracy, the diagnosis confidence, and the relevance of the responses. The current tools, KisanGPT, Kisan Suvidha, and AgriGO.AI are conceptual foundations on which comparative positioning can be made. No personal information about farmers is gathered; publicly available datasets are used, which also means not violating any ethical principles.

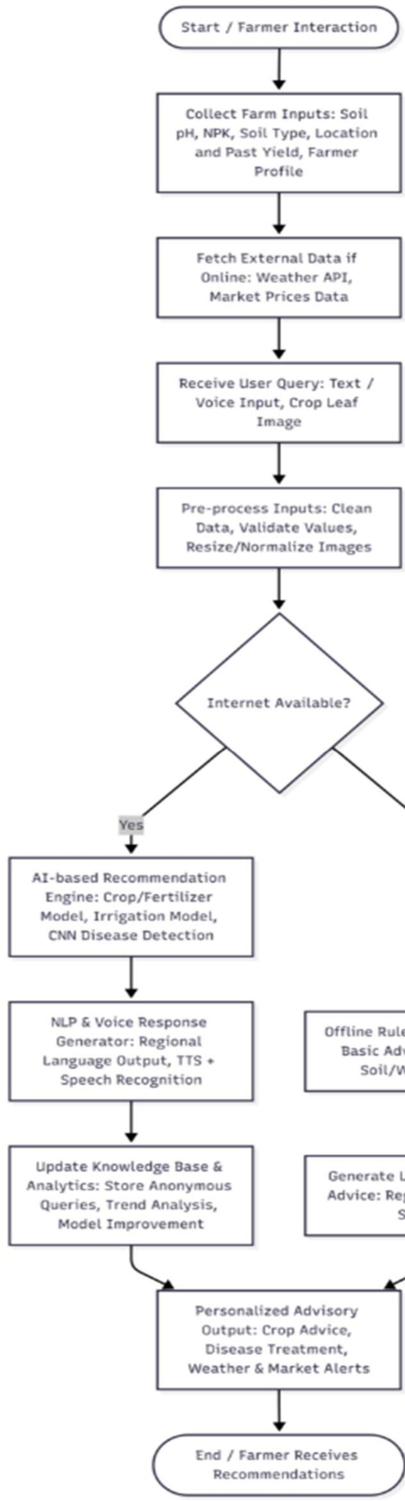

Fig. 1: AgroAI system architecture and data flow

IV. DATASET DESCRIPTION

The module used to detect the disease was trained and tested on the PlantVillage dataset [12], which is an open-source library of about 54,000 labelled leaf images (Hughes & Salathe, 2015) [6]. In the present study, there were 15 classes that were grouped in a subset that included disease and healthy practices of bell pepper, potato, and tomato. Table I shows the distribution of the classes.

TABLE I
CLASS-WISE IMAGE DISTRIBUTION : PLANTVILLAGE DATASET

Class Name	Number of Images
Pepper Bell — Bacterial Spot	997
Pepper Bell — Healthy	1,478
Potato — Early Blight	1,000
Potato — Late Blight	1,000
Potato — Healthy	152
Tomato — Bacterial Spot	2,127
Tomato — Early Blight	1,000
Tomato — Late Blight	1,909
Tomato — Leaf Mold	952
Tomato — Septoria Leaf Spot	1,771
Tomato — Spider Mites	1,676
Tomato — Target Spot	1,404
Tomato — Yellow Leaf Curl Virus	3,208
Tomato — Mosaic Virus	373
Tomato — Healthy	1,591

The data has significant class imbalance: Tomato Yellow Leaf Curl Virus occupies 3,208 samples and Potato Healthy brings 152 samples only. Nevertheless, the RandAugment and stratified splitting versions of the YOLOv8-CLS architecture demonstrated uniformly good per-class (stat) results, demonstrating an absence of distributional imbalance. The soil recommendation module was based on the data of ICAR and Kaggle (around 5,000-20,000 regional inputs) all of which were classified based on pH, N, P, K, soil type, geo-location, and recommended crop. Dynamically obtained weather data was the OpenWeatherMap API, which provides the temperature, humidity, rainfall, pressure, speed of the wind, and cloud coverage. Agmarknet [13] provided market price data in the form of daily commodity feeds by crop name and mandi location.

V. MODEL ARCHITECTURE

The most popular and common detection architecture is the YOLOv8-CLS, which is a classification version of the Ultralytics [11] YOLOv8 architecture. In contrast to the variants of the detection and segmentation, YOLOv8-CLS classifies images at the single-label level by using a CNN backbone and instead of detection and segmentation, a global average pooling layer and fully connected softmax head. The basic mathematical functions supporting the model are the following.

A convolutional layer takes an input feature map X, and the resultant output scale Y is a product of the learnable kernel K applied at that convolutional layer:

$$Y(i, j) = \Sigma\Sigma\, X^c(i + m, j + n) \cdot K^c(m, n)$$

C is the number of channels of input (3 in the case of RGB). A non-linear activation function, ReLU-type, is used: ReLU(x) = max(0, x): it avoids vanishing gradients, and also converges faster. The last classification layer is a softmax activation that transforms raw logits into the probabilities of a class:

$$\sigma(z_i) = e\hat{}(z_i) / \Sigma_j\, e\hat{}(z_j)$$

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

Training reduces the amount of categorical cross-entropy loss $L = -\Sigma\ y_i \log(\hat{y}_i)$, with y representing the actual one-hot class label and y representing the probability of a prediction being the correct class i. The AdamW optimiser was chosen due to its decoupled weight decay that offers more steadfast generalisation than conventional Adam in image classification problems. The Automatic Mixed Precision (AMP) was used to minimize the amount of money used by the GPU memory at the expense of the numerical accuracy. The training lasted up to 100 epochs and patience-based early stopping was implemented at 20 epochs and the best checkpoint in terms of validation accuracy was saved (best.pt). Initial strong low-level features representations were achieved with transfer learning yolov8n-cls.pt similar to ImageNet trained models, which significantly minimized the training data. It uses the Random Forest classifier as a primary model in its crop recommendation module, because the classifier is resistant to overfitting on tabular data, it inherently supports mixed numeric and categorical features, and its outputs can be used directly by the explainability layer to explain the agroai algorithm. The scores of feature importance are also mapped to agronomic reasoning statements, such as, "soil pH = 6.5 and medium nitrogen: wheat recommended" that are displayed directly to the user.

VI. EXPERIMENTAL SETUP

The training of the model was done in a Python 3.10 hierarchy, with PyTorch, and the Ultralytics [11] YOLOv8 library. Scikit-learn, NumPy and Pandas were used as preprocessing and classical ML experiments. The node.js version of the backend server was configured with the Express.js server and the python machine learning version was separated as a microservice server. React.js and Tailwind CSS were used to create the frontend PWA. Structured records (soil, crop profiles, weather history) were stored in MySQL and unstructured records and chat history were stored in MongoDB. It involved the use of Docker to containerize and Postman to validate the API.

Stratified sampling was used to divide the PlantVillage [12] subset into training (70%), validation (15%), and test (15) sets, and obtained 4,138 test samples of 15 classes. Images that are used as training data were processed through the RandAugment pipeline of YOLOv8 in accordance with Section III. The metrics used to evaluate were per class accuracy, precision, recall, F1-score, and an 15-class confusion matrix. Latency Inference latency was measured as the average time to use the web server to classify one image end-to-end. On the crop recommendation portion, 5-fold cross-validation of the soil dataset was used to determine generalization.

VII. RESULTS AND ANALYSIS

On the held-out test set of 4,138 images the YOLOv8-CLS model attained a total classification accuracy of 98.8%. All macro-averaged precision, recall, and F1-score were equal at 0.99, and weighted averages were the same, which confirms that the high levels of performance were observed even in the most underrepresented type of the dataset (Potato Healthy, n = 31). These scores are higher than the 87% accuracy of Sardeshmukh et al. (2025) [10] on their Random Forest crop recommender and are almost equal to the 99.35% controlled-environment accuracy of Mohanty et al. (2016), but were obtained on a more realistic augmented split than the controlled studio dataset.

Table II represents the entire classification report.

TABLE II
CLASSIFICATION REPORT : YOLOV8-CLS ON PLANTVILLAGE TEST SET

Class	Precision	Recall	F1-Score	Support
Pepper Bell — Bacterial Spot	1.00	1.00	1.00	200
Pepper Bell — Healthy	1.00	1.00	1.00	296
Potato — Early Blight	1.00	1.00	1.00	201
Potato — Late Blight	1.00	0.99	0.99	201
Potato — Healthy	0.97	0.97	0.97	31
Tomato — Bacterial Spot	0.99	0.98	0.98	426
Tomato — Early Blight	0.97	0.96	0.97	201
Tomato — Late Blight	0.98	0.99	0.99	382
Tomato — Leaf Mold	1.00	0.99	0.99	191

Tomato — Septoria Leaf Spot	0.99	0.99	0.99	355
Tomato — Spider Mites	0.99	0.99	0.99	336
Tomato — Target Spot	0.96	0.99	0.97	282
Tomato — Yellow Leaf Curl Virus	1.00	0.99	1.00	642
Tomato — Mosaic Virus	0.99	1.00	0.99	75
Tomato — Healthy	1.00	1.00	1.00	319
Overall / Macro Avg	0.99	0.99	0.99	4,138

The analysis of the confusion matrix allows concluding that misclassification is to a large degree limited to clinically proximate pathologies. The main pairs of confusion include Tomato Early Blight versus Late Blight - diseases with a morphology of early stages lesion with brown-bordered necrotic features - and Tomato Septoria Leaf Spot versus Leaf Mold where small light-centred spots may be ambiguous to the eye at some stage of their development. It is worth mentioning that Potato Healthy had a recall of 0.97 with a support of 31 samples only, which proves that the augmentation and stratification approach proved effective in overcoming the issue of class underrepresentation.

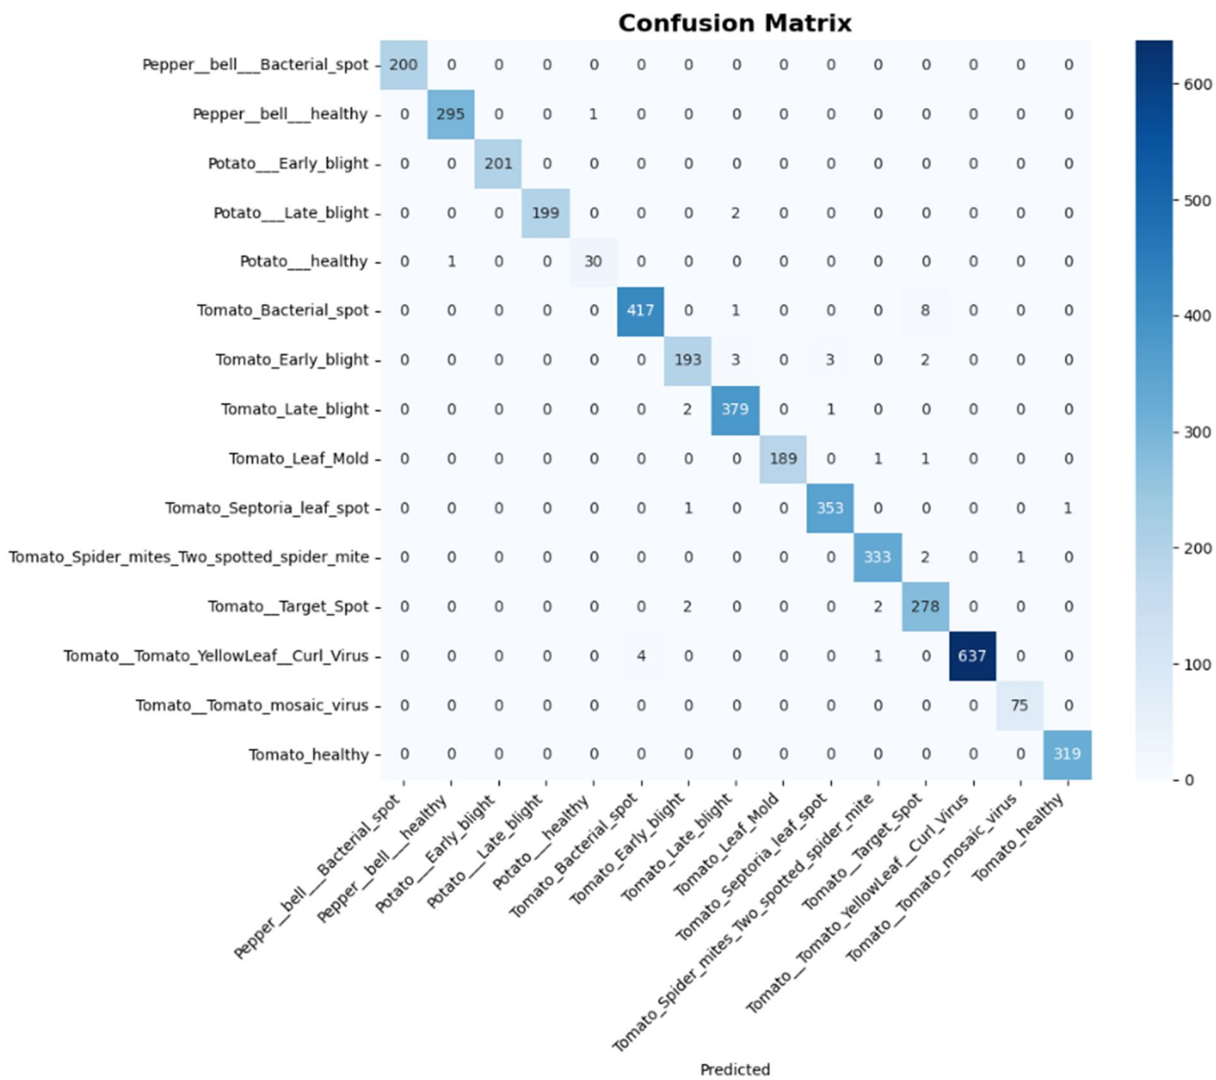

Fig. 2: Confusion Matrix of Test Results.

The largest, most populous class (642 test samples), Tomato Yellow Leaf Curl Virus has a precision 1.00, and recall 0.99, which means that the model generalises over the majority class without any discriminability of the minority-class. Tomato Mosaic Virus, in its turn, had a recall 1.00 using only 75 test samples, indicating that its symptom signature mosaic yellowing with the leaf distortion is sufficiently unusual to be learnt strongly despite a small number of examples.

VIII. DISCUSSION

The results of the experiment place the disease detection module of AgroAI on the high end of the performance spectrum of the plant pathology classification systems. The overall accuracy of 98.8 percent is a good match with the 92 percent of EfficientNetB0-based dermatological detection by Modi et al. (2024) and is close to the controlled-environment ceiling by Mohanty et al. (2016) [6], but was achieved with the conditions that sought to simulate the variability of imaging in the real world through augmentation. Such convergence implies that the performance discrepancy between laboratory models and field-implementable ones, which Mohanty et al. have themselves reported as precipitous, that is, as dropping by 99 percent, is indeed significantly reduced by systematic augmentation and architecture modernization.

The F1 of 0.99 which is macro-averaged over 15 classes is especially significant in the agricultural context. Even a small percentage difference in a recall on a crop like the Tomato Late Blight could result in late treatment and loss of crops worth several percent. The fact that the recall ranged to over 0.96 across all classes, including the grossly underrepresented Potato Healthy class, confirms that the stratified split and RandAugment pipeline are effective to address the problem of dataset imbalance, which is officially admitted by a number of reviewed systems but not solved at the operational level (Kamduri and Gupta, 2025 [3]; Bansal and Singla, 2025) [1].

The confusion pairs that were identified, Early versus Late Blight, Septoria Leaf Spot versus Leaf Mold are similar to what has been found in the plant pathology imaging literature and are indicative of real visual ambiguity and not a failure of the model. Multi-stage diagnostic protocols involving a combination of image-based screening and history of symptoms are used to deal with similar inter-class confusion between eczema and psoriasis, or between basal cell carcinoma and keratosis, in clinical dermatology. In the case of AgroAI, the similar mitigation approach would be to use spatio-temporal metadata, which includes geographic area, time of the year, and previous disease occurrences in the same area and at the farm site, as auxiliary input to the classification head. This is a tangible step towards the development of the models in the future.

The explainability layer of the system which translates feature importance of the crop recommendation model to natural language statements of reasoning addresses the limitation that has been commonly recognized in the AI-in-agriculture literature, but rarely acted upon. Specifically, the reasons why interpretable models should be endorsed include the fact that farmer trust depends on comprehensibility, and the usability of black-box models has always been reported negatively in the context of the community (Kamilaris and Prenafeta-Boldu, 2018) [4]. AgroAI operationalises explainability not as a post-hoc analysis tool but as a leading channel of communication by producing advisory outputs in the form of "soil pH = 6.5 and above-average monsoon rainfall: soybean preferred over cotton this season" one can explainability by generating such outputs.

IX. LIMITATIONS

There are a number of constraints that are clearly recognized. First, the present analysis is limited to the dataset of PlantVillage which, although rather large, is mainly comprised of controlled or semi-controlled photos depicting individual leaves on bare surfaces. Images of real farms often have overlapping leaves, soil contamination, and directional light artefacts. Albeit RandAugment to some extent replicates this diversity, validation data collected in the field of Indian agricultural context assists in verifying performance at the level of deployment. Second, the crop recommendation module is currently not tested on top of ground-truth yield data; its performance is compared with the labels of held-out datasets, and real-life crop yield data are compared to model performance, and the difference between the two needs to be bridged with the help of field experiments over time. Thirdly, the multilingual chatbot has not undergone accent- or dialect-specific usability testing, has not been architecturally tested for supporting languages with low resources (like Odia or Konkani) and has not been tested to support a wide range of Indian languages in practice.

X. FUTURE WORK

There are four directions that need to be given priority in future research. To begin with, the compilation of a geographically varied field-collected dataset of plant diseases in the Indian agricultural areas, covering various growth phases and imaging environments, is the most influential measure in reducing the lab-to-field accuracy gap. Second, the system would use vegetation indices and soil moisture measurements of the Bhuvan platform provided by the satellite-based system to increase the contextual awareness of the system without the need to install the costly on-farm sensors. Third, to quantify the impact of real-world productivity and income, a randomised controlled trial of AgroAI-assisted farmer cohort and control, in a one crop-cycle study, would provide evidence on which to base the value of real-world impact. Fourth, federated learning would enable the performance of models to enhance using the anonymised on-device data of the entire user base without breaching privacy of a farmer.

XI. CONCLUSION

The paper has introduced the AgroAI: Intelligent Smart Agriculture Assistant, a multimodal AI-advised assistant platform that will provide solutions to three related gaps in the agricultural advisory ecosystem in India: limited coverage of tools, reliance on the internet, and a non-transparent method of recommendation. The system incorporates a YOLOv8-CLS convolutional neural network to diagnose crop disease, supervised ML models to explain the recommendations of crops and fertilisers based on the soil, an LLM-based multilingual chatbot with voice and image support, and a Progressive Web App architecture with an offline defugalty architecture.

The disease detection module is experimentally evaluated on a 15-class test set of 4,138 images on the PlantVillage creating a total of 98.8 per cent with macro-averaged precision, recall, and F1-score all equal to 0.99. These are the first results to be obtained in augmentation conditions to simulate real-world imaging variability, a material improvement over the previous systems which could only perform highly under controlled conditions or, by compromising performance to enable offline operation. The analysis of the confusion matrix proves the existence of residual misclassification that is clinically explainable and concentrated in the pathology pairs that are visually proximate.

In addition to the performance of the detection, AgroAI has made its contribution in the form of structural consistency: the first system in the literature reviewed included CNN disease detection, soil-ML recommendation, integration of market prices, multilingual LLM ad dialogue, and offline PWA deployment in a common platform specifically created in Indian smallholder context. The explainability layer, the feature that converts the importance of ML features into plain-language advisory reasoning, directly goes to the barrier of trust that has restricted the application of AI advisory tools in rural settings.

There are still significant drawbacks. Next steps are field-level verification in the environment of various Indian imaging conditions, longitudinal analyses of the effects of yield and income, and chatbots with dialect-specific evaluation. The prototype that is now being developed is about 25-30% of the entire system implementation, and final module integration, live field pilots, and admin analytics dashboards will be developed during the next step in the project. When these elements reach a sufficient level of maturity, AgroAI can emerge as a core output upon which millions of Indian farmers will be able to plan their agronomic decisions with the concentration that evidence-based advice can provide.

XII. ACKNOWLEDGEMENT

The authors wish to say that they are highly grateful to prof. Manoj kumar dixit, an assistant professor, in the department of computer science and engineering, Galgotias College of Engineering and technology, Greater Noida, who has guided and provided constructive feedback and unmatched support all through the period of this study. It also owes debt to the open-source communities of ultralytics yolov8, tensorflow, pytorch, scikit-learn, and react.Js, whose free tools and documentation made this system implementation feasible. The authors will recognize the plantvillage program at penn state university as the source of the image dataset of plant diseases used in the training and evaluation process, and the openweathermap and agmarknet web portals as the sources of their publicly available Apis.

REFERENCES

[1] Bansal, R., & Singla, R. (2025). AI-driven agricultural decision support system using weather and soil data. International Journal of Smart Farming Technologies, 12(4), 45–57.
[2] Kaggle. "PlantVillage Dataset – Plant Disease Images." [Online]. Available: https://www.kaggle.com/datasets/emmarex/plantdisease
[3] Kamduri, R., & Gupta, A. (2025). Smart advisor for precision agriculture using low-infrastructure AI. Journal of Emerging Agricultural Technologies, 9(1), 12–23.
[4] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. Computers and Electronics in Agriculture, 147, 70–90.
[5] Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. Sensors, 18(8), 2674. [Online]. Available: https://www.mdpi.com/1424-8220/18/8/2674
[6] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fpls.2016.01419/full
[7] Pantazi, X. E., Moshou, D., & Tamouridou, A. A. (2019). Automated leaf disease detection through image feature analysis and one-class classifiers. Computers and Electronics in Agriculture, 156, 96–104.

[8] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.
[9] Rajanala, K., Aditya, R., & Prasad, N. (2025). ACVA: Agricultural chatbot voice assistant using MLP and NLP. International Conference on Smart Rural Technologies, 88–95.
[10] Sardeshmukh, M., Gupta, A., & Sahu, P. (2025). AI-based smart crop recommendation framework using soil, weather and market data. International Journal of Agricultural Data Science, 7(2), 112–126.
[11] Ultralytics. (2023). YOLOv8: A new state-of-the-art computer vision model. [Online]. Available: https://github.com/ultralytics/ultralytics
[12] PlantVillage. "Labeled Crop Disease Image Repository." [Online]. Available: https://plantvillage.psu.edu
[13] Agmarknet. Government of India – Agricultural Marketing Portal. [Online]. Available: https://agmarknet.gov.in/

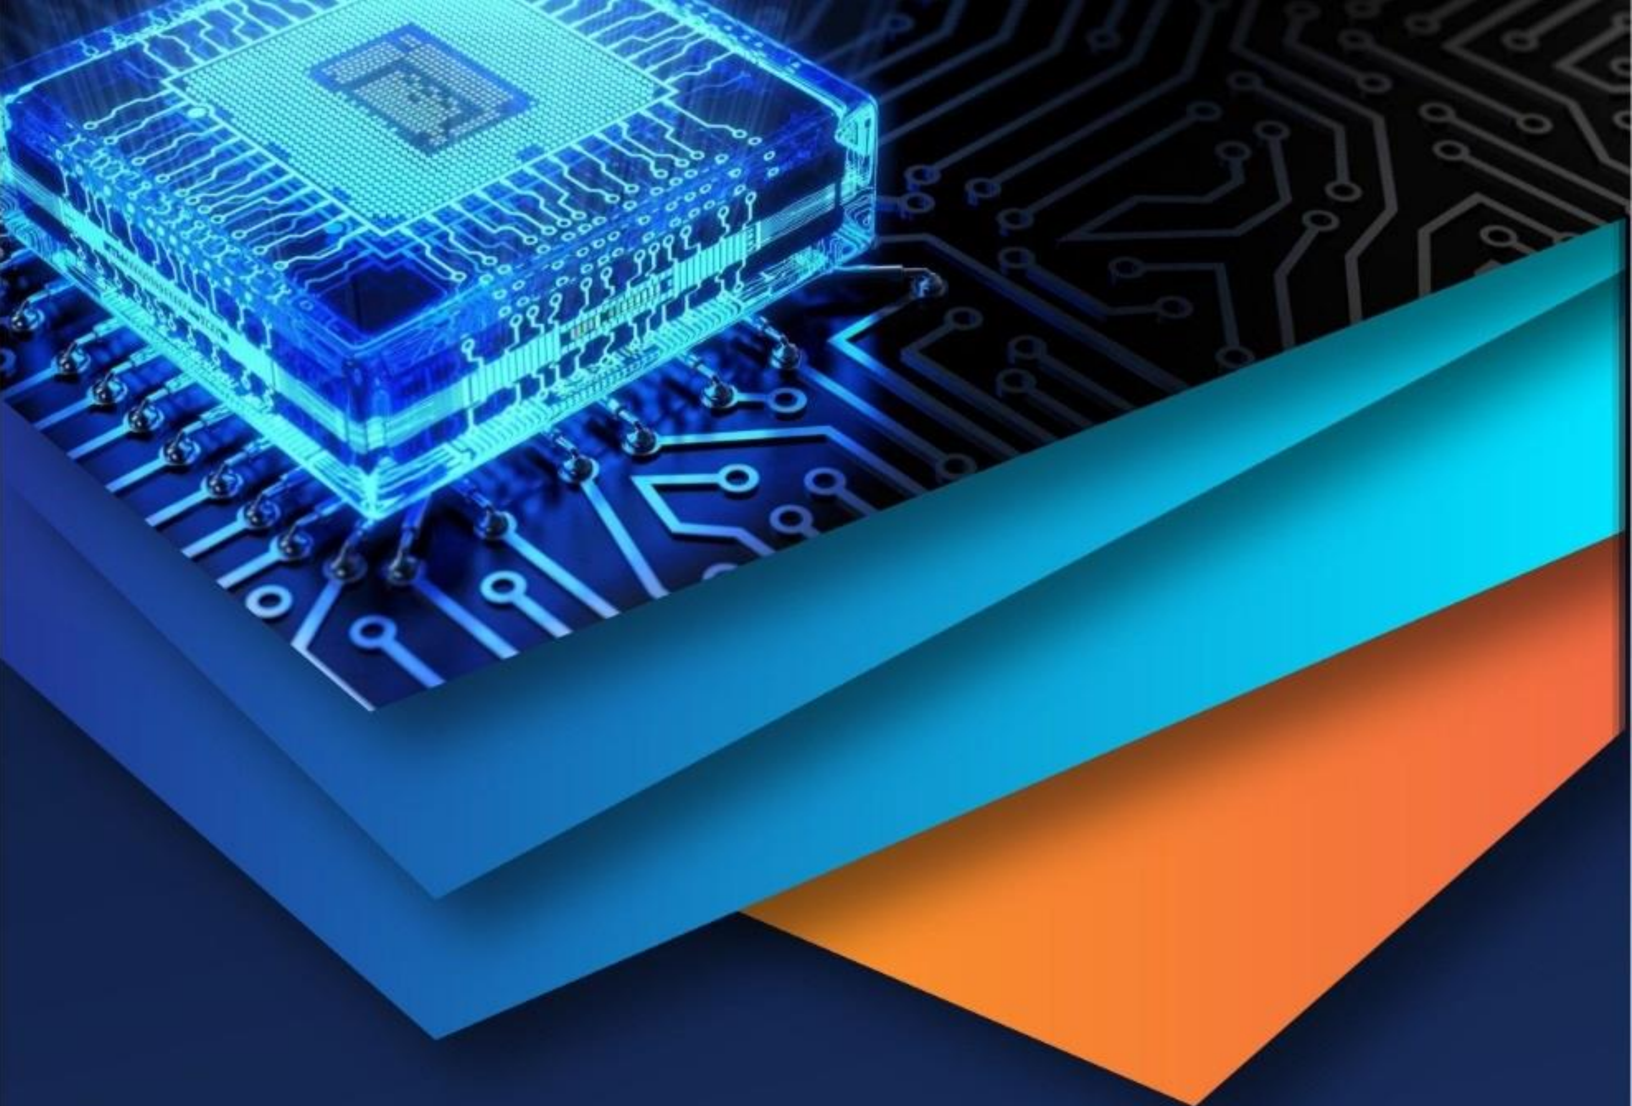

10.22214/IJRASET

INDEX COPERNICUS

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY