AGRISENSE: AN AI-DRIVEN AGRICULTURE ASSISTANCE SYSTEM FOR SOIL ANALYSIS, CROP RECOMMENDATION, AND PLANT DISEASE DETECTION

Antony Sagaya Surya.S[1], Aswin Balaji.A[2], Amarjith.M[3], Mr. A. T. Barani Vijayakumar[4]

[123]UG Student, Department of Computer Science and Engineering, Dhanalakshmi Srinivasan University, Tamilnadu, India

Assistant Professor, Department of Computer Science and Engineering, Dhanalakshmi Srinivasan Institute of Technology, Trichy-621112, Tamil Nadu, India

---***---

Abstract - *Agrisense is an AI-powered agricultural support system that empowers farmers through three intelligent modules: soil analysis, crop recommendation, and plant disease detection. The soil analysis module employs Convolutional Neural Networks (CNNs) to classify soil into Sandy, Clay, Loamy, and Black soil types based on visual features extracted from uploaded images. The crop recommendation engine maps identified soil types to suitable crops and fertilizers. The plant disease detection module uses deep learning with Transfer Learning (CNN + ResNet-50) to analyze plant leaf images, identify diseases, and suggest treatments and preventive measures for more than 10 plant species. A Gemini-powered multilingual chatbot enables farmers to interact via voice or text in their native regional language. The web-based interface Agrisense integrates all modules, supporting image upload, automated analysis, result visualization, and recommendation delivery in real time.*

Keywords: Agrisense, Soil Analysis, Crop Recommendation, Plant Disease Detection, CNN, Deep Learning, Transfer Learning, ResNet-50, Gemini AI, Multilingual Chat bot, Agricultural Support System, Image Classification, Flask, Web Application.

1. INTRODUCTION

Agriculture is the backbone of the global economy, providing livelihoods for over 2.5 billion people worldwide. Despite its critical importance, the agricultural sector faces persistent challenges including soil degradation, improper crop selection, rampant plant diseases, and limited access to expert agronomic knowledge — particularly in developing nations like India. These challenges result in significant annual crop losses, reduced yield, and economic hardship for farming communities that depend on agriculture as their primary source of income.

Artificial Intelligence (AI) and Deep Learning have emerged as transformative technologies capable of addressing these challenges. Convolutional Neural Networks (CNNs) have demonstrated exceptional performance in image-based classification tasks, enabling automated and accurate analysis of soil samples and plant leaf images. Transfer learning techniques further enhance model accuracy by leveraging features learned from large-scale image datasets such as ImageNet, even when domain-specific training data is limited.

Existing agricultural AI tools are largely fragmented addressing only one problem such as disease detection or soil classification and typically lack multilingual support. This limits their adoption by farmers who are not proficient in English. Most systems are also not integrated into userfriendly web platforms accessible on mobile devices, creating accessibility barriers for rural users. Traditional farm advisory services require costly agronomist visits and are not available to small and marginal farmers in real time. The absence of a unified, low-cost, multilingual, AI-powered platform remains a critical gap in the agricultural technology ecosystem.

To address these gaps, this paper proposes Agrisense, a unified AI-powered agricultural support system combining soil analysis, crop recommendation, and plant disease detection in a single web application, with a Gemini-powered multilingual chatbot. The system employs a CNN for soil classification and ResNet-50 transfer learning for plant disease detection. The Gemini API powers the multilingual chatbot, supporting Tamil, Telugu, Hindi, Kannada, Malayalam, and English via voice and text input. The web-based Agrisense interface integrates all modules, enabling seamless image upload, automated analysis, result visualization, and recommendation delivery for farmers across diverse linguistic and geographic contexts.

2. LITERATURE REVIEW

The field of agricultural AI has experienced significant growth with the proliferation of deep learning and computer vision technologies. The following studies from literature form the foundation for the proposed Agrisense system. Table 1 presents a structured summary of the key works reviewed.

Table 1: *Inferences from Literature Survey*

S. No	Author	Title	Year	Methodology	Inference	Merits
1	Diana Susan Joseph et al.	AI in precision agriculture	2024	Hybrid ML	Combines weather & soil for yield prediction	Good adaptability
2	Tripathi & Yadav	AI in sustainable agriculture	2022	Survey	Summarizes AI advancements	Covers multiple agri domains
3	Anuja Bhargava et al.	Plant disease detection using deep learning	2024	CNN	Detects multiple diseases from leaf images	High accuracy in controlled datasets
4	Adhikari & Verma	Soil classification and fertility prediction	2020	Decision Tree	Soil types mapped to crop suitability	Helps in crop recommendation
5	David J. Ritcher et al.	Accelerated CNN Plant disease classification	2025	ResNet-50	Effective for multi-disease detection	Automates disease identification

From the literature review, it is evident that while individual modules for soil classification, crop recommendation, and plant disease detection have been explored in isolation, no unified platform integrates all three alongside a multilingual conversational interface. Most platforms are English-only and do not support voice interaction. Agrisense addresses these gaps through comprehensive preprocessing, transfer learning, and an end-to-end integrated deployment platform with voice-enabled multilingual support.

3. MATERIALS AND METHODS

3.1 Dataset

For soil classification, images of four soil types — Sandy, Clay, Loamy, and Black soil — were collected from Kaggle and Roboflow Universe, totaling approximately 2,400 labeled images (600 per class). Sandy soil images were sourced from arid and semi-arid region datasets, featuring characteristic grainy textures and light brown color profiles. Clay soil samples exhibit smooth, dense textures with darker brown or reddish tones. Loamy soil images show mixed granular textures with moderate moisture retention patterns. Black soil samples, common in Maharashtra and Tamil Nadu, display characteristic dark coloring and high clay content signatures.

For plant disease detection, the PlantVillage dataset was used, comprising over 54,000 labeled leaf images across 14 crop species and 26 disease classes. Class imbalance was addressed through weighted sampling during training. Rare disease classes such as grape leaf blight and corn common rust were upsampled to prevent model bias toward more common classes like tomato early blight. Both datasets were split 80% training, 10% validation, and 10% testing.

3.2 Pre-processing Pipeline

All images were resized to 224×224 pixels and normalized to [0,1]. Gaussian and median filtering reduced noise and enhanced feature clarity for consistent model input. Data augmentation — rotation (±20°), horizontal flipping, zoom (±10%), and brightness adjustment (±15%) — improved generalization across varied field imaging conditions. The augmentation pipeline was validated to ensure no augmented samples leaked into the validation or test sets.

After pre-processing, model training stability improved significantly compared to raw image inputs. Normalization reduced gradient vanishing issues during CNN training. Data augmentation increased effective training set size from 1,920 to approximately 9,600 images, providing the model with diverse viewpoints, orientations, and lighting conditions representative of real-world field photography.

3.3 Soil Classification CNN

The soil classification module uses a custom CNN with four convolutional blocks (Conv2D, Batch Normalization, ReLU, MaxPooling), two Dense layers with Dropout (0.4), and a Softmax output. Trained for 30 epochs using Adam optimizer (lr=0.001, cosine decay), batch size 32, categorical cross-entropy loss. The architecture was designed to be lightweight enough for deployment on standard web server hardware without GPU requirements at inference time.

3.4 Plant Disease Detection (ResNet-50)

Transfer learning with pre-trained ResNet-50 is used for plant disease detection. The base layers are frozen while the custom head — Global Average Pooling, Dense(512, ReLU), Dropout(0.5), Dense(26, Softmax) — is trained. A second fine-tuning phase unfreezes the top 30 ResNet layers at lr=0.0001. This two-phase approach allows the model to adapt ImageNet-learned features to the agricultural domain while avoiding catastrophic forgetting of foundational visual representations.

3.5 Crop Recommendation Engine

A rule-based knowledge system maps identified soil types to crops and fertilizers. Sandy soil: groundnut, watermelon, cassava. Clay soil: rice, wheat, cotton. Loamy soil: vegetables, maize, sugarcane. Black soil: soybean, sunflower, jowar. Irrigation guidance is also provided based on water retention characteristics of each soil type. The knowledge base was compiled in consultation with agronomic literature and validated against regional agricultural extension guidelines for Tamil Nadu, Maharashtra, and Andhra Pradesh.

3.6 Multilingual Chatbot

Powered by the Gemini API, the chatbot accepts voice or text input in Tamil, Telugu, Hindi, Kannada, Malayalam, and English. Voice is captured via the Web Speech API, converted to text, and sent to Gemini with a structured agricultural context prompt. Responses are delivered in the user-selected language. The system provides expert-level guidance on soil health, pest management, irrigation scheduling, fertilizer application, and real-time market price information for major crops..

3.7 Training Configuration

Table 2: *Training Configuration*

Parameter	Soil CNN	Disease (ResNet-50)
Input Size	224×224×3	224×224×3
Epochs	30	25 + 15 fine-tune
Batch Size	32	32
Optimizer	Adam	Adam
Learning Rate	0.001	0.001 / 0.0001
Loss	Categorical Cross-Entropy	Categorical Cross-Entropy
Dropout	0.4	0.5

4. SYSTEM ARCHITECTURE

The Agrisense system follows a modular client-server architecture designed for scalability, maintainability, and ease of deployment. The web frontend (HTML5, CSS3, JavaScript) handles user interaction including image upload, language selection, voice input, and result display. The Flask REST API backend routes incoming requests to the appropriate deep learning model or Gemini API endpoint based on request type.

The backend exposes three primary REST API endpoints: /api/soil for soil classification, /api/disease for plant disease detection, and /api/chat for multilingual chatbot queries. Each endpoint receives a multipart form-data POST request, performs preprocessing and inference, and returns a structured JSON response containing prediction results, confidence scores, and actionable recommendations. The backend is stateless, enabling horizontal scaling and future cloud deployment.

MySQL stores user sessions, uploaded image metadata, prediction history, and chatbot interaction logs. The database schema includes three primary tables: users (session data), predictions (image path, model output, timestamp), and chat_logs (query, response, language, timestamp). This data supports future model retraining and agricultural trend analysis at the regional level. Results are returned as JSON and rendered on the frontend as structured visual cards.

5. MODULE DESCRIPTION

5.1 Soil & Crop Planning

The user uploads a soil sample image through the Agrisense web interface. The interface supports drag-and-drop and file browser uploads on both desktop and mobile devices, accepting JPEG, PNG, and WEBP formats. After OpenCV-based pre-processing including resizing to 224×224, normalizing to [0,1], and applying Gaussian filtering, the CNN classifies the soil as Sandy, Clay, Loamy, or Black. The system then retrieves crop and fertilizer recommendations from the knowledge base and displays them instantly as structured visual cards optimized for mobile screen readability.

Fig. 2: Agrisense — Soil & Crop Planning Interface

5.2 Plant Disease Detection

The farmer uploads a plant leaf image through the Disease Detection & Cure panel. After pre-processing, the fine-tuned ResNet-50 identifies the disease with confidence scoring across 10+ plant species and 26 disease classes. Outputs include disease name, confidence percentage, chemical and organic treatment options, and preventive measures. The interface also suggests when to consult a local agronomist for severe or ambiguous cases. The module processes images within 1.8 seconds on average, enabling farmers to receive diagnostic guidance while still in the field.

5.3 Multi-Language Agri Chatbot

Farmers ask agricultural queries via voice or text in their preferred regional language through the Multi-Language Agri Chatbot panel. The chatbot supports Tamil, Telugu, Hindi, Kannada, Malayalam, and English. Users select their language from a dropdown, then type or speak their question. The Gemini API processes the query with domain-specific agricultural context and provides expert-level responses on soil health, pest management, irrigation scheduling, fertilizer dosage, and live market prices. Voice input is triggered by a dedicated microphone button, captured via the Web Speech API, and converted to text before transmission.

6. SYSTEM WORKFLOW

The end-to-end workflow of Agrisense follows a six-step pipeline. Step 1: The user uploads a soil or leaf image through the Agrisense web interface using drag-and-drop or file browser on desktop or mobile, accepting JPEG, PNG, and WEBP formats. Step 2: The Flask backend performs preprocessing using OpenCV — resizing to 224×224, normalizing to [0,1], and applying Gaussian filtering to reduce noise and enhance feature clarity for consistent model input.

Step 3: The pre-processed image is routed to the appropriate model. For soil analysis, the custom CNN performs inference. For disease detection, the fine-tuned ResNet-50 is invoked. Both models return a predicted class and confidence score within 1.4–1.8 seconds. Step 4: Based on the prediction, recommendations are generated. For soil, suitable crops, fertilizers, and irrigation methods are retrieved from the knowledge base. For disease, treatment protocols and preventive measures are fetched from the advisory database.

Step 5: Results are rendered on the Agrisense frontend as structured visual cards showing the prediction, confidence bar, and recommendations. All results are optimized for mobile screen readability. Step 6: The farmer may ask follow-up questions via the multilingual Chabot. The Chabot accepts voice or text in the selected regional language and provides Gemini-powered expert responses. All interactions are logged in MySQL for future system improvement and agricultural trend analysis.

7. IMPLEMENTATION DETAILS

7.1 Web Application — Agrisense

Agrisense is built with HTML5, CSS3, and vanilla JavaScript. The interface features a dark green theme reflecting the agricultural context and ensures readability in outdoor sunlight. A scrolling news ticker at the top displays real-time agricultural news and market watch updates. The application is structured into three functional panels on a single page: Soil & Crop Planning, Disease Detection & Cure, and Multi-Language Agri Chatbot — eliminating navigation complexity for first-time users. The responsive layout adapts to all screen sizes from 320px mobile to 4K desktop displays.

7.2 Backend Architecture

The Flask backend exposes three REST API endpoints: /api/soil for soil classification, /api/disease for disease detection, and /api/chat for chatbot queries. Each endpoint receives a multipart form-data POST request, performs preprocessing and inference, and returns a JSON response with prediction, confidence score, and recommendations. The backend is stateless, enabling horizontal scaling and future cloud deployment on AWS, Google Cloud Run, or Heroku. TensorFlow SavedModel format is used for optimized model serving.

7.3 Database Design

MySQL stores user sessions, uploaded image metadata, prediction history, and chatbot interaction logs. The schema includes three primary tables: users (session data), predictions (image path, model output, timestamp), and chat_logs (query, response, language, timestamp). Uploaded images are stored temporarily on the server only for the duration of inference and are automatically deleted after results are returned. No personal identification data is collected. All API endpoints are protected against SQL injection through parameterized queries and against cross-site scripting through input sanitization.

7.4 Security and Privacy

Data privacy is a core design principle of Agrisense. The MySQL database stores only anonymized prediction metadata — soil type, disease class, timestamp — without any linkage to individual user identities. The voice input feature records audio only during the active recording session explicitly triggered by the user. No background audio capture occurs, and recorded audio is converted to text locally via the Web Speech API before transmission. Raw audio is never stored or transmitted to external servers, ensuring compliance with user privacy expectations in rural farming contexts.

From an ethical standpoint, Agrisense is designed to augment — not replace — the expertise of professional agronomists and extension officers. All recommendations are presented as data-driven guidance rather than definitive prescriptions. The interface includes clear disclaimers advising users to validate critical decisions with local agricultural experts, particularly for large-scale farming operations involving significant pesticide or fertilizer investments.

8. MODEL PERFORMANCE

8.1 Performance Metrics

The soil CNN achieved 94.5% accuracy (precision: 0.93, recall: 0.92, F1: 0.92). The ResNet-50 disease model achieved 96.2% accuracy (precision: 0.95, recall: 0.94, F1: 0.94) across 26 disease classes. Crop recommendation accuracy was validated at 91.8% against ground truth agronomic guidelines. Chatbot response latency averaged under 2 seconds across all six supported languages.

Table 3: Module Performance Metrics

Module	Metric	Value
Soil Classification (CNN)	Accuracy	94.5%
Soil Classification (CNN)	Precision	0.93
Soil Classification (CNN)	Recall	0.92
Soil Classification (CNN)	F1-Score	0.92
Disease Detection (ResNet-50)	Accuracy	96.2%
Disease Detection (ResNet-50)	Precision	0.95
Disease Detection (ResNet-50)	Recall	0.94
Disease Detection (ResNet-50)	F1-Score	0.94
Crop Recommendation	Accuracy	91.8%
Chatbot (Gemini)	Avg. Response	< 2 sec

8.2 Real-Time Performance

Soil classification inference averaged 1.4 seconds per image on a standard CPU server. Disease detection averaged 1.8 seconds per image. Chatbot response latency remained under 2 seconds across all languages. The web interface loaded within 2.5 seconds on broadband and 4.1 seconds on 4G mobile networks, confirming real-time suitability for field deployment. Memory footprint of the CNN soil model is 18MB and ResNet-50 disease model is 94MB, both within acceptable limits for server deployment.

8.3 Error Analysis and Failure Cases

A systematic error analysis was conducted on the test set predictions for both models. The soil classification CNN showed the highest confusion between Clay and Loamy soil samples, which share overlapping visual features including similar color tones and granular surface textures under certain lighting conditions. This confusion accounted for approximately 60% of all soil classification errors. Future work will address this by incorporating additional texture feature extraction techniques such as Local Binary Patterns (LBP) and Gabor filters.

For disease detection, the most common failure cases involved early-stage diseases where symptoms were subtle and not yet visually distinctive. The model also showed reduced confidence on images captured under very bright direct sunlight, which can wash out leaf color and reduce the visibility of disease lesions. Images taken in low-light or highly shadowed conditions similarly showed lower confidence scores. These findings inform the preprocessing recommendations provided to end users within the Agrisense interface, including guidance on optimal image capture conditions.

9. TOOLS & TECHNOLOGIES

Programming Language: Python 3.10

Deep Learning Framework: TensorFlow 2.x / Keras

Transfer Learning Model: ResNet-50 (pre-trained on ImageNet)

Preprocessing: OpenCV 4.x, NumPy, Albumentations

Web Framework: Flask (Backend REST API)

Frontend: HTML5, CSS3, JavaScript

Database: MySQL

AI Chatbot: Gemini API (Google)

Voice Input: Web Speech API

Visualization: Matplotlib, Seaborn

Dataset Sources: Roboflow Universe, Kaggle, PlantVillage

10. DISCUSSION

The Agrisense system demonstrates strong performance across all three core modules. The soil CNN 94.5% accuracy reflects the robustness of the architecture trained with comprehensive data augmentation. The disease detection model 96.2% accuracy confirms that ResNet-50 transfer learning significantly improves performance even with moderate dataset sizes compared to training from scratch. Both results exceed previously reported benchmarks for comparable agricultural image classification systems on the PlantVillage and Roboflow datasets.

The integration of a Gemini-powered multilingual chatbot is a key innovation distinguishing Agrisense from all reviewed existing platforms. Language barriers have historically limited agricultural technology adoption in rural India. By supporting six regional languages through both voice and text, Agrisense removes this barrier and makes AI-powered guidance genuinely accessible to smallholder farmers with limited English literacy or keyboard proficiency. This is especially significant for elderly farmers and women farmers who form a substantial portion of the agricultural workforce in southern India.

The unified web platform Agrisense provides a seamless experience by integrating all modules under a single interface, eliminating the need for farmers to use multiple disconnected tools. Mobile compatibility ensures access from smartphones in the field. The system's three-module integration — soil, disease, and crop recommendation — within a single freely accessible platform represents a significant advancement over existing single-domain solutions.

A key strength of Agrisense over existing solutions is its zero-cost deployment model. Most commercial agricultural AI platforms such as Plantix and AgroStar operate on subscription-based or freemium models that are often economically inaccessible to smallholder farmers operating on thin margins. By designing Agrisense as a freely deployable open platform for agricultural institutions, Krishi Vigyan Kendras, and state government departments, the system can reach millions of farmers as a public service without imposing any direct cost on the end user.

Voice-enabled interaction in regional languages is particularly transformative for digital agricultural inclusion. Research consistently shows that voice interfaces dramatically lower the digital literacy barrier for first-time technology users. Farmers who have never typed on a keyboard can access expert agricultural guidance simply by speaking into their smartphone — a capability that has significant potential to democratize access to AI-powered farming intelligence across rural India.

11. SOCIAL IMPACT AND SIGNIFICANCE

India's agricultural sector provides livelihoods for approximately 58% of the rural population. Farmers face severe information asymmetry — lacking timely access to agronomic expertise, market intelligence, and disease diagnostics. Agrisense bridges this gap by delivering AI-powered guidance through a mobile-accessible web platform available to any farmer with a smartphone and internet connection.

The multilingual support feature is of particular social significance. India has 22 officially recognized languages. The majority of smallholder farmers in Tamil Nadu, Andhra Pradesh, Karnataka, Kerala, and northern states are not proficient in English. By supporting Tamil, Telugu, Kannada, Malayalam, Hindi, and English through voice and text, Agrisense ensures that language is never a barrier to accessing intelligent agricultural guidance.

Early and accurate plant disease detection can prevent crop losses of 20–40% annually in developing countries. The Agrisense disease detection module, with 96.2% accuracy and 1.8-second inference time, enables farmers to identify diseases within minutes of symptom appearance — dramatically reducing the time between detection and treatment compared to traditional methods that rely on waiting days or weeks for an expert visit. This speed advantage directly translates to reduced crop damage and improved yield for smallholder farming families.

Proper soil management is foundational to sustainable agriculture. Inappropriate crop selection leads to poor yield and unnecessary fertilizer expenditure. The Agrisense soil analysis module provides instant, data-driven classification and crop-fertilizer recommendations without requiring expensive laboratory soil testing — democratizing access to soil intelligence for smallholder farmers who cannot afford formal testing services that typically cost INR 500–1,500 per sample.

12. COMPARATIVE ANALYSIS

Several agricultural AI platforms exist globally including Plantix, AgroStar, and the ICRISAT AI Sowing App. Plantix focuses exclusively on disease detection and does not provide soil analysis or crop recommendation. AgroStar provides market and advisory services but does not incorporate image-based AI analysis. The ICRISAT Sowing App uses weather and sensor data for sowing recommendations but does not support image-based soil or disease analysis. None of the reviewed commercial platforms offer integrated multilingual voice interaction in Indian regional languages.

Table 4: Feature Comparison with Existing Platforms

Feature	Agrisense	Plantix	AgroStar	ICRISAT App
Soil Classification	Yes	No	No	Partial
Crop Recommendation	Yes	No	Yes	Yes
Disease Detection	Yes (10+ crops)	Yes	No	No
Multilingual Support	Yes (6 lang.)	Partial	Partial	No
Voice Input	Yes	No	No	No
Mobile Compatible	Yes	Yes	Yes	Yes
Free Access	Yes	Freemium	Freemium	Limited
Integrated Platform	Yes	No	No	No

13. DEPLOYMENT STRATEGY

Agrisense is currently deployed as a local web server application for testing and validation purposes. The Flask backend and MySQL database run on a single server instance suitable for small-scale pilot deployment with up to

100 concurrent users. For production deployment at national scale, a cloud-based containerized architecture is planned using Docker and Kubernetes. The soil and disease models will be served via TensorFlow Serving, enabling GPU-accelerated inference and horizontal scaling during peak agricultural seasons.

A three-phase deployment roadmap is defined. Phase 1 — Pilot (Month 1–3): Deploy to 50 volunteer farmers in 5 villages in Tamil Nadu in collaboration with a local agricultural university. Collect real-field image data and farmer feedback to improve model accuracy for local conditions. Phase 2 — Regional Launch (Month 4–8): Expand to 500 farmers across Tamil Nadu and Andhra Pradesh. Retrain models with Phase 1 field data. Launch Android mobile application. Phase 3 — National Scale (Month 9–18): Cloud deployment supporting 50,000+ users, additional regional languages, and expanded soil and disease coverage.

14. LIMITATIONS

The soil classification model covers four broad soil categories. Regional variations such as red soil common in Tamil Nadu and laterite soil in coastal regions are not yet covered. The disease detection model was primarily trained on PlantVillage laboratory images; real-field images with occlusions, shadows, and multiple overlapping diseases may show reduced confidence. The system currently requires an active internet connection, limiting usability in remote areas with poor connectivity. The crop recommendation engine uses a rule-based approach that does not account for local weather patterns, seasonal variations, or market price fluctuations in real time.

15. CONCLUSION

This paper presents Agrisense, a comprehensive AI-powered agricultural support system that integrates soil analysis, crop recommendation, plant disease detection, and a multilingual chatbot into a unified web platform. The CNN-based soil classifier achieves 94.5% accuracy and the ResNet-50 disease detection model achieves 96.2% accuracy across 14 crop species and 26 disease classes. These results confirm that deep learning and transfer learning are highly effective for agricultural image classification tasks in real-world deployment scenarios.

The Gemini-powered multilingual chatbot supports voice and text interaction in six regional languages, making the system highly accessible to farmers across diverse linguistic backgrounds. The web-based Agrisense interface facilitates seamless image upload, automated analysis, and actionable recommendations, supporting faster and more informed agricultural decision-making. The system's zero-cost, open-platform deployment model makes it uniquely positioned to serve smallholder farmers across India who cannot afford commercial agricultural AI subscriptions.

Overall, Agrisense demonstrates high reliability, scalability, and practical value as a tool for advancing intelligent and inclusive agriculture in developing regions. The integration of soil, disease, and crop modules within a single voice-enabled, multilingual platform addresses a critical gap in existing agricultural AI systems and represents a meaningful contribution toward digital agricultural transformation in India and similar developing economies.

16. FUTURE WORK

The plant disease detection model will be extended to support more plant species, including tropical and region-specific crops such as banana, turmeric, cardamom, and areca nut. Expanding the soil class taxonomy to include red soil and laterite soil common in Tamil Nadu and coastal regions is also planned. Integration of real-time weather data from public APIs and LoT-based soil sensor readings will enable dynamic, location-specific crop recommendations accounting for seasonal and environmental variability.

The multilingual chatbot will be fine-tuned on domain-specific agricultural datasets to improve response accuracy for technical and region-specific queries. Support for additional regional languages including Bengali, Marathi, and Punjabi is planned. A dedicated Android and iOS mobile application will be developed to improve field accessibility and provide offline functionality in areas with limited internet connectivity. Offline inference using Tensor Flow Lite models will be integrated to enable disease detection and soil classification without internet access.

Clinical-style validation trials in partnership with agricultural universities and government extension services such as Krishi Vigyan Kendras are planned to evaluate the system's real-world impact on crop yield, disease incidence, and farmer income. These trials will provide empirical evidence to support broader deployment and policy adoption of AI-based agricultural decision support systems across India and other developing nations. Integration with the PM-KISAN farmer database is being explored to enable personalized recommendations based on registered land holdings and historical crop patterns.

Acknowledgements

The authors sincerely thank the faculty and staff of the Department of Computer Science and Engineering, Dhanalakshmi Srinivasan University, for their continuous guidance and support. Special thanks to the project guide for valuable feedback, technical insights, and encouragement throughout the development process. The authors also acknowledge the open-source community behind TensorFlow, Keras, Flask, and OpenCV, and Google's Gemini API team for access to the language model powering the chatbot. This project was completed as part of the undergraduate final year project requirements for the academic year 2025–2026.

REFERENCES

[1] Anuja Bhargava, Aasheesh Shukla, Om Prakash Goswami (2024). Plant disease detection using deep learning: A comprehensive review of methods and benchmarks. Journal of Agricultural Informatics, 15(2), 45–62.

[2] Adhikari, R., & Verma, R. (2020). Soil classification using machine learning and digital image processing. International Journal of Agricultural Science and Research, 10(3), 112–124.

[3] David J. Ritcher, Md Ilias Bappi, Shivani Sanjay Kolekar (2025). Accelerated CNN-based plant disease classification using ResNet-50 transfer learning. IEEE Transactions on Agricultural Engineering, 12(1), 88–97.

[4] Tripathi, A., & Yadav, P. (2022). Artificial intelligence in sustainable agriculture: A comprehensive survey. Computers and Electronics in Agriculture, 198, 107–125.

[5] Diana Susan Joseph, Pranav, Kaustubh Chackradeo (2024). AI in precision agriculture: Hybrid machine learning models for yield prediction. Smart Agriculture Technology, 7, 100–115.

[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.

[7] Kamilaris, A., & Prenafeta-Boldú, F.X. (2018). Deep learning in agriculture: A survey. Computers and Electronics in Agriculture, 147, 70–90.

[8] Ferentinos, K.P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318.

[9] Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2021). Machine learning applications for precision agriculture: A comprehensive review. IEEE Access, 9, 4843–4873.

[10] Google DeepMind (2024). Gemini: A family of highly capable multimodal models. Technical Report, Google Research.