

AgriVeda AI: An Intelligent Smart Agriculture Platform Using Machine Learning and Explainable AI

PRAKRUTHI H V

PostGraduate Student, Dept. of CSE
Rajeev Institute of Technology, Hassan
Hassan, India
Prakruthivenkatesh92@gmail.com

Dr. H N Prakash

Professor, Dept. of CSE
Rajeev Institute of Technology, Hassan
Hassan, India
Prakash_hn@yahoo.com

Abstract - Agriculture sustains the livelihoods of more than 700 million people across India and contributes roughly 18% of the national GDP. Despite this central role, farmers continue to face compounding difficulties: late detection of crop diseases, unreliable yield estimates, scarce expert guidance, and unpredictable market conditions. This paper introduces AgriVeda AI, an integrated smart agriculture web platform powered by Artificial Intelligence and Machine Learning. The system consolidates ten functional modules within a single browser-accessible dashboard: dual-mode crop disease detection (tabular field-condition analysis using Random Forest and leaf-image analysis using MobileNetV2 with Grad-CAM heatmaps), SHAP-explainable yield prediction, soil-driven crop recommendation, FAO-model-based smart irrigation planning, a Retrieval-Augmented Generation (RAG) Farm Assistant, real-time weather intelligence, live market prices, government scheme discovery, and an AI-powered shopping list. Performance evaluation on held-out test data yields 94% accuracy for CNN-based disease detection across 38 disease classes, 98.2% classification accuracy for crop recommendation, and an R-squared value of 0.98 for yield prediction. The cross-module shopping list, which automatically aggregates purchase recommendations from disease, irrigation, and crop-recommendation outputs, represents a practical advance toward holistic precision agriculture for smallholder farmers in Karnataka, India.

Key Words: Machine Learning, Deep Learning, Convolutional Neural Network, Random Forest, Grad-CAM, SHAP, Natural Language Processing, Retrieval-Augmented Generation, Flask, React

1. INTRODUCTION

Agriculture contributes approximately 18% of India's Gross Domestic Product and sustains over 700 million livelihoods across the country. Despite this foundational importance, the sector suffers from persistent technological deficiencies that reduce productivity and constrain farmer income. Conventional farming practice relies heavily on manual field observation, delayed consultation with agronomists, and restricted access to timely market and weather data—conditions that collectively result in significant crop losses and inefficient use of inputs such as water, fertiliser, and pesticides. The growing convergence of Artificial Intelligence (AI), Machine Learning (ML), and affordable web technologies opens a genuine pathway toward precision agriculture[1]. However, the digital tools currently available to Indian farmers remain fragmented, each addressing a single

sub-problem in isolation—disease identification, weather forecasting, or mandi price lookup—without the contextual integration required for holistic farm management. Small and marginal farmers, who represent over 86% of all Indian farmholders [1], cannot practically navigate multiple disconnected applications under time and connectivity constraints.

This paper presents AgriVeda AI, a unified smart agriculture platform that brings ten AI-powered modules together within one dashboard accessible from any standard web browser. The backend is implemented as a Flask (Python) REST API while the frontend is built with React, making the system deployable locally or on cloud infrastructure without additional hardware. Four specific contributions distinguish AgriVeda AI from prior work: (1) a dual-mode disease detection engine combining tabular Random Forest classification with CNN image analysis augmented by Grad-CAM visual explanations; (2) SHAP-based per-prediction feature importance for yield forecasting, translating model outputs into actionable farm management advice; (3) a RAG-based Farm Assistant covering over 100 topics across eight agricultural knowledge domains; and (4) a cross-module shopping list that automatically consolidates purchase recommendations generated by the disease, yield, irrigation, and crop-recommendation modules. The remainder of this paper is structured as follows. Section II reviews related work. Section III describes the system architecture. Section IV details the methodology of each functional module. Section V covers implementation tools. Section VI presents experimental results. Section VII concludes with future directions.

2. RELATED WORK

Plant disease detection using deep learning gained widespread attention after Mohanty et al. [2] showed that a CNN trained on the PlantVillage image repository could identify 26 disease categories with 99.35% accuracy under laboratory conditions. Ferentinos [3] extended this work across a broader disease taxonomy using AlexNet and GoogleNet architectures. A consistent limitation of these early systems, however, was that their prediction outputs provided no visual guidance to the end-user regarding which regions of a leaf drove the classification decision. Selvaraju et al. [4] addressed this interpretability gap through Gradient-weighted Class Activation Mapping (Grad-CAM), which generates heatmaps localising the network's attention onto the input image.

AgriVeda AI adopts Grad-CAM to overlay disease-focus regions directly onto uploaded leaf photographs, giving

farmers and extension workers a transparent basis for accepting or querying a diagnosis[2].

Jeong et al. [5] applied Random Forest and Gradient Boosting regressors to FAO multi-country datasets and demonstrated that ensemble methods substantially outperform traditional linear models on heterogeneous agricultural data. Pantazi et al. [6] combined remote sensing imagery with ML to predict wheat yield at the sub-field level. A recurring limitation in this body of work is that predictions are delivered as single scalar values without explanatory context, making it difficult for farmers to identify which management factors are responsible for a shortfall. AgriVeda AI addresses this by integrating SHAP TreeExplainer [7], which decomposes each yield prediction into per-feature Shapley value contributions, visualised as waterfall charts and accompanied by plain-language management insights. The application of Retrieval-Augmented Generation (RAG) to domain-specific question answering was formalised by Lewis et al. [8].

Kamilaris and Prenafeta-Boldu [9] surveyed deep learning applications in agriculture and identified NLP-based advisory systems as a priority research direction. The Farm Assistant module in AgriVeda AI realises this direction through TF-IDF cosine-similarity retrieval from curated knowledge bases, with optional large language model enhancement via the GROQ API, producing grounded, citation-backed responses that are more reliable than purely generative approaches. Compared with existing platforms such as PlantVillage, AgroAI, KrishiNet, and FASAL, AgriVeda AI is, to the best of the authors' knowledge, the first unified system to offer dual-mode disease detection, SHAP-explainable yield prediction, crop recommendation, AI chatbot with RAG, live market prices, and government scheme discovery within a single production-ready web application.

3. SYSTEM ARCHITECTURE

AgriVeda AI follows a six-layer architecture as illustrated in Fig. 1. The layers are: (1) React Frontend — User Interface; (2) Flask Backend — REST API layer; (3) Trained ML/DL Models (.pkl, .h5); (4) Training Datasets (Kaggle, FAO, PlantVillage); (5) External APIs for live data; and (6) User-facing Output. All inter-layer communication uses HTTP/JSON.

The frontend is a React 18 single-page application (SPA) built with Vite 5, featuring a dark-themed sidebar navigation providing access to all ten modules. The Flask 3.0 backend exposes RESTful endpoints under the /api/ namespace. Models are loaded once at startup using joblib (scikit-learn models) and keras.models.load_model (CNN), minimizing per-request inference latency. A session-based in-memory history store tracks prediction results across modules to enable the cross-module shopping list integration. This paper presents AgriVeda AI, a unified smart agriculture platform that brings ten AI-powered modules together within one dashboard accessible from any standard web browser. The backend is implemented as a Flask (Python) REST API while the frontend is built with React, making the system deployable locally or on cloud infrastructure without additional hardware. Four specific contributions distinguish AgriVeda AI from prior work:

- (1) a dual-mode disease detection engine combining tabular Random Forest classification with CNN image analysis augmented by Grad-CAM visual explanations;
- (2) SHAP-based per-prediction feature importance for yield forecasting, translating model outputs into actionable farm management advice;
- (3) a RAG-based Farm Assistant covering over 100 topics across eight agricultural knowledge domains; and
- (4) a cross-module shopping list that automatically consolidates purchase recommendations generated by the disease, yield, irrigation, and crop-recommendation modules. The remainder of this paper is structured as follows. Section II reviews related work. Section III describes the system architecture. Section IV details the methodology of each functional module.




Figure 1: AgriVeda AI complete data flow diagram showing all six layers from user input to AI-generated outputs.

The frontend is a React 18 single-page application (SPA) built with Vite 5, featuring a dark-themed sidebar navigation providing access to all ten modules. The Flask 3.0 backend exposes RESTful endpoints under the /api/ namespace. Models are loaded once at startup using joblib (scikit-learn models) and keras.models.load_model (CNN), minimizing per-request inference latency. A session-based in-memory history store tracks prediction results across modules to enable the cross-module shopping list integration.

4. FUNCTIONAL MODULES

A. Crop Disease Detection

The disease detection module employs a dual-mode architecture to handle both tabular field conditions and uploaded leaf images. In tabular mode, six observational parameters are collected: crop type, ambient humidity (%), temperature (degrees Celsius), recent rainfall (mm), leaf color (yellowing/browning/dark spots/wilting), and spot/lesion size (<5mm, 5-15mm, >15mm). These features are encoded using

LabelEncoder and classified by a Random Forest Classifier ($n_estimators=200$) trained on a curated field condition dataset. The model outputs a disease label, confidence score, and top-3 predictions with probabilities. In image mode, the uploaded leaf photograph is resized to 224x224 pixels, normalized to [0,1], and fed to a fine-tuned MobileNetV2 CNN pre-trained on ImageNet and fine-tuned on the PlantVillage dataset [10] comprising 54,306 images across 38 disease classes. The final fully connected layer uses softmax activation for 38-class output. Grad-CAM [4] is applied to the last convolutional layer to generate a heatmap overlaid on the original image, with red regions indicating disease focus and blue indicating healthy tissue. An infection severity percentage is computed as the proportion of high-activation pixels (threshold: top 40%) to total leaf-area pixels, classified as: Mild (<30%), Moderate (30-60%), or Severe (>60%). Both chemical and organic/biological treatment recommendations are retrieved from a disease-treatment lookup database and displayed alongside the diagnosis.

B. Crop Yield Prediction with SHAP Explainability

The yield prediction module uses a Random Forest Regressor ($n_estimators=100$, $random_state=42$) trained on the FAO Global Crop Yield dataset covering 100+ countries and multiple decades (1990-2023). Input features include: crop type, country/region, year, average annual rainfall (mm/year), average temperature (degrees Celsius), and pesticide usage (tonnes). Categorical variables are encoded using LabelEncoder fitted on training data.

Model performance on the 20% held-out test set: R-squared = 0.98, Mean Absolute Error (MAE) = 892 kg/ha, Root Mean Squared Error (RMSE) = 1,247 kg/ha. Post-prediction, SHAP TreeExplainer computes per-feature Shapley values, visualized as a waterfall chart showing absolute feature contributions (in kg/ha) relative to the baseline average yield. Each feature card provides a specific farm insight — for example, flagging suboptimal pesticide dosage if the pesticide feature shows a large negative SHAP contribution. Model performance on the 20% held-out test set: R-squared = 0.98, Mean Absolute Error (MAE) = 892 kg/ha, Root Mean Squared Error (RMSE) = 1,247 kg/ha. For each user query, cosine similarity retrieval identifies the top-k most relevant text chunks, which are concatenated and returned with source citations.

C. Crop Recommendation

The crop recommendation module accepts seven soil and climate parameters: Nitrogen (N, kg/ha), Phosphorus (P, kg/ha), Potassium (K, kg/ha), soil pH (0-14), rainfall (mm), temperature (degrees Celsius), humidity (%), and optional region name. A Random Forest Classifier ($n_estimators=200$) trained on the Kaggle

Crop Recommendation Dataset (2,200 samples, 22 crop classes) achieves 98.2% test accuracy. The system returns the top-5 recommended crops with confidence percentages and agronomic justification for each recommendation, enabling farmers to make informed planting decisions.

D. AI Farm Assistant (RAG-NLP)

The Farm Assistant implements a Retrieval-Augmented Generation (RAG) pipeline. Four curated agricultural knowledge bases are indexed using TF-IDF vectorization: Soil Health Guide, Pest Management Guide, Irrigation Techniques

Guide, and Fertiliser Usage Guide. For each user query, cosine similarity retrieval identifies the top-k most relevant text chunks, which are concatenated and returned with source citations[9]. Optional integration with the GROQ API (Llama-3 70B) generates fluent, LLM-enhanced responses grounded in the retrieved context.

The chatbot interface includes eight navigable category tabs (Farming Methods, Soil & Fertilizers, Irrigation, Pests & Diseases, Crops & Stats, Govt. Schemes, Animal Husbandry, Technology) with 100+ pre-loaded sample questions, enabling farmers to browse topics without formulating search queries. The system is available 24/7 without requiring internet connectivity beyond the initial page load. Model performance on the 20% held-out test set: R-squared = 0.98, Mean Absolute Error (MAE) = 892 kg/ha, Root Mean Squared Error (RMSE) = 1,247 kg/ha. Post-prediction, SHAP TreeExplainer computes per-feature Shapley values, visualized as a waterfall chart showing absolute feature contributions (in kg/ha) relative to the baseline average yield.

E. Smart Irrigation Planning

The irrigation module collects seven farm parameters: crop type, soil type, crop growth stage, land area (ha), current soil moisture (%), ambient temperature (degrees Celsius), and expected rainfall (mm). A Random Forest Regressor (R-squared = 0.95) predicts the water requirement in litres. The system additionally computes: irrigation frequency (days between sessions), optimal schedule time (morning 6-8AM or evening 5-7PM), and crop/soil type classification. Historical irrigation plans are stored in an in-memory session log.

F. Weather Intelligence, Market Prices & Government Schemes

The Weather Intelligence module integrates the Open-Meteo API to retrieve real-time current conditions and 7-day forecasts for any geographic location. Users may search by city name or select from a Karnataka Districts quick-select dropdown. The displayed data includes temperature, humidity, wind speed (km/h), precipitation (mm), UV index, and an Agricultural Advisory generated from rule-based logic applied to retrieved conditions.

The Market Prices module fetches live vegetable mandi prices from the AGMARKNET[5] Government of India API, presenting average modal price, minimum price, maximum price, and a mandi-wise breakdown table for any selected commodity and state. The Government Schemes module presents 35 national and state agricultural schemes (including PM-KISAN, PMFBY, Kisan Credit Card, Soil Health Card Scheme, and PMKSY) with eligibility criteria, benefit amounts, deadlines, and direct Apply links."

5. IMPLEMENTATION AND TOOLS

AgriVeda AI is implemented as a full-stack web application. Table I summarizes the complete technology stack. The Flask REST API backend is implemented in Python 3.11. The ML models are trained and serialized using scikit-learn 1.4 (joblib) and TensorFlow 2.13/Keras. Image preprocessing uses OpenCV 4.9. SHAP analysis uses the shap 0.45 library with TreeExplainer.

External API communication uses the Python requests library. The React 18 frontend is built with Vite 5 and uses custom CSS for the dark agricultural theme.

TABLE I

AGRIVEDA AI TECHNOLOGY STACK

Component	Technology / Library
Backend	Python 3.11, Flask 3.0, Gunicorn
ML	scikit-learn 1.4, shap 0.45, joblib
DL	TensorFlow 2.13, Keras, OpenCV 4.9
Data	Pandas 2.1, NumPy 1.26
NLP/RAG	NLTK, TF-IDF, GROQ API (Llama-3)
Frontend	React 18, Vite 5, JavaScript ES6+
APIs	Open-Meteo, AGMARKNET, Kaggle Hub
Dev Tools	VS Code, Git, npm, Node.js 20

The application is launched with two commands: python app.py (Flask backend, port 5000) and npm run dev (React frontend). Models are loaded once at startup. The full source project is organized into backend/ (Flask + ML) and frontend/ (React + Vite) directories under the agri-smart-hub root. Cloud deployment is supported via Heroku, AWS EC2, or Railway with minimal configuration changes.

6. RESULTS AND DISCUSSION

A. Model Performance

Table II summarizes the performance of all trained ML/DL models. The Random Forest Crop Recommender achieves the highest classification accuracy at 98.2% on the 20% held-out test set. The CNN disease detector achieves 94% test accuracy across 38 disease classes on the PlantVillage dataset. The yield prediction model achieves R-squared of 0.98, indicating that 98% of variance in crop yield is explained by the six input features.

TABLE II

MODEL PERFORMANCE SUMMARY

Model	Algorithm	Dataset	Metric
Yield Prediction	RF Regressor	FAO Global (100+ countries)	R2=0.98 MAE=892
Disease (Image)	MobileNetV2	PlantVillage (54,306)	Acc=94%
Disease	RF Classifier	Field conditions	Acc=87%

Model	Algorithm	Dataset	Metric
(Tabular)		DB	
Crop Recommendation	RF Classifier	Kaggle (2,200)	Acc=98.2 %
Irrigation Planning	RF Regressor	Agri irrigation DB	R2=0.95

B. Disease Detection Validation

Testing across 20 leaf images covering tomato, pepper, potato, and squash crops confirmed accurate disease identification. For a pepper leaf image (Fig. 3), the CNN+Grad-CAM system diagnosed Bacterial Spot with 53% confidence, correctly identifying the bacterial lesion regions as the primary disease focus in the heatmap. The severity computation returned 41.8% leaf area affected (Moderate Infection).

The Grad-CAM overlay correctly highlighted the darker lesion zones in the central and basal regions of the leaf, consistent with expert agronomist assessment. Tabular mode testing for Tomato with yellowing leaves, small lesions, humidity=75%, temperature=25 degrees Celsius, and rainfall=30mm correctly identified Early Blight (22% confidence top prediction, followed by Septoria Leaf Spot at 20% and Bacterial Spot at 14%).

C. Yield Prediction Validation

For Potatoes in India with 800mm annual rainfall, 25 degrees Celsius average temperature, 50 tonnes pesticides, and 1 ha land area, the system predicted 11,360 kg/ha (Fig. 4). For Cassava under identical conditions, the prediction was 20,305 kg/ha — consistent with FAO benchmark yields for these crops in Indian conditions.

The SHAP analysis (Fig. 5) revealed that pesticide overuse (50 tonnes) was the largest negative contributor (-5,805 kg/ha, 51.1%), generating a specific farm insight recommending consultation with an agronomist to review application timing and dosage. Crop type selection contributed the largest positive impact (+9,054 kg/ha, 79.7%), validating the model's sensitivity to crop-specific yield potential.

D. Crop Recommendation Validation

For Hassan district, Karnataka (N=80, P=40, K=40 kg/ha, pH=6.5, Rainfall=800mm, Temperature=25 degrees Celsius, Humidity=60%), the system recommended Coffee (50%), Rice (25%), Jute (23%), Maize (1%), and Pigeonpeas (1%). This accurately reflects the known crop suitability of Hassan district, which is a major coffee-growing region in Karnataka. Testing for coastal Karnataka conditions (high humidity, high rainfall) correctly elevated Rice and Coconut to top recommendations.

E. System Usability and Integration

The updated dashboard (Fig. 9) incorporates a live weather widget displaying real-time conditions for the farmer's detected location (Dakshina Kannada: 28 degrees Celsius,

Overcast, 78% humidity, 5.9 km/h wind), demonstrating seamless API integration within the main dashboard.

The Shopping List module successfully auto-populated with 10 relevant items (Mancozeb 75% WP fungicide, Knapsack Sprayer, Safety Gloves, Soil Test Kit, etc.) for a Tomato crop on 2 acres, with an estimated total of Rs. 16,920, correctly sourcing recommendations from the Disease Detection, Yield Prediction, Irrigation Planning, and Crop Recommendation modules[10].

F. Comparison with Existing Systems

Table III compares AgriVeda AI against representative existing smart agriculture tools across key feature dimensions. AgriVeda AI is the only system offering all six major feature categories in a single platform, demonstrating its comprehensive scope.

TABLE III

FEATURE COMPARISON WITH EXISTING SYSTEMS

Feature	PlantVil.	AgroAI	Krishi Net	FASAL	AgriVeda
Disease Detection	Yes	Yes	No	No	Yes (Dual)
Yield Prediction	No	Yes	Yes	Yes	Yes+SHAP
Crop Recommendation	No	No	Yes	No	Yes
Weather + Advisory	No	Yes	No	Yes	Yes
AI Chatbot / RAG	No	No	No	No	Yes (RAG)
Market Prices	No	No	No	No	Yes (Live)
Govt. Schemes	No	No	No	No	Yes (35)
Explainability	No	No	No	No	SHAP+ GradCAM

VII.CONCLUSION

This paper presented AgriVeda AI, a comprehensive smart agriculture platform integrating ten AI-powered modules within a unified Flask-React web application. The system delivers strong technical performance: 98.2% crop recommendation accuracy, R-squared of 0.98 for yield prediction, 94% CNN disease detection accuracy, and real-time integration with Open-Meteo and AGMARKNET APIs. The inclusion of SHAP explainability for yield predictions and Grad-CAM visualization for disease detection significantly enhances the transparency and interpretability of AI outputs — critical requirements for farmer adoption and trust.

The cross-module integration demonstrated through the AI Shopping List — which automatically consolidates outputs from disease detection, yield prediction, irrigation planning, and crop recommendation — represents a novel contribution to holistic precision agriculture platforms. The modular Flask-React architecture ensures independent maintainability and scalability of each component.

Future work will focus on: (1) multilingual support (Kannada, Hindi, Tamil) for broader rural accessibility; (2) mobile application development with offline-capable model inference; (3) IoT sensor integration for real-time soil NPK and moisture data input; (4) expansion of the disease detection model to additional crop species and regional disease variants; (5) time-series yield forecasting incorporating historical farm data; and (6) federated learning for privacy-preserving model improvement across distributed farm data. The AgriVeda AI platform demonstrates the practical feasibility of integrated AI-powered smart agriculture and represents a significant contribution toward digitally empowered Indian farming.

ACKNOWLEDGMENT

The authors thank the Department of Computer Science and Engineering, Rajeev Institute of Technology, for providing computational resources. We acknowledge the open-source communities behind scikit-learn, TensorFlow, React, Flask, and the FAO, Kaggle, and PlantVillage dataset contributors. We also thank the farming communities of Karnataka for inspiring this research and validating the platform's practical utility.

REFERENCES

- [1] Ministry of Agriculture & Farmers Welfare, Government of India, "Agriculture Census 2015-16," New Delhi, India, 2019.
- [2] S. P. Mohanty, D. P. Hughes, and M. Salathe, "Using deep learning for image-based plant disease detection," *Front. Plant Sci.*, vol. 7, p. 1419, Sep. 2016.
- [3] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Comput. Electron. Agric.*, vol. 145, pp. 311-318, Feb. 2018.
- [4] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations

- from deep networks via gradient-based localization," in Proc. IEEE ICCV, pp. 618-626, 2017.
- [5] J. H. Jeong, J. P. Resop, N. D. Mueller, D. H. Fleisher, K. Yun, E. E. Butler, D. J. Timlin, K.-J. Shim, J. S. Gerber, V. R. Reddy, and D. Kim, "Random forests for global and regional crop yield predictions," PLoS ONE, vol. 11, no. 6, p. e0156571, 2016.
- [6] X. E. Pantazi, D. Moshou, T. Alexandridis, R. L. Whetton, and A. M. Mouazen, "Wheat yield prediction using machine learning and advanced sensing techniques," Comput. Electron. Agric., vol. 121, pp. 57-65, Feb. 2016.
- [7] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. NIPS, vol. 30, pp. 4765-4774, 2017.
- [8] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in Proc. NeurIPS, vol. 33, pp. 9459-9474, 2020.
- [9] A. Kamilaris and F. X. Prenafeta-Boldu, "Deep learning in agriculture: A survey," Comput. Electron. Agric., vol. 147, pp. 70-90, Apr. 2018.
- [10] D. P. Hughes and M. Salathe, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv:1511.08060, 2015.