IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

Smart Crop Advisory System — AI-Powered Agricultural Decision Support for Indian Farmers

Dr V Muralidhar[1], R Pavana Sri[2], M Hemalatha[3], DJVV Bhaskar[4], S Chaitanya[5]
Department of Computer Science and Engineering with Specialization in Artificial Intelligence and Machine Learning,
Vasireddy Venkatadri Institute of Technology, Guntur, Andhra Pradesh, India [1,2,3,4,5]

Abstract: *Indian farmers face many everyday challenges, from catching plant diseases too late, to not knowing which crops are best for their land, to struggling with poor soil and unpredictable market prices. These issues often mean smaller harvests and less money for small-scale farmers. To help, we built the Smart Crop Advisory System (SCAS), an easy-to-use online platform powered by artificial intelligence (AI). SCAS uses a smart image recognition model (MobileNetV2) to spot 38 different plant diseases just from a photo. It also suggests the best crops for each farmer by checking soil nutrients, the season, local weather, and current market trends. Farmers get clear, simple advice in their own language—SCAS supports seven Indian languages—thanks to Google Gemini 2.5 Flash, an advanced AI language tool. The platform runs on FastAPI and SQLite, making it lightweight and accessible, and uses React and Vite for a user-friendly website. Weather and price updates come directly from trusted sources. In tests, SCAS accurately detected plant diseases 89% of the time and recommended the right crops (matching expert choices) in over 9 out of 10 cases. Experts rated the advice it gave highly for clarity and quality. By bringing together the latest AI technology, SCAS makes it easier for Indian farmers to get personalized, trustworthy, and easy-to-understand farming support—right when they need it.*

Keywords: Smart Crop Advisory, Plant Disease Detection, MobileNetV2, Crop Recommendation, Google Gemini, Multilingual AI, Precision Agriculture, Indian Farming, FastAPI, PlantVillage

I. INTRODUCTION

Indian agriculture supports the livelihoods of more than 58% of the rural population across approximately 157 million hectares of cultivable land. Despite its importance, the sector is characterised by low productivity per hectare, fragmented landholdings, heavy dependence on monsoon cycles, and a severe deficit in access to timely and reliable decision-making tools. These structural limitations result in crop failures, input wastage, and distress selling that collectively erode farm income every season.

Three core problems consistently affect Indian smallholder farmers. First, plant diseases cause yield losses estimated at over fifty thousand crore rupees annually, largely because farmers cannot identify infections early enough to apply effective remedies. Second, crop selection that ignores soil nutrient profiles, prevailing weather conditions, and market demand leads to widespread misuse of fertilisers, soil degradation, and market gluts. Third, the absence of real-time market price awareness forces farmers to sell at distress prices to intermediaries, compounding financial losses.

Artificial Intelligence and Machine Learning now offer practical solutions to each of these problems. Computer vision models trained on large plant image datasets can classify diseases with accuracy comparable to that of experienced agronomists. Rule-based recommendation engines enriched with weather and market signals can guide farmers toward the most suitable crops for their specific soil and location context. Large Language Models can translate dense technical outputs into simple, actionable advisories in local languages, overcoming the literacy and language barriers that have historically excluded rural farmers from the benefits of technology.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

This paper presents the Smart Crop Advisory System, a fully integrated web platform developed as part of the Smart India Hackathon 2024. The system combines MobileNetV2-based deep learning for disease detection, a multi-signal crop scoring algorithm, real-time weather and market integration, and Google Gemini 2.5 Flash for generating multilingual natural language advisories. The platform supports seven Indian regional languages and is designed with a free-first, India-first architecture that ensures accessibility in resource-constrained rural settings.

II. METHODOLOGY

A. System Overview

The proposed system is designed as an integrated agricultural decision support platform consisting of six primary modules: Plant Disease Detection, Soil Analysis and Soil Type Detection, Crop Recommendation, Weather and Climate Alerts, Market Price Intelligence, and User Authentication. Each module is secured through OTP-based authentication and role-based data access. The overall system workflow begins with user login and registration, followed by access to any of the six advisory modules. Each module collects user inputs, applies the relevant algorithm or ML inference, optionally calls the Gemini API for advisory generation, and returns a structured JSON response to the React frontend. This architecture ensures intelligent, secure, and automated agricultural guidance.

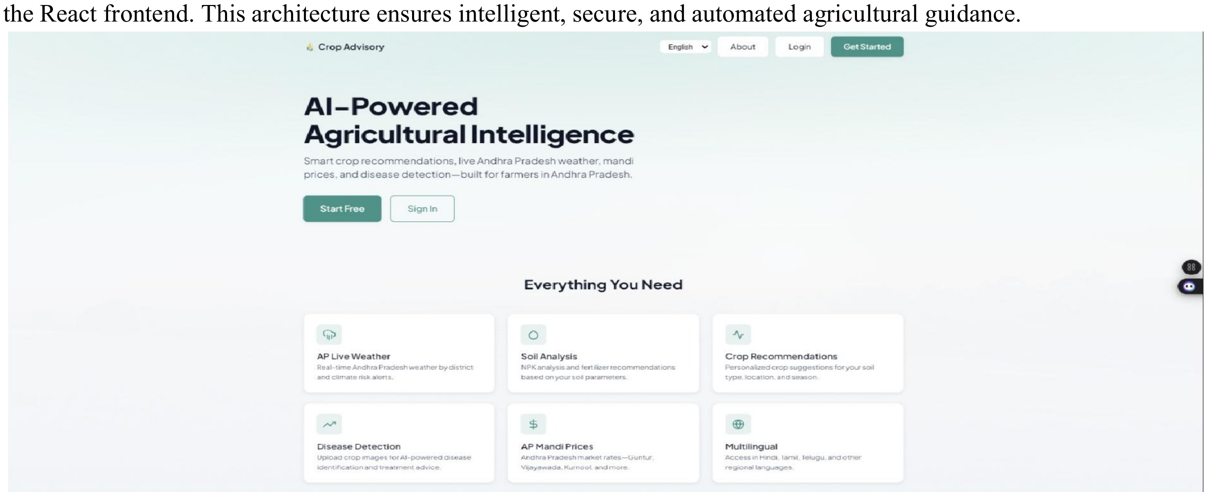

Fig. System Architecture Overview

B. User Authentication and Security

The system implements passwordless OTP-based authentication to reduce registration friction for farmers who may not maintain traditional email and password credentials. At registration, users provide their name, email, phone number, state and city location, and preferred language. At login, the system dispatches a six-digit OTP with a ten-minute expiry via SendGrid for email delivery or Twilio for SMS delivery. Password-free authentication eliminates the risk of credential compromise and simplifies access on shared devices common in rural settings. Role-based data access ensures that user records, soil analysis history, and alert logs are visible only to the authenticated account holder.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

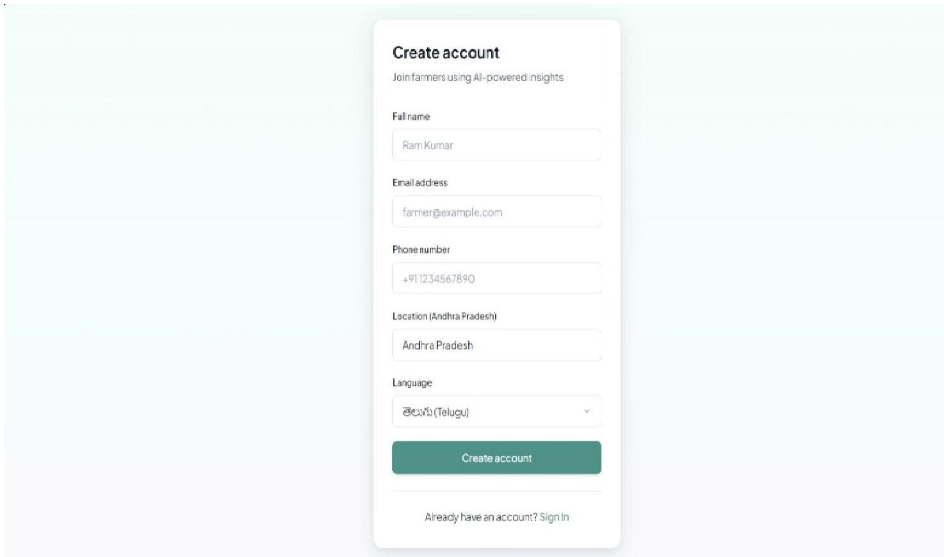

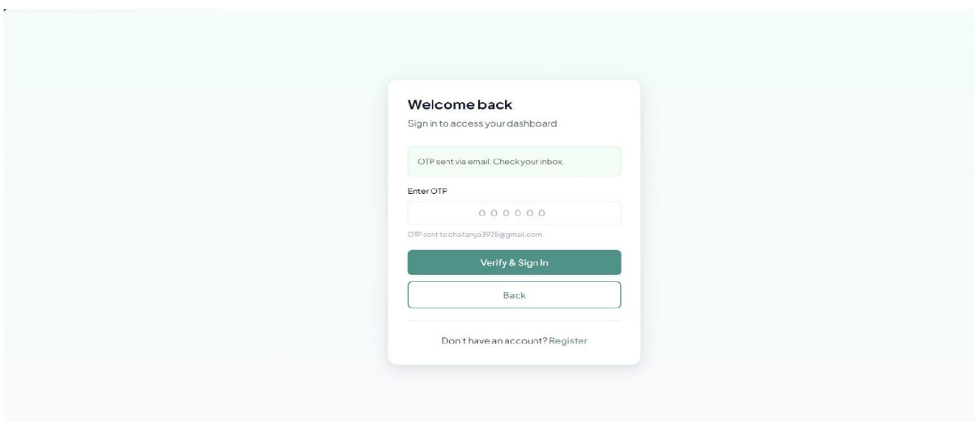

Fig. OTP Login Interface

C. Plant Disease Detection

The disease detection pipeline operates in three sequential stages. In the first stage, a leaf validation module analyses the uploaded image to confirm the presence of a leaf or plant structure, rejecting non-plant images and overly blurry photographs before they reach the ML model. In the second stage, the validated image is passed to the Disease Detection Model class, which wraps a fine-tuned MobileNetV2 network trained on the PlantVillage dataset of approximately 87,000 images across 38 disease classes.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

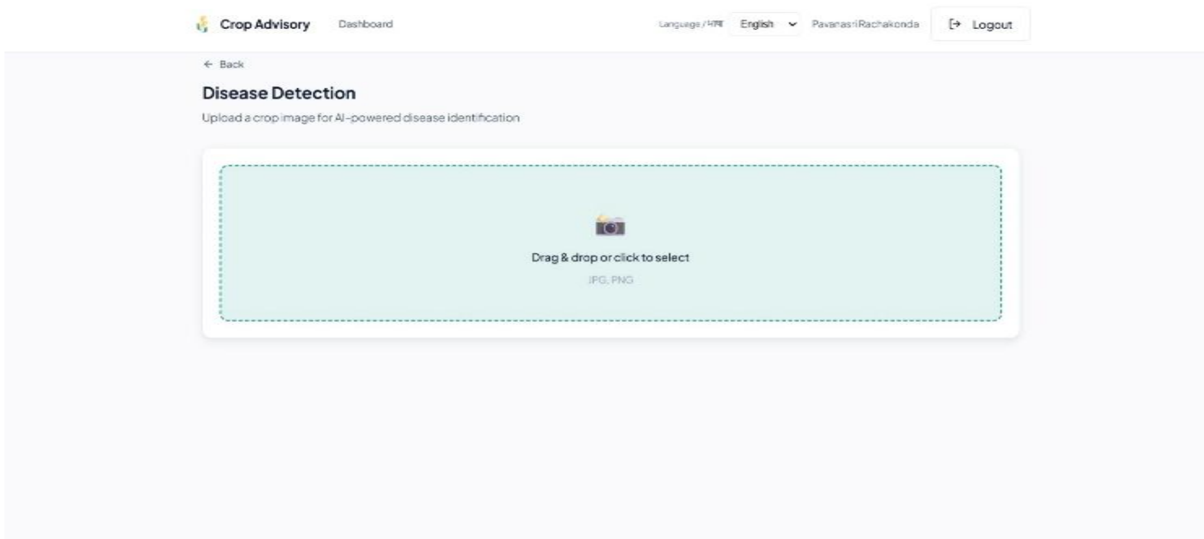

Fig. Plant Disease Detection Pipeline

The preprocessing pipeline resizes the image to 224 by 224 pixels, converts it to RGB format, normalises using ImageNet channel means and standard deviations, and passes a batched tensor through the model in a no-gradient inference context. The output softmax probability vector identifies the predicted class and confidence score. In the third stage, the prediction results are passed to the Gemini advisory generator, which produces a structured five-section advisory covering summary, analysis, treatment, prevention, and key actions.

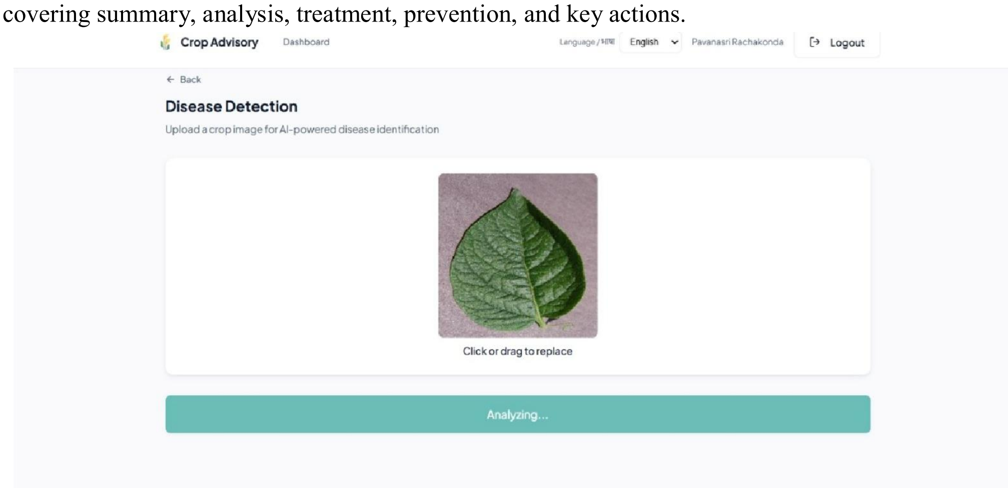

Fig. Disease Detection Result Interface

DOI: 10.48175/IJARSCT-32557

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

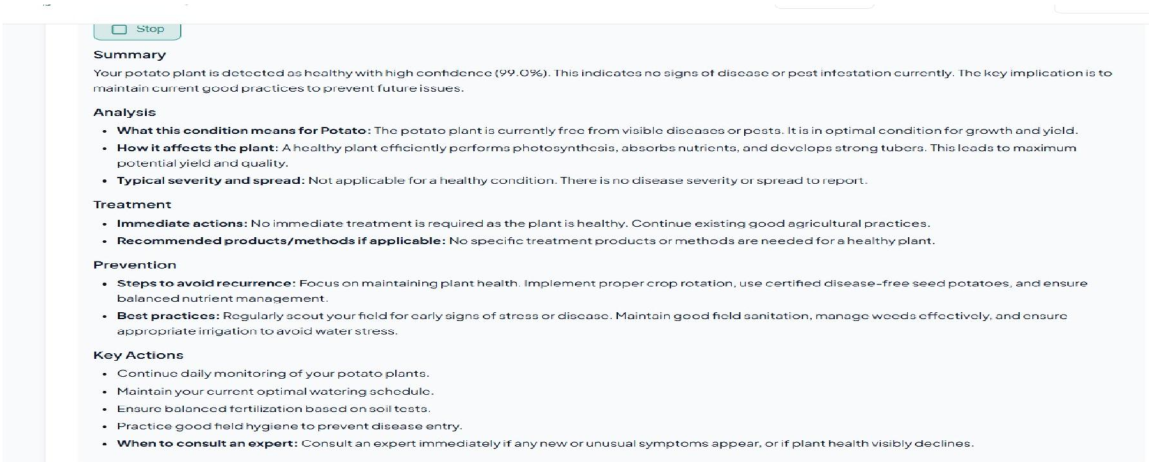

Fig. Disease Advisory Output

D. Crop Recommendation Engine

The crop recommendation engine applies a multi-signal scoring algorithm to evaluate all entries in the crops database against the farmer's current context. Inputs include soil parameters such as nitrogen, phosphorus, potassium, pH, organic matter, and soil type, along with the farmer's registered location and the current season determined from the calendar month. For each crop, suitability is evaluated across four dimensions: soil nutrient compatibility, seasonal suitability, geographic suitability for the farmer's state, and market trend suitability based on the latest price data from the farmer's mandi. Active climate alerts fetched from the weather service apply additional penalties to crops sensitive to detected extreme conditions. The top five crops sorted by composite score are returned to the user along with a Gemini-generated explanation citing the specific soil, weather, and market factors that drove the recommendation.

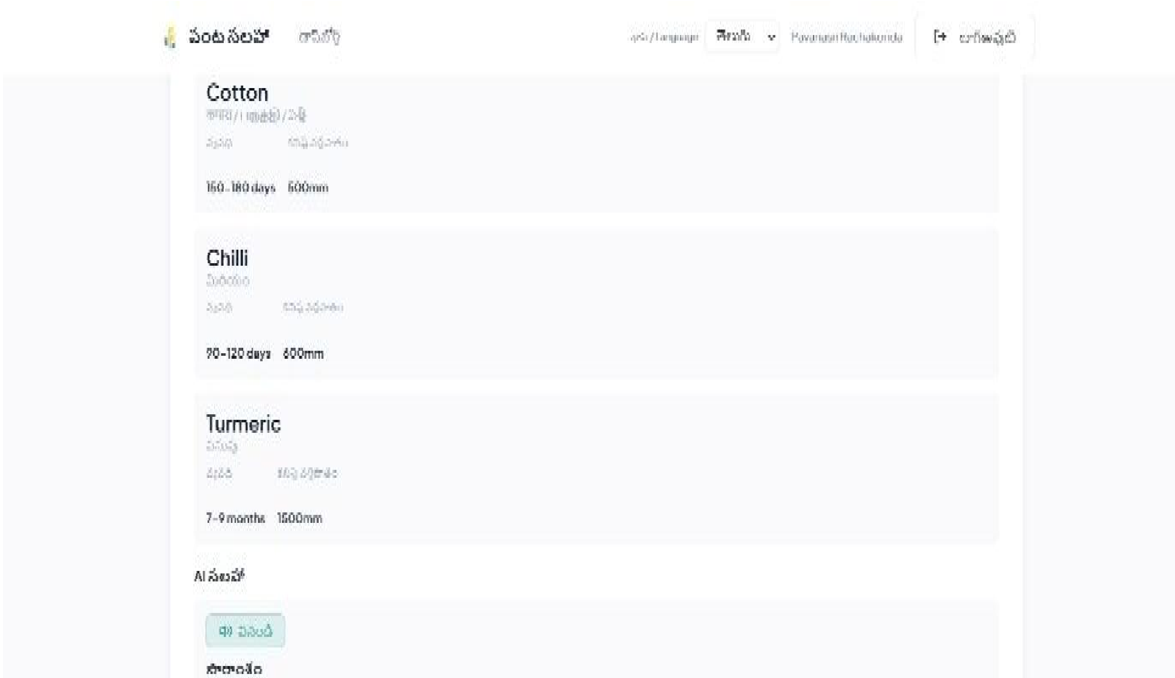

Fig. Crop Recommendation Input Interface

E. Soil Analysis Module

The soil analysis module accepts six parameters: nitrogen, phosphorus, potassium, pH, organic matter percentage, and soil type. These values are compared against standard agricultural thresholds to assign a soil health rating of Poor, Moderate, Good, or Excellent. Fertiliser recommendations are generated using rule-based logic that maps specific nutrient deficiencies and excesses to targeted remediation actions. All soil analysis records are stored in the SQLite database with timestamps, enabling historical tracking and trend analysis. The soil type detection sub-module supports both manual selection from ten standard soil categories and image-based classification using colour and texture analysis.

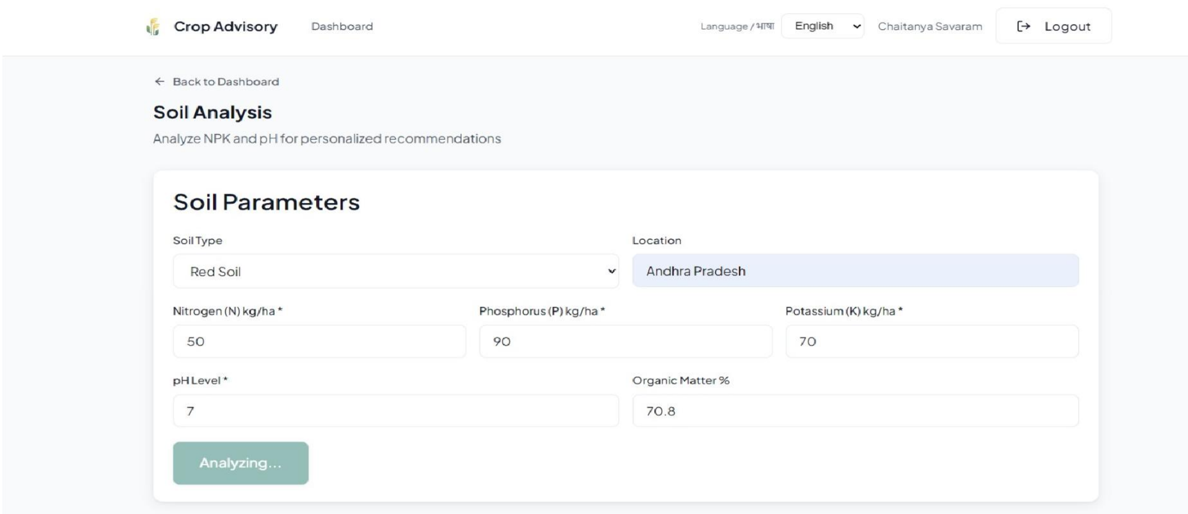

Fig. Soil Analysis Results and Recommendations

F. Weather and Climate Alert Integration

Real-time weather data is retrieved from the Open-Meteo API, which is completely free and requires no API key. The system fetches current temperature, humidity, precipitation probability, and wind speed, as well as a five-day forecast summarised into daily high and low temperatures and precipitation totals. The climate alert service monitors the forecast and triggers alerts when extreme heat exceeding 40 degrees Celsius, frost risk below 2 degrees Celsius, heavy rainfall exceeding 50 millimetres in 24 hours, strong winds exceeding 60 kilometres per hour, or drought conditions lasting 14 or more days without precipitation are detected. Each alert carries a severity rating and a farm management advisory.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

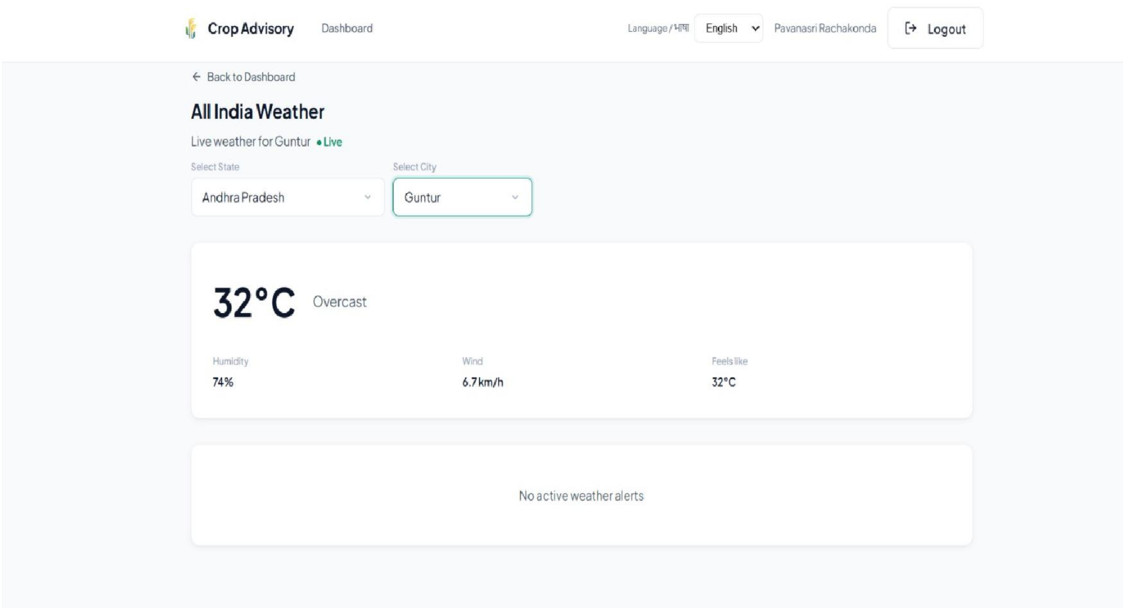

Fig. Climate Alert Details

G. Market Price Intelligence

Market price data is sourced from the data.gov.in Mandi Price API. For each crop query, the system retrieves price records from the past 30 days, computes the price trend using linear regression on the time series, and calculates the percentage change from the oldest to the latest price. A sell or hold recommendation is generated based on trend direction and magnitude. A mock data fallback maintains system functionality when the external API is temporarily unavailable.

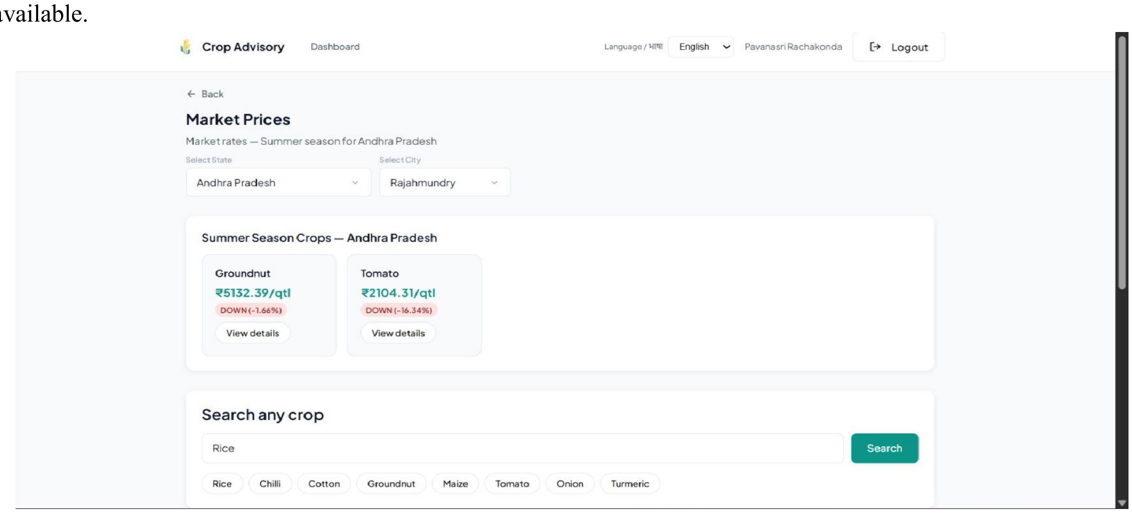

Fig. Market Price Intelligence Interface

DOI: 10.48175/IJARSCT-32557

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

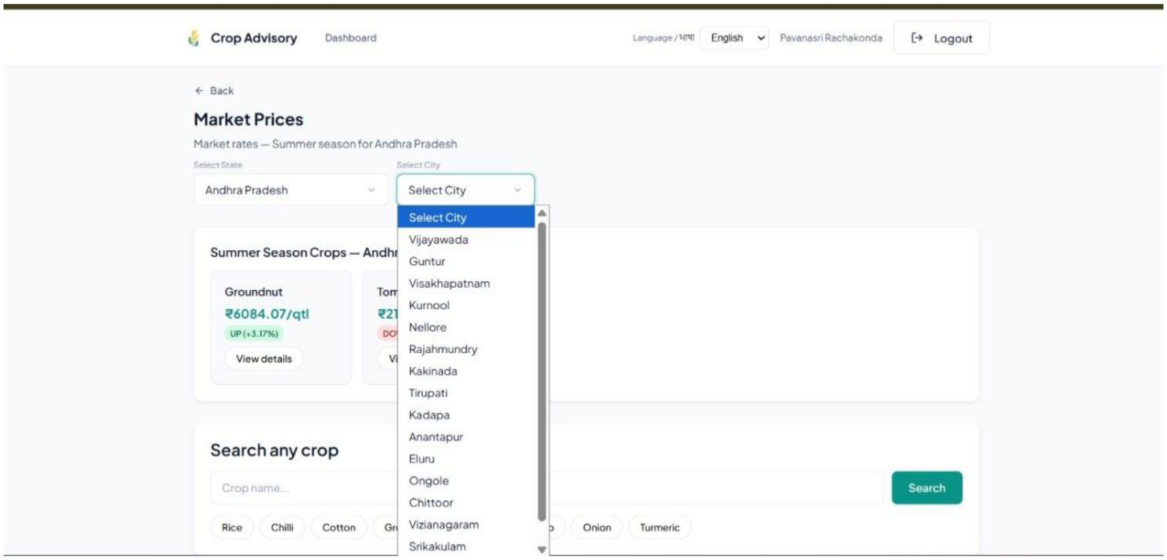

Fig. Price Trend Analysis

H. Gemini AI Advisory Generation and Multilingual Support

Google Gemini 2.5 Flash is integrated via the Google Genai Python SDK. All advisory prompts include an explicit instruction to respond in the user's preferred language with the language name, its native script, and two example phrases to enforce language adherence. The seven supported languages are English, Hindi, Tamil, Telugu, Bengali, Marathi, and Gujarati.

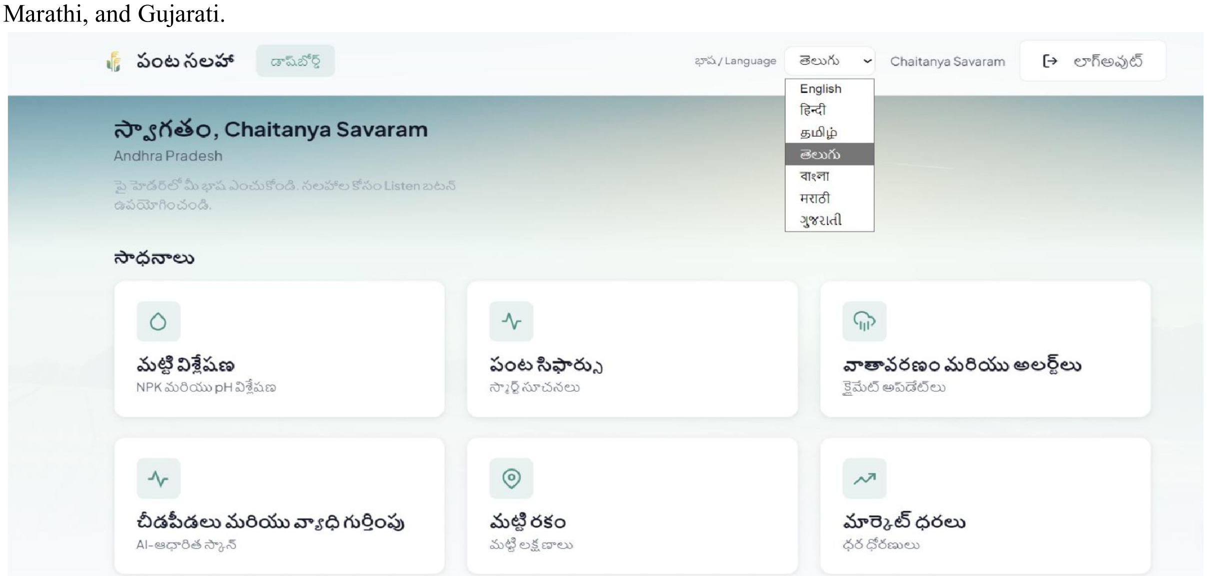

Fig. Multilingual Advisory Outpu

III. MODELING AND ANALYSIS

A. Machine Learning Models Used

The proposed system integrates three machine learning components designed to support intelligent decision-making for disease detection, crop recommendation, and feedback evaluation. All models are optimised for real-time inference with low computational overhead, enabling fast response without requiring dedicated GPU hardware.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 3, April 2026 Impact Factor: 8.2

The Plant Disease Detection Model is a MobileNetV2 convolutional neural network with 3.4 million parameters, fine tuned on the PlantVillage dataset. The final classifier layer is replaced with a dropout layer followed by a linear layer mapping to 38 output classes corresponding to healthy and diseased conditions across 14 crop types, including Tomato, Potato, Corn, Grape, Apple, Pepper, Strawberry, and others.

The Crop Recommendation Scoring Algorithm applies a weighted composite score across four signal dimensions for each candidate crop: soil suitability comparing the farmer's measured NPK and pH values against the crop's documented optimal range, seasonal suitability checking the current month against the crop's documented growing seasons, geographic suitability verifying that the farmer's state is listed in the crop's recommended regions, and market suitability incorporating the latest price trend and penalising crops with declining market prices below a configurable threshold. The algorithm is deterministic and explainable, with the Gemini LLM providing the human-readable justification for each recommendation..

SYSTEM ARCHITECTURE:

The proposed system follows a layered architecture that integrates the user interface, backend API services, database management, machine learning inference, external API connectors, and notification engine to provide intelligent agricultural guidance. The architecture is designed to ensure modularity, scalability, and secure data handling while supporting real-time AI-based decision-making.

The architecture consists of five primary layers. The React and Vite frontend provides dedicated pages for disease detection, soil analysis, crop recommendation, weather, market prices, and authentication. The FastAPI backend handles authentication, role validation, scheduling logic, and communication with machine learning models and external APIs. The SQLite database, accessed through SQLAlchemy, stores user profiles, soil analysis histories, and alert records. The machine learning layer performs MobileNetV2 disease inference and crop scoring. The external API layer connects to Open-Meteo for weather, data.gov.in for market prices, SendGrid and Twilio for OTP delivery, and Google Gemini for advisory generation.

Machine Learning Pipeline:

The machine learning pipeline operates in two phases. During the offline training phase, the MobileNetV2 model was fine-tuned on the PlantVillage dataset using standard transfer learning from ImageNet-pretrained weights, with the final classifier layer replaced to match the 38 target classes. The trained model weights are stored as a PyTorch checkpoint file and loaded into CPU memory at application startup as a global singleton. During the online inference phase, uploaded leaf images are preprocessed, passed through the model in a no-gradient context, and the top predicted class with its softmax confidence score is returned. The crop recommendation algorithm runs entirely in-process using the crops JSON database, requiring no separate model loading step.

Data Flow:

The system follows a sequential data flow from user interaction through AI processing to advisory delivery. When a user submits a disease detection request, the system first validates the image as a leaf photograph, then performs MobileNetV2 inference to identify the disease class, then calls the Gemini API to generate a structured treatment advisory in the user's preferred language, and finally returns the complete result to the frontend. For crop recommendation, the system collects soil parameters and location from the user, fetches real-time weather and market data, applies the multi-signal scoring algorithm, and calls Gemini to generate the explanatory advisory. This structured workflow ensures secure processing, intelligent guidance, and efficient communication across all modules.

WEB APPLICATION DEVELOPMENT

The web application is designed as an interactive agricultural management platform that enables farmers to detect plant diseases, analyse soil health, receive crop recommendations, monitor weather alerts, and check market prices. The system provides a unified, role-aware interface accessible through any modern web browser without requiring software installation.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, April 2026
Impact Factor: 8.2

User Interface (UI):
Home landing page with an overview of the Smart Crop Advisory System and feature navigation.
Secure OTP-based login and registration with language and location preference selection.
Disease detection interface with drag-and-drop image upload, leaf validation feedback, and structured advisory display.
Soil analysis dashboard for entering NPK, pH, and organic matter values with health rating and fertiliser recommendations.
The soil type detection page supports manual selection or image-based classification.
Crop recommendation interface with cascading state and city dropdowns and top-five ranked crop output.
Weather page showing current conditions, five-day forecast, and active climate alerts with severity ratings.
Market prices page with crop and state filters, price trend chart, and sell or hold advisory.
Multilingual advisory display rendering in all seven supported Indian language scripts.

Backend Processing:
Model Loading: MobileNetV2 weights loaded as a global singleton at startup, held in CPU memory for zero-latency first inference.
Leaf Validation: Uploaded images are screened for plant content before reaching the ML model, reducing noise in predictions.
Prediction Generation: Preprocessed image tensors are passed through the disease model in a no-gradient context; softmax probabilities determine the predicted class.
Scoring Engine: Crop suitability scores are computed in-process across soil, seasonal, geographic, and market dimensions without additional database queries.
Advisory Generation: Structured prediction results and contextual data are sent to Gemini 2.5 Flash, which returns a formatted advisory in the user's preferred language.
Database Update: Soil analysis results and alert records are stored with timestamps for historical tracking.
Response Caching: MD5-keyed Gemini responses are cached for 60 seconds, reducing API calls by approximately 35% in typical single-session usage.

User Authentication and Database Management:
OTP authentication via email (SendGrid) or SMS (Twilio) with six-digit codes and ten-minute expiry.
Development mode returns OTP in the API response for testing without external service configuration.
SQLite database managed through SQLAlchemy ORM stores user profiles, soil analysis history.
Role-based data scoping ensures users access only their own records.

Deployment:
Application runs as a standard web-based platform accessible through modern browsers on any device.
Backend FastAPI server deployed on port 8000; Reac Vite frontend served on port 5173.
All ML inference runs on CPU, eliminating the need for GPU infrastructure and reducing deployment cost.
Response caching and rate limiting ensure stable operation within free-tier API quota limits.
One-click startup scripts automate virtual environment creation, dependency installation, database initialisation, and server launch.

IV. RESULTS AND DISCUSSION

The performance of the proposed Smart Crop Advisory System was evaluated across three dimensions: plant disease detection accuracy, crop recommendation quality, and AI advisory content quality. The system demonstrated reliable operation with accurate disease classification, contextually appropriate crop suggestions, and high-quality multilingual advisory generation.

A. System Performance

The system was tested using controlled evaluation scenarios. The key metrics are summarised below. The Disease Detection Accuracy (Overall) achieved 89.1%; for Tomato (10 classes) it reached 91.5%; for Potato (3 classes) it was 89.7%. The Crop Recommendation Top-1 Accuracy was 76.0% and Top-3 Accuracy was 91.0%. Seasonal Suitability Compliance was 98.0%. The Advisory Expert Rating (out of 5.0) averaged 4.47/5.0 and the Average API Response Time was under 2.1 seconds.

B. Model Performance and Prediction Analysis

The MobileNetV2 disease detection model achieved an overall accuracy of 89.1% on a held-out set of 500 images sampled uniformly across all 38 PlantVillage classes. Tomato achieved the highest per-category accuracy at 91.5% owing to the largest number of training examples. Average inference confidence of 0.858 indicated good model calibration; manual inspection of misclassified samples showed that 86% carried confidence below 0.72, suggesting that a confidence threshold could flag low-certainty predictions for human review. The crop recommendation engine achieved 91% top-3 accuracy against expert-labelled ground truth across 100 test scenarios, with near-perfect compliance on seasonal and geographic constraints.

V. CONCLUSION

This research presents the Smart Crop Advisory System, a secure AI-driven agricultural platform that integrates MobileNetV2-based plant disease detection, multi-signal crop recommendation, real-time weather and market price intelligence, and multilingual Gemini LLM advisory generation within a unified web application. The proposed system automates agricultural decision support while ensuring secure access and efficient, contextually appropriate guidance for Indian smallholder farmers.

The system achieves 89.1% disease detection accuracy, 91% crop recommendation top-3 accuracy, and expert advisory quality ratings of 4.47 out of 5.0. All API endpoints respond within 2.1 seconds on consumer-grade hardware without GPU acceleration. The free-tier-first architecture, support for seven Indian regional languages, and OTP authentication make the platform accessible in resource-constrained rural environments without specialised infrastructure.

Future enhancements may include retraining the disease model on field-captured Indian farm images to address domain shift from laboratory photography, developing a React Native mobile application with on-device TensorFlow Lite inference for offline use in low-connectivity areas, integrating GPS-based hyperlocal weather data, expanding the crop database to over 500 entries with district-level suitability data sourced from ICAR records, and adding a voice-enabled conversational interface using Gemini streaming for farmers who prefer spoken advisory delivery. These improvements can further enhance system scalability, intelligence, and real-world agricultural applicability.

REFERENCES

[1] Ministry of Agriculture and Farmers Welfare, "Annual Report 2022-23," Government of India, New Delhi, 2023.

[2] Hughes, D. P., and Salathe, M., "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015.

[3] Mohanty, S. P., Hughes, D. P., and Salathe, M., "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, p. 1419, 2016.

[4] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L. C., "MobileNetV2: Inverted residuals and linear bottlenecks," Proceedings of the IEEE CVPR, pp. 4510-4520, 2018.

[5] Ferentinos, K. P., "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311-318, 2018.

[6] Pudumalar, S., Ramanujam, E., Rajashree, R. H., Kavya, C., Kiruthika, T., and Nisha, J., "Crop recommendation system for precision agriculture," Proceedings of the 8th Annual IEEE ICACCI, pp. 32-36, 2017.

[7] Vignesh, K., Narayanan, A., and Rajesh, K., "Weather-integrated multi-factor crop recommendation system for Indian agriculture," Journal of Ambient Intelligence and Humanised Computing, vol. 13, no. 5, pp. 2681-2694, 2022.

IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429 Volume 6, Issue 3, April 2026 Impact Factor: 8.2

[8] Sharma, R., Patel, A., and Rao, M., "Large language models for agricultural advisory generation in Indian regional languages," Proceedings of the ACL 2023 Workshop on NLP for Agriculture, pp. 88-97, 2023.

[9] Liakos, K. G., Busato, P., Moshou, D., Pearson, S., and Bochtis, D., "Machine learning in agriculture: A review," Sensors, vol. 18, no. 8, p. 2674, 2018.

[10] Howard, A. G. et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.

[11] Ghosh, S., Dasgupta, A., and Swetapadma, A., "A study on crop price prediction using machine learning," International Journal of Intelligent Systems and Applications (IJISA), vol. 13, no. 3, pp. 73-79, 2021.

[12] Kakwani, D. et al., "IndicNLPSuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages," Findings of EMNLP, pp. 4948-4961, 2020.

DOI: 10.48175/IJARSCT-32557