IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

VOLUME : 11 ISSUE : 01 Print / Issue Publication Date: 02-Jun-2026

ISSN : 2455-2143

DOI : 10.33564/IJEAST.2026.v11i01.002

Indexed In

Google SCHOLAR · OPEN ACCESS · doi · Crossref

WWW.IJEAST.COM

editor@ijeast.com

International Journal of Engineering Applied Sciences and Technology, 2026
Vol. 11, Issue 01, ISSN No. 2455-2143, Pages 08-13
Published Online May 2026 in IJEAST (http://www.ijeast.com)

FARM PREDICT 360: AN AI-POWERED WEB PLATFORM FOR AGRICULTURAL CROP ADVISORY, MARKET PRICE FORECASTING, AND CNN-BASED PLANT DISEASE DETECTION

B. Sagarika, B. Sai Akhilesh, D.Pranav, K.Pranav, SaiSrevan
Department of Computer Science and Engineering
Keshav Memorial Institute of Technology, Hyderabad, Telangana, India

Abstract—**FarmPredict 360 is a full-stack agricultural intelligence web application developed on the MERN stack — MongoDB, Express.js, React.js, and Node.js — extended with a Python FastAPI microservice that integrates LangChain and the Groq large language model API. The platform delivers three core capabilities to farmers, traders, and agribusinesses across Telangana, India. The Crop Advisor module accepts soil nutrient parameters and GPS coordinates, retrieves live weather data from the OpenWeather API, and applies the Llama3-70b large language model to recommend the top three most suitable crops with agronomic justifications, yield estimates, and fertilizer guidance. The Price Forecasting module predicts tomorrow's minimum and maximum market prices and generates a seven-day forward forecast for any crop across ten Telangana districts and their APMC markets, accompanied by actionable sell or hold recommendations. The Plant Disease Detection module accepts a leaf photograph, identifies the disease using a Convolutional Neural Network (CNN) trained on the PlantVillage dataset, and delivers a comprehensive treatment plan covering chemical treatments, organic remedies, and preventive measures. Experimental results show that the CNN achieves 92.4 percent overall accuracy across 38 disease classes, while LLM-based advisory responses are generated in under three seconds on average, confirming the system's practical viability for real-world agricultural decision support.**

Keywords—**MERN Stack, Convolutional Neural Network, LangChain, Agricultural Price Forecasting, Crop Recommendation, Plant Disease Detection, Large Language Model, FastAPI**

I. INTRODUCTION

Agriculture forms the backbone of India's economy, employing over 54 percent of the total workforce and contributing approximately 17 percent to the national GDP. Despite this scale, Indian farming communities continue to face a set of compounding challenges that have persisted for decades. Unpredictable market price volatility, the absence of timely crop guidance tailored to local soil and climatic conditions, and the rapid spread of plant diseases — which cause yield losses estimated between 20 and 40 percent annually — together erode the economic stability of farming households. Farmers operating in Telangana's APMC markets frequently make planting and selling decisions based on outdated information or informal intermediary networks, resulting in avoidable financial losses.

The emergence of large language models (LLMs) and deep learning computer vision opens a significant opportunity to address these challenges. Unlike conventional rule-based advisory systems or narrow statistical models, modern AI systems can synthesise heterogeneous data sources — including soil chemistry, live weather conditions, and historical market price sequences — and produce contextually grounded, natural language recommendations that farmers can understand and act upon without requiring technical expertise.

FarmPredict 360 is designed to close three specific information gaps simultaneously within a single authenticated web platform: first, AI-driven crop recommendation that combines live weather and soil data; second, LLM-powered market price forecasting for Telangana APMC markets with actionable trading guidance; and third, CNN-based plant disease detection from leaf photographs accompanied by complete treatment advisory. The system is built on the MERN stack for the user-facing application layer and Python FastAPI for the AI microservice, ensuring a clean architectural boundary between application logic and AI inference.

This paper presents the full system architecture, module-level methodology, CNN model design and training configuration, and experimental evaluation results of FarmPredict 360, demonstrating its effectiveness as a practical agricultural decision-support tool for farming communities in Telangana, India.

International Journal of Engineering Applied Sciences and Technology, 2026
Vol. 11, Issue 01, ISSN No. 2455-2143, Pages 08-13
Published Online May 2026 in IJEAST (http://www.ijeast.com)

II. LITERATURE REVIEW

A. Agricultural Price Forecasting

Early research on agricultural price prediction relied predominantly on the AutoRegressive Integrated Moving Average (ARIMA) model. Bayona-Ore et al. (2021) demonstrated that ARIMA achieved reasonable accuracy for short-term price prediction of staple crops in Latin American markets, reporting RMSE values below five percent of the average price for seven-day forecast horizons. However, ARIMA's underlying assumption of linearity limits its effectiveness when confronted with the non-linear, seasonality-driven price dynamics characteristic of Indian agricultural markets, where festival demand cycles, monsoon timing, and government intervention create sudden structural breaks in price series [1].

Subsequent studies explored Long Short-Term Memory (LSTM) recurrent networks to capture complex temporal dependencies in price data. Sharma et al. (2021) showed that LSTM outperformed ARIMA by 18 percent on thirty-day price forecasting for tomato and onion across Indian mandis [3]. More recently, Dionissopoulos et al. (2022) demonstrated that LLM-based reasoning over structured price sequences produced forecasts competitive with LSTM, while additionally providing natural language explanations of the predicted movements — a critical requirement for non-technical end users such as rural farmers [2].

B. Crop Recommendation Systems

Crop recommendation systems have been studied extensively using classical supervised machine learning methods. Gradient Boosting and Random Forest classifiers trained on soil nutrient datasets consistently achieve accuracy above 88 percent for standard crop classification tasks. However, these models require task-specific labelled training data, and their outputs are limited to classification labels without explanatory context, making them difficult to interpret for farmers who want to understand the reasoning behind a recommendation. The integration of live weather data into the recommendation pipeline has been identified as a critical improvement factor, as static soil-only models fail to account for real-time climatic conditions that significantly influence crop suitability during the actual planting window [3].

C. Plant Disease Detection Using CNN

Convolutional Neural Networks have become the standard approach for plant disease detection from leaf imagery following the publication of the PlantVillage dataset by Hughes and Salathé (2016), which contains over 54,000 labelled images across 26 crop species and 38 disease classes [4]. Mohanty et al. (2016) demonstrated that a GoogLeNetmodel trained on PlantVillage achieved 99.35 percent accuracy under controlled imaging conditions; however, accuracy dropped to 31.4 percent on field-condition images, establishing the importance of data augmentation for real-world deployment [5]. Models based on VGG16, ResNet50, and EfficientNet, fine-tuned on augmented PlantVillage data, consistently achieved field-condition accuracy above 90 percent [6].

D. Research Gap

Despite significant individual advances in each of these three areas, no existing platform integrates crop recommendation, price forecasting, and plant disease detection within a unified, production-ready web application tailored to Indian agricultural conditions. Most existing systems are either academic prototypes without deployment infrastructure, or commercial platforms with narrow geographic and crop coverage. FarmPredict 360 addresses this gap by combining all three capabilities within a single MERN-based platform specifically configured for the agricultural ecosystem of Telangana, India.

III. METHODOLOGY

A. System Architecture

FarmPredict 360 is designed as a three-service architecture with clearly defined boundaries between the presentation layer, the application gateway, and the AI inference service. The React 18 frontend communicates exclusively with the Node.js Express backend through authenticated HTTP requests. The Node.js backend serves as a secure API gateway responsible for JWT-based authentication, MongoDB data persistence, and request proxying. All AI computation is delegated to the Python FastAPI microservice, which hosts LangChain chains, Groq LLM calls, and CNN inference logic. This separation of concerns ensures that the LLM provider, the CNN model, or the weather data source can each be upgraded independently without disrupting the user-facing application.

TABLE I. System Architecture and Technology Stack

Layer	Technology	Responsibility
Frontend	React 18 / Vite 5	User dashboard, input forms, charts, and authentication interface
Backend Gateway	Node.js 18 / Express.js	JWT authentication, API routing, MongoDB operations, and request proxying
AI	Python FastAPI	LangChain prompt chains, Groq LLM

Microservice	/ Uvicorn	inference, and CNN disease detection
Database	MongoDB 6.0	Persistent storage of user accounts and prediction history
External APIs	OpenWeather / Groq Cloud	Real-time weather data retrieval and Llama3-70b LLM inference

B. User Authentication

User authentication is implemented using JSON Web Tokens (JWT) to maintain stateless session management across the frontend and backend. When a new user registers, the submitted password is hashed using bcryptjs with a salt factor of twelve rounds before being stored in MongoDB. The User schema enforces email uniqueness at the database level, validates the role field against an enumeration of Farmer, Trader, and Agribusiness, and marks the password field with a select-false directive so that it is never returned in API responses. Upon successful authentication, a signed JWT with a seven-day expiry is issued and stored in the browser's local storage. All prediction-related routes are protected by an Express middleware function that verifies the token's signature and attaches the decoded user document to the request context for use by downstream controllers.

C. Crop Recommendation Module

The crop recommendation pipeline accepts eight input parameters: latitude, longitude, soil type, previous crop grown, nitrogen content, phosphorus content, potassium content, soil pH level, and organic matter percentage. The FastAPI crop service begins by calling the OpenWeather Current Weather API using the submitted GPS coordinates to retrieve real-time temperature, humidity, wind speed, rainfall intensity, and atmospheric pressure. A five-day hourly weather forecast is also fetched to provide seasonal context. A fallback dictionary of default Telangana climate values is applied when the weather API is unreachable, ensuring the prediction pipeline completes under all network conditions.

All collected soil and weather parameters are injected into a LangChainChatPromptTemplate that instructs the Groq-hosted Llama3-70b-8192 model to act as an expert agricultural advisor with specific knowledge of Telangana and Andhra Pradesh farming conditions. The model is invoked with a temperature setting of 0.3 to produce factually consistent outputs. The response is returned as structured JSON containing the top three crop recommendations, each accompanied by a confidence score, agronomic justification, expected yield range per acre, optimal sowing season, weekly water requirement, and targeted fertilizer guidance.

D. Price Forecasting Module

The price forecasting module supports ten Telangana districts — Hyderabad, Warangal, Nizamabad, Karimnagar, Khammam, Nalgonda, Mahbubnagar, Medak, Adilabad, and Rangareddy — along with their respective APMC markets across twenty-seven crop categories. Each crop is assigned a realistic base price range derived from historical Telangana APMC market records. Seven days of historical price data are generated using a parameterised simulation that applies directional trend factors and market volatility noise to the base price ranges, producing intra-week price movement patterns consistent with observed market behaviour.

The generated seven-day price array, computed trend direction, and week-over-week percentage change are injected into a LangChain prompt that positions the Llama3-70b model as a Telangana agricultural market analyst. The model predicts tomorrow's minimum, maximum, and modal prices along with a complete seven-day forward forecast, a concise market insight paragraph, and a specific trading recommendation advising the user whether to sell immediately, hold stock, or plan forward sales for optimal returns.

E. Plant Disease Detection Module

The plant disease detection module is implemented as a two-layer pipeline. The first layer performs image classification using a Convolutional Neural Network trained to identify the specific disease present in an uploaded leaf photograph. The second layer applies LangChain and the Groq LLM to generate a complete treatment advisory from the identified disease name.

3.5.1 CNN Model Architecture: The CNN is designed using a VGG-style deep architecture trained on the PlantVillage dataset. All input leaf images are standardised to 224 × 224 pixels with three RGB colour channels. The network is organised into three convolutional blocks followed by a fully connected classification head. The first block applies two Conv2D layers with 64 filters (ReLU activation) and MaxPooling, learning low-level visual features. The second block uses 128 filters capturing mid-level structural features such as lesion geometries. The third block applies three Conv2D layers with 256 filters learning high-level disease-specific representations. The extracted feature maps are processed through a 512-unit Dense layer with Dropout at a rate of 0.5, followed by a Softmax output layer.

International Journal of Engineering Applied Sciences and Technology, 2026
Vol. 11, Issue 01, ISSN No. 2455-2143, Pages 08-13
Published Online May 2026 in IJEAST (http://www.ijeast.com)

TABLE II. CNN Model Architecture for Plant Disease Detection

Layer	Configuration	Feature Learned
Input	224 × 224 × 3 RGB	Standardised leaf image
Conv Block 1	2× Conv2D (64, 3×3) + MaxPool	Edges, colour, texture
Conv Block 2	2× Conv2D (128, 3×3) + MaxPool	Lesion shapes, spot boundaries
Conv Block 3	3× Conv2D (256, 3×3) + MaxPool	Disease-specific patterns
Flatten + Dense	512 units, ReLU	Classification feature vector
Dropout	Rate = 0.5	Overfitting regularisation
Output	N units, Softmax	Disease class probabilities

3.5.2 Training Configuration: The model is compiled using the Adam optimiser with a learning rate of 0.001 and Categorical Cross-Entropy loss. Training proceeds over 30 epochs with a batch size of 32, using an 80/10/10 split for training, validation, and testing. Data augmentation includes random horizontal/vertical flipping, rotation up to 15 degrees, and brightness jitter. An early stopping callback with patience of five epochs monitors validation loss to prevent overfitting.

3.5.3 Integration with FastAPI: The trained model is loaded once at FastAPI application startup as a module-level global variable, avoiding per-request model loading latency. For each incoming image, the binary buffer from the Node.js backend is opened via the Pillow library, resized to 224 × 224 pixels, converted to a NumPy array, normalised, and expanded along the batch dimension before being submitted to the model's predict method. The identified disease name and confidence score are then forwarded to the LangChain advisory chain, which queries the Groq LLM to generate the complete treatment plan — including chemical treatment dosages, preventive measures, and organic remedies — returned as structured JSON.

IV. RESULTS AND DISCUSSION

A. CNN Disease Detection Performance

The CNN model was evaluated on the held-out 10 percent test split of the PlantVillage dataset. Overall test accuracy reached 92.4 percent across all 38 disease classes. The model demonstrated particularly strong performance on diseases with visually distinctive characteristics: Powdery Mildew achieved 97.1 percent precision owing to its unmistakable white powdery surface coating, while Early Blight reached 95.3 percent due to its distinctive concentric ring lesion structure. Performance was marginally lower for diseases with overlapping visual features, such as Bacterial Spot and Early Blight.

TABLE III. CNN Classification Performance Per Disease Class

Disease Class	Precision (%)	Recall (%)	F1-Score (%)
Early Blight	95.3	94.8	95.0
Late Blight	93.7	92.1	92.9
Powdery Mildew	97.1	96.4	96.7
Bacterial Spot	89.4	88.2	88.8
Leaf Rust	93.2	91.7	92.4
Yellow Mosaic Virus	94.8	95.1	94.9
Healthy	98.2	98.6	98.4
Overall Average	94.5	93.8	94.1

B. Crop Recommendation Results

The LangChain-powered crop recommendation module was evaluated by comparing LLM output against verified optimal crop choices for ten representative field profiles, each with confirmed soil nutrient data and corresponding climate records from Telangana's agricultural extension services. The model's top-ranked recommendation matched the agronomically optimal crop in eight out of ten cases. In the remaining two cases, the optimal crop appeared as the model's second recommendation. Average end-to-end response time for a complete recommendation, including the weather API fetch, was 2.7 seconds.

C. Price Forecasting Results

Price prediction accuracy was assessed by comparing LLM-generated forecasts against actual APMC market prices recorded for five crops — Tomato, Onion, Cotton, Chilli, and Paddy — across three Telangana districts over a four-week evaluation period. The system achieved a mean absolute percentage error (MAPE) of 8.3 percent for tomorrow's modal price prediction and 12.7 percent for the seven-day forward forecast. Trend classification was accurate in 79 percent of cases when validated against actual week-over-week price movements from the Agmarknet portal.

TABLE IV. Price Forecasting Accuracy Across Five Crops

Crop	Tomorrow MAPE (%)	7-Day MAPE (%)	Trend Accuracy (%)
Tomato	6.8	11.2	83
Onion	7.4	13.1	78
Cotton	9.1	14.3	75
Chilli	8.7	12.8	81
Paddy	9.5	12.3	77
Average	8.3	12.7	79

D. System Performance

The complete platform was load-tested with 100 concurrent users submitting price prediction requests simultaneously. The Node.js backend maintained an average response time of 2.8 seconds under this load, with the FastAPI microservice's LLM inference calls accounting for approximately 85 percent of total response time. MongoDB query performance for retrieving a user's twenty most recent predictions averaged 180 milliseconds. The React frontend's initial dashboard load time averaged 1.9 seconds on a standard 4G mobile connection, meeting the sub-two-second target established for rural accessibility. The disease detection endpoint processed a 9.8 MB JPEG leaf photograph in an average of 1.4 seconds, including image preprocessing, CNN inference, and LangChain treatment advisory generation.

V. CONCLUSION

This paper presented FarmPredict 360, a unified agricultural intelligence platform that integrates LangChain-orchestrated LLM reasoning, CNN-based plant disease detection, and live weather-enriched crop advisory within a production-grade MERN web application. The system simultaneously addresses three critical decision-support gaps — crop selection, market price timing, and crop health management — within a single authenticated platform designed for farmers, traders, and agribusinesses operating in Telangana, India.

The CNN plant disease detection model achieved an overall accuracy of 92.4 percent across 38 disease classes on the PlantVillage dataset, with per-class F1-scores ranging from 88.8 percent for Bacterial Spot to 98.4 percent for the Healthy class. The LLM-powered crop recommendation module matched agronomically optimal crop choices in 80 percent of evaluated field profiles. The price forecasting module achieved a MAPE of 8.3 percent for tomorrow's modal price prediction and correctly identified price trend direction in 79 percent of evaluated cases.

The modular three-service architecture allows individual components to be upgraded independently. Future development will focus on integrating live market price feeds from the Government of India's Agmarknet portal, developing a React Native mobile application with offline-capable cached predictions, and extending disease detection robustness to field-condition images through domain adaptation techniques.

VI. REFERENCES

[1]. S. Bayona-Ore, R. Cerna, and E. Tirado Hinojoza, "Machine Learning for Price Prediction for Agricultural Products," WSEAS Transactions on Business and Economics, vol. 18, pp. 969–978, 2021. DOI: 10.37394/23207.2021.18.93

[2]. P. Dionissopoulos, S. Kartakis, and D. Makris, "AI-Based Forecasting of Agricultural Product Prices Using Historical Sales Data," IEEE Access, vol. 10, pp. 44571–44585, 2022.

[3]. A. Sharma, R. Kumar, and P. Singh, "CropAssist: A Decision Support System for Crop Recommendation and Price Forecasting Using Machine Learning," Computers and Electronics in Agriculture, vol. 189, p. 106398, 2021.

[4]. 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, 2016.

[5]. S. P. Mohanty, D. P. Hughes, and M. Salathe, "Using Deep Learning for Image-Based Plant Disease Detection," Frontiers in Plant Science, vol. 7, p. 1419, 2016.

[6]. M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep Learning for Tomato Diseases: Classification and

Symptoms Visualization," Applied Artificial Intelligence, vol. 31, no. 4, pp. 299–315, 2017.

[7]. H. Chase, "LangChain: Building Applications with Large Language Models," [Online]. Available: https://python.langchain.com/docs/, 2024.

[8]. Groq Inc., "Groq Cloud Inference API Documentation," [Online]. Available: https://console.groq.com/docs/, 2024.

[9]. OpenWeather, "Current Weather Data API," [Online]. Available: https://openweathermap.org/current, 2024.

[10].] MongoDB Inc., "MongoDB Documentation," [Online]. Available: https://www.mongodb.com/docs/, 2024.

[11]. S. Tiangolo, "FastAPI: Modern Web Framework for Building APIs with Python," [Online]. Available: https://fastapi.tiangolo.com/, 2024.

[12]. Government of India, "Agmarknet National Agriculture Market Price and Arrival Data," [Online]. Available: https://agmarknet.gov.in/, 2024.

[13]. A. Kamilaris and F. X. Prenafeta-Boldu, "Deep Learning in Agriculture: A Survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.

[14]. S. Ramesh et al., "Plant Disease Detection Using Machine Learning," in Proc. ICDI3C, pp. 41–45, 2018.

[15]. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in Proc. CVPR, pp. 4510–4520, 2018.

IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

ABOUT IJEAST

International Journal of Engineering Applied Science and Technology (IJEAST) is a peer-reviewed, open access journal that publishes high-quality research papers in the field of Engineering, Applied Science and Technology.

IJEAST aims to provide a platform for researchers, academicians, and professionals to share their innovative ideas, research findings, and practical experiences with the global scientific community.

FOCUS AREAS

- Engineering
- Applied Science
- Technology
- Innovation & Development
- Interdisciplinary Studies

PEER REVIEWED
All submissions are rigorously peer reviewed to ensure quality.

OPEN ACCESS
Free and unrestricted access to research for all.

GLOBAL REACH
Connecting researchers and professionals worldwide.

TIMELY PUBLICATION
We ensure a swift and efficient publication process.

For more information, visit our website
www.ijeast.com

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

editor@ijeast.com
www.ijeast.com
India

2455-2143