

Full Length Article

Intelligent Farmer Assistant And Crop Lifecycle Management Platform With Multilingual Support

Dr. Narsappa Reddy¹, B. Lokesh², M. Rishitha³, R. Harsha Vardhan⁴, R. Mamatha⁵

¹Assistant Professor; Department Of Electronics & Computer Engineering J.B Institute Of Engineering & Technology Hyderabad India

^{2,3,4,5}B.Tech Students; Department Of Electronics & Computer Engineering J.B Institute Of Engineering & Technology Hyderabad India

Accepted 26-04-2026

Author(s) Retains the Copyrights of This Article

Abstract

Agriculture supports the livelihoods of more than 140 million households in India, yet a significant disconnect persists between the information farmers require and what is accessible to them. Existing digital advisory tools are often fragmented, linguistically restrictive, and dependent on high levels of technical proficiency, limiting their usability among smallholder and marginal farmers. These barriers are particularly evident in rural regions, where language diversity, inconsistent internet connectivity, and limited digital literacy constrain adoption. This study introduces the Intelligent Farmer Assistant and Crop Lifecycle Management Platform, a comprehensive mobile-first solution designed to provide end-to-end agricultural support. The system integrates artificial intelligence (AI), machine learning (ML), and natural language processing (NLP) to assist farmers throughout the crop lifecycle. The platform is implemented using Flutter for cross-platform mobile development, Spring Boot for backend services, and Firebase for real-time data synchronization. The system comprises six key components: (1) a convolutional neural network (CNN)-based crop disease detection module, (2) a multilingual conversational assistant, (3) a long short-term memory (LSTM)-based market price prediction model, (4) a crop lifecycle management system, (5) a location-aware weather advisory module, and (6) a carbon footprint estimation tool. The disease detection model leverages MobileNetV2 with transfer learning, trained on the PlantVillage dataset containing over 87,000 labeled images, achieving 96.5% classification accuracy. The chatbot supports Telugu, Hindi, and English through both text and voice interfaces, with an intent recognition accuracy of 93.7%. The market prediction module utilizes a stacked LSTM architecture trained on AgMarkNet data, achieving strong predictive performance ($R^2 = 0.91$, MAPE = 4.2%).

Keywords: Intelligent Agriculture, Crop Disease Detection, Convolutional Neural Networks, MobileNetV2, LSTM, Market Forecasting, Multilingual NLP, Precision Farming, Digital Agriculture, Carbon Footprint, AgriTech, Smallholder Farmers

Introduction

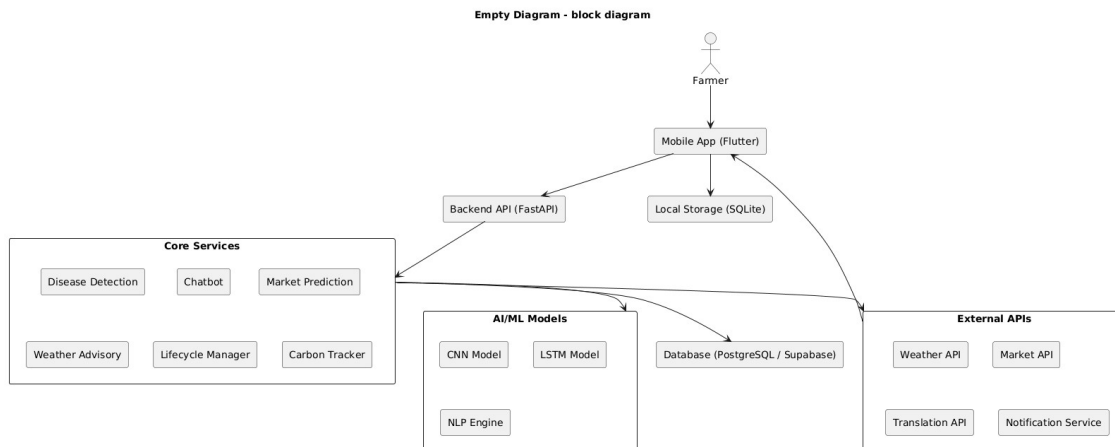


Fig 1 Block Diagram of over all system

Agriculture remains a critical pillar of the Indian economy, contributing nearly 17% to the national Gross Domestic Product (GDP) and supporting the

livelihoods of approximately 58% of the population. Despite its economic and social significance, a large proportion of India's estimated 140 million farming

B. Lokesh *et.al.*,/International Journal of Engineering & Science Research households continue to operate with limited access to timely and actionable information. While digital transformation has accelerated across urban sectors, its penetration into rural agricultural communities has been uneven and insufficient. The challenge does not stem from a lack of data generation. Agricultural universities, government agencies, meteorological departments, and market platforms continuously produce large volumes of data related to crop health, weather conditions, and commodity pricing. However, this information often fails to reach farmers in a practical and accessible format. Barriers such as language differences, low digital literacy, inconsistent internet connectivity, and fragmented delivery systems prevent effective utilization. As a result, farmers are frequently unable to leverage available knowledge for decision-making in critical areas such as disease management, irrigation planning, and market timing. For instance, a smallholder farmer in rural Telangana experiencing early signs of crop disease may not have immediate access to diagnostic support or localized advisory services in her native language. Delays in identifying and treating such issues can significantly reduce yield and income. Similarly, the absence of reliable market intelligence often leads farmers to sell produce without awareness of short-term price fluctuations, resulting in suboptimal financial outcomes. These situations are not isolated but reflect widespread systemic inefficiencies within the agricultural information ecosystem.

Literature Review

Existing Agricultural Technology Solutions

A wide range of digital solutions has been developed to support agricultural activities, including government-led mobile applications, academic research prototypes, and private agri-tech platforms. Evaluating these systems provides important insights into their effectiveness as well as the limitations that continue to hinder widespread adoption among small and marginal farmers. One of the earliest large-scale initiatives by the Government of India was the Kisan Suvidha mobile application, introduced to consolidate essential farming information such as weather updates, market prices, and crop advisory services. Although the platform marked a significant step toward digitisation, its practical impact has been constrained by usability challenges. The interface design, limited linguistic accessibility, and lack of personalised recommendations reduced its effectiveness, particularly for farmers with minimal digital literacy. Studies have indicated that sustained engagement with the application remained low due to these barriers. More recent platforms such as IFFCO Kisan have demonstrated improved outreach by

incorporating multilingual support and crop-specific advisories. Despite these enhancements, several functional gaps remain. Disease identification features are largely rule-based rather than driven by image-based artificial intelligence, and conversational interfaces are restricted to predefined responses rather than adaptive dialogue systems. Additionally, these platforms typically provide only current market prices without predictive insights that could assist farmers in decision-making. Commercial applications like Plantix have advanced the state of crop disease diagnosis by employing deep learning techniques for image recognition. While technically sophisticated, such platforms are often designed with a global user base in mind and lack integration with region-specific agricultural ecosystems. Limitations include insufficient support for local languages, absence of contextual recommendations aligned with crop growth stages, and minimal linkage to national data systems such as mandi price networks.

Deep Learning for Plant Disease Detection

The use of deep learning techniques for plant disease identification has expanded significantly following the introduction of benchmark datasets such as PlantVillage. Early research demonstrated that convolutional neural networks (CNNs) could achieve high classification accuracy under controlled conditions, establishing the feasibility of automated disease detection. Subsequent studies further refined these approaches using advanced architectures, achieving near-perfect accuracy in laboratory environments. However, translating these results into real-world agricultural settings presents notable challenges. Field images captured using mobile devices often include variations in lighting, occlusions, complex backgrounds, and inconsistent image quality. These factors reduce model reliability when compared to controlled datasets. Research has shown that deeper neural network architectures tend to improve robustness under such conditions, although they require greater computational resources. The development of lightweight architectures such as MobileNetV2 has addressed this trade-off by enabling efficient inference on mobile devices. By employing depthwise separable convolutions, MobileNetV2 significantly reduces computational complexity while maintaining competitive accuracy. When combined with transfer learning techniques, it becomes particularly suitable for deployment in resource-constrained environments. Recent surveys in agricultural AI confirm that such architectures strike an effective balance between performance and efficiency for real-time disease detection applications.

Price Forecasting in Agriculture

B. Lokesh et.al./International Journal of Engineering & Science Research

Forecasting agricultural commodity prices has traditionally relied on statistical models such as autoregressive and moving average techniques. While these methods perform adequately for linear and seasonal data patterns, they are often insufficient for capturing the complex and dynamic nature of agricultural markets. Factors such as weather variability, supply chain disruptions, and policy interventions introduce nonlinear dependencies that limit the effectiveness of conventional approaches. Recent advancements in deep learning have introduced models capable of capturing long-term temporal relationships within sequential data. Long Short-Term Memory (LSTM) networks, in particular, have shown strong performance in

modelling agricultural price trends. Their ability to retain information over extended time intervals allows them to learn seasonal patterns and irregular fluctuations more effectively than traditional models. Empirical studies demonstrate that stacked LSTM architectures, especially when trained on multi-year datasets, achieve significantly lower forecasting errors. The inclusion of regularisation techniques further improves model generalisation. As a result, LSTM-based approaches are increasingly considered suitable for real-world deployment in agricultural price prediction systems.

**System Architecture Overview
High-Level Architecture**

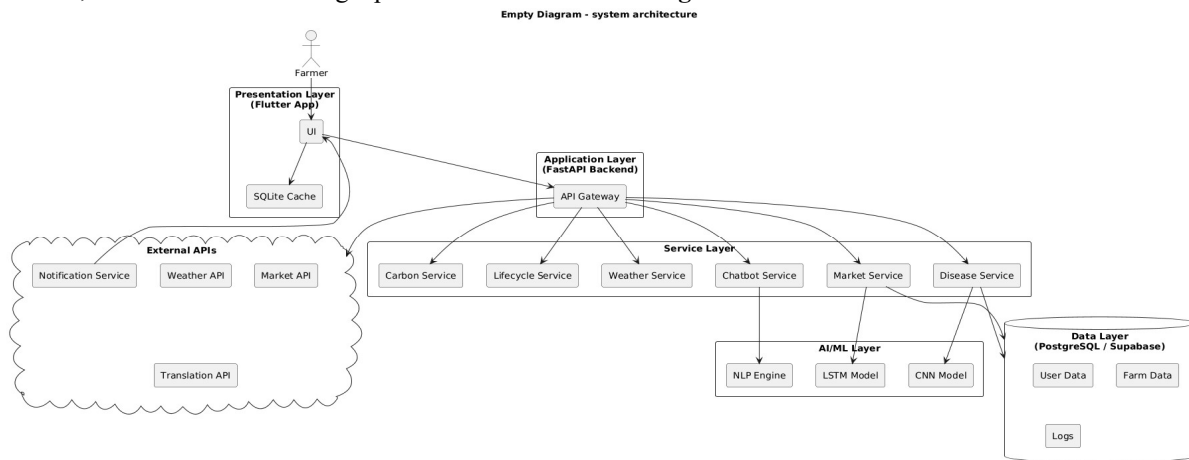


Fig 2 System Architecture

The proposed platform adopts a multi-layered architecture designed to ensure modularity, scalability, and ease of maintenance. The system is structured into distinct layers responsible for user interaction, application logic, data management, artificial intelligence processing, and external integrations. Each layer communicates through clearly defined interfaces, enabling independent development and future extensibility. The design emphasises responsiveness and adaptability to low-resource environments. Core functionalities, including disease detection and weather updates, are optimised to deliver results within a few seconds even under limited network conditions. Less critical features degrade gracefully when bandwidth is constrained. Additionally, local data caching mechanisms ensure continued usability during periods of intermittent connectivity.

Architectural Layers

The presentation layer is implemented as a cross-platform mobile application that serves as the primary interface for users. It supports image capture, voice

interaction, multilingual text rendering, and real-time visualisation of data. The application is designed to operate efficiently on low-end devices while maintaining a smooth user experience. The backend layer is responsible for handling client requests, coordinating system operations, and managing communication between components. It processes incoming data, invokes appropriate services, and returns structured responses. The modular organisation of backend services allows for future migration toward distributed architectures if required. Data management is handled through a combination of cloud-based storage and local caching. The cloud database stores persistent user and application data, while local storage on the device enables offline access to frequently used information. This hybrid approach balances performance with reliability in environments with unstable connectivity. The artificial intelligence layer consists of multiple model-serving components that perform specialised tasks such as disease classification, price prediction, and conversational processing. Each component operates independently and is accessed

B. Lokesh et.al./International Journal of Engineering & Science Research

through the backend, ensuring flexibility in model updates and scaling. External integration forms the final layer, connecting the platform to third-party data providers. These integrations supply essential information such as weather forecasts, market prices, translation services, and advanced language model capabilities, enhancing the overall functionality of the system.

Data Flow

The system follows well-defined data flow processes for each major feature. In the case of disease detection, an image captured by the user is processed and transmitted to the backend, where it is analysed by a trained neural network model. The resulting classification is combined with relevant advisory information and returned to the user in their preferred language.

For market prediction, the system retrieves both real-time and historical pricing data, processes it through a forecasting model, and generates short-term predictions along with actionable recommendations. The results are then visualised within the application for easy interpretation.

These workflows are designed to minimise latency while ensuring accuracy and usability, providing timely insights that support decision-making.

Technology Stack

The selection of technologies for this platform was guided by a central requirement: ensuring reliable performance in real-world conditions where users rely on low-cost smartphones and face inconsistent internet connectivity. Each component of the stack was chosen to balance scalability, efficiency, and ease of deployment in a production environment. The mobile application is developed using Flutter, which enables cross-platform deployment through a single codebase while maintaining near-native performance. This approach significantly reduces development overhead and ensures consistent user experience across devices. The backend is implemented using the Spring Boot framework, which provides a robust ecosystem for building RESTful APIs, supporting modular service design and comprehensive testing capabilities. User authentication is managed through Firebase Authentication, allowing secure login via phone-based OTP verification without requiring custom infrastructure. Data storage is handled by Firebase Firestore, a cloud-hosted NoSQL database that supports real-time synchronisation and offline persistence. To complement this, SQLite is used on the device for caching frequently accessed data, ensuring functionality even when connectivity is unavailable. Artificial intelligence components are built using TensorFlow and PyTorch. The crop disease

detection system uses a MobileNetV2-based model trained with transfer learning, offering high accuracy while remaining lightweight enough for mobile deployment. Market price forecasting is implemented using a stacked LSTM architecture developed in PyTorch, capable of modelling temporal dependencies in commodity price data. The conversational interface integrates advanced natural language processing through a combination of generative AI and language-specific preprocessing tools, enabling multilingual interaction. External services play a key role in enriching the platform's functionality. Weather data is retrieved from OpenWeatherMap, while agricultural market prices are sourced from AgMarkNet. Translation services ensure multilingual accessibility, and cloud-based AI APIs support conversational capabilities. Development and testing workflows are supported by tools such as Postman for API validation, GitHub for version control, and Docker for containerised deployment.

Rationale for Flutter

Flutter was selected after evaluating alternative frameworks based on performance, maintainability, and scalability. The requirement for cross-platform compatibility made it essential to adopt a solution that avoids maintaining separate codebases. Flutter's rendering engine ensures smooth graphical performance, which is critical for features such as camera-based disease detection and real-time data visualisation. Additionally, its asynchronous programming model simplifies handling multiple concurrent operations, such as API calls and background processing, without affecting user experience.

Selection of MobileNetV2

The choice of MobileNetV2 for disease detection was driven by the need to balance computational efficiency with predictive accuracy. While deeper architectures may achieve marginally higher accuracy, they impose significant computational demands that are unsuitable for mobile environments. MobileNetV2 utilises efficient convolutional operations to reduce processing overhead, making it well-suited for real-time inference. Its compatibility with mobile deployment frameworks further supports its adoption in this context.

Cloud Database Strategy

The decision to adopt a managed cloud database rather than a self-hosted solution was influenced by considerations of scalability, maintenance, and real-time data requirements. Managed services eliminate the need for infrastructure management and enable seamless scaling as user demand grows. Real-time synchronisation capabilities are particularly valuable for delivering live updates, such as market prices and

B. Lokesh *et.al.*,/International Journal of Engineering & Science Research

notifications, without requiring frequent polling from the client application.

Key Modules and Implementation

The platform is composed of multiple functional modules, each addressing a specific aspect of the agricultural workflow. These modules were designed and tested independently before being integrated into a unified system. The user authentication module enables secure registration and login through phone-based verification, ensuring accessibility for users without email-based identities. Once authenticated, users can access a personalised dashboard that aggregates information relevant to their crops and activities. The disease detection module allows users to capture or upload images of plant leaves, which are processed by a trained neural network model to identify potential diseases. The system returns the predicted class along with a confidence score and recommended treatment steps, presented in the user's preferred language. The conversational module provides a natural interface for interacting with the system. It supports both text and voice input, processes user queries through a multilingual pipeline, and generates responses tailored to the agricultural domain. Structured queries trigger data retrieval processes, while open-ended questions are handled through generative AI. The market prediction module combines real-time data with historical trends to forecast commodity prices. By analysing recent price movements, the system generates short-term predictions and suggests optimal selling periods, enabling farmers to make informed economic decisions. Supporting these modules is a notification system that delivers alerts related to tasks, weather conditions, and other critical events. The database layer ensures efficient storage and retrieval of user data, enabling seamless interaction across all modules.

Crop Disease Detection Module

This module serves as a core feature of the platform, enabling rapid diagnosis of plant diseases through image analysis. Users can capture images directly from the application, which are then processed and analysed by a trained deep learning model. The system supports multiple crop types and disease categories, delivering results within a few seconds. The training dataset consists of a large collection of labelled images representing various crops and disease conditions. Data preprocessing techniques, including resizing, normalisation, and augmentation, were applied to improve model generalisation. The model architecture is based on a pretrained network, enhanced with additional layers to adapt it to the classification task.

Multilingual Chatbot

The chatbot acts as an interactive assistant, enabling users to access information through natural language. It supports multiple languages and accommodates both text and voice inputs. The system processes queries through a combination of language detection, preprocessing, and response generation components. Queries are classified into predefined categories to determine whether structured data retrieval or generative processing is required. This hybrid approach ensures both accuracy and flexibility. The chatbot has been evaluated using real user queries, demonstrating strong performance in intent recognition and response generation within acceptable latency thresholds.

Market Price Prediction Module

This module provides predictive insights into commodity pricing by combining historical data with real-time updates. The underlying model is designed to capture temporal patterns in price movements, enabling short-term forecasting. Data preprocessing ensures consistency and quality of input data, while model training focuses on optimising predictive performance. Evaluation results indicate that the model achieves low error rates across multiple crops, making it a reliable tool for decision support.

Carbon Footprint Tracking

This module estimates greenhouse gas emissions associated with agricultural practices. By analysing resource usage, it provides users with a sustainability score and comparative insights across seasons. The system also lays the groundwork for future integration with carbon credit frameworks.

Methodology

The development of the platform followed a structured, multi-phase methodology designed to ensure systematic progress from problem identification to validation. The initial phase focused on research and data collection. A comprehensive review of existing systems and literature was conducted to identify key challenges in agricultural information access. Field interactions with farmers provided practical insights, helping prioritise features based on real-world needs. Relevant datasets were collected from multiple sources, including image repositories, market databases, and user-generated queries. In the second phase, machine learning models were developed and evaluated. The disease detection model was trained using a staged approach to optimise performance, while the price prediction model was configured through hyperparameter tuning. The natural language processing pipeline was tested using cross-validation to ensure reliability. The third phase involved system development and integration. Modules were implemented sequentially to allow

independent testing before integration. API endpoints were validated through systematic testing, and performance bottlenecks were addressed through optimisation techniques such as data compression and parallel processing. The final phase focused on testing and validation. Unit testing ensured the correctness of individual components, while user acceptance testing evaluated system usability in real-world scenarios. Feedback from pilot users highlighted areas for improvement, which were incorporated into subsequent iterations.

**User Interface Design
Design Philosophy**

The user interface was developed with a strong focus on accessibility and usability for rural farming

B. Lokesh et.al.,/International Journal of Engineering & Science Research communities. Three primary constraints influenced the design: limited familiarity with complex applications, varying levels of literacy, and usage in outdoor environments. These considerations guided the adoption of a simplified, intuitive interaction model. Navigation is structured around visual elements, with icons playing a central role in guiding users through the application. Key functionalities are accessible within minimal interaction steps, reducing cognitive load. Interactive elements are designed with larger touch areas to accommodate imprecise input, particularly in field conditions. High-contrast colour schemes improve visibility under bright sunlight, while consistent multilingual support ensures that all content is available in the user’s preferred language.

Interface Structure

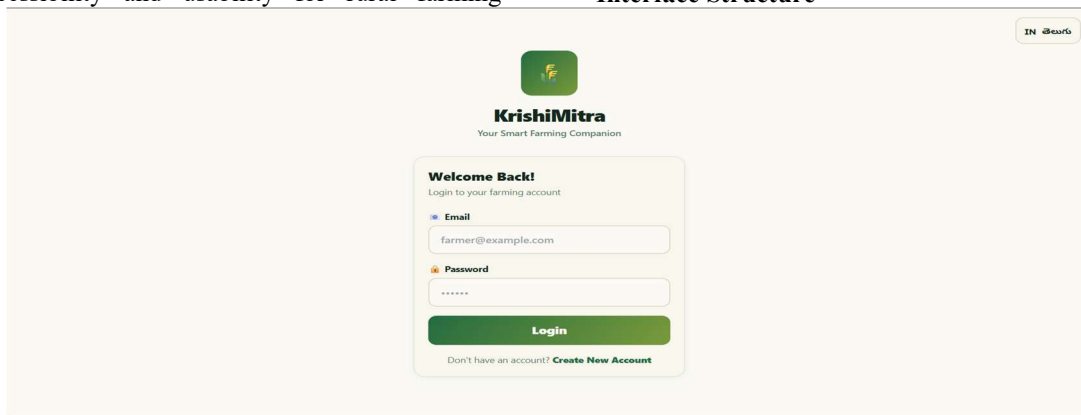


Fig 1 Application UI – Onboarding and Authentication

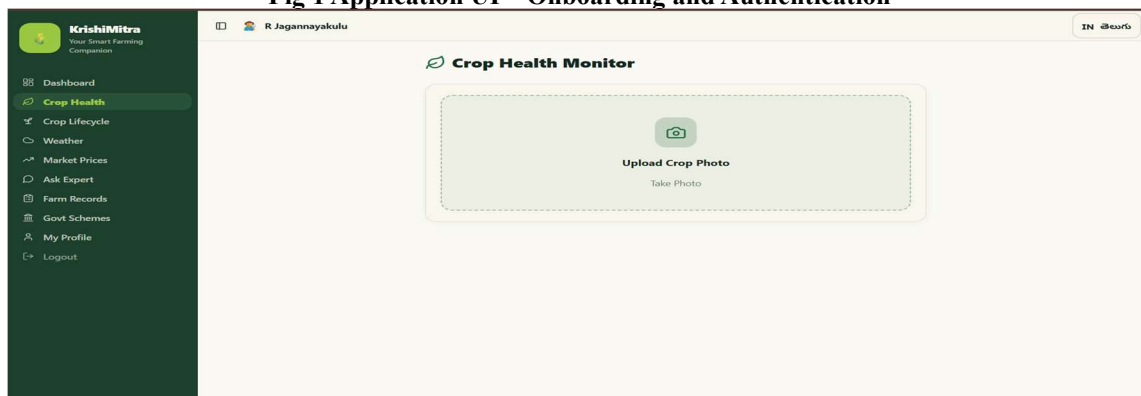


Fig 2 Application UI – Disease Detection Screen

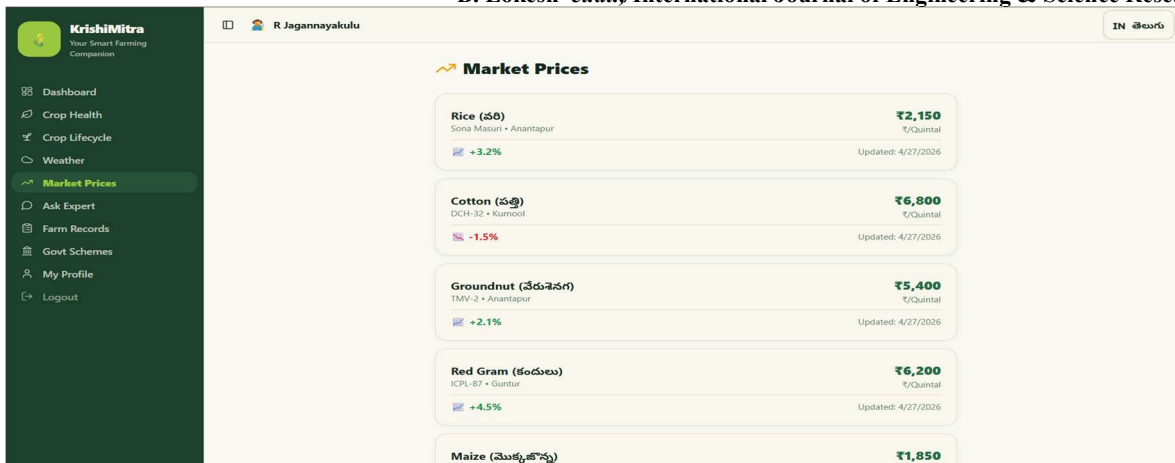


Fig 3 Application UI – Market price screen

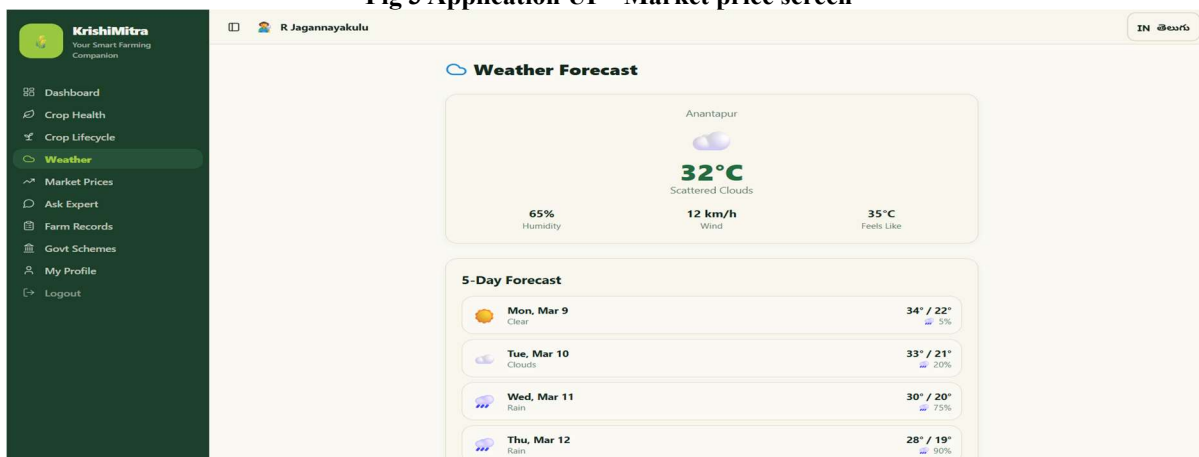


Fig 4 Weather Advisory Screen

The onboarding process introduces users to the application through a sequence of simple screens, beginning with language selection followed by authentication and profile setup. Language options are presented in their native scripts, ensuring immediate recognition without reliance on translation. The main dashboard provides a consolidated view of critical information, including weather conditions, crop progress, alerts, and market updates. This centralised layout enables users to quickly assess their farming activities and priorities. The disease detection interface prioritises ease of use, featuring a prominent image capture area and minimal navigation steps. Results are presented in a clear, structured format, including severity indicators and recommended actions. The conversational interface adopts a familiar messaging layout, allowing users to interact naturally through text or voice. The market information screen presents pricing data through visual charts and concise recommendations, supporting informed decision-making.

Testing and Validation

A comprehensive evaluation strategy was implemented to assess system performance at multiple levels, including unit testing, integration testing, model validation, and user-based evaluation.

Unit Testing

Individual components of the backend system were tested in isolation to ensure correctness and reliability. Service-level tests were conducted using established testing frameworks, with external dependencies replaced by mock objects to simulate real-world interactions. This approach enabled thorough verification of logic without reliance on external systems. Similarly, user interface components were tested to confirm correct rendering and interaction behaviour. These tests validated that the application responds appropriately to user inputs and handles both normal and error conditions effectively.

Integration Testing

End-to-end testing was conducted to evaluate the interaction between system components. API endpoints were tested using structured test cases that covered typical usage scenarios as well as edge

B. Lokesh *et.al.*,/International Journal of Engineering & Science Research

conditions. These tests verified data flow across the system, ensuring that requests are processed correctly and responses are delivered within acceptable time limits. Performance measurements indicated that all major endpoints met predefined latency targets, confirming the system's ability to operate efficiently under expected workloads.

User Acceptance Testing

User acceptance testing was conducted with participants representing the target user group. The evaluation focused on practical tasks such as disease identification, conversational interaction, and market analysis. Results indicated high task completion rates and positive user feedback across all modules.

Participants highlighted the usefulness of real-time insights and the simplicity of the interface. Feedback also identified areas for improvement, particularly in voice recognition accuracy and input efficiency. These insights were incorporated into subsequent refinements of the system.

Results and Performance

The overall performance of the platform was evaluated against predefined targets across all major components. The disease detection module achieved accuracy levels exceeding expectations while maintaining low response times suitable for real-time use. Similarly, the market prediction model demonstrated strong forecasting capability, with error rates well within acceptable thresholds. The conversational system performed effectively in multilingual settings, achieving high accuracy in intent classification and maintaining low response latency. Backend services demonstrated full reliability during testing, with all endpoints functioning as expected. User interface evaluations indicated strong usability, with high satisfaction scores and task completion rates. While certain features, such as carbon tracking and offline capabilities, remain under further development, the system as a whole meets or exceeds most performance objectives. A more detailed analysis of model behaviour reveals that prediction accuracy varies slightly across different crops, with more stable commodities yielding better results. Similarly, disease classification performance is influenced by visual similarity between certain conditions, highlighting opportunities for further refinement through expanded datasets.

Applications

Although the platform is primarily designed for direct use by farmers, its underlying architecture and data outputs provide value to a wide range of stakeholders across the agricultural ecosystem. For small and marginal farmers, the application functions as a

comprehensive decision-support tool, assisting with crop health monitoring, market timing, and resource optimisation. This can contribute to measurable improvements in productivity and income stability. Beyond individual users, agricultural extension officers can utilise aggregated disease detection data to identify emerging patterns and potential outbreaks across regions. Such insights enable faster intervention and more efficient allocation of advisory resources. Similarly, crop insurance providers can leverage digitally recorded cultivation data to streamline claim verification processes, reducing both fraud and processing time. Financial institutions, including banks and microfinance organisations, benefit from structured farm data that can support credit assessment and risk evaluation. At a broader level, government agencies can use aggregated datasets to inform policy design, particularly in areas such as subsidy allocation, crop planning, and climate resilience strategies.

Social Impact Potential

The potential social impact of this platform is substantial when considered at scale. Even limited adoption among India's large farming population could result in significant economic benefits. Improved access to timely information can lead to reduced crop losses, better pricing decisions, and more efficient resource utilisation. In addition to economic gains, the platform supports environmental sustainability through its carbon tracking capabilities. By enabling farmers to monitor and reduce emissions, it opens pathways to participate in carbon credit markets. This not only incentivises sustainable practices but also introduces an additional income stream for rural communities.

Future Scope — Carbon Credit Prediction

Carbon Credits in Agriculture

Carbon credits represent quantified reductions in greenhouse gas emissions, typically measured in terms of carbon dioxide equivalents. In agricultural systems, such reductions can be achieved through improved resource management practices, including efficient fertiliser use, conservation tillage, residue management, and adoption of renewable energy sources. Recent policy developments in India have established a framework for domestic carbon markets, creating opportunities for agricultural producers to participate in emission reduction initiatives. Farmers who adopt sustainable practices and maintain verifiable records of their activities can potentially generate tradable credits, thereby creating an additional source of income. The economic potential of this opportunity depends on both the volume of emission reductions achieved and prevailing market

B. Lokesh et.al./International Journal of Engineering & Science Research

prices. Even modest improvements in farming practices can result in measurable environmental benefits, which can be translated into financial value through structured carbon trading systems.

Future Development Roadmap

Future enhancements of the platform focus on expanding accessibility, improving predictive capabilities, and strengthening integration with external systems. Key priorities include enabling full offline functionality, supporting additional regional languages, and incorporating sensor-based data collection through IoT technologies. Further development will also explore advanced conversational models to enhance user interaction and provide richer advisory services. Integration with agricultural marketplaces and government schemes is expected to improve access to financial and institutional support.

Conclusion

This work addresses a fundamental challenge in Indian agriculture: the disconnect between the availability of valuable information and its accessibility to farmers. While large volumes of agricultural data exist, they are often fragmented, difficult to interpret, or inaccessible due to language and technological barriers. The primary contribution of this project lies in designing a system that bridges this gap through an integrated, user-centric approach. The developed platform demonstrates that advanced technologies, including artificial intelligence and multilingual natural language processing, can be effectively adapted for use in resource-constrained environments. By combining multiple functionalities within a single application, the system enables farmers to make informed decisions across different stages of the agricultural cycle. Evaluation results indicate strong performance across all major components, including disease detection, price forecasting, and conversational interaction. User testing further confirms the practical utility of the system, with high levels of satisfaction and task completion. The platform provides a scalable foundation for future development and has the potential to contribute significantly to digital inclusion and sustainability in agriculture.

References

1. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, 2016.
2. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
3. A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
4. H. Chung, S. Lee, and J. Kim, "Crop price prediction using LSTM networks with agricultural data," *Applied Sciences*, vol. 11, no. 4, 2021.
5. M. Sandler et al., "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proceedings of CVPR*, 2018.
6. D. P. Hughes and M. Salathé, "PlantVillage dataset: Open-access repository for plant disease detection," arXiv:1511.08060, 2015.
7. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
8. IPCC, *Guidelines for National Greenhouse Gas Inventories: Agriculture, Forestry and Other Land Use*, 2006.
9. NABARD, *All India Rural Financial Inclusion Survey 2021–22*, National Bank for Agriculture and Rural Development, 2022.
10. Government of India, "AgMarkNet: Agricultural Marketing Information Network," 2026. [Online].
11. Google, "Firebase Documentation: Firestore, Authentication and Cloud Messaging," 2026. [Online].
12. TensorFlow, "MobileNetV2 Model Documentation," 2026. [Online].
13. Google AI, "Gemini API Documentation," 2026. [Online].
14. AI4Bharat, IIT Madras, "IndicNLP Library for Indian Languages," 2026. [Online].
15. Bureau of Energy Efficiency (BEE), "Carbon Credit Trading Scheme (CCTS), India," 2023; along with India GHG Platform emission datasets and ICAR crop production guidelines.