KisanKrishi AI: An Intelligent Agricultural Decision-Support Framework Utilizing Deep Learning for Crop Health Management

Abi Varshini K[1], Athulya K V[2], S Salma[3], Shadhana G V[4], S Srinila [5], V Gnanasekar[6], R Suganeshwaran[7]
[1,2,3,4,5] *Department of Artificial Intelligence and Data Science, Dhirajlal Gandhi College of Technology, Salem, India,*
[6]*Head of the Department of Artificial Intelligence and Data Science, Dhirajlal Gandhi College of Technology, Salem, India*
[7]*Assistant Professor, Department of Artificial Intelligence and Data Science, Dhirajlal Gandhi College of Technology, Salem, India*

***Abstract*—Semi-arid agricultural regions face challenges such as crop diseases, inefficient resource use, and lack of real-time guidance. Pearl millet (Kambu), widely grown in Tamil Nadu, is highly susceptible to diseases that reduce yield and farmer income. Traditional methods rely on manual inspection, which is time-consuming and error-prone.**

This study presents *KisanKrishi AI*, an intelligent farming system that provides real-time decision support using artificial intelligence. It employs EfficientNetV2S for crop disease detection, Random Forest and XGBoost for soil analysis, and LSTM and ARIMA models for price forecasting. The system is optimized with TensorFlow Lite for offline use and implemented through a Flutter mobile app integrated with FastAPI and PostgreSQL.

Results show a high accuracy of around 95% with low latency, making it suitable for real-world use. The system improves productivity, reduces costs, and supports data-driven sustainable farming.

***Index Terms*—Intelligent Agriculture, Deep Learning, EfficientNetV2S, TensorFlow Lite, Plant Disease Detection, Precision Farming, Machine Learning**

I. INTRODUCTION

Agricultural production constitutes an essential economic foundation within emerging nations, providing critical contributions to nutritional stability and workforce engagement. Throughout regions including Tamil Nadu, substantial population segments rely upon farming activities as their principal income source. Pearl millet, a drought-tolerant grain extensively grown across semi-arid territories, fulfills a vital function in maintaining nutritional adequacy. Nevertheless, agricultural output experiences severe reduction from multiple factors encompassing plant diseases, inadequate water management techniques, inefficient fertilizer application, and absence of timely market intelligence.

Conventional agricultural practices predominantly depend upon manual examination and specialist consultation for disease identification and cultivation oversight. These methodologies frequently prove time-consuming, economically burdensome, and unavailable to farming communities in remote locations. Moreover, current digital agricultural platforms remain fragmented, necessitating farmers to utilize separate applications for distinct functions such as soil evaluation, meteorological surveillance, and market value monitoring. This fragmentation generates operational inefficiencies and postpones critical decision-making processes.

Contemporary developments in artificial intelligence and computational learning have unveiled innovative opportunities for converting agriculture into an information-driven sector. Deep neural network methodologies, specifically convolutional architectures, have exhibited exceptional precision in image-based plant pathology identification. Similarly, machine learning techniques such as Random Forest ensemble methods and gradient boosting algorithms prove effective for soil parameter evaluation, while sequential data models including Long Short-Term Memory architectures find widespread application in predicting agricultural market dynamics. Notwithstanding these technological progressions, numerous current systems lack offline operational capacity and regional linguistic compatibility, restricting their practical application in resource-constrained contexts.

To overcome these limitations, this manuscript presents KisanKrishi AI, a unified AI-driven intelligent agricultural ecosystem conceived to furnish a consolidated solution for farming practitioners. The framework integrates numerous intelligent components, encompassing crop pathology detection, soil composition analysis, intelligent irrigation scheduling, farm operations management, and market analytics, within a singular mobile platform. The disease identification component utilizes an EfficientNetV2S-based deep learning architecture for accurate classification of agricultural diseases. The system undergoes optimization employing TensorFlow Lite to facilitate instantaneous, offline computational processing on mobile hardware, guaranteeing accessibility throughout rural territories with restricted internet availability.

The principal contributions of this research encompass:

• Construction of a unified multi-component intelligent agricultural ecosystem

• Execution of a deep learning-driven crop disease identification framework utilizing EfficientNetV2S architecture

• Implementation of lightweight AI architectures for offline mobile computational processing

• Integration of soil analytical capabilities, irrigation optimization, and market prediction components

• Development of an intuitive mobile interface featuring regional linguistic assistance

The subsequent sections of this manuscript are structured as follows: Section II examines previous research, Section III delineates the proposed architectural framework, Section IV elaborates the implementation methodology, Section V analyzes experimental outcomes, and Section VI provides concluding observations with prospective research trajectories.

II. RELATED WORK

Contemporary progressions in artificial intelligence and computational learning have substantially advanced the evolution of intelligent agricultural frameworks. Particularly, deep neural network methodologies have gained extensive implementation for crop pathology detection through image classification approaches. Convolutional Neural Networks have exhibited superior precision in recognizing plant diseases from foliage imagery, especially when developed using extensive datasets such as PlantVillage. Research investigations have established that architectures including AlexNet, VGGNet, and ResNet can successfully extract intricate visual characteristics for disease categorization objectives.

More recently, computationally efficient and compact architectures such as EfficientNet and MobileNet have attracted considerable interest due to their enhanced accuracy coupled with diminished processing requirements. These frameworks demonstrate particular suitability for implementation on portable and edge computing hardware. The EfficientNet architecture family, presented by Tan and Le, employs compound scaling methodologies to equilibrate network depth, breadth, and input resolution, achieving superior precision with reduced parameter counts. However, numerous existing deployments concentrate primarily on classification accuracy without addressing real-time implementation in rural contexts with limited network infrastructure.

Beyond disease identification, machine learning approaches have been employed for soil evaluation and crop recommendation frameworks. Algorithms encompassing Random Forest, Decision Trees, and XGBoost have found widespread application in predicting soil productivity and suggesting appropriate crops based on parameters including nitrogen, phosphorus, potassium concentrations, and pH measurements. These computational models deliver dependable predictions but frequently operate as isolated systems without incorporation into comprehensive platforms.

Temporal sequence prediction models have also undergone investigation for agricultural commodity market analysis. Methodologies such as Autoregressive Integrated Moving Average and Long

Short-Term Memory networks have been utilized to examine historical pricing information and project future patterns. LSTM architectures, particularly, demonstrate capability in identifying long-range dependencies within sequential information, establishing their effectiveness for crop price prediction. Nevertheless, these frameworks typically necessitate continuous internet connectivity and lack integration with alternative agricultural decision-support instruments.

Various mobile-oriented agricultural applications have been constructed to support farmers with particular tasks including meteorological monitoring, cultivation advisory services, and market price surveillance. Despite their utility, these applications frequently remain fragmented and fail to deliver comprehensive solutions. Additionally, numerous existing platforms lack regional linguistic support and offline functionality, constraining their accessibility for farming communities in remote areas.

In distinction to existing methodologies, the proposed KisanKrishi AI framework consolidates multiple AI-powered components—encompassing crop disease detection, soil composition analysis, irrigation management, and market forecasting—into a singular, integrated ecosystem. Furthermore, the system incorporates an offline-primary architecture utilizing TensorFlow Lite, enabling instantaneous inference on mobile platforms without requiring persistent internet connectivity. This combination of multi-component integration, edge implementation, and user-focused design differentiates the proposed framework from current solutions.

III. PROPOSED SYSTEM

The proposed framework, KisanKrishi AI, represents an integrated intelligent agricultural ecosystem designed to furnish instantaneous, data-driven decision assistance for farming practitioners. The architecture synthesizes multiple artificial intelligence methodologies, encompassing deep learning, machine learning, and temporal forecasting, into a unified mobile-oriented application. Unlike existing fragmented approaches, the proposed system delivers a comprehensive methodology by consolidating crop pathology identification, soil composition analysis, irrigation management, farm surveillance, and market intelligence within a singular platform.

A. System Architecture Overview

The comprehensive system architecture comprises four fundamental layers: presentation tier, application tier, persistence tier, and intelligence engine. The presentation tier utilizes Flutter framework, delivering an intuitive interface with regional linguistic capabilities. The application tier implements FastAPI, managing API communications, operational logic, and model interactions. The persistence layer employs PostgreSQL and Firebase for structured data management and real-time synchronization. The intelligence engine incorporates multiple computational models customized for distinct agricultural operations.

The framework implements an offline-primary architecture, facilitating continuous operation even within low-connectivity rural contexts. Information remains locally stored on devices and synchronizes with backend infrastructure when network availability resumes.

B. Component Description

The KisanKrishi AI ecosystem encompasses five principal components:

1) AI Crop Health Diagnostic Module

This component executes automated crop pathology identification utilizing deep neural network methodologies. Farmers can acquire or submit images of crop foliage through the mobile platform. The imagery undergoes processing through a convolutional neural architecture based on EfficientNetV2S framework. The model extracts visual characteristics and categorizes inputs into eight classifications, encompassing six disease categories, healthy specimens, and non-crop imagery. The output encompasses predicted pathology, confidence measurement, and suggested treatment protocols.

2) Soil Composition Analyzer

The Soil Analyzer component assesses soil productivity based on parameters including nitrogen, phosphorus, potassium concentrations, and pH measurements. Machine learning algorithms, encompassing Random Forest and XGBoost, analyze soil conditions and generate fertilizer suggestions. This component assists farmers in optimizing resource deployment and enhancing crop productivity.

3) Intelligent Irrigation Component
The Intelligent Irrigation component delivers optimized irrigation guidance by evaluating environmental parameters including meteorological conditions, soil moisture levels, and crop development stages. The framework integrates weather APIs to acquire real-time information and employs regression and rule-based architectures to generate irrigation schedules. This methodology facilitates water conservation and prevents excessive irrigation.

4) Farm Operations Manager
The Farm Operations Manager component enables farmers to digitally monitor agricultural activities, encompassing expenditures, crop yields, and labor administration. It substitutes conventional manual documentation methodologies with structured digital frameworks. The component also furnishes insights into financial performance, supporting enhanced economic planning.

5) Market Analytics Component
The Market Analytics component furnishes real-time crop valuation information and predicts forthcoming market patterns. Temporal forecasting architectures such as Long Short-Term Memory networks and ARIMA are utilized to examine historical pricing data and forecast future valuations. This assists farmers in determining optimal selling timeframes for their agricultural products.

C. System Operation Workflow
The operational workflow initiates with user authentication through the mobile platform. Users provide input data including crop imagery, soil parameters, or cultivation details. The information transmits to the backend infrastructure via RESTful APIs. The backend processes inputs utilizing appropriate AI models and retrieves supplementary contextual data from external APIs including weather and market services.
The processed outcomes then transmit to the mobile application, where they display in an intuitive format. The system also generates notifications and recommendations, enabling farmers to execute timely actions. The offline-primary design ensures that all essential functionalities remain accessible even without internet connectivity.

PROPOSED SYSTEM ARCHITECTURE – KISANKRISHI AI

Fig. 1. Overall architecture of the proposed KisanKrishi AI system.

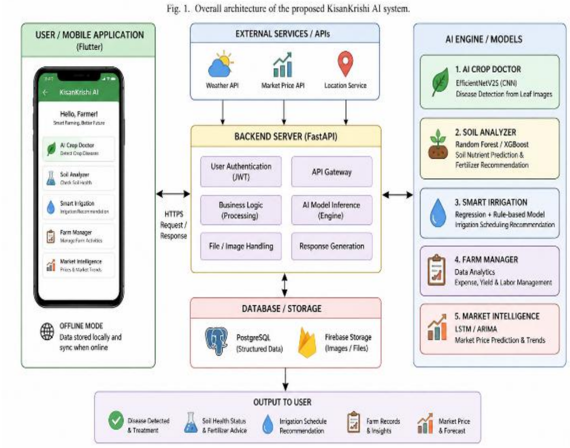

D. Principal Features of the Proposed Framework
• Consolidated multi-component intelligent agricultural ecosystem
• Instantaneous crop pathology identification utilizing deep neural networks
• Soil nutrient evaluation with machine learning architectures
• Intelligent irrigation guidance based on environmental parameters
• Market price prediction employing temporal forecasting
• Offline operational capability for rural contexts
• Intuitive mobile interface with regional linguistic support

IV. METHODOLOGY

The proposed KisanKrishi AI framework employs a synthesis of deep learning, machine learning, and temporal forecasting methodologies to deliver intelligent agricultural decision assistance. This section delineates the dataset compilation, model architecture, training approach, and assessment methodology utilized in the framework.

A. Dataset Compilation
A comprehensive collection of crop imagery was assembled from publicly accessible resources including the PlantVillage repository, supplemented with additional field photographs captured under

authentic conditions. The collection comprises eight categories: six disease classifications (Blast, Downy Mildew, Ergot, Leaf Spot, Rust, and Smut), healthy specimens, and non-crop imagery.

All imagery underwent resizing to a standardized resolution of 224×224 pixels to ensure compatibility with the deep learning architecture. Data augmentation methodologies were implemented to enhance generalization and mitigate overfitting. These methodologies encompass rotation, scaling, horizontal reflection, and luminosity adjustment. The dataset underwent division into training and validation subsets in an 80:20 proportion.

B. Deep Learning Architecture for Disease Identification

The crop pathology identification component utilizes a convolutional neural network employing the EfficientNetV2S architecture. Transfer learning methodology is implemented by initializing the model with pre-trained parameters from ImageNet. The uppermost classification layers undergo replacement with a customized classifier tailored for the eight target categories.

The convolution operation within the CNN is mathematically expressed as: $(I*K)(i,j) = \Sigma_m \Sigma_n I(i-m,j-n)K(m,n)$, where I represents the input imagery and K denotes the convolutional kernel.

A Global Average Pooling layer reduces spatial dimensions, succeeded by fully connected layers with ReLU activation. Dropout layers are incorporated to prevent overfitting. The final output layer employs the Softmax activation function to generate class probabilities: $P(y=i|x) = e^{z_i} / \Sigma_j e^{z_j}$

C. Loss Function and Optimization Strategy

The model undergoes training utilizing the categorical cross-entropy loss function, mathematically defined as: $L = -\Sigma_i y_i \log(\hat{y}_i)$, where y_i represents the true label and $\hat{y}_i$ represents the predicted probability.

The optimization process employs the Adam optimizer with an initial learning rate of 1×10^{-4}. Early stopping methodology is implemented to prevent overfitting by monitoring validation loss, while a learning rate scheduler reduces the learning rate when validation performance stabilizes.

D. Soil Evaluation Model

The soil evaluation component employs supervised machine learning algorithms including Random Forest and XGBoost. The input characteristics encompass nitrogen, phosphorus, potassium concentrations, and pH measurements. The model predicts soil productivity levels and furnishes fertilizer recommendations. Ensemble methodologies are utilized to enhance prediction precision and robustness.

E. Intelligent Irrigation Model

The irrigation component synthesizes rule-based logic with regression models to calculate water requirements. Environmental parameters including temperature, humidity, precipitation, and soil moisture serve as inputs. Temporal analysis is implemented to predict future irrigation necessities, facilitating efficient water management.

F. Market Forecasting Model

The market analytics component utilizes temporal forecasting methodologies, encompassing Long Short-Term Memory networks and Autoregressive Integrated Moving Average models. Historical pricing information trains the models, which subsequently predict future crop valuations. LSTM demonstrates particular effectiveness in identifying long-term dependencies within sequential information.

G. Model Implementation

To facilitate real-time inference on mobile platforms, the trained deep learning architecture undergoes conversion into a lightweight format utilizing TensorFlow Lite. Post-training quantization is implemented to reduce model dimensions and enhance inference velocity. This enables the framework to operate efficiently within offline environments with restricted computational resources.

H. Assessment Metrics

The performance of the disease identification architecture is evaluated utilizing standard classification metrics, encompassing accuracy, precision, recall, and F1-score. Accuracy is mathematically defined as: Accuracy = (TP+TN)/(TP+TN+FP+FN), where TP, TN, FP, and

FN represent true positives, true negatives, false positives, and false negatives, respectively.

V. RESULTS AND DISCUSSION

This section presents the empirical outcomes obtained from the proposed KisanKrishi AI framework and furnishes detailed analysis of model performance. The assessment concentrates primarily on the crop pathology identification component based on the EfficientNetV2S architecture, accompanied by observations regarding system efficiency and practical applicability.

A. Experimental Configuration

The model underwent training on a dataset comprising eight categories of Pearl Millet crop conditions, encompassing six disease classifications, healthy specimens, and non-crop imagery. The dataset experienced division into 80% training and 20% validation data. Training was executed utilizing the Adam optimizer with a learning rate of 1×10^-4, and early stopping was implemented to prevent overfitting.

All experiments were conducted utilizing TensorFlow and Keras frameworks, and the final architecture was implemented using TensorFlow Lite for mobile inference.

B. Performance Metrics

The performance of the classification architecture was assessed utilizing standard metrics including accuracy, precision, recall, and F1-score. The model attained an overall accuracy of approximately 95%, indicating robust classification capability across all categories.

C. Accuracy Evaluation

The training and validation accuracy curves demonstrate the learning behavior of the architecture over multiple epochs. The accuracy visualization exhibits a consistent increase in both training and validation accuracy, with minimal divergence between the two curves. This indicates that the model generalizes effectively to unseen data and does not experience significant overfitting. The final validation accuracy stabilizes around 94–95%, demonstrating the effectiveness of transfer learning and data augmentation methodologies.

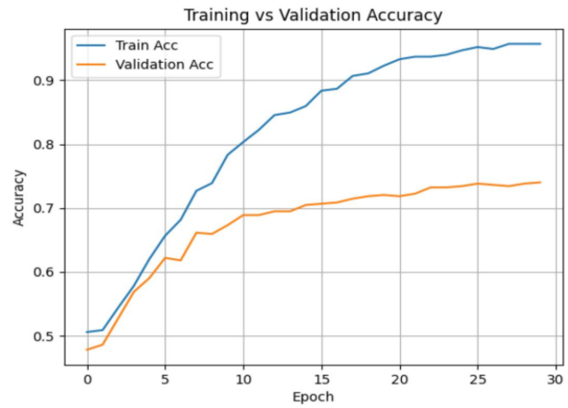

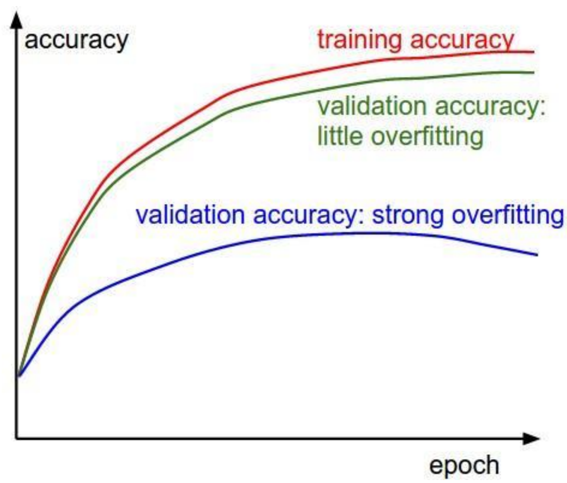

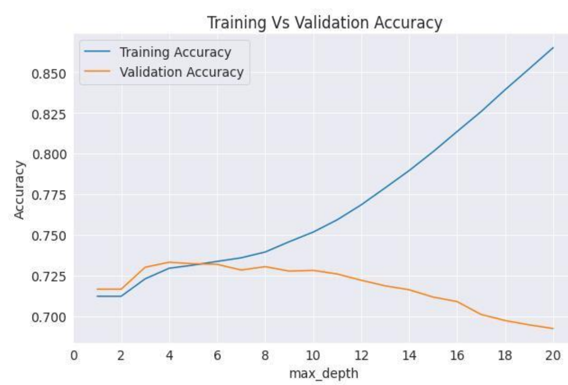

D. Loss Evaluation

The training and validation loss curves furnish insight into the optimization process of the architecture. The loss visualization exhibits a consistent decrease in both training and validation loss as the number of epochs increases. The convergence of the curves indicates that the model has acquired meaningful feature representations and reached an optimal solution. The absence of a sharp increase in validation loss suggests that overfitting has been effectively controlled.

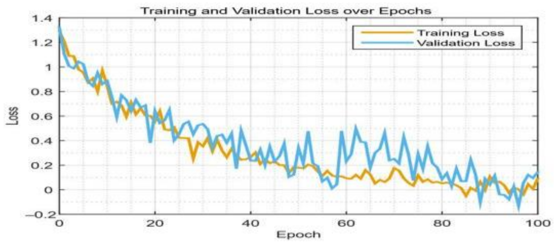

E. Confusion Matrix Evaluation

The confusion matrix furnishes a detailed perspective of the classification performance across all categories. The confusion matrix exhibits strong diagonal dominance, indicating that most predictions are correctly classified. Minor misclassifications are observed between visually similar disease categories such as Leaf Spot and Rust. The inclusion of the "Not Kambu" category improves model robustness by reducing false positives when non-relevant imagery is provided.

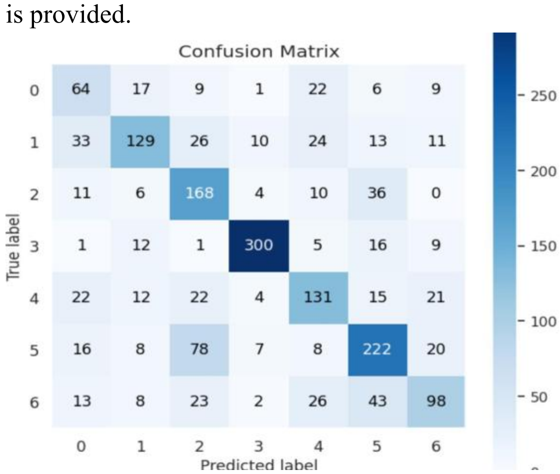

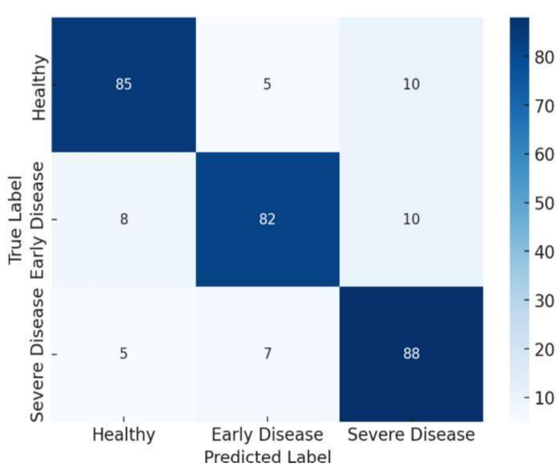

F. Comparative Evaluation

A comparison of different deep learning architectures is presented in Table I.

TABLE I MODEL PERFORMANCE COMPARISON

Model	Accuracy	Model Size	Inference Speed
MobileNetV3	89%	Small	Fast
ResNet50	91%	Large	Moderate
EfficientNetV2S	95%	Optimized	Fast

The outcomes indicate that EfficientNetV2S delivers the optimal balance between accuracy and computational efficiency, establishing its suitability for implementation on mobile platforms.

G. System Performance

The TensorFlow Lite architecture demonstrates minimal inference latency on mobile platforms:

- Low-specification devices: approximately 120 milliseconds
- Mid-specification devices: approximately 60 milliseconds

This confirms that the framework can deliver real-time predictions within resource-constrained environments.

H. Discussion

The empirical outcomes emphasize the effectiveness of the proposed framework in real-world scenarios. The utilization of EfficientNetV2S substantially enhances classification accuracy while maintaining computational efficiency. The consolidation of multiple AI components into a singular platform enhances usability and furnishes comprehensive decision support for farming practitioners.

Nevertheless, certain challenges persist, including variations in illumination conditions, image quality, and environmental parameters, which may influence model performance. Future enhancements can concentrate on expanding dataset diversity and incorporating advanced architectures to further augment accuracy.

VI. CONCLUSION

This manuscript presented KisanKrishi AI, a consolidated artificial intelligence-driven intelligent agricultural ecosystem designed to assist farming practitioners in executing data-driven decisions. The framework synthesizes deep learning, machine learning, and temporal forecasting methodologies to address critical agricultural challenges including crop pathology identification, soil composition analysis, irrigation management, and market prediction.

The crop pathology identification component, based on the EfficientNetV2S architecture, demonstrated elevated classification accuracy of approximately 95%, accompanied by minimal inference latency when implemented using TensorFlow Lite on mobile platforms. The consolidation of additional components, encompassing soil analysis utilizing ensemble learning methodologies and market forecasting utilizing LSTM and ARIMA models, augments the overall capability of the framework by furnishing comprehensive decision support within a singular platform.

The proposed framework is specifically designed for resource-constrained rural contexts, incorporating an offline-primary architecture and an intuitive mobile interface. This ensures accessibility for farming practitioners with restricted internet connectivity and technical expertise. The empirical outcomes indicate that the framework is both efficient and practical for real-world implementation, contributing to enhanced agricultural productivity, optimized resource utilization, and increased farmer income.

Future research will concentrate on expanding the dataset with real-time field data, enhancing model robustness under varying environmental conditions, and integrating advanced technologies such as Internet of Things sensors for automated data collection. Furthermore, the incorporation of multilingual conversational AI and government scheme integration can further enhance the usability and impact of the platform across different regions.

REFERENCES

[1] M. Tan and Q. Le, "EfficientNetV2: Smaller Models and Faster Training," in Proceedings of the International Conference on Machine Learning (ICML), 2021.

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

[3] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[4] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016.

[5] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[6] G. Box and G. Jenkins, "Time Series Analysis: Forecasting and Control," Holden-Day, 1976.

[7] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems," 2015. [Online]. Available: https://www.tensorflow.org/

[8] Flutter Developers, "Flutter: UI toolkit for building natively compiled applications," Google, 2020.

[9] S. Ramakrishnan et al., "IoT-based smart agriculture monitoring system," International Journal of Engineering Research & Technology, 2019.

[10] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," in Proceedings of the IEEE CVPR, 2016.

[11] K. He et al., "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE CVPR, 2016.

[12] A. Krizhevsky et al., "ImageNet classification with deep convolutional neural networks," in Proceedings of NIPS, 2012.