

Crop Vision AI: An AI-Powered Smart Farming Assistant Using CNN and Machine Learning

1. Dr. Akhilesh Tayade

2. Aniket Pravin Wankhade 3. Dweep D. Gadge 4. Vedant Shaileshrao Khanke

5. Rushikesh Ganesh Pilantre 6. Swaraj Arvind Kalane 7. Ajay Rajesh Pawar

Department of Computer Science and Engineering

P. R. Pote Patil College of Engineering & Management, Amravati

ABSTRACT

Agricultural productivity is influenced by a wide range of interdependent factors, including the timely detection of crop diseases, accurate yield forecasting, appropriate fertilizer application, and awareness of prevailing weather conditions. In many farming regions, decisions on these matters are still guided largely by inherited experience and visual inspection rather than by structured, data-driven analysis, which can lead to delayed responses, avoidable crop losses, and inefficient use of agricultural inputs. This work introduces Crop Vision AI, a unified, AI-driven farming-assistance platform designed to bring several independent decision-support functions together under a single, farmer-accessible interface. The platform applies a

Convolutional Neural Network (CNN) to identify visible crop diseases from uploaded leaf photographs, employs a Random Forest ensemble model to project expected crop yield from agronomic and environmental variables, and uses a Decision Tree classifier to recommend suitable fertilizers based on crop and soil attributes. A live Weather API supplies location-specific meteorological updates that can be viewed alongside the AI-generated recommendations. Rather than functioning as isolated utilities, these components are designed to operate cooperatively within one workflow, reducing the need for farmers to consult multiple disconnected tools. The architecture is intentionally modular, allowing new crops, disease classes, regional datasets, and additional agricultural services to be incorporated as the platform matures. This document expands on the original project report by elaborating the background, methodology, implementation considerations, and evaluation approach in greater depth, while preserving the overall structure and intent of the initial submission.

Keywords: *Crop Vision AI, Smart Farming, Convolutional Neural Network, Machine Learning, Crop Disease Detection, Random Forest, Decision Tree, Fertilizer Recommendation, Weather API, Precision Agriculture, Agricultural Decision Support.*

1. INTRODUCTION

1.1 Background

Agriculture continues to be one of the most demanding domains for real-time, evidence-based decision-making. A single growing season may require a farmer to recognize the early symptoms of an infection on a leaf, judge whether the coming rainfall will

support or damage a standing crop, choose among several fertilizer options for a particular soil type, and estimate what the eventual harvest might look like. Each of these judgments has traditionally depended on years of accumulated field experience, informal advice from neighbouring growers, or occasional visits from agricultural extension officers, none of which are consistently available at the moment a decision is actually needed. As smartphone penetration and mobile connectivity have expanded into rural areas, there is now a practical opportunity to place a portion of this decision-making support directly into the hands of the farmer through software rather than through a scarce human expert.

Artificial intelligence and machine learning are particularly well suited to this setting because they can encode patterns learned from large collections of prior examples: thousands of labelled leaf photographs, historical yield records tied to weather and soil variables, or documented fertilizer outcomes across different crop-soil combinations. Once trained, such models can generate a recommendation in seconds, a speed that manual expert consultation cannot match. Crop Vision AI is motivated by this gap between the availability of AI techniques and their limited integration into a single, farmer-facing product.

1.2 Existing Challenges

A review of currently available digital agricultural tools shows that most are narrow in scope. A disease-identification app typically does not forecast yield; a yield-estimation tool typically does not recommend fertilizer; and a weather application typically has no awareness of the crop being grown at all. This fragmentation forces the farmer to move between several unrelated applications, each with its own interface, login, and data format, which adds friction precisely at the moments when a fast decision is most valuable.

Beyond fragmentation, each underlying task carries its own technical difficulty. Disease recognition from images depends on the availability of a sufficiently large and visually diverse training set, since diseases can appear differently depending on crop variety, growth stage, and image quality. Yield estimation depends on the availability of clean historical records that combine weather, soil, and crop information, which are not always maintained consistently at a local level. Fertilizer recommendation depends on an accurate understanding of both the crop's nutritional requirements and the specific deficiencies of the soil in question, information that is not always tested or recorded. These constraints do not prevent the construction of a

useful system, but they do shape the design choices made in Crop Vision AI, particularly its emphasis on modularity and incremental data expansion.

1.3 Proposed Crop Vision AI

Crop Vision AI is put forward as a single integrated assistant rather than a collection of separate scripts. A farmer interacts with one interface to upload a leaf photograph, enter basic agronomic details, and select a location; the platform then coordinates the CNN, Random Forest, Decision Tree, and Weather API components internally and returns a combined set of results. The intention is that the farmer experiences one coherent interaction rather than several disconnected ones, even though different machine learning techniques are working behind the interface to answer different questions.

Workflow: Farmer Input → Data/Image Processing → AI/ML Models → Disease, Yield and Fertilizer Results → Weather Information → Decision Support

1.4 Motivation

The motivation for this project stems from the observation that many small and mid-sized farms do not have consistent access to agricultural extension officers or agronomists, particularly during time-sensitive periods such as the early onset of disease symptoms or the approach of unfavourable weather. A tool that can offer an immediate, if preliminary, opinion on these matters has the potential to reduce the cost of delay, even where it is not intended to fully replace expert judgment. There is also a broader motivation connected to data collection: as more farmers use the platform, the resulting usage patterns and, where consented, image submissions can help refine future versions of the underlying models.

1.5 Scope of the Project

The current scope of Crop Vision AI covers four coordinated services: image-based disease identification for a defined set of crops and disease classes, numerical yield prediction based on available agronomic and environmental parameters, rule-based fertilizer recommendation derived from crop and soil inputs, and live weather retrieval tied to farmer location. The scope explicitly excludes, at this stage, functions such as automated pesticide dosage calculation, direct integration with farm machinery, and financial or market-price forecasting, all of which are identified as candidates for later expansion rather than part of the present implementation.

2. PROBLEM STATEMENT

Farmers are routinely expected to make decisions across several agricultural dimensions at once, even though the information required for each dimension typically originates from a different source and is rarely presented together. Disease identification depends on visual inspection, yield estimation depends on historical and environmental data, fertilizer selection depends on soil and crop characteristics, and weather awareness depends on external meteorological services. When these information streams remain separate, farmers must either consult several distinct tools or rely on informal judgment, both of which introduce delay and inconsistency.

The problem is compounded where access to trained agronomists or extension staff is limited, seasonal, or geographically distant, and where existing software solutions are designed around a single narrow function rather than the farmer's overall workflow. There is, therefore, a clear need for a consolidated platform capable of accepting crop images and agricultural parameters as input and returning disease, yield, fertilizer, and weather information together, so that a single interaction can support several related decisions rather than one.

Crop Vision AI is proposed in direct response to this gap. By centralizing disease detection, yield prediction, fertilizer recommendation, and weather monitoring within one AI-based platform, the system is intended to shorten the path between observing a problem in the field and receiving actionable, data-informed guidance.

3. OBJECTIVES

The overarching objective of this work is to design and demonstrate an AI-powered smart farming assistant that consolidates several forms of agricultural decision support within one accessible platform. The specific objectives pursued are listed below.

To design a centralized platform that brings together crop disease detection, yield prediction, fertilizer recommendation, and weather monitoring in one workflow.

To implement a Convolutional Neural Network capable of classifying crop leaf images into disease categories or a healthy class.

To implement a Random Forest regression model that predicts expected crop yield from agronomic and environmental parameters.

To implement a Decision Tree classifier that recommends suitable fertilizers based on available crop and soil inputs.

To integrate a Weather API that supplies real-time, location-specific meteorological information.

To design an interface that presents AI-generated agricultural information in a clear and easy-to-use format, suitable for users with limited technical background.

To evaluate the proposed system with respect to accuracy, usability, reliability, and overall usefulness to the intended farming audience.

To structure the underlying dataset pipeline and model training process so that additional crops, disease classes, and regional data can be incorporated without redesigning the platform.

To consider basic data-handling and privacy safeguards for farmer-submitted images and agronomic information.

4. LITERATURE REVIEW

Research into automated plant disease diagnosis has consistently pointed to the effectiveness of deep learning for image-based classification tasks. Mohanty et al. examined the use of deep convolutional networks for identifying plant diseases directly from leaf photographs and reported that such models could achieve strong classification performance when trained on

sufficiently large image collections, establishing an early foundation for CNN-based diagnostic tools. [1] Building on the need for accessible training data, Hughes and Salathé compiled and released an open repository of plant health images specifically intended to support the development of mobile disease-diagnostic applications, underscoring how public datasets can accelerate progress in this area. [2]

Ferentinos extended this line of work by evaluating several deep learning architectures for plant disease detection and diagnosis across a broad range of crop species, concluding that convolutional architectures generalize well to agricultural image classification when trained appropriately. [3] Too et al. carried out a comparative study of fine-tuned deep learning models for plant disease identification, finding that transferring features learned on large general-purpose image datasets to the agricultural domain could improve recognition accuracy while reducing the amount of task-specific data required. [4] Sladojevic et al. similarly reported strong recognition performance for deep neural networks applied to leaf-image classification, reinforcing the broader consensus that convolutional approaches are well matched to this problem. [12]

On the prediction side of agricultural analytics, ensemble learning methods have proven useful for handling the many interacting variables that influence crop outcomes. Breiman's original formulation of Random Forests demonstrated that combining a large number of decision trees, each trained on a randomly sampled subset of data and features, produces a model that is both more accurate and more resistant to overfitting than any single tree. [5] Decision Trees themselves, as formalized by Quinlan, remain attractive for recommendation-style problems precisely because their internal logic can be traced and explained, a property that is valuable in an agricultural context where a farmer may want to understand why a particular fertilizer was suggested. [6]

At a broader level, several survey papers have documented how artificial intelligence and connected technologies are reshaping agricultural practice. Kamilaris and Prenafeta-Boldú reviewed a wide range of deep learning applications across agricultural tasks and noted the growing dependence of such systems on the availability and quality of domain-specific datasets. [7], [10] Wolfert et al. examined the role of big data in smart farming and argued that the real value of digital agriculture emerges when disparate data streams, such as weather, soil, and crop information, are combined rather than analysed in isolation, a conclusion that directly supports the integrated design adopted by Crop Vision AI. [8] Khan et al. further explored deep learning-based plant disease detection and classification pipelines, reporting consistent gains in accuracy from architectural refinements and improved data preprocessing. [9]

Dataset characteristics have also been shown to materially affect model reliability. Barbedo investigated how dataset size and variety influence the effectiveness of deep learning and transfer learning for plant disease classification, concluding that models trained on narrow or imbalanced datasets tend to perform poorly when exposed to field conditions that differ from their training distribution, a finding that directly informs the dataset dependency limitation discussed later in this report. [11] Chawla et al.

reviewed the broader landscape of machine learning applications in smart agriculture, covering not only disease detection but also yield forecasting and resource optimization, and observed that most deployed systems still address individual tasks rather than offering a combined decision-support experience. [13] Institutional publications reinforce this observation at a policy level: the Food and Agriculture Organization has documented the expanding role of digital technologies in agriculture and rural development, while the World Bank has highlighted data-driven farming as a mechanism for improving smallholder productivity and resilience. [14], [15]

Taken together, the reviewed literature supports the technical feasibility of each individual component proposed in this work, namely CNN-based disease recognition, Random Forest yield estimation, and Decision Tree fertilizer recommendation, while also repeatedly identifying the absence of a single, farmer-oriented platform that brings these functions together. Crop Vision AI is positioned to address precisely this integration gap rather than to introduce an entirely new algorithmic technique.

5. PROPOSED METHODOLOGY

5.1 Proposed System Overview

Crop Vision AI is organized around a shared input stage and four cooperating analytical modules. When a farmer submits a leaf photograph, that image is routed to the image-processing pipeline that feeds the CNN classifier. Agronomic details entered by the farmer, such as crop type, soil characteristics, and rainfall estimates, are simultaneously routed to the Random Forest yield model and the Decision Tree fertilizer module. The farmer's selected location is used to query the Weather API independently of the other three modules. Once each module returns its result, the platform assembles the outputs into a single response view rather than presenting them as separate, disconnected screens.




Fig. 1. Workflow of the proposed Crop Vision AI system.

5.2 CNN-Based Crop Disease Detection

The disease-detection module accepts a leaf image, resizes and normalizes it, and passes it through a convolutional neural network trained to distinguish between a defined set of disease categories and a healthy-leaf class. The convolutional layers progressively learn low-level visual patterns such as edges and colour gradients in early layers and more complex structures such as lesion shapes and discoloration patterns in deeper layers. Pooling layers reduce spatial dimensionality between convolutional blocks, and dropout layers are included to reduce the likelihood of the network memorizing the training set rather than learning generalizable features. A final fully connected layer with a softmax activation produces a probability distribution across the possible disease classes, and the class with the highest probability is returned as the predicted result. CNN was selected for this task because convolutional architectures are able to learn discriminative visual features directly from raw pixel data without requiring manually engineered descriptors. [1], [3], [4], [12]

Workflow: Leaf Image → Preprocessing → CNN → Feature Extraction → Classification → Disease Result

5.3 Random Forest for Crop Yield Prediction

The yield-prediction module estimates expected crop output using a Random Forest regression ensemble. Random Forest was chosen because it combines many individually trained decision trees, each built from a bootstrapped sample of the training data and a randomly restricted subset of available features, and averages their individual predictions into a single, more stable estimate. This ensembling reduces the variance associated with any single decision tree and allows the model to capture nonlinear relationships between agronomic variables without requiring those relationships to be specified in advance. [5]

Workflow: Agricultural Data → Preprocessing → Random Forest → Yield Prediction

5.3.1 Yield Prediction Inputs

Inputs to the yield model include crop type, seasonal rainfall, average temperature, soil texture or fertility indicators, irrigation availability, and, where available, historical yield figures for the same field or region. Categorical inputs such as crop type and soil class are encoded numerically, and continuous variables such as rainfall and temperature are scaled to comparable ranges before being supplied to the trained model, so that no single variable dominates the learning process purely because of its numeric scale.

5.3.2 Yield Prediction Output

The Random Forest model returns a single numerical estimate of expected yield per unit area, derived from averaging the predictions of its constituent trees. Alongside the point estimate, the relative importance of each input feature can be extracted from the trained ensemble, which offers some insight into which agronomic factors most strongly influenced a particular prediction. Farmers are advised to treat this figure as a planning aid rather than a guaranteed outcome, since actual yield will still depend on conditions that arise after the prediction is made, such as unexpected pest pressure or extreme weather events.

5.4 Decision Tree for Fertilizer Recommendation

The fertilizer-recommendation module uses a single Decision Tree classifier trained to map crop type and soil attributes, such as nitrogen, phosphorus, and potassium levels along with soil pH, to a recommended fertilizer category. The tree is grown by repeatedly splitting the training data on the attribute that most effectively separates the fertilizer classes at each step, using an impurity measure such as Gini index or information gain, and splitting stops once a stopping criterion, such as a minimum node size or maximum depth, is reached to reduce the risk of overfitting. The resulting tree structure can be traced from root to leaf, which allows the recommendation to be explained in terms of the specific soil and crop conditions that led to it. [6]

Workflow: Crop & Soil Information → Input Processing → Decision Rules → Fertilizer Recommendation

5.5 Weather API Integration

The weather module retrieves current conditions and short-term forecast information from a third-party Weather API based on the location supplied by the farmer, either through manual entry or device-based geolocation. Retrieved data typically includes temperature, humidity, precipitation probability, and wind conditions, all of which can influence decisions such as the timing of irrigation, pesticide spraying, or harvesting. To reduce the number of external API calls and improve responsiveness, weather responses for a given location can be cached for a short interval before being refreshed, and the module is designed to degrade gracefully by displaying the most recently retrieved data if the external service is temporarily unavailable. **5.6 System Architecture**

5.7 Model Development and Evaluation

The development pipeline for each model follows a broadly similar sequence: data collection, cleaning and preprocessing, model training, validation, testing, and finally integration into the shared platform. For the CNN classifier, a labelled collection of crop leaf images is divided into training, validation, and testing subsets, with the validation subset used to tune hyperparameters such as learning rate and network depth before final testing. Agricultural datasets used for the Random Forest yield model are similarly partitioned, and cross-validation is used to obtain a more reliable estimate of prediction error than a single train-test split would provide. The Decision Tree fertilizer model is evaluated on held-out crop and soil records to confirm that its recommendations remain consistent with expected agronomic practice. The integrated platform as a whole can be assessed using classification accuracy and confusion-matrix analysis for the disease module, error metrics such as mean absolute error for the yield module, recommendation consistency for the fertilizer module, and general measures of usability and system response time for the platform as a whole.

5.8 Data Preprocessing and Augmentation

Image data intended for the CNN module is preprocessed through resizing to a fixed input resolution, pixel-value normalization, and augmentation techniques such as random rotation, horizontal flipping, brightness adjustment, and slight zooming. These augmentation steps artificially expand the effective size and

diversity of the training set, which helps the network generalize better to leaf images captured under varying lighting conditions, camera angles, and background clutter that are typical of real field photography rather than controlled laboratory imagery. Tabular data used by the Random Forest and Decision Tree modules is cleaned to handle missing values, checked for inconsistent units, and, where necessary, normalized or encoded so that it is compatible with the respective model's input requirements.

5.9 System Implementation Considerations

The platform is conceived as a layered system consisting of a frontend interface accessible through a web browser or mobile application, a backend service layer responsible for coordinating requests to the CNN, Random Forest, Decision Tree, and Weather API components, and a data storage layer for persisting user records, historical predictions, and model artifacts. Python-based machine learning frameworks are well suited to training and serving the CNN and ensemble models, while a lightweight web framework can expose these models through application programming interfaces that the frontend consumes. Structuring the system in this layered manner allows individual components, such as replacing the disease-detection model with an improved version, to be updated without requiring changes to the rest of the platform.

5.10 Security and Privacy Considerations

Because Crop Vision AI processes farmer-submitted images and agronomic details, basic safeguards are considered as part of the design rather than as an afterthought. These include transmitting data over encrypted connections, restricting access to stored images and personal details to authorized platform components, and avoiding retention of identifying information beyond what is required to deliver and improve the service. Farmers are also expected to be informed about how their submitted data may be used, particularly if it is intended to contribute to future model retraining.

6. RESULTS AND DISCUSSION

6.1 Results

The platform is designed to generate four coordinated outputs from a single farmer interaction: a predicted disease category or healthy status for the submitted leaf image, an estimated crop yield, a recommended fertilizer category, and current weather information for the specified location. The CNN module classifies the uploaded image into the most probable disease class, the Random Forest module returns a numerical yield estimate, the Decision Tree module returns a fertilizer recommendation with its underlying rule path, and the Weather API supplies locationspecific conditions. Presenting these four outputs together, rather than through separate applications, is the primary functional result of the proposed design.

6.2 Discussion

The combination of three distinct machine learning paradigms, a convolutional network for image classification, an ensemble regressor for numerical prediction, and a single interpretable tree for rule-based recommendation, illustrates how different algorithmic strengths can be matched to different agricultural

subproblems within one platform rather than forcing a single model to address all of them. The CNN component contributes strong performance on unstructured image data, the Random Forest component contributes robustness against noisy or incomplete numerical inputs, and the Decision Tree component contributes interpretability that supports farmer trust in the fertilizer recommendation. The external Weather API contributes real-time situational context that none of the trained models could otherwise provide. Overall system usefulness remains closely tied to the quality and representativeness of the underlying training data, and validation under genuine field conditions, rather than curated benchmark datasets alone, would be an important next step in assessing practical performance.

6.3 Integrated User Workflow

In a typical interaction, a farmer opens the platform, uploads a photograph of an affected or healthy leaf, enters relevant agronomic details such as crop type and soil information, and specifies or allows detection of their location. The leaf image is routed to the CNN model while the numerical inputs are simultaneously routed to the Random Forest and Decision Tree models, and the Weather API is queried for the specified location. Once all three model responses and the weather data are available, the platform consolidates them into a single results view for the farmer to review, avoiding the need to revisit separate screens or tools for each piece of information.

6.4 Expected Benefits

By consolidating four previously separate agricultural functions into one workflow, the platform is expected to reduce the time and effort required to gather comparable information from disconnected sources. Farmers can review disease status, yield expectations, fertilizer guidance, and weather conditions in one place, which may improve both the speed and consistency of resulting field decisions, particularly for users who would otherwise lack convenient access to several specialized tools or human experts.

6.5 Comparative Perspective on Existing Systems

Relative to single-purpose applications documented in the literature, such as disease-only diagnostic tools or weather-only forecasting services, Crop Vision AI does not claim superior accuracy in any individual task; rather, its distinguishing characteristic is the breadth of decision support offered within one interface. Where a farmer might otherwise need to consult a disease-identification app, a separate agronomic advisory service, and a general weather application, the proposed platform is intended to reduce this to a single point of interaction, which is the primary practical advantage highlighted by this project relative to the fragmented tools described in the reviewed literature. [7], [8], [13]

7. LIMITATIONS

7.1 Dataset Dependency

The accuracy of every model in the platform is bounded by the size, quality, and diversity of the datasets used to train it. Datasets that are small, imbalanced across classes, or captured under narrow conditions may cause the trained models to perform

noticeably worse on inputs that differ from their training distribution.

7.2 Disease Coverage

The CNN classifier can only recognize disease categories that were represented during training. Diseases that are new, rare, region-specific, or visually similar to categories already in the training set may be misclassified or not recognized at all, which limits the diagnostic scope of the current implementation.

7.3 Prediction Uncertainty

Crop yield is influenced by numerous interacting factors, including weather variability, soil condition, farming practices, pest pressure, and crop variety, not all of which can be captured in a fixed set of input parameters. The yield figure produced by the Random Forest model should therefore be interpreted as an informed estimate rather than a guaranteed outcome.

7.4 Weather Dependency

The weather module depends entirely on the continued availability and accuracy of an external Weather API and on stable network connectivity at the point of use. Interruptions to either the external service or the farmer's internet connection will directly affect the availability of this component.

7.5 Expert Validation

Outputs from the disease-detection and fertilizer-recommendation modules are intended to function as decision-support information rather than as definitive agronomic instruction. For consequential decisions, particularly those involving significant financial outlay or the risk of crop loss, confirmation from a qualified agricultural expert remains advisable.

7.6 Computational and Infrastructure Constraints

Training and, to a lesser extent, serving the CNN model require adequate computational resources, and deployment in regions with limited internet bandwidth may affect the responsiveness of image upload and analysis. These constraints are relevant to how widely and how quickly the platform can be made available in low-connectivity rural settings.

7.7 Overall Limitation

The limitations described above reflect the current developmental stage of Crop Vision AI rather than fundamental barriers to its underlying approach. Further dataset expansion, model refinement, controlled field testing, and structured expert validation are identified as necessary steps before the platform could be considered ready for large-scale deployment.

8. FUTURE SCOPE

Several directions are identified for extending Crop Vision AI beyond its present form. The set of supported crops and disease categories can be broadened by incorporating additional, geographically diverse image datasets, which would also help address some of the dataset-dependency limitations noted above. Direct mobile-camera integration would allow farmers to capture and submit leaf images without needing a separate device or filetransfer step, improving convenience in field conditions.

Integration with low-cost Internet of Things soil sensors could supply the yield-prediction and fertilizer-recommendation

modules with real-time soil moisture, temperature, and nutrient readings, reducing reliance on manually entered or estimated values. Multilingual interface support would extend accessibility to farmers who are more comfortable in regional languages than in English, which is particularly relevant given the platform's intended rural user base. Improved weather-based recommendations, such as alerts tied specifically to spraying or irrigation windows, could turn the weather module from a passive information display into a more active advisory feature.

The incorporation of explainable AI techniques, such as visual attention maps for the CNN classifier or clearer featureimportance summaries for the Random Forest model, could help farmers and agricultural advisors understand why a particular disease, yield estimate, or fertilizer recommendation was produced, which may in turn increase trust in the system's outputs. Finally, formal integration with local agricultural extension services and expert consultation channels is identified as a way to combine the speed of automated recommendations with the accountability of human expert review.

9. CONCLUSION

This report has presented an expanded description of Crop Vision AI, an integrated smart farming assistant that combines CNNbased crop disease detection, Random Forest yield prediction, Decision Tree fertilizer recommendation, and real-time weather retrieval within a single platform. Rather than treating these as separate tools, the proposed design routes farmer-submitted images and agronomic details through a coordinated pipeline so that disease, yield, fertilizer, and weather information can be reviewed together in one interaction. The approach is grounded in established techniques documented across the reviewed literature, and its principal contribution lies in the integration of these techniques into one farmer-facing workflow rather than in the introduction of an entirely new algorithm.

While the system's practical usefulness remains dependent on the quality of its underlying datasets, the robustness of its trained models, the continued availability of external services such as the Weather API, and eventual validation under real field conditions, the work described here provides a structured foundation on which a more comprehensive AI-assisted farming platform can be developed. Continued dataset expansion, model refinement, and collaboration with agricultural experts are identified as the most direct paths toward realizing that broader goal.

ACKNOWLEDGEMENT

The authors wish to express their sincere gratitude to Dr. Akhilesh Tayade for his continued guidance and constructive feedback throughout the development of this project. The authors also thank the Department of Computer Science and Engineering, P. R. Pote Patil College of Engineering & Management, Amravati, for providing the resources and environment necessary to carry out this work.

10. 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, Art. no. 1419, 2016.
- [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:1511.08060*, 2015.
- [3] K. P. Ferentinos, "Deep Learning Models for Plant Disease Detection and Diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
- [4] E. C. Too, Y. Li, S. Njuki, and L. Yingchun, "A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.
- [5] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, pp. 5–32, 2001.
- [6] J. R. Quinlan, "Induction of Decision Trees," *Machine Learning*, vol. 1, pp. 81–106, 1986.
- [7] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep Learning in Agriculture: A Survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
- [8] S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, "Big Data in Smart Farming – A Review," *Agricultural Systems*, vol. 153, pp. 69–80, 2017.
- [9] M. A. Khan et al., "Plant Disease Detection and Classification Using Deep Learning," *IEEE Access*, 2020.
- [10] A. Kamilaris, "Deep Learning in Agriculture: A Survey," *Computers and Electronics in Agriculture*, 2018.
- [11] J. G. A. Barbedo, "Impact of Dataset Size and Variety on the Effectiveness of Deep Learning and Transfer Learning for Plant Disease Classification," *Computers and Electronics in Agriculture*, 2018.
- [12] S. Sladojevic et al., "Deep Neural Networks Based Recognition of Plant Diseases by Leaf Image Classification," *Computational Intelligence and Neuroscience*, 2016.
- [13] M. P. S. Chawla et al., "Machine Learning Applications in Smart Agriculture: A Review," *International Journal of Advanced Computer Science and Applications*, 2021.
- [14] Food and Agriculture Organization, "Digital Technologies in Agriculture and Rural Areas," *FAO Publications*, 2019.
- [15] World Bank, "Agriculture, Food and Rural Development: Digital Agriculture and Data-Driven Farming," *World Bank Reports*, 2021.