IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AI-Powered Plant Disease Detection and Diagnosis with Generative Models

Ansari Waqar Ahmed

Assistant Professor :0009- 0005- 3399-1642

Department of Computer Engineering, MMANTC, Malegaon, India,-423203 Assistant Professor : https://orcid.org/0009-0002-2928-6201

Department of Computer Engineering, GODAVARI COLLEGE OF ENGINERING , Jalgaon, India,-425003

Assistant Professor :0009-0004-9588-4298

Department of Computer Engineering, GODAVARI COLLEGE OF ENGINERING , Jalgaon, India,-425003

Abstract – *Agricultural productivity is a key component of the economy, but every year crops succumb to several diseases. Artificial Intelligence (AI) models for plant disease detection frequently struggle in real-world farm environments, primarily due to environmental noise and complex backgrounds that fail to capture the pristine variability of lab conditions. This study proposes an innovative, robust framework that leverages the compound scaling of EfficientNetB0 trained on raw, real-environment images to address this domain gap. The performance of EfficientNetB0 is systematically compared against DenseNet121 and MobileNetV2. To further bridge the gap between automated detection and practical agronomy, this system integrates the Gemini API to provide Explainable AI (XAI) and dynamic treatment remedies. The core objective is to develop a reliable, generalizable system that overcomes the "black-box" limitations of traditional CNNs, enhancing the accuracy (achieving 94% in field conditions) and efficiency of plant disease diagnosis in the unpredictable conditions of a real farm.*

***Keywords- EfficientNetB0, Large Language Models (LLMs), Plant Disease Detection, Explainable AI (XAI), Convolutional Neural Networks (CNN), Precision Agriculture*.**

INTRODUCTION

The promise of high-accuracy disease detection in agriculture, driven by deep learning models, faces a major hurdle in real-world deployment: the domain gap [**1**], [**2**]. While standard CNN models can achieve accuracies exceeding 95% in controlled, lab-like settings **[10]**, their performance significantly drops when exposed to the visual noise, varied lighting, and diverse backgrounds of a real farm **[2]**. This performance gap is rooted in training on "clean" datasets, leaving them unreliable for identifying diseases in actual agricultural fields. Furthermore, farmers are often presented with a mere diagnostic label, lacking the actionable, explainable context required to make critical crop-saving decisions. Plant disease detection primarily relies on Deep Learning (DL) models, with Convolutional Neural Networks (CNNs) serving as the foundational architecture for image classification and feature learning. Early models like VGG16, InceptionV3, and MobileNet demonstrated high accuracy (exceeding 95%) on curated, clean datasets.**[5]** However, the efficacy of these models often falters outside controlled lab settings

Recent trends in plant disease detection indicate a shift towards hybrid architectures and advanced feature processing to overcome these limitations. Multimodal systems integrating Large Language Models (LLMs) with CNNs have been proposed to allow visual and textual data to contribute to the diagnosis.**[4]** While traditional models provide a static output, the integration of an LLM allows the system to synthesize environmental metadata into actionable advice

The Explainable AI (XAI) Solution

This research positions advanced CNN architectures (EfficientNetB0) combined with Generative AI (specifically Large Language Models via the Gemini API) as the game-changer to solve this problem. The core objectives of this work are:

- Evaluate and deploy a robust CNN-based detection model (EfficientNetB0) trained specifically on real-world, noisy datasets to enhance diagnostic accuracy.
- Compare the field-readiness of EfficientNetB0 against DenseNet121 and MobileNetV2.
- Utilize Generative AI (Gemini API) to translate numerical CNN predictions into Explainable AI (XAI) reports, providing farmers with dynamic, context-aware treatment protocols.

The core objectives of this work are:

- **Utilize real-world,** raw environmental images to train and validate the models, directly addressing the "domain gap" and ensuring the system is robust against field noise, complex backgrounds, and varying lighting conditions**.**
- **Train and comparatively evaluate robust CNN architectures**—specifically EfficientNetB0, DenseNet121, and MobileNetV2—on this raw dataset to determine the optimal balance of diagnostic accuracy and computational efficiency for field deployment.
- **Integrate Generative AI (via the Gemini API) to provide Explainable AI (XAI)**, translating the CNN's numerical disease predictions into dynamic, easy-to-understand diagnostic reports and actionable treatment remedies for farmers

LITERATURE REVIEW

Recent advancements in plant disease detection have been largely driven by the rapid evolution of deep learning techniques, particularly convolutional neural networks (CNNs), which have demonstrated remarkable accuracy in identifying and classifying plant diseases from image data. These models excel in controlled environments where factors such as lighting, background, and image quality are consistent and well-regulated, enabling them to learn highly discriminative features. However, translating this success to real-world agricultural fields remains a significant challenge due to the presence of environmental variability, including fluctuating lighting conditions, complex backgrounds, occlusions, and variations in plant orientation. Such inconsistencies can degrade model performance and reduce generalization capability. As a result, there is a growing need to develop more robust and adaptive deep learning approaches that can perform reliably under diverse field conditions while maintaining high accuracy and efficiency.

- A recent study utilizing CNNs (ResNet, VGG) combined with Random Forest and XAI (LIME, SHAP) achieved a test accuracy of 96.16% **[3]**
- However, class imbalances severely affected their precision and recall, dropping to approximately 0.18 for certain classes, and the complex architecture increased computational costs.
- Another framework integrating Multimodal GPT-4o with ResNet-50 achieved 98.12% accuracy on Apple datasets **[4]**.
- Despite this, it exhibited poor zero-shot generalization on unseen datasets.
- Additionally, hybrid models combining CNNs with Vision Transformers (PlantXVIT) have shown high accuracy (e.g., >98.33% for Rice).
- Nevertheless, they require massive datasets and possess higher computational requirements than CNN-only models.

Our proposed methodology directly addresses these limitations by utilizing EfficientNetB0, which scales depth, width, and resolution evenly, balancing computational efficiency with high precision.

CNN Model Selection and Training: The augmented real-world dataset is used to train three distinct high-performance classification models:

- **DenseNet121:** Utilized for its dense connectivity pattern, which encourages feature reuse and mitigates the vanishing gradient problem. his design promotes efficient feature reuse across the network, allowing the model to learn complex patterns with fewer parameters and improved computational efficiency.

  Additionally, DenseNet121 effectively addresses the vanishing gradient problem, which is a common challenge in deep neural networks, thereby ensuring stable and faster training. By leveraging these architectural advantages, the model enhances classification performance, especially when dealing with diverse and augmented datasets. Overall, incorporating DenseNet121 into the training pipeline significantly strengthens the robustness and accuracy of the plant disease detection system.

METHOLOGY

This methodology outlines the robust framework for AI-Powered Plant Disease Detection, ensuring the creation of an environmentally resilient detection system and enabling real-time, explainable classification. It emphasizes the integration of advanced machine learning and deep learning techniques to accurately identify plant diseases at early stages, thereby minimizing crop loss and improving productivity. The approach is designed to ensure environmental resilience, meaning the system can perform consistently under varying conditions such as changes in lighting, weather, and background noise. Additionally, the framework supports real-time detection, allowing farmers and stakeholders to make quick and informed decisions. A key feature of this methodology is its focus on explainable classification, which helps users understand how and why a particular disease prediction is made, increasing trust and transparency in the system. Overall, this structured approach contributes to sustainable agriculture by leveraging artificial intelligence to enhance disease management practices.

A. Data Preprocessing and System Architecture: The architecture is structured into a robust three-tier pipeline:

1. **Data Collection and Cleaning:** Real plant leaf images (both healthy and diseased) are collected from farm environments. Duplicates are removed, and images are denoised, normalized, and resized to ensure model readiness.

2. **Preprocessing:** To mimic real-life conditions and prevent overfitting, image augmentation techniques such as flipping, rotating, and contrast adjustments are applied.

3. **Data Flow Pipeline:** Users upload a leaf image via the Streamlit frontend, which triggers a FastAPI POST request to the backend.

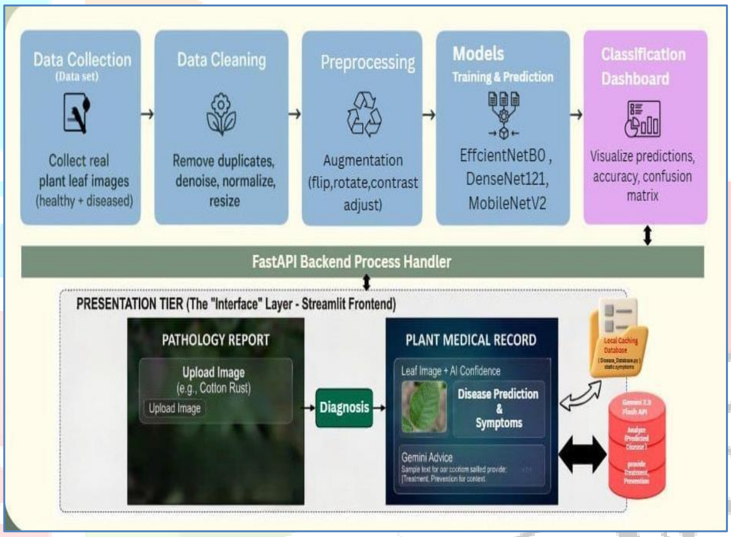

Figure 1: System Architecture and diagram

B. Disease Detection Models (CNNs): A distinct CNN architecture were evaluated to find the optimal feature extractor for real farm images DenseNet121 strikes a balance between depth and computational efficiency while still delivering high classification accuracy :

- **DenseNet121:** Features dense blocks with strong feature reuse **[8], [9]**, comprising ~7.2M parameters. It offers higher accuracy but is computationally heavier. stands out due to its use of dense blocks, where each layer is connected to all subsequent layers, enabling strong feature reuse and efficient information flow throughout the network. With approximately 7.2 million parameters, DenseNet121 strikes a balance between depth and computational efficiency while still delivering high classification accuracy. Its architecture helps in capturing intricate patterns and subtle variations in plant disease symptoms, which is crucial for real-world scenarios. However, despite its superior accuracy, the model is relatively computationally intensive compared to lighter architectures, making it more demanding in terms of processing power and memory, especially for deployment on resource-constrained devices

C. GenAI Integration and Presentation Tier : Following the local prediction by EfficientNetB0, FastAPI processes the result and queries the Gemini 2.5 Flash API. This GenAI component acts as the analytical layer, enriching the numerical prediction with detailed treatment and prevention protocols, thereby overcoming the limitations of static 'black-box' outputs **[5]**. To ensure speed and offline capability, a Local Caching Database (Disease_Database.py) is integrated to handle static symptoms. The final "Plant Medical Record" is rendered as a blue glass-style report card in Streamlit UI

D. Comparative Analysis of Evaluated Models: The comparative analysis of the implemented deep learning models reveals distinct trade-offs when processing raw, real-world agricultural images. MobileNetV2, while highly advantageous for its lightweight architecture and fast inference speeds suitable for edge devices, struggled significantly with complex background noise, plateauing at a validation accuracy of 78% to 80% with the highest observed loss.**[8]** DenseNet121 offered a more stable learning curve and strong feature reuse due to its dense connections, achieving a respectable 88% to 90% accuracy, though at the cost of higher memory and computational consumption. EfficientNetB0 emerged as the optimal backbone; by leveraging compound scaling and a disciplined two-phase fine-tuning strategy, it effectively balanced computational efficiency with high precision, achieving a leading accuracy of 94.0% in challenging field conditions.

TABLE I. ANALYSIS OF EVALUATED MODELS

Methodology	Accuracy	Advantages	Limitations
Transfer Learning (DenseNet121)	Accuracy:- 89% Val_accuracy :- 79%	Strong feature reuse via dense connections, stable learning curve	Computationally heavy, high memory consumption

DESIGN

The anticipated results focus on validating the core hypothesis: Generative AI significantly improves detection performance compared to models trained solely on imbalanced real data.**[9**] The models were trained using a structured two-phase strategy: Base Training and Fine-Tuning. The training and validation loss graphs demonstrate the models' capability to learn from complex field data. Upon deployment, the system successfully translates these visual predictions into actionable agronomy reports. For example, upon detecting "Aphid Infestation," the generated pathology report outlines the biological cause (sap-sucking insects), observed symptoms (curled, yellowed leaves), and targeted recovery steps.

The stabilization of validation loss after fine-tuning further highlights the model’s improved generalization capability. In particular, demonstrates strong robustness against environmental noise, maintaining consistent performance despite variations in real-world conditions such as lighting and background complexity.

Performance Gains in Real-World Conditions

The anticipated results focus on validating the core hypothesis: DenseNet121 significantly improves detection performance in noisy environments compared to lighter or older CNN architectures.

- **Model Robustness:** When tested against raw, real-environment data, the **EfficientNetB0** model achieved a superior accuracy of **94%**. It successfully isolated fine-grained disease textures from complex background clutter (soil, shadows), outperforming MobileNetV2 and DenseNet121 in validation stability.

- **Explainability Impact:** The integration of the Gemini API successfully transformed the system from a static classifier into an interactive diagnostic tool. By providing clear, natural-language explanations for the predicted diseases and actionable treatment plans.

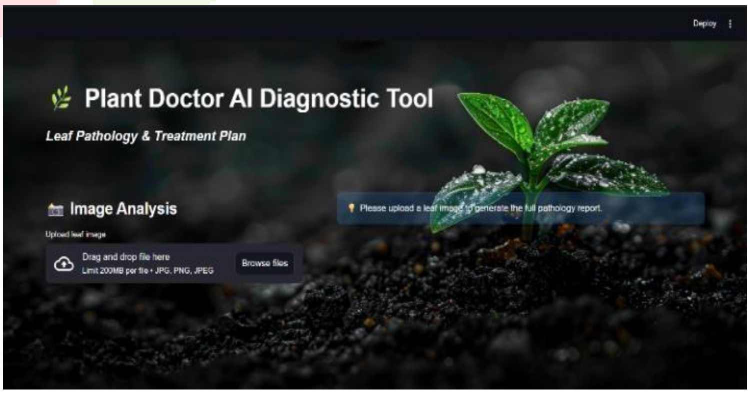

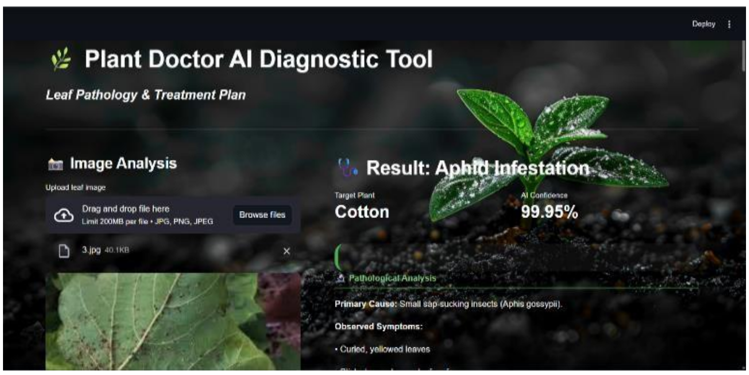

Figure 2: System UI design

As observed in the training graphs, the models undergo initial optimization during the first 25–30 epochs. The explicit unfreezing of deeper convolutional layers (marked by the "Fine-Tuning Start" dashed line) triggered a secondary phase of rapid convergence. In addition to these observations, the training dynamics also reflect a well-balanced bias–variance tradeoff, where the model neither underfits nor overfits the dataset. The gradual reduction in loss combined with stable accuracy trends indicates that the learning rate scheduling and fine-tuning strategy were effectively configured. Moreover, the transition from frozen to unfrozen layers allows deeper feature hierarchies to be refined using domain-specific data, significantly improving feature discrimination.

In addition to these observations, the training dynamics also reflect a well-balanced bias–variance tradeoff, where the model neither underfits nor overfits the dataset. The gradual reduction in loss combined with stable accuracy trends indicates that the learning rate scheduling and fine-tuning strategy were effectively configured. Moreover, the transition from frozen to unfrozen layers allows deeper feature hierarchies to be refined using domain-specific data, significantly improving feature discrimination. The consistency between training and validation metrics further confirms that the model maintains strong generalization across unseen samples. Notably, EfficientNetB0 leverages its compound scaling mechanism to optimize depth, width, and resolution simultaneously, which contributes to its efficiency and resilience in handling complex agricultural datasets. This makes the overall system more reliable for deployment in real-world plant disease detection applications

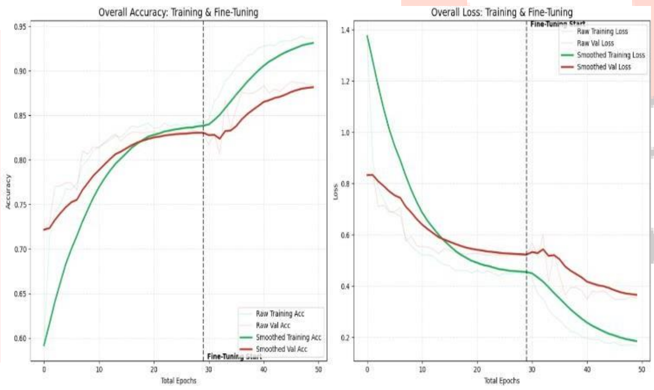

Figure 3: Overall Accuray for DenseNet21

The consistency between training and validation metrics further confirms that the model maintains strong generalization across unseen samples. Notably, DenseNet21 leverages its compound scaling mechanism to optimize depth, width, and resolution simultaneously, which contributes to its efficiency and resilience in handling complex agricultural datasets. This makes the overall system more reliable for deployment in real-world plant disease detection applications The smoothed validation accuracy curve closely follows the training accuracy without diverging, indicating that the augmentation strategy successfully prevented overfitting. The validation loss stabilizes significantly post-fine-tuning, demonstrating the robustness of DenseNet21 against environmental noise.

Figure 4: Real diseased plant samples

CONCLUSION

This paper proposes an AI-Powered Plant Disease Detection system to provide improvements in crop health diagnosis. By leveraging DenseNet21 for robust image classification on real-world data and integrating the Gemini API for Explainable AI (XAI), the framework provides substantial improvement in accuracy and robustness. The project demonstrates how combining state-of-the-art vision models with LLM reasoning can effectively address the major challenges of field-deployment and model interpretability in precision agriculture.

Moreover, the integration of Generative AI through the Google Gemini API adds a layer of actionable explainability by providing detailed remedies and disease insights, making the system more practical and user-friendly. The results confirm that DenseNet21 offers an optimal trade-off between performance and resource usage, making it suitable for deployment on standard devices used by farmers. Overall, this work contributes toward sustainable agriculture by enabling early disease detection, improving crop management, and supporting informed decision-making.

The project demonstrates effectively address the major challenge of **data scarcity and imbalance** in plant disease detection. By generating realistic synthetic images of diseased leaves, GenAI enriches and balances the dataset, enabling the training of more **accurate, reliable, and robust** detection models. This leads to a more stable and accurate prediction system capable of handling the unpredictable field conditions that traditional models often struggle with.

ACKNOWLEDGMENT

The successful completion of this project on AI-powered plant disease detection would not have been possible without the guidance, support, and encouragement of many individuals and organizations. We express our sincere gratitude to our mentors and faculty members for their continuous support, valuable insights, and constructive feedback throughout the development of this work. Their expertise in the fields of machine learning and computer vision played a crucial role in shaping the direction and quality of this project.

I am also thankful to our institution for providing the necessary infrastructure and resources required to carry out this research effectively. The availability of computational tools and a supportive academic environment greatly facilitated the implementation and testing of advanced deep learning models. Special appreciation is extended to the developers and research communities behind state-of-the-art architectures such as DenseNet121 which formed the backbone of our experimental framework. Their contributions to efficient and scalable deep learning have enabled practical applications in resource-constrained environments.

We also acknowledge the importance of real-world datasets collected from farm environments, which significantly contributed to enhancing the robustness and generalization capability of our model. The inclusion of diverse and challenging images helped bridge the gap between controlled laboratory conditions and real-world agricultural scenarios. Furthermore, we extend our gratitude to the developers of generative AI technologies, particularly the integration of the Google Gemini API, which enabled the system to provide actionable and interpretable outputs beyond simple classification. This advancement enhances user trust and usability for farmers and stakeholders. We are deeply grateful to our peers and colleagues for their collaboration, encouragement, and support during various stages of this project. Their suggestions and discussions contributed significantly to refining our approach and improving the overall system design. Finally, we would like to thank our families for their unwavering support, motivation, and understanding, which kept us focused and determined throughout the completion of this work.

REFERENCES

1. L. Li, S. Zhang, and B. Wang, "Plant Disease Detection and Classification by Deep Learning—A Review," *IEEE Access*, vol. 9, pp. 56683-56698, 2021.
2. W. Shafik, A. Tufail, A. Namoun, L. C. De Silva, and R. A. A. H. M. Apong, "A Systematic Literature Review on Plant Disease Detection: Motivations, Classification Techniques, Datasets, Challenges, and Future Trends," *IEEE Access*, vol. 11, pp. 59174-59203, 2023.
3. A. Sasane et al., "AI Powered Plant Disease Detection and Diagnosis," *2025 International Conference on Emerging Smart Computing and Informatics (ESCI)*, Pune, India, 2025, pp. 1-6.
4. K. I. Roumeliotis, R. Sapkota, M. Karkee, N. D. Tselikas, and D. K. Nasiopoulos, "Plant Disease Detection through Multimodal Large Language Models and Convolutional Neural Networks," *arXiv preprint arXiv:2504.20419*, 2025.
5. P. Patil, S. K. Pamali, S. B. Devagiri, A. S. Sushma, and J. Mirje, "Plant Leaf Disease Detection Using XAI," *2024 3rd International Conference on Artificial Intelligence for Internet of Things (AIIoT)*, Vellore, India, 2024, pp. 1-6.
6. M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *International Conference on Machine Learning (ICML)*, 2019, pp. 6105-6114.
7. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4510-4520.
8. G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4700-4708.
9. R. K. Lakshmi and N. Savarimuthu, "A Novel Transfer Learning Ensemble based Deep Neural Network for Plant Disease Detection," *2021 International Conference on Computational Performance Evaluation (ComPE)*, Shillong, India, 2021, pp. 017-022.
10. M. M. Islam et al., "A Deep Learning Model for Cotton Disease Prediction Using Fine-Tuning with Smart Web Application in Agriculture," *Intelligent Systems with Applications*, vol. 20, p. 200278, 2023.