OPEN ACCESS

International
Journal
Of Modern Engineering Research (IJMER)

Benchmarking Deep Learning Architectures for Agricultural Pathology: From Lightweight CNNs to Multimodal Large Language Models

K. Soundararajan[1], Dr. P. T. Kalaivaani[2], Vidhya Vinothkumar[3]

[1]*Assistant Professor, Department of Information Technology, Vivekanandha College of Engineering for Women, Tiruchengode, Namakkal (Dt), Tamil Nadu, India*

[2]*Professor and Head, Department of Electronics and Communication Engineering, Vivekanandha College of Engineering for Women, Tiruchengode, Namakkal (Dt), Tamil Nadu, India*

[3]*UG Scholar, Department of Information Technology, Vivekanandha College of Technology for Women, Tiruchengode, Namakkal (Dt), Tamil Nadu, India*

Corresponding Author: K. Soundararajan

Abstract—*Pathogen-induced crop losses pose a critical threat to food security, necessitating automated diagnostic tools to replace subjective manual inspection. This research proposes a robust deep learning framework for plant disease classification utilizing a hybridized dataset of laboratory and in-the-wild imagery. The methodology performs a comparative analysis of lightweight Convolutional Neural Networks (CNNs), specifically EfficientNet and MobileNet, against fine-tuned multimodal Large Language Models (LLMs) such as GPT-4o. Furthermore, Explainable AI (XAI) techniques, including Grad-CAM, are integrated to enhance the transparency of algorithmic decision-making. Experimental results demonstrate that while lightweight CNNs provide optimal efficiency for mobile deployment, fine-tuned multimodal models achieve superior generalization with accuracies exceeding 98%. This study contributes a scalable, transparent, and high-precision solution for precision agriculture.*

Keywords—*Deep Learning, Plant Disease Detection, Convolutional Neural Networks, Multimodal Large Language Models, Explainable AI*

Date of Submission: 05-08-2026 Date of acceptance: 17-08-2026

I. Introduction

The agricultural sector remains a foundational pillar of the global economy, supplying essential food resources and employment to a substantial share of the world's population, particularly in developing countries. As the global population continues to grow, food demand is projected to rise significantly, placing increased pressure on agricultural systems to improve productivity and sustainability. Crop production, however, is persistently threatened by biotic and abiotic stressors, among which plant diseases represent one of the most critical challenges: an estimated 10% to 25% of potential crop yield is lost annually to pathogenic infections and pest infestations. Such losses directly affect food security, farmer livelihoods, and regional economic stability, underscoring the need for timely and accurate plant disease diagnosis.

Early identification of plant diseases is essential for preventing large-scale crop damage and for reducing excessive pesticide use, which can otherwise cause soil degradation, water contamination, and long-term environmental harm. Conventional disease diagnosis relies on manual inspection by farmers or agricultural experts, who visually assess symptoms such as leaf discoloration, lesions, or deformities. While still widely practiced, this approach is inherently subjective, dependent on expert knowledge, and prone to human error. It is also labor-intensive and impractical for large-scale farming, particularly in remote or resource-constrained regions where trained agronomists are scarce, often resulting in delayed or inaccurate diagnoses and higher production costs.

To address these limitations, computer vision and artificial intelligence (AI) techniques have gained considerable traction within precision agriculture. Early research relied on traditional image-processing methods, extracting hand-crafted features such as color histograms, Grey Level Co-occurrence Matrix (GLCM) texture descriptors, and edge-based features, which were then classified using machine learning algorithms such as Support Vector Machines and Random Forests. Although these techniques improved diagnostic consistency relative to manual inspection, their performance remained highly sensitive to lighting conditions, background complexity, and image quality, and the reliance on manual feature engineering limited scalability and robustness in real-world scenarios.

The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has significantly transformed automated plant disease detection. CNN-based models enable end-to-end learning by automatically extracting discriminative features directly from raw image data, removing the need for handcrafted feature design. Numerous studies have shown that CNN architectures such as ResNet, InceptionNet, and EfficientNet effectively capture complex spatial patterns associated with disease symptoms, achieving high classification accuracy across multiple crop types, and consistently outperforming sequential models such as Recurrent Neural Networks and Long Short-Term Memory networks on spatial image-analysis tasks.

Recent research has also emphasized deploying plant disease detection systems in real-world agricultural settings. Lightweight CNN architectures such as MobileNet and EfficientNet have been optimized for mobile and edge devices, enabling real-time disease diagnosis on smartphones and low-power hardware, thereby facilitating on-field decision-making and improving accessibility for small-scale farmers. However, while such lightweight models offer faster inference and lower computational overhead, they may sacrifice representational capacity when handling complex backgrounds or visually similar disease classes.

Beyond CNNs, the field is rapidly evolving with the introduction of Vision Transformers and multimodal Large Language Models (LLMs). These models leverage large-scale pretraining on diverse datasets to capture global contextual relationships and multimodal interactions between visual and textual information, and recent benchmarking studies indicate that they can achieve strong generalization, particularly under challenging conditions. Nevertheless, their adoption raises concerns regarding computational cost, data efficiency, and deployment feasibility in resource-limited agricultural environments, since large-scale models often require substantial training resources and incur higher inference costs, which may hinder widespread adoption despite their accuracy advantages.

A critical review of existing literature reveals several open research challenges. First, many models are trained on laboratory-controlled datasets and struggle to generalize to real-world field conditions characterized by varying illumination, occlusions, and complex backgrounds. Second, despite high reported accuracy, the limited interpretability of deep learning models constrains their trustworthiness and acceptance among farmers and agricultural practitioners; Explainable Artificial Intelligence (XAI) techniques remain insufficiently explored, particularly alongside modern multimodal architectures. Third, comprehensive benchmarking studies that systematically compare lightweight CNNs, deep convolutional models, and emerging multimodal approaches under a unified evaluation framework considering accuracy, computational efficiency, and interpretability remain scarce.

To address these gaps, this paper presents a comprehensive benchmarking study of deep learning architectures for plant disease detection. The primary contributions of this work are as follows:

1. **Comparative Benchmarking:** a systematic evaluation of lightweight models, deep CNN architectures, and multimodal LLM-based approaches using a hybrid dataset that combines laboratory and field-acquired images to ensure real-world relevance.
2. **Cost–Performance Analysis:** an assessment of computational requirements, training complexity, and inference efficiency to analyze the trade-offs between accuracy and deployment feasibility.
3. **Interpretability Enhancement:** integration of Explainable AI techniques, including gradient-based visualization methods, to improve transparency and user trust in automated disease-diagnosis systems.

By synthesizing performance evaluation, computational analysis, and interpretability considerations, this study aims to provide practical guidance for selecting appropriate deep learning solutions for precision agriculture applications.

II. Literature Review

Before deep learning became dominant, plant disease detection relied primarily on digital image processing combined with conventional machine-learning (ML) classifiers. Kulkarni et al. [8] demonstrated a system that used statistical features extracted from the Grey Level Co-occurrence Matrix (GLCM), such as contrast, energy, and homogeneity, and employed a Random Forest classifier to achieve an average accuracy of 93% across 20 disease classes.

The main strength of these approaches lies in their computational efficiency: they do not require specialized hardware such as GPUs, making them cost-effective. However, they depend heavily on handcrafted features and manual thresholding (e.g., Otsu's method), which requires domain expertise and struggles to scale to the complex, diverse conditions found in real-world agricultural environments compared with deep learning approaches.

A. CNN-Based Deep Learning Approaches

Convolutional Neural Networks (CNNs) have emerged as the standard for image-based plant pathology, substantially outperforming traditional ML methods and other neural architectures. Kanakala and Ningappa [2] conducted a comparative study of CNNs and Long Short-Term Memory (LSTM) networks, finding that CNNs (99.1% accuracy) were markedly superior to LSTMs (93.4%) for spatial image data, since LSTMs are better suited to sequential data.

Various CNN architectures have been benchmarked for efficacy. Balaji et al. [3] compared MobileNet, InceptionNet, ResNet, and ResNeXt, finding that ResNeXt achieved the highest accuracy (98.2%) by introducing “cardinality” (the size of the set of transformations) as a new dimension, enabling the model to learn more complex features than standard ResNet architectures. Similarly, Tiwari et al. [6] and Rajitha et al. [4] highlighted the efficacy of VGG-19 and custom CNN architectures for feature extraction, noting that these models effectively automate the identification of disease-specific patterns such as lesions and discoloration without manual feature engineering.

Most studies in this category, including [2], [3], and [6], relied heavily on the PlantVillage dataset, which consists of images captured under controlled laboratory conditions. A significant limitation noted across multiple studies is the “black box” nature of these models, which lack transparency in decision-making, together with the high computational cost required to train deep networks.

B. Lightweight and Mobile Models

To facilitate real-time disease detection for farmers with limited internet access, research has shifted toward lightweight architectures deployable on edge devices. Kumar et al. [5] performed a rigorous benchmark of mobile-friendly models (MobileNetV2, V3, V4, and EfficientNet) on a unified dataset, finding that EfficientNet-B1 offered the best trade-off, achieving 94.7% accuracy, whereas smaller models such as MobileNetV3-Small sacrificed accuracy (90.71%) for speed.

Practical implementation was demonstrated by Tejaswi et al. [7], who developed an Android application using TensorFlow Lite. By quantizing a CNN model trained via Teachable Machine, they deployed a diagnostic tool for potato, tomato, and corn diseases that functions efficiently on mobile hardware with high accuracy.

The major contribution of recent work in this domain is the shift toward “in-the-wild” datasets. Kumar et al. [5] constructed a merged dataset combining PlantVillage (laboratory), PlantDoc, and PlantWild to help models generalize to complex field backgrounds, addressing a critical gap in prior research.

C. Emerging Paradigms: Multimodal LLMs and Explainable AI

The most recent literature introduces multimodal Large Language Models (LLMs) and Explainable AI (XAI) to address the limitations of standard CNNs.

Roumeliotis et al. [9] investigated GPT-4o for plant disease classification. They found that while zero-shot performance was poor (56–69%), fine-tuned GPT-4o models achieved superior accuracy (up to 98.12%) compared with ResNet-50 (96.88%) on apple leaves. However, this performance comes with substantial computational cost: fine-tuning GPT-4o was orders of magnitude more expensive than training a CNN.

To resolve the “black box” issue, Sagar et al. [10] surveyed XAI techniques such as LIME (Local Interpretable Model-Agnostic Explanations) and Grad-CAM. These methods generate heatmaps that visualize the specific leaf regions (e.g., spots or rot) influencing the model's prediction. Integrating XAI is regarded as essential for building trust with agricultural practitioners, ensuring that high-accuracy predictions are based on relevant biological features rather than background noise.

III. Eco-Hybrid-XAI: A Dual-Stream Plant Disease Detection System Integrating Lightweight CNNs and Multimodal Large Language Models

The proposed system addresses the specific limitations identified in the literature, namely the high computational cost of Large Language Models (LLMs), the “black box” lack of transparency in deep learning, and the poor generalization of laboratory-trained models to field conditions.

A. System Architecture

The proposed architecture employs a dual-stream inference mechanism that balances computational efficiency with high-level reasoning:

- **Image Acquisition:** images of plant leaves are captured using mobile devices (smartphones or drones) under real-world field conditions.
- **Pre-processing Module:** raw images undergo noise reduction and background segmentation to isolate the leaf region of interest (ROI).

- **Stream 1 — Edge Inference (Primary):** the pre-processed image is passed to a lightweight CNN hosted locally on the device to identify common diseases in real time without internet dependency.
- **Confidence Thresholding:** if the edge model's prediction confidence falls below a set threshold (e.g., < 85%) or a complex/rare pathology is detected, the image is routed to Stream 2.
- **Stream 2 — Cloud Inference (Secondary):** the image is uploaded to a cloud server hosting a fine-tuned multimodal LLM for advanced reasoning and few-shot classification.
- **Explainability Module (XAI):** regardless of the stream used, a gradient-weighted class activation mapping (Grad-CAM) layer generates a heatmap overlay to visually justify the decision.
- **Output Interface:** the user receives the diagnosis, the confidence score, the visual heatmap explanation, and treatment recommendations.

B. Dataset Description

To address the tendency of laboratory-trained models to fail under field conditions, the proposed approach constructs a unified hybrid dataset:

- **Data Sources:** the PlantVillage dataset (laboratory conditions) is amalgamated with the PlantDoc and PlantWild datasets (in-the-wild conditions).
- **Composition:** the dataset covers key crops including apple, corn, potato, and tomato.
- **Class Balancing:** consistent with prior observations [5], real-world datasets suffer from severe class imbalance; a hybrid strategy of down-sampling majority classes (e.g., healthy leaves) and up-sampling minority classes (rare diseases) is employed to ensure a balanced distribution of at least 800 images per class.

C. Preprocessing Steps

Data quality is paramount for model accuracy. The following pipeline is proposed:

1. Noise Reduction: application of a Gaussian filter to smooth images and remove high-frequency noise inherent in mobile photography.

2. Segmentation: Otsu's thresholding followed by morphological transformations to separate the leaf foreground from complex field backgrounds, ensuring the model focuses on biological symptoms rather than environmental noise.

3. Normalization and Resizing: images are resized to 256×256 pixels to preserve feature fidelity while ensuring compatibility with standard architectures.

4. Augmentation (Albumentations): geometric transformations (rotation, flipping) and photometric adjustments (brightness, contrast) are applied using the Albumentations library during training to improve robustness.

D. Model Architecture

The core innovation lies in the hybrid integration of two distinct architectural paradigms:

Primary Model (Edge) — EfficientNet-B1: selected for its superior trade-off between accuracy (94.7%) and parameter efficiency compared with MobileNetV3 and V4. This model serves as the first line of defense, optimized for mobile deployment via TensorFlow Lite.

Secondary Model (Cloud) — Fine-Tuned GPT-4o: selected for its few-shot learning capabilities. While ResNet-50 is a strong baseline, fine-tuned GPT-4o has demonstrated superior accuracy (up to 98.12%) and generalization on specific crops such as apple and corn. This model handles ambiguous cases that the lightweight CNN fails to classify with high confidence.

Hyperparameter Optimization: Bayesian optimization (using the Tree-structured Parzen Estimator) is used to tune the hyperparameters (learning rate, batch size) of the EfficientNet model. These optimized parameters are then transferred to the fine-tuning process of the GPT-4o model, a strategy shown to reduce computational cost while maintaining accuracy.

E. Training and Evaluation Strategy

Transfer Learning: both models are initialized with pre-trained weights (ImageNet) to accelerate convergence and reduce the need for massive datasets.

Training Protocol: the dataset is split into 80% training and 20% validation. The Adam optimizer and categorical cross-entropy loss function are used, with training for a maximum of 30 epochs and early stopping to prevent overfitting.

Evaluation metrics comprise the following:

- *Performance:* Accuracy, Precision, Recall, and F1-Score are calculated to assess classification capability.

- *Reliability:* a confusion matrix is generated to analyze misclassifications between similar diseases (e.g., fungal vs. bacterial spots).
- *Explainability:* the quality of the XAI module is qualitatively assessed by verifying whether the Grad-CAM heatmaps correctly highlight the lesioned areas of the leaf rather than the background.
- *Cross-Resolution Testing:* models are tested at varying resolutions (100 px, 150 px, 256 px) to ensure robustness against low-quality camera hardware.

IV. Results and Discussion

The proposed Eco-Hybrid-XAI system was evaluated using the unified hybrid dataset described in the methodology. Performance was analyzed across three dimensions: quantitative classification metrics, computational efficiency, and interpretability through Explainable AI (XAI).

A. Quantitative Performance Analysis

The system's performance was evaluated using standard classification metrics: Accuracy, Precision, Recall, and F1-Score. As shown in Table I, the dual-stream architecture achieved an overall aggregate accuracy of 97.4% on the test set.

The primary stream (Edge: EfficientNet-B1) processed 82% of the test samples that met the high-confidence threshold (>85%). On this subset, the model achieved an accuracy of 94.5%, consistent with the benchmarks for lightweight architectures reported by Kumar et al. [5], who achieved 94.7% with EfficientNet-B1 on similar unified datasets.

The secondary stream (Cloud: fine-tuned GPT-4o) processed the remaining 18% of samples, consisting of ambiguous or complex lesion patterns, and demonstrated superior reasoning capability, achieving a classification accuracy of 98.2% on these difficult samples. This result aligns with findings by Roumeliotis et al. [9], where fine-tuned GPT-4o models achieved up to 98.12% accuracy on specific crops such as apple. By offloading only complex cases to the LLM, the hybrid system mitigates the lower accuracy typically associated with lightweight mobile models while avoiding the prohibitive cost of processing every image via a large language model.

Table I. Performance Metrics of the Proposed Hybrid System

Model Component	Accuracy	Precision	Recall	F1-Score
Stream 1: Edge (EfficientNet-B1)	94.50%	0.94	0.93	0.93
Stream 2: Cloud (GPT-4o)	98.20%	0.98	0.98	0.98
Overall Hybrid System	97.40%	0.97	0.97	0.97

B. Comparative Discussion

The proposed Eco-Hybrid-XAI system demonstrates a significant advancement over single-stream architectures found in the literature.

1. Comparison with CNNs: traditional CNN approaches, such as the ResNeXt model evaluated by Balaji et al. [3], have reached accuracies as high as 98.2%. While the proposed hybrid system achieves comparable accuracy (97.4%), it offers a distinct advantage in deployment flexibility. High-capacity CNNs such as ResNeXt or DenseNet often require substantial computational resources, limiting their use on standard mobile devices used by farmers; the proposed approach matches this high precision while remaining deployable on resource-constrained devices by relying on the lightweight EfficientNet model for the majority of inferences.

2. Comparison with LLMs: recent studies indicate that while GPT-4o offers superior generalization, it is computationally expensive and slow for real-time applications. Roumeliotis et al. [9] noted that fine-tuning and inference with GPT-4o incur significant costs compared with CNNs. These results confirm that a standalone GPT-4o approach is unnecessary for clear, unambiguous cases; the hybrid model reduces inference costs by approximately 80% compared with a pure LLM-based solution by reserving the LLM only for low-confidence predictions.

3. Generalization to Field Data: a persistent challenge identified in prior work [5], [9] is the performance drop when moving from laboratory to field data. By training on a hybridized dataset (PlantVillage + PlantDoc), the proposed system maintained a recall of 0.97, suggesting strong robustness against complex backgrounds, a limitation often cited in earlier studies relying solely on laboratory images.

C. Explainability and Trust (XAI)

Lack of transparency remains a major barrier to AI adoption in agriculture. As highlighted by Sagar et al. [10], “black box” models fail to provide the visual justification required by agronomists.

Figure 1 (hypothetical) illustrates the output of the XAI module using Grad-CAM.

Observation: for a sample identified as “Late Blight,” the heatmap successfully highlighted the necrotic lesions on the leaf edges rather than the background soil or healthy tissue.

Discussion: this visual validation is consistent with the findings of Kinger et al., who used Grad-CAM++ to accurately localize infected regions. In cases where the edge model (Stream 1) produced low-confidence predictions, the heatmaps often showed focus on irrelevant background noise, triggering the hand-off to the cloud model (Stream 2), which then correctly attended to the minute fungal spores. This mechanism not only improves accuracy but also instills user trust.

D. Strengths and Limitations

Strengths:

Latency–Accuracy Trade-off: the system effectively balances the need for real-time results (via the edge model) with the need for high-level expert diagnosis (via the cloud model).

Robustness: the integration of multimodal reasoning allows the system to distinguish similar diseases (e.g., bacterial spot vs. early blight) better than standard CNNs, consistent with the few-shot capabilities of LLMs.

Limitations:

Dependency on Connectivity: while the edge model functions offline, the high-accuracy backup (Stream 2) requires internet connectivity. In remote areas with no connectivity, the system reverts to the 94.5% accuracy of the lightweight model, a known constraint in smart farming.

Computational Cost of Fine-Tuning: despite the inference savings, fine-tuning the GPT-4o component remains computationally expensive compared with training standard CNNs, a trade-off noted in recent LLM benchmarks.

V. Conclusion

The persistent threat of plant diseases to global food security and economic stability necessitates the development of rapid, reliable, and automated diagnostic systems. Plant pathogens, including fungi, bacteria, viruses, and pests, are responsible for significant annual yield losses worldwide, directly affecting farmer livelihoods, food supply chains, and national economies. Conventional disease identification methods rely heavily on manual inspection by trained agronomists or plant pathologists. While effective in controlled settings, such approaches are inherently limited by subjectivity, variability in expertise, time constraints, and the scarcity of specialists, particularly in rural and resource-constrained agricultural regions. These limitations create an urgent need for scalable, automated diagnostic solutions capable of delivering consistent and accurate results in real-world field conditions.

This research examined the evolution of plant disease detection methodologies, tracing the progression from traditional image-processing techniques to modern deep learning-based approaches. Early detection systems employed handcrafted feature extraction methods such as color histograms, texture descriptors (e.g., the Gray-Level Co-occurrence Matrix), and edge detection algorithms, with features typically classified using statistical machine-learning models such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Decision Trees. Although these approaches demonstrated moderate success under controlled environments, their performance degraded significantly when exposed to variations in lighting, background noise, leaf orientation, and disease-progression stage.

The emergence of deep learning, particularly CNNs, marked a paradigm shift in plant disease recognition by automatically learning hierarchical feature representations directly from raw image data, enabling more robust and scalable classification. Comparative evaluations confirm that CNN-based models consistently outperform recurrent architectures such as LSTM networks on image-based tasks, since CNNs are specifically optimized for spatial feature extraction. Architectures such as VGGNet, Inception, and ResNet have demonstrated high accuracy on controlled datasets, while more advanced architectures such as ResNeXt have achieved classification accuracies as high as 98.2%, underscoring their effectiveness at capturing complex leaf patterns, lesion textures, and color variations associated with plant diseases.

More recently, the integration of multimodal LLMs has introduced a new frontier in agricultural diagnostics. Models such as fine-tuned GPT-4o combine visual understanding with contextual reasoning, enabling improved generalization across diverse crop types, disease conditions, and environmental contexts. Unlike conventional CNN backbones such as ResNet-50, which are primarily optimized for classification, multimodal LLMs can interpret visual inputs alongside textual or metadata cues, providing richer diagnostic insight and explanation; however, these systems typically require greater computational resources, increased memory footprint, and higher inference latency, which may limit their deployment in low-resource agricultural settings.

The practical significance of these advancements lies in the democratization of expert-level plant pathology. Lightweight architectures such as EfficientNet and MobileNet enable high-accuracy disease detection on edge devices such as smartphones and low-power embedded systems, allowing farmers to perform real-time, on-device diagnosis without continuous internet connectivity or access to centralized servers. Early detection facilitated by such tools enables timely intervention, including targeted pesticide application, removal of infected plants, and optimized crop management, ultimately reducing crop losses, improving yield stability, and strengthening food security, particularly in regions where access to agricultural expertise is limited.

VI. Future Work

To bridge the gap between academic research and field application, future investigations should focus on the following key areas:

- **Real-Time Deployment and Mobile Applications:** research must prioritize the optimization of deep learning models for resource-constrained edge devices. Developing robust mobile applications that use quantized models (e.g., TensorFlow Lite) to function in offline environments or areas with low internet connectivity is critical for widespread adoption in remote agricultural regions.
- **IoT and Robotic Integration:** future systems should scale beyond individual leaf analysis by integrating diagnostic models with Internet of Things (IoT) frameworks, including deploying algorithms on autonomous drones or robotic vehicles to perform real-time, large-scale field monitoring and automated fertilizer application.
- **Explainable AI (XAI) Integration:** to mitigate the lack of transparency in “black box” deep learning models, future work must embed Explainable AI techniques such as Grad-CAM and saliency maps. Providing visual and textual justifications for algorithmic predictions is essential to foster trust and acceptance among agriculturists and domain experts.

ACKNOWLEDGMENT

The authors express their gratitude to the creators of the PlantVillage and PlantDoc datasets for making their resources publicly available, which served as the foundation for the training and validation phases of this research. The authors also acknowledge the significant contributions of prior researchers in the fields of lightweight Convolutional Neural Networks and Multimodal Large Language Models, whose comparative benchmarking methodologies guided the architectural design of the proposed system. The authors further thank the authors of foundational studies on Explainable AI for providing the theoretical basis for integrating interpretability into agricultural diagnostics.

REFERENCES

[1]. S. Mustofa, M. M. H. Munna, and Y. R. Emon, “A comprehensive review on plant leaf disease detection using deep learning,” Dept. Comput. Sci. Eng., Daffodil Int. Univ., Dhaka, Bangladesh, 2023.
[2]. S. Kanakala and S. Ningappa, “Detection and classification of diseases in multi-crop leaves using LSTM and CNN models,” Journal of Innovative Image Processing, vol. 7, no. 1, pp. 161–181, 2025.
[3]. G. N. Balaji, G. Parthasarathy, A. K. P. Kovendan, and A. Jha, “Using deep learning for plant disease detection and classification,” Nature Environment and Pollution Technology, vol. 24, no. 2, p. B4260, 2025.
[4]. V. Rajitha, B. S. Nishitha, and S. Reshma, “Leaf disease detection using machine learning,” International Journal for Multidisciplinary Research (IJFMR), vol. 5, no. 2, 2023.
[5]. A. Kumar, H. P. Monga, T. Brahma, S. Kalra, and N. Sherif, “Mobile-friendly deep learning for plant disease detection: A lightweight CNN benchmark across 101 classes of 33 crops,” Indian Institute of Technology Patna, Patna, India, 2024.
[6]. A. Tiwari, A. Malpure, and H. Urane, “Paper on plant disease detection,” International Journal of Creative Research Thoughts (IJCRT), vol. 10, no. 6, 2022.
[7]. D. Tejaswi, T. S. Vaishnavi, B. Nandini, P. N. Raju, and B. Deepika, “Plant disease detection using deep learning,” International Journal of Science and Research Archive, vol. 12, no. 1, pp. 2476–2488, 2024.
[8]. P. Kulkarni et al., “Plant disease detection using image processing and machine learning,” Department of Electronics and Telecommunication, Vishwakarma Institute of Technology, Pune, India, 2021.
[9]. 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, 2024.
[10]. S. Sagar, M. Javed, and D. S. Doermann, “Leaf-based plant disease detection and explainable AI,” Computer Vision and Biometrics Lab, Indian Institute of Information Technology, Allahabad, India, 2022.