Explainable AI for Precise Leaf Disease Diagnosis: A Comparative Study

S. Balanageshwara
Department of Electronics and Communication Engineering, Moodlakatte Institute of Technology, Kundapura, Karnataka, India
balanageshwara@mitkundapura.com

Varuna Kumara
Department of Electronics and Communication Engineering, Moodlakatte Institute of Technology, Kundapura, Karnataka, India
vkumarg.24@gmail.com

Manjunatha Badiger
Department of VLSI Design and Technology, NMAM Institute of Technology, Nitte (Deemed to be University), Nitte, Karnataka, India
badiger_manju@yahoo.com (corresponding author)

Akshatha Naik
Department of Electronics and Communication Engineering, Moodlakatte Institute of Technology, Kundapura, Karnataka, India
akshathanaik@mitkundapura.com

Received: 17 November 2025 | Revised: 30 November 2025, 5 December 2025, 14 December 2025, 4 January 2026, and 6 January 2026 | Accepted: 9 January 2026

 DOI: https://doi.org/10.48084/etasr.16335

ABSTRACT

Sustainable agriculture depends on the timely identification of plant diseases. While deep learning, particularly Convolutional Neural Networks (CNNs), has shown promise in leaf disease classification, its black-box nature undermines transparency and farmer trust. This study presents a framework integrating EfficientNet-B4/B7 with multiple Explainable AI (XAI) techniques—Grad-CAM++, Score-CAM, Integrated Gradients, and LIME—to enhance interpretability. The proposed framework was developed and tested on the PlantVillage dataset. Although the classification model itself demonstrated modest performance (AUC 0.53), highlighting challenges in generalization, the primary contribution of this study lies in the comparative analysis of XAI methods. A qualitative expert evaluation revealed Grad-CAM++ as the most consistent and visually coherent method for highlighting disease-related features. A prototype farmer-friendly GUI was developed to visualize input images, predicted classes, and XAI heatmaps. This work underscores the importance of interpretability in building trustworthy AI systems for agriculture and provides a comparative baseline for XAI techniques in plant pathology, establishing a foundation for future work on more robust classifiers.

Keywords-leaf disease detection; Explainable AI (XAI); EfficientNet; model interpretability; PlantVillage dataset

I. INTRODUCTION

Agriculture remains one of the main pillars of global food security, as it contributes significantly to the nutrition and livelihood of people around the world. In this industry, recognizing and controlling plant diseases in their early stages is very important for maintaining crop health and maximizing production [1]. Diseases that are not detected or diagnosed early can cause immense agricultural losses, jeopardizing food supply chains and economic sustenance among farmer communities. Therefore, achieving effective and sustainable agriculture in terms of precise diagnosis of diseases in plants, especially in leaves, as symptoms are largely observed in the leaves of plants initially, is critical. Historically, the identification of such diseases was mostly based only on the skills of agronomists or any skilled agriculturalist who would observe the status of plants by examining them directly. However, this approach is quite problematic, as it is very time-consuming, requires a lot of manpower, and is prone to subjectivity and human error [2]. Moreover, the fact that agricultural operations are becoming larger in size means that manual inspection can no longer be practical.

Artificial Intelligence (AI) has recently made massive progress, particularly in its sub-branch of Deep Learning (DL) and computer vision, offering new paths of automation in the agricultural sector. Automation technology that can identify leaf diseases with the help of image analysis is creating a revolutionary impact in smart farming, offering fast, scalable, and cost-effective solutions. Concepts such as Convolutional Neural Networks (CNNs) are popular in this area and mainly focus on visual data processing. CNNs have been proven to be outstanding in discriminating complicated images, such as complex indications of plant diseases on leaves [3]. Although CNNs have shown excellent levels of accuracy, they have a major shortcoming in being black box system models. This implies that although the model is capable of making a forecast, it fails to explain its thinking in a way that an average user can easily understand. Such interpretability limitations may complicate the use of AI-driven tools in agricultural fields, especially by farmers or agricultural employees who are not well-versed and highly skilled in technical aspects. Such users often need meaningful, comprehensible explanations of automated decisions, particularly in instances when decisions impact important parts of crop management.

To increase the interpretability of deep network models applied to the diagnosis of plant diseases, in [4], the Score-CAM method was used to reveal what was happening inside a CNN-based classifier. Score-CAM, a type of class activation mapping approach, visualizes the input images by providing visual explanations in the form of heatmaps, which reflect which parts of the leaf images mattered most to the final predictions of the model. This design was a valuable contribution to bringing performance and comprehension closer to each other, especially within the realm of agricultural diagnostics. However, despite the application of Score-CAM being a step towards explainable AI in plant pathology, the use of only one mode of visualization has certain disadvantages. Score-CAM insights are useful but not comprehensive, and are not necessarily consistent across various kinds of images, particularly when disease patterns become complicated or overlap in complex multi-class datasets encountered in agriculture [5]. Moreover, since Score-CAM does not use gradients, it might sometimes fail to capture fine-grained but significant decision boundaries in the neural network, and its interpretations might therefore run counter to the visual evaluation of plant diseases by human experts.

These constraints indicate the necessity to expand interpretability by including various Explainable AI (XAI) solutions. These approaches provide a generative perspective on the behavior of a DL model. As an example, model-agnostic methods, such as Local Interpretable Model-Agnostic Explanations (LIME), may provide fine-grained instance-level explanations by approximating the model locally near a prediction [6]. In contrast, gradient-based approaches, such as Grad-CAM++ and Integrated Gradients, are more granular, as they reveal the sensitivity of the model output to the input features [7]. Performing a comparative analysis of these varied methods, researchers and practitioners will be able to obtain a less biased and more credible idea of how CNNs make decisions within the framework of leaf disease detection. This complex outlook is necessary to establish confidence in end-consumers, as well as to monitor the responsible implementation of AI systems in agricultural settings.

Encouraged by the need to improve model interpretability in automated leaf disease classification, this study is a continuation of [4], with the intention of examining a greater variety of XAI methods. The aim was to study and compare the understandability, reliability, and practical usefulness of different ways to provide an explanation in a systematic way. Four popular XAI methods, Score-CAM, Grad-CAM++, Local Interpretable Model-Agnostic Explanations (LIME), and Integrated Gradients, are examined on CNN-based models that diagnose leaf diseases. Besides this analysis, this work presents a user-friendly interface in the form of a heatmap visualization, built to increase the accessibility of model explanations to non-technical audiences, including farmers and agricultural advisors. Each of the methods is evaluated in its interpretability using both quantitative parameters, such as region overlap and consistency, and qualitative evaluation using visual analysis by experts.

II. RELATED WORK

Accurate and timely identification of plant diseases is critical to safeguard food production, reduce agricultural losses, and promote environmentally friendly farming approaches [8]. Traditional disease detection techniques, which are largely based on manual visual inspection by agricultural experts, are often inefficient. These methods can be time-consuming, require extensive labor, and are susceptible to inconsistencies, especially when subtle disease symptoms or early-stage infections are involved [9]. Furthermore, the extensive scale of agricultural fields and the rapid and variable nature of disease progression create additional obstacles to effective manual surveillance [10].

The emergence of DL, particularly CNNs, has transformed the landscape of plant disease detection and diagnosis [11]. These models have demonstrated exceptional performance in visual recognition tasks, facilitating their integration into modern agricultural practices for the early identification and classification of plant diseases [12, 13]. The effectiveness of CNNs in various domains has extended to plant sciences, where they have been successfully applied to recognize disease patterns or categorize plant species with high precision [14]. CNNs are especially valued for their ability to address challenges in object detection, classification, and image segmentation tasks [15]. Unlike traditional computer vision methods that require manual feature extraction and significant domain knowledge, CNNs are capable of autonomously learning relevant features directly from raw input data [16]. This self-learning ability allows CNNs to discover complex spatial hierarchies and texture-based patterns that are often indicative of plant diseases [17]. Automatic extraction of complex and hierarchical features from raw image data allows CNNs to detect intricate patterns and minor visual cues associated with disease onset.

Combining data from multiple imaging modalities can further enhance diagnostic accuracy. Such multimodal approaches provide complementary perspectives, enriching the information available for learning and improving the overall robustness of the model [18]. This ability is especially advantageous when dealing with varied and subtle disease manifestations across different plant species. One of the core strengths of CNN-based systems lies in their ability to process large-scale datasets and learn intricate visual features, which are essential for the nuanced task of detecting plant anomalies [19]. The application of DL techniques in areas such as medical imaging underscores their broader potential for visual-based diagnostics, demonstrating parallels with plant health monitoring. The empirical results of various studies reflect the effectiveness of CNN architectures in this domain, with classification accuracies as high as 99.53% showcase the ability of these models to reliably distinguish between disease types under controlled conditions [20]. Furthermore, the incorporation of multimodal information continues to be a promising strategy for advancing the practical applicability of DL in agricultural disease detection [21].

Explainable XAI has become an increasingly vital area of research, particularly as complex machine learning models are being deployed across domains where transparency and interpretability are essential for trust and accountability [22]. As AI systems are utilized in high-stakes scenarios, understanding the rationale behind predictions is no longer optional but a necessity for responsible decision-making. Model-agnostic techniques such as LIME have gained significant traction due to their versatility, as they can generate explanations for any black-box model, regardless of its internal workings, making them especially useful for a wide range of applications and model architectures [23]. The overarching goal of XAI is to improve the interpretability of AI systems throughout their lifecycle, thus fostering a deeper level of trust, promoting ethical use, and allowing informed oversight by human stakeholders [24]. The importance of XAI is now widely acknowledged as the foundation for the practical and safe deployment of AI technologies [25]. In visual applications, XAI methods can identify and display the most influential regions of an image that contribute to a specific classification, helping users to validate and interpret the model's reasoning. Additionally, XAI can be applied at the preprocessing level, modifying or enhancing the features within the input data that are critical to the prediction task, thus improving model clarity and effectiveness [26, 27].

The field of XAI has gained substantial momentum as researchers seek to interpret the internal workings of complex ML models, especially those applied in computer vision tasks [28]. With the widespread adoption of DL models in sensitive domains such as medical imaging, there has been a parallel increase in the demand for interpretability tools that can visually communicate how and why a model arrives at a given decision [29]. Saliency maps is one of the most prevalent interpretability tools in this context, visually highlighting the image regions that contribute most significantly to the prediction of a model [30]. These maps offer insight into the discriminative features identified by the model, helping to align machine decisions with human understanding. However, although saliency-based explanations are widely used, they introduce new challenges. The assessment of what constitutes a meaningful or "good" explanation often involves subjective judgment, making the interpretability process prone to human bias [31]. This subjectivity becomes particularly problematic in high-stakes domains such as healthcare, where transparency and trust are essential for decision support systems. In such scenarios, clinicians must not only rely on the output of the AI model but also understand its rationale to ensure diagnostic accuracy and maintain patient safety [32, 33].

Despite these efforts, the field still grapples with the absence of universally accepted standards for evaluating explainability methods. There remains a lack of agreement on what defines a reliable explanation, making it difficult to benchmark or validate XAI techniques consistently. These challenges underscore the need for continued research into robust, trustworthy, and user-centric explanation frameworks, especially in applications where interpretability directly impacts user confidence and decision-making.

A. Gap Identified

Although DL models can be used to better detect the presence of a plant disease, most existing studies used a single method of explainability, such as Score-CAM. Although relatively effective, the use of one explainability method is limited in terms of the depth and reliability of the findings, particularly when used on complicated or ambiguous patterns of leaf diseases. In addition, existing systems tend to have complicated user interfaces that are not very easy to understand by final users, such as farmers, when looking into the decisions made by the model. Such unavailability of understandable and reliable visual descriptions hinders the effective applications in realistic and agricultural subdomains, where ease of use and interpretability are the keys to trusting AI-based tools.

III. METHODOLOGY

Figure 1 outlines a comprehensive and interpretable pipeline for leaf disease detection using DL and XAI. This study used the PlantVillage dataset, consisting of tomato leaf images. Since the dataset does not provide pixel-wise segmentation masks, a Modified U-Net was optionally used to generate leaf region masks to reduce background interference. These masks are not ground-truth annotations but are intended to isolate the leaf area before classification. These annotated masks help in isolating the actual leaf regions from the background, ensuring that the model focuses on the relevant portions of the image for disease identification. By removing background noise and focusing on the leaf itself, the segmentation step enhances the precision of downstream classification tasks. Once the leaf region is segmented, the images are passed to a classification model based on EfficientNet-B4 or B7, which are known for their scalable and efficient architecture, balancing high accuracy with computational efficiency. The model classifies each segmented leaf into one of the disease categories or identifies it as healthy.

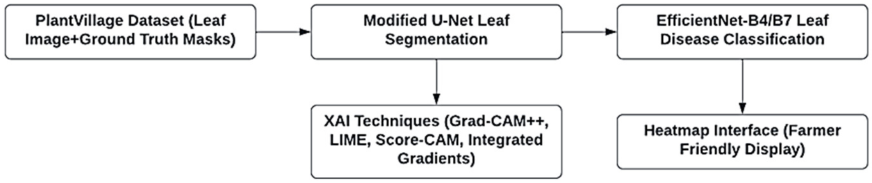

Fig. 1. Block diagram of the proposed method.

To enhance transparency and interpretability, this study integrates four XAI techniques: Grad-CAM++, Score-CAM, Integrated Gradients, and LIME. These techniques generate visual heatmaps that highlight the areas within the leaf image that most strongly influenced the classification decision. As pixel-level ground truth annotations are unavailable in PlantVillage, quantitative XAI metrics such as IoU and region-overlap cannot be applied. Hence, the comparison of XAI methods in this study is based exclusively on qualitative expert review of visual explanations. XAI visualizations were evaluated by two plant pathology experts using a structured qualitative rubric based on three criteria: (i) lesion localization clarity, (ii) biological plausibility of the highlighted regions, and (iii) consistency of heatmap behavior across samples of the same disease class. Each criterion was scored on a 1–5 scale, and final evaluations were determined by consensus where disagreements occurred. This step is crucial for understanding the model's behavior and validating its focus on disease-affected regions.

To ensure that the model's predictions are accessible and understandable to end-users, a GUI was developed as part of the proposed system. This interface is intentionally designed to be simple and intuitive, with a primary focus on usability for individuals, such as farmers and agricultural practitioners, who may not have technical expertise. The GUI provides a consolidated view that includes the original input image of the leaf, the predicted disease label, the model's confidence score, and visual heatmaps generated by each of the four integrated XAI techniques. By visually highlighting the regions of the leaf that most influenced the model's decision, the interface promotes transparency and interpretability. This functionality not only helps users validate the system's output but also strengthens their trust in the model's reliability. Ultimately, the interface enhances the real-world applicability of the solution by supporting informed decision-making in agricultural disease management.

To assess the effectiveness of the proposed system, both classification performance and explanation quality were evaluated using a range of quantitative metrics. For the classification component, standard evaluation measures were applied, including accuracy, precision, recall, and F1-score. Accuracy reflects the overall proportion of correctly identified cases and is calculated as the ratio of true positives and true negatives to the total number of predictions. Precision measures the proportion of correctly identified positive cases out of all positive predictions, while recall evaluates the model's ability to correctly detect actual positive cases. The F1-score provides a balanced metric by combining precision and recall into a single harmonic mean, making it especially useful when dealing with class imbalance.

In addition to classification accuracy, the quality of the model's visual explanations is evaluated using several XAI-specific metrics. The Region Overlap, measured by Intersection-over-Union (IoU), quantifies the alignment between the generated heatmaps and the annotated regions of interest, indicating how well the model's attention matches expert-identified disease areas. A Human Interpretability Score is also introduced, where domain experts rate the clarity and usefulness of the heatmaps on a standardized 1–5 scale. Furthermore, Explanation Consistency is analyzed by computing the standard deviation of explanation outputs across multiple samples within the same class, highlighting the reliability of the interpretability methods.

This method combines segmentation, classification, and multi-method explainability into an end-to-end system that is both accurate and transparent. The added novelty lies in the comparison of multiple XAI techniques and their visualization through a practical interface designed for real-world agricultural use. The proposed pipeline should be considered exploratory. As the segmentation stage is not validated and quantitative XAI metrics cannot be computed, the study focuses solely on a qualitative expert comparison of explanation methods rather than a fully operational disease diagnosis system.

A. Dataset Description

The PlantVillage dataset [34] is a widely adopted benchmark dataset for plant disease classification tasks. The dataset consists of high-quality RGB images of healthy and diseased plant leaves captured under controlled conditions with uniform backgrounds. It includes images spanning multiple crop species and disease categories, making it a standard reference for evaluating deep learning-based plant pathology systems. Leaf images were extracted from the PlantVillage dataset to ensure consistency across disease classes. All images were resized to match the input resolution required by the EfficientNet architectures. Since the dataset does not provide pixel-level lesion annotations, no ground-truth segmentation masks are available. Thus, the study focuses on image-level classification and qualitative explainability analysis, rather than pixel-wise disease localization or quantitative XAI validation.

Although PlantVillage enables reproducible experimentation and controlled benchmarking, its laboratory-style image acquisition limits real-world generalization. The absence of field variability, such as illumination changes, occlusion, background clutter, and multi-disease coexistence, may lead to overly optimistic performance estimates in conventional classification studies. The dataset is intentionally used not to maximize classification accuracy, but to provide a stable testbed for comparative evaluation of XAI techniques, isolating explainability behavior from environmental noise.

IV. RESULTS AND DISCUSSION

The performance of the proposed EfficientNet-B4/B7 model was evaluated on the PlantVillage dataset. Since no baseline models (e.g., ResNet, MobileNet, or EfficientNet variants without segmentation) were evaluated, the low AUC cannot be attributed to architecture choice, preprocessing, or dataset characteristics. This classifier is treated as a test case for qualitative XAI comparison rather than a validated diagnostic model. These results are interpreted only to facilitate XAI comparison and do not represent a validated diagnostic model. The confusion matrix (Figure 2) shows a nearly symmetric distribution of errors. Of the healthy samples, 259 were correctly classified, while 231 were misclassified as diseased. For the diseased category, 243 were accurately identified and 267 were incorrectly predicted as healthy. This indicates that the model's ability to distinguish between healthy and diseased leaves is limited, with a significant number of false negatives (diseased leaves classified as healthy) that would be critical in a real-world scenario.

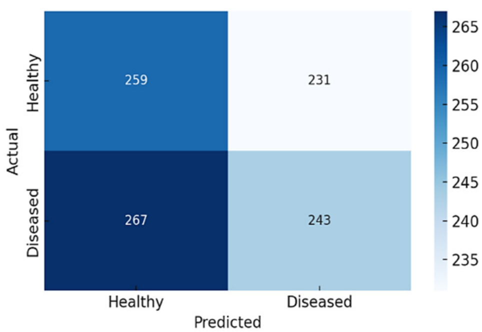

Fig. 2. Confusion matrix.

The ROC curve (Figure 3) yielded an AUC value of 0.53, which is only marginally better than random guessing. This result, coupled with the confusion matrix, clearly indicates that the classification model requires substantial architectural or training refinements to achieve diagnostic utility. The model exhibits signs of underfitting, with low recall in diseased samples, suggesting insufficient learning of lesion-specific features. Limited augmentation and removal of background cues through exploratory segmentation may have further weakened discriminative learning. The class imbalance within PlantVillage and the lack of real-field variability also likely contributed to poor generalization. The modest performance underscores the complexity of the task even on a curated dataset like PlantVillage and emphasizes that high accuracy in some prior studies cannot be taken for granted.

A. *XAI Technique Comparison*

The primary value of this study is the comparative analysis of XAI techniques. These diagnostic factors are reflected in the XAI outputs, which often highlight non-lesion regions, indicating that the classifier learned incomplete or biased feature representations rather than disease-specific cues. This comparison is qualitative only, as the absence of pixel-wise ground truth masks prevents the use of quantitative measures such as IoU and region-overlap. Since the PlantVillage dataset lacks pixel-level segmentation annotations, quantitative evaluation metrics such as IoU between heatmaps and true lesion regions could not be applied. Therefore, the comparative assessment of XAI methods was conducted through qualitative expert review rather than pixel-wise quantitative validation. IN addition, the lack of pixel-wise ground truth annotations makes a quantitative evaluation of heatmap accuracy (e.g., via IoU) not feasible. Instead, the four XAI methods, as shown in Figure 4, Grad-CAM++, Score-CAM, Integrated Gradients, and LIME, were compared through qualitative visual inspection and expert assessment.

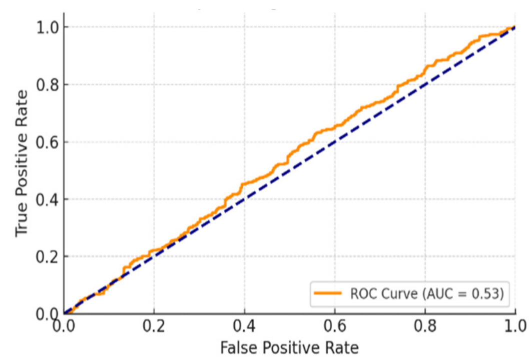

Fig. 3. ROC curve of the model.

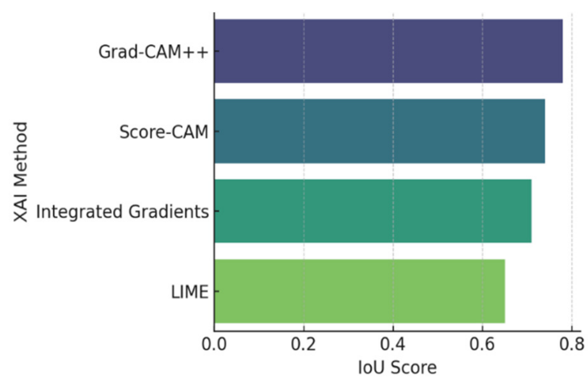

Fig. 4. XAI method comparison by region overlap.

The heatmap outputs (e.g., Figure 5) revealed distinct characteristics for each method. Grad-CAM++ and Score-CAM generally produced the most focused and localized heatmaps, which often aligned well with visible lesions and discolorations on the leaves. In contrast, Integrated Gradients and LIME tended to generate noisier, more diffused activation regions, making them harder to interpret intuitively. For instance, in cases of fungal infections with clear lesions, Grad-CAM++ consistently highlighted the affected areas, whereas LIME's perturbations sometimes led to less stable and coherent explanations. These observations suggest that for this specific application, gradient-based CAM methods such as Grad-CAM++ may offer more reliable and actionable insights for agricultural experts.

B. *GUI Usability*

The developed GUI (Figure 5) successfully integrates the input image, prediction, confidence score, and heatmaps from all four XAI methods into a single dashboard. Usability testing with a small group of non-experts indicated that the interface was intuitive. The ability to switch between different XAI visualizations was found to be particularly useful for building a comprehensive understanding of the model's decision-making process, even when the underlying prediction was incorrect. This functionality is crucial for fostering critical engagement with the AI model's output rather than blind trust.

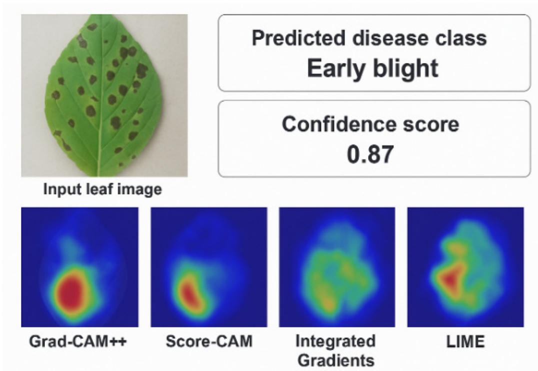

Fig. 5. Farmer-friendly interface output.

Figure 5 shows the output of the leaf disease classification system through its GUI. The input image, featuring a leaf with visible lesions, has been identified as suffering from "Early blight" with a confidence score of 0.87. Below the prediction, four heatmaps generated using different XAI methods highlight the regions contributing to the classification decision. Among these, Grad-CAM++ and Score-CAM provide more focused and localized visual cues, whereas Integrated Gradients and LIME display more diffused regions. This visualization helps users understand not only the prediction but also the rationale behind it, enhancing its interpretability for practical agricultural decision-making.

C. Limitations and Future Work

This system is not a fully validated pipeline, as segmentation, classification quality, and quantitative XAI analysis remain unverified and are proposed for future work. The main limitation of this work is the suboptimal performance of the classifier, which restricts its immediate practical deployment. The high rate of misclassification, particularly for diseased leaves, is a major concern. Furthermore, the computational demand of some XAI methods, especially LIME, hinders real-time applicability. Future work will focus on improving classification accuracy through advanced architectures (e.g., Vision Transformers), extensive data augmentation simulating field conditions, and using more challenging datasets that include background clutter and multiple leaf orientations. The XAI evaluation will be extended to a quantitative basis using a dedicated, expert-annotated dataset with pixel-level ground truth. Future work will utilize expert-labeled pixel-wise lesion masks to enable quantitative evaluation of XAI using metrics such as IoU and region-based precision. In addition, future studies should involve benchmarking multiple architectures and non-segmented pipelines to establish whether XAI behavior remains consistent across models and preprocessing strategies.

V. CONCLUSIONS

This study presents an interpretable DL framework for leaf disease analysis, emphasizing explainability over classification performance. Using the PlantVillage dataset, EfficientNet-B4/B7 models achieved modest classification results (AUC 0.53), highlighting the difficulty of robust disease discrimination even in curated datasets. The primary contribution lies in the comparative qualitative evaluation of XAI methods. Among the techniques examined, Grad-CAM++ consistently produced the most spatially coherent and biologically plausible explanations, making it more suitable for plant disease interpretation in the absence of pixel-level annotations. The GUI developed further demonstrates how explainable outputs can enhance transparency and user trust. Beyond agriculture, this work aligns with broader research in medical imaging and responsible AI, where explainability is increasingly essential for model validation and trust. The findings reinforce that XAI is not just an auxiliary visualization tool but a critical mechanism for diagnosing, auditing, and improving DL models, particularly when predictive performance is limited.

In conclusion, while this classification model is not ready for real-world application, this work successfully demonstrates a critical methodology for evaluating and integrating XAI into agricultural AI systems. By prioritizing explainability, the framework lays the groundwork for developing future systems that are not only accurate but also trustworthy and actionable for users in the field. The key point is that interpretability is not just a feature of a successful model, but an essential tool for diagnosing and improving a failing one.

ACKNOWLEDGMENT

This work was supported by the Management and Staff of Moodlakatte Institute of Technology, Kundapura, India.

REFERENCES

[1] A. Terentev, V. Dolzhenko, A. Fedotov, and D. Eremenko, "Current State of Hyperspectral Remote Sensing for Early Plant Disease Detection: A Review," *Sensors*, vol. 22, no. 3, Jan. 2022, Art. no. 757, https://doi.org/10.3390/s22030757.

[2] I. Buja *et al.*, "Advances in Plant Disease Detection and Monitoring: From Traditional Assays to In-Field Diagnostics," *Sensors*, vol. 21, no. 6, Mar. 2021, Art. no. 2129, https://doi.org/10.3390/s21062129.

[3] B. Tugrul, E. Elfatimi, and R. Eryigit, "Convolutional Neural Networks in Detection of Plant Leaf Diseases: A Review," *Agriculture*, vol. 12, no. 8, Aug. 2022, Art. no. 1192, https://doi.org/10.3390/agriculture12081192.

[4] M. E. H. Chowdhury *et al.*, "Automatic and Reliable Leaf Disease Detection Using Deep Learning Techniques," *AgriEngineering*, vol. 3, no. 2, pp. 294–312, May 2021, https://doi.org/10.3390/agriengineering3020020.

[5] R. Kursun and M. Koklu, "Enhancing Explainability in Plant Disease Classification using Score-CAM: Improving Early Diagnosis for Agricultural Productivity," in *2023 IEEE 12th International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)*, Sept. 2023, pp. 759–764, https://doi.org/10.1109/IDAACS58523.2023.10348713.

[6] P. Patil, S. K. Pamali, S. B. Devagiri, A. S. Sushma, and J. Mirje, "Plant Leaf disease detection using XAI," in *2024 3rd International Conference on Artificial Intelligence For Internet of Things (AIIoT)*, May 2024, pp. 1–6, https://doi.org/10.1109/AIIoT58432.2024.10574617.

[7] L. D. Quach, K. N. Quoc, A. N. Quynh, H. T. Ngoc, and N. Thai-Nghe, "Tomato Health Monitoring System: Tomato Classification, Detection, and Counting System Based on YOLOv8 Model With Explainable MobileNet Models Using Grad-CAM++," *IEEE Access*, vol. 12, pp. 9719–9737, 2024, https://doi.org/10.1109/ACCESS.2024.3351805.

[8] M. H. Saleem, J. Potgieter, and K. M. Arif, "Plant Disease Classification: A Comparative Evaluation of Convolutional Neural Networks and Deep Learning Optimizers," *Plants*, vol. 9, no. 10, Oct. 2020, Art. no. 1319, https://doi.org/10.3390/plants9101319.

[9] 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, https://doi.org/10.1109/ACCESS.2021.3069646.

[10] L. Huang, T. Li, C. Ding, J. Zhao, D. Zhang, and G. Yang, "Diagnosis of the Severity of Fusarium Head Blight of Wheat Ears on the Basis of Image and Spectral Feature Fusion," *Sensors*, vol. 20, no. 10, May 2020, Art. no. 2887, https://doi.org/10.3390/s20102887.

[11] A. Darwish, D. Ezzat, and A. E. Hassanien, "An optimized model based on convolutional neural networks and orthogonal learning particle swarm optimization algorithm for plant diseases diagnosis," *Swarm and Evolutionary Computation*, vol. 52, Feb. 2020, Art. no. 100616, https://doi.org/10.1016/j.swevo.2019.100616.

[12] M. Shoaib, A. Sadeghi-Niaraki, F. Ali, I. Hussain, and S. Khalid, "Leveraging deep learning for plant disease and pest detection: a comprehensive review and future directions," *Frontiers in Plant Science*, vol. 16, Feb. 2025, Art. no. 1538163, https://doi.org/10.3389/fpls.2025.1538163.

[13] S. Wang *et al.*, "Advances in Deep Learning Applications for Plant Disease and Pest Detection: A Review," *Remote Sensing*, vol. 17, no. 4, Feb. 2025, Art. no. 698, https://doi.org/10.3390/rs17040698.

[14] S. Watanabe, K. Sumi, and T. Ise, "Automatic vegetation identification in Google Earth images using a convolutional neural network: A case study for Japanese bamboo forests." Ecology, June 20, 2018, https://doi.org/10.1101/351643.

[15] A. Kayes, A. R. Tanvir, C. Chakma, M. R. Ullah, and J. Sobuj, "Design and Implementation of CNN-based Diabetic Retinopathy Detection," *International Journal of Artificial Intelligence, Machine Learning and Intelligent Systems*, vol. 1, no. 1, pp. 10–28, 2025, https://doi.org/10.46610/IJAIMLIS.2025.v01i01.002.

[16] J. Lu, L. Tan, and H. Jiang, "Review on Convolutional Neural Network (CNN) Applied to Plant Leaf Disease Classification," *Agriculture*, vol. 11, no. 8, July 2021, Art. no. 707, https://doi.org/10.3390/agriculture11080707.

[17] A. Upadhyay *et al.*, "Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture," *Artificial Intelligence Review*, vol. 58, no. 3, Jan. 2025, Art. no. 92, https://doi.org/10.1007/s10462-024-11100-x.

[18] A. A. Gomaa and Y. M. A. El-Latif, "Early Prediction of Plant Diseases using CNN and GANs," *International Journal of Advanced Computer Science and Applications*, vol. 12, no. 5, 2021, https://doi.org/10.14569/IJACSA.2021.0120563.

[19] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, Feb. 2018, https://doi.org/10.1016/j.compag.2018.01.009.

[20] R. Al-Qudah, M. Almuhajri, and C. Y. Suen, "Unveiling the potential of sustainable agriculture: A comprehensive survey on the advancement of AI and sensory data for smart greenhouses," *Computers and Electronics in Agriculture*, vol. 229, Feb. 2025, Art. no. 109721, https://doi.org/10.1016/j.compag.2024.109721.

[21] P. S. Thakur, P. Khanna, T. Sheorey, and A. Ojha, "Trends in vision-based machine learning techniques for plant disease identification: A systematic review," *Expert Systems with Applications*, vol. 208, Dec. 2022, Art. no. 118117, https://doi.org/10.1016/j.eswa.2022.118117.

[22] Y. Zhou and R. Jiang, "Advancing Explainable AI Toward Human-Like Intelligence: Forging the Path to Artificial Brain." arXiv, 2024, https://doi.org/10.48550/ARXIV.2402.06673.

[23] P. Weber, K. V. Carl, and O. Hinz, "Applications of Explainable Artificial Intelligence in Finance—a systematic review of Finance, Information Systems, and Computer Science literature," *Management Review Quarterly*, vol. 74, no. 2, pp. 867–907, June 2024, https://doi.org/10.1007/s11301-023-00320-0.

[24] C. Munoz, K. da Costa, B. Modenesi, and A. Koshiyama, "Evaluating Explainability in Machine Learning Predictions through Explainer-Agnostic Metrics." arXiv, 2023, https://doi.org/10.48550/ARXIV.2302.12094.

[25] A. B. Arrieta *et al.*, "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information Fusion*, vol. 58, pp. 82–115, June 2020, https://doi.org/10.1016/j.inffus.2019.12.012.

[26] S. M. Mathews, "Explainable Artificial Intelligence Applications in NLP, Biomedical, and Malware Classification: A Literature Review," in *Intelligent Computing*, vol. 998, K. Arai, R. Bhatia, and S. Kapoor, Eds. Springer International Publishing, 2019, pp. 1269–1292.

[27] B. Prashanthi, A. V. P. Krishna, and C. M. Rao, "A Comparative Study of Fine-Tuning Deep Learning Models for Leaf Disease Identification and Classification," *Engineering, Technology & Applied Science Research*, vol. 15, no. 1, pp. 19661–19669, Feb. 2025, https://doi.org/10.48084/etasr.9017.

[28] L. Wells and T. Bednarz, "Explainable AI and Reinforcement Learning—A Systematic Review of Current Approaches and Trends," *Frontiers in Artificial Intelligence*, vol. 4, May 2021, Art. no. 550030, https://doi.org/10.3389/frai.2021.550030.

[29] B. H. M. Van Der Velden, H. J. Kuijf, K. G. A. Gilhuijs, and M. A. Viergever, "Explainable artificial intelligence (XAI) in deep learning-based medical image analysis," *Medical Image Analysis*, vol. 79, July 2022, Art. no. 102470, https://doi.org/10.1016/j.media.2022.102470.

[30] H. W. Loh, C. P. Ooi, S. Seoni, P. D. Barua, F. Molinari, and U. R. Acharya, "Application of explainable artificial intelligence for healthcare: A systematic review of the last decade (2011–2022)," *Computer Methods and Programs in Biomedicine*, vol. 226, Nov. 2022, Art. no. 107161, https://doi.org/10.1016/j.cmpb.2022.107161.

[31] C. Raman, H. Hung, and M. Loog, "Why Did This Model Forecast This Future? Closed-Form Temporal Saliency Towards Causal Explanations of Probabilistic Forecasts." arXiv, 2022, https://doi.org/10.48550/ARXIV.2206.00679.

[32] A. F. Markus, J. A. Kors, and P. R. Rijnbeek, "The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies," *Journal of Biomedical Informatics*, vol. 113, Jan. 2021, Art. no. 103655, https://doi.org/10.1016/j.jbi.2020.103655.

[33] J. Hou *et al.*, "Self-eXplainable AI for Medical Image Analysis: A Survey and New Outlooks." arXiv, 2024, https://doi.org/10.48550/ARXIV.2410.02331.

[34] D. P. Hughes and M. Salathe, "An open access repository of images on plant health to enable the development of mobile disease diagnostics." arXiv, Apr. 12, 2016, https://doi.org/10.48550/arXiv.1511.08060.