

Explainable Deep Learning for Cotton Leaf Disease Classification Using ResNet50 and Grad-CAM

D.Y. Tayade¹, Dr.D.N. Besekar²

¹Assistant Professor, Department of Computer, P.N.College Pusad

²Professor, Department of Computer Science, Shri.Shivaji College Akola

Abstract—Cotton is one of India’s most important commercial crops, yet its yield is significantly affected by foliar diseases such as Curl Virus, Fusarium Wilt, Bacterial Blight, Powdery Mildew, Target Spot, and pest-induced damage. Early identification is essential for controlling disease spread, but traditional manual diagnosis performed by farmers is slow, expertise-dependent, and often inaccurate. To address this challenge, this study proposes an explainable deep learning framework for the automatic classification of cotton leaf diseases using a ResNet50-based Convolutional Neural Network (CNN) integrated with Gradient-weighted Class Activation Mapping (Grad-CAM). A merged dataset of 5,561 cotton leaf images collected from the Akola and Yavatmal districts and supplemented with publicly available samples was used for experimentation. The images were categorized into seven classes: Bacterial Blight, Curl Virus, Fusarium Wilt, Powdery Mildew, Target Spot, Pest Damage, and Healthy leaves.

A two-stage transfer learning strategy was adopted, involving frozen feature extraction followed by selective fine-tuning of deeper layers for disease-specific representation learning. The proposed ResNet50 model achieved a test accuracy of 99.2%, demonstrating strong generalization performance.

The core contribution of this research lies in integrating Grad-CAM to provide visual explanations for model predictions. The generated heatmaps highlight disease-relevant regions such as lesions, vein distortions, fungal growth, and pest damage, thereby improving transparency and trust. The proposed framework represents a step toward Explainable AI (XAI) in agriculture and is suitable for integration into mobile or web-based decision-support systems for early disease diagnosis.

Index Terms—Cotton leaf diseases, Deep learning, Explainable AI, Grad-CAM, CNN, ResNet50, Transfer learning.

I. INTRODUCTION

Cotton (*Gossypium hirsutum*) is one of India’s most important commercial crops and plays a vital role in the agricultural economy, particularly in Maharashtra. However, cotton cultivation is highly susceptible to foliar diseases such as Curl Virus, Bacterial Blight, Fusarium Wilt, Powdery Mildew, Target Spot, and pest-induced damage. These diseases cause significant yield losses, degrade fiber quality, and increase production costs. Traditional disease diagnosis relies on manual field inspection by farmers or agricultural experts, which is time-consuming, subjective, and often inaccurate, especially during early infection stages.

Recent advances in deep learning, particularly Convolutional Neural Networks (CNNs), have demonstrated strong performance in plant disease classification by automatically learning discriminative features such as texture variations, lesion patterns, vein discoloration, and leaf deformation. Despite their high accuracy, CNN-based systems are often criticized for their black-box nature, providing little insight into the reasoning behind predictions. This lack of interpretability limits trust and adoption in agriculture, where incorrect predictions can have serious economic consequences.

Explainable Artificial Intelligence (XAI) techniques aim to address this limitation by making deep learning decisions more transparent. Gradient-weighted Class Activation Mapping (Grad-CAM) is a widely used XAI method that highlights image regions contributing most to a model’s prediction. By generating visual heatmaps, Grad-CAM enables verification of whether CNNs focus on biologically

meaningful disease symptoms rather than background artifacts.

In this work, an explainable deep learning framework is proposed for cotton leaf disease classification using a ResNet50-based CNN integrated with Grad-CAM. A dataset of 5,561 cotton leaf images collected from the Akola and Yavatmal districts of Maharashtra and supplemented with publicly available samples was used to classify seven disease categories. A two-stage transfer learning strategy was adopted to improve performance on limited agricultural data. The key contribution of this study lies in combining high classification accuracy with visual interpretability, enabling reliable and transparent disease diagnosis suitable for real-world agricultural decision-support systems.

II. LITERATURE REVIEW

Deep learning has seen rapid adoption in agricultural informatics due to its superior ability to learn discriminative features from plant images. Early studies primarily focused on traditional machine learning algorithms such as SVMs, Random Forests, and K-Nearest Neighbours, which required handcrafted features such as colour histograms, GLCM texture descriptors, and shape-based measurements. While these methods achieved moderate success, they lacked robustness against real-world variability in lighting, leaf orientation, and background clutter.

The introduction of Convolutional Neural Networks (CNNs) revolutionized plant disease classification. Mohanty et al. (2016) demonstrated the capability of CNNs to classify 14 crop species and 26 diseases with high accuracy. Similarly, Ferentinos (2018) evaluated deep learning models on plant leaves and reported significant improvement compared to classical approaches. These foundational works paved the way for numerous experiments in crop-specific disease recognition.

In the context of cotton, several researchers explored CNN-based disease identification. Patil et al. (2020) used AlexNet and VGG16 to classify cotton leaf diseases, reporting satisfactory performance but facing limitations in computational cost and overfitting. Later studies incorporated lightweight architectures such as

MobileNet and EfficientNet, enabling faster deployment on mobile devices. However, many of these works lacked focus on explainability, leading to limited adoption in field applications.

A growing body of literature emphasizes the role of transfer learning in improving accuracy, particularly for agricultural datasets that are typically small in size. Models pretrained on ImageNet, including ResNet, Inception, and DenseNet, have been widely used due to their strong feature extraction capabilities. Researchers such as Too et al. (2019) systematically compared CNN backbones for plant disease classification, highlighting the superior generalization ability of deep pretrained networks when fine-tuned appropriately.

Despite accuracy improvements, a major challenge persisted: deep learning models behave as black boxes, offering no clarity on how predictions are made. This lack of transparency poses risk in sensitive domains like agriculture where wrong predictions may misguide farmers. To address this, researchers began incorporating Explainable Artificial Intelligence (XAI) techniques. Selvaraju et al. (2017) introduced Gradient-weighted Class Activation Mapping (Grad-CAM), which visually highlights important image regions contributing to model output. Since then, Grad-CAM has been applied in medical imaging, food quality inspection, and increasingly in agricultural disease diagnosis.

In plant pathology, XAI-based works are still emerging. Barbedo (2019) discussed the importance of interpretability in agricultural AI, while Chouhan et al. (2020) applied Grad-CAM to tomato leaf disease classification, showing how heatmaps improved trustworthiness of the predictions. However, literature on cotton leaf diseases with Grad-CAM-based explainability remains very limited, creating a gap that this research intends to fill. Zhao et al. (2020) integrated an attention mechanism into a CNN architecture for tomato disease recognition, achieving 96.81% accuracy by focusing on disease-relevant regions. Cengil et al. (2021) proposed a hybrid approach where features extracted from pretrained CNNs such as AlexNet, ResNet50, and VGG16 were fused and classified using machine learning algorithms, achieving up to 98.3% accuracy.

This study distinctly contributes by integrating ResNet50 with Grad-CAM visualization on a real-world cotton leaf dataset collected from the Akola and Yavatmal districts. Unlike earlier works, the focus here is not only on classification accuracy but also on visual interpretability, enabling validation of model behaviour using disease-specific symptoms. The heatmaps produced in this research help confirm whether CNNs are learning meaningful lesion patterns, making the system reliable, transparent, and suitable for end-user deployment.

III. DATASET DESCRIPTION

The dataset used in this study consists of a total of 5561 cotton leaf images, compiled from two primary sources

(1) Field images collected from Akola and Yavatmal districts in Maharashtra, India, and

(2) Publicly available datasets, including Kaggle repositories and open-source agricultural image collections.

This combination of real-field images and curated samples ensures diversity in environmental conditions, disease severity, background noise, and leaf orientation.

3.1 Dataset Composition

Field images were captured using Canon EOS 1500D DSLR camera under natural lighting conditions across multiple farms during the cropping seasons of 2023–2025. The images were categorized into seven classes representing commonly observed cotton leaf conditions: Bacterial Blight, Curl Virus, Fusarium Wilt, Powdery Mildew, Target Spot, Pest Damage, and Healthy leaves. The dataset was split into training, validation, and test sets as summarized in Table 1.

The dataset includes seven distinct cotton leaf classes that represent commonly observed diseases in Maharashtra cotton fields

Class	Train	Validation	Test	Total
Bacterial_Blight	798	99	101	998
Curl_Virus	773	96	98	967
Fusarium_Wilt	775	96	98	969
Healthy	780	97	99	976
Pest	440	55	55	550
Powdery_Mildew	440	55	55	550
Target Spot	440	55	55	550
TOTAL	4446	553	562	5561

Table 1: Distribution of images per class across training, validation, and test splits.

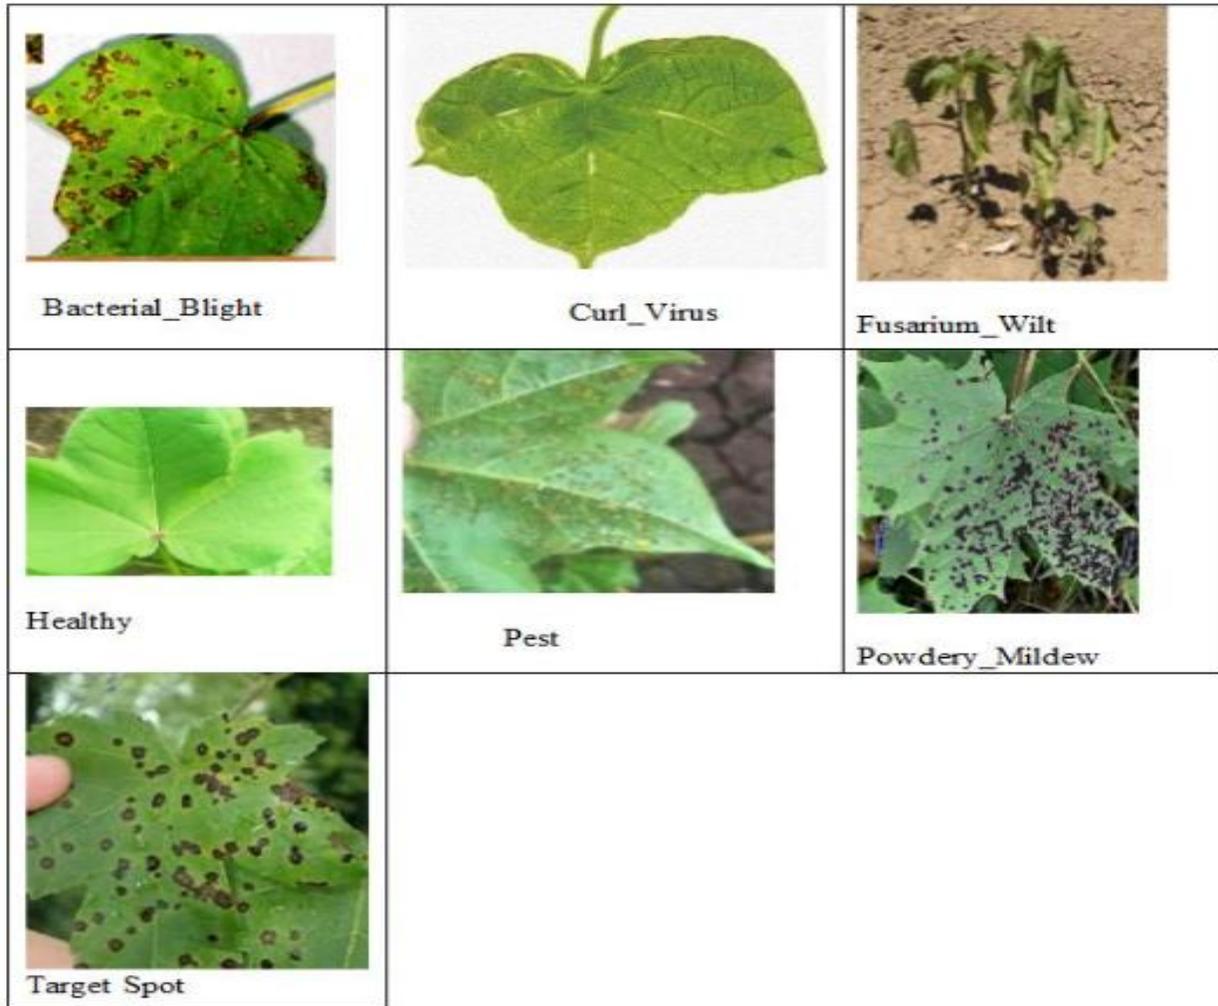


Figure 1: Representative images from the cotton leaf dataset across seven classes.

3.2 Challenges in the Dataset

The dataset presents typical real-world agricultural challenges, including non-uniform lighting, complex backgrounds, overlapping leaves, and subtle visual differences between disease classes. These challenges motivate the use of robust deep learning models combined with explainability techniques to ensure reliable and interpretable predictions.

Addressing these complexities requires models that are not only accurate but also interpretable—thus motivating the integration of Grad-CAM in this study.

IV. METHODOLOGY

The proposed explainable cotton leaf disease classification framework consists of five stages: image pre-processing, data augmentation, CNN-based feature extraction, two-stage transfer learning, and Grad-CAM-based explainability. The overall workflow of the system is illustrated in Figure 2, integrating automated disease classification with visual interpretation to enhance transparency and reliability.

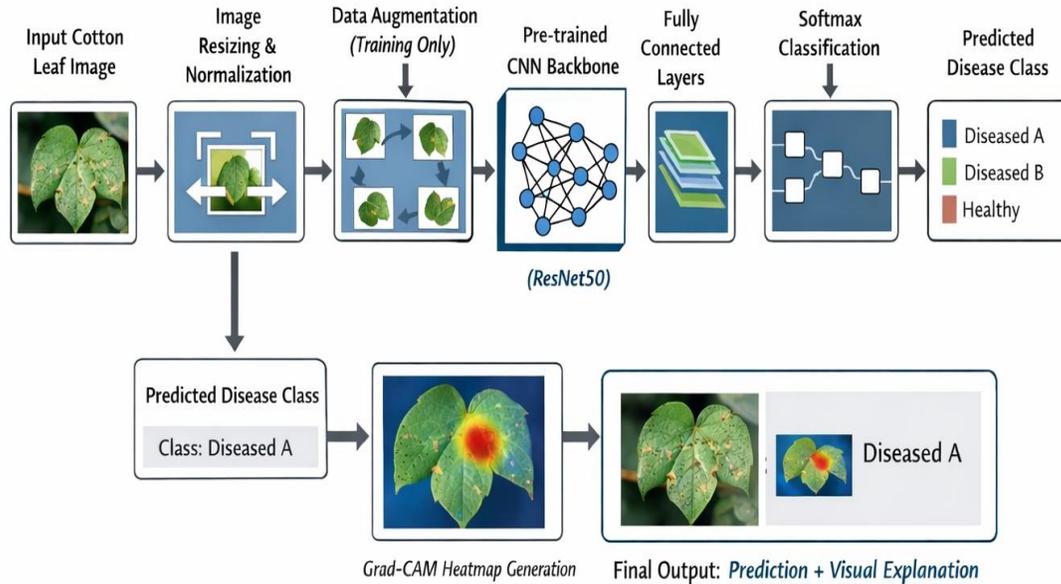


Figure 2. Block diagram of the proposed cotton leaf disease classification framework using ResNet50 and Grad-CAM.

4.1 System Pipeline

A cotton leaf image is resized to 224×224 pixels and normalized before being processed by a ResNet50 CNN pre-trained on ImageNet. During training, data augmentation is applied to improve generalization. Deep feature representations extracted by the CNN are passed to a SoftMax classifier to predict the disease class. Grad-CAM is subsequently applied to the final convolutional layer to generate class-discriminative heatmaps highlighting disease-relevant regions.

4.2 Preprocessing and Data Augmentation

All images were resized to $224 \times 224 \times 3$ and normalized to the range $[0, 1]$. One-hot encoding was used for multi-class labelling. To reduce overfitting and improve robustness under real-field conditions, augmentation techniques including random flipping, rotation ($0-20^\circ$), zoom (0.1), and contrast adjustment (0.1) were applied during training.

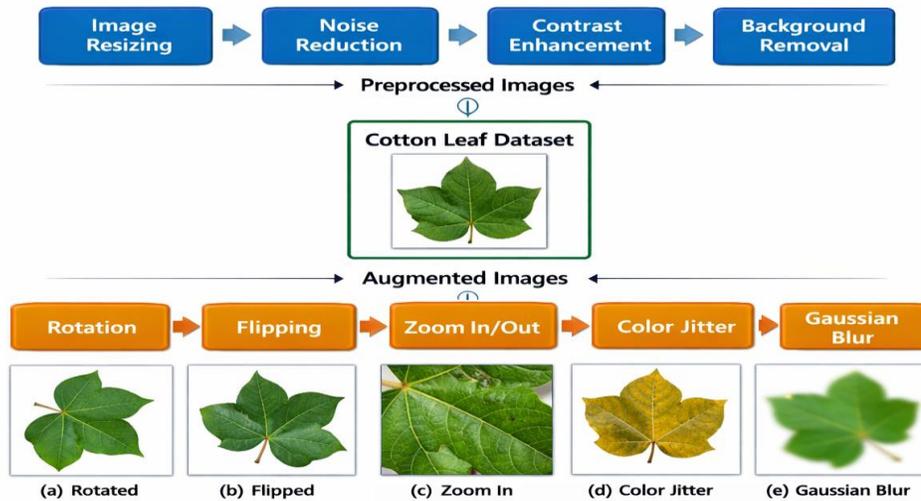


Figure 3 Preprocessing and data augmentation workflow applied to cotton leaf images.

4.3 Model Selection and Feature Extraction

ResNet50 was selected due to its residual learning mechanism, which enables effective training of deep networks and strong feature extraction. The model was initialized with ImageNet-pretrained weights and configured with `include_top = False`, allowing it to function as a feature extractor capable of learning disease-specific patterns such as lesion textures, vein discoloration, fungal growth, and pest-induced damage.

4.4 Two-Stage Transfer Learning Strategy

A two-stage transfer learning approach was adopted. In the first stage, all convolutional layers were frozen and a custom classification head consisting of Global Average Pooling, Dropout (0.25), and a SoftMax layer was trained using the Adam optimizer (learning rate = $1e-3$). In the second stage, the top 30–40 layers of ResNet50 were selectively unfrozen and fine-tuned with a reduced learning rate ($1e-4$), along with Early Stopping and ReduceLRonPlateau, enabling task-specific feature adaptation while preserving general representations.

4.5 Evaluation Metrics

Model performance was evaluated using accuracy, precision, recall, F1-score, and confusion matrix

analysis to assess both overall classification quality and per-class performance.

4.6 Explainability Using Grad-CAM

To address the black-box nature of CNNs, Grad-CAM was integrated to generate visual explanations. Grad-CAM heatmaps highlight the image regions that most influenced the predicted class, ensuring that model decisions are based on biologically meaningful disease symptoms rather than background artifacts. Detailed mathematical formulation is provided in the following section.

V. CNN ARCHITECTURE

5.1 ResNet50 Backbone

ResNet50 is a deep CNN architecture consisting of 50 layers and employs residual connections to mitigate the vanishing gradient problem. This design enables effective learning of complex visual features while maintaining training stability. The model contains approximately 23.6 million parameters and accepts input images of size $224 \times 224 \times 3$, making it well-suited for distinguishing visually similar cotton leaf diseases.

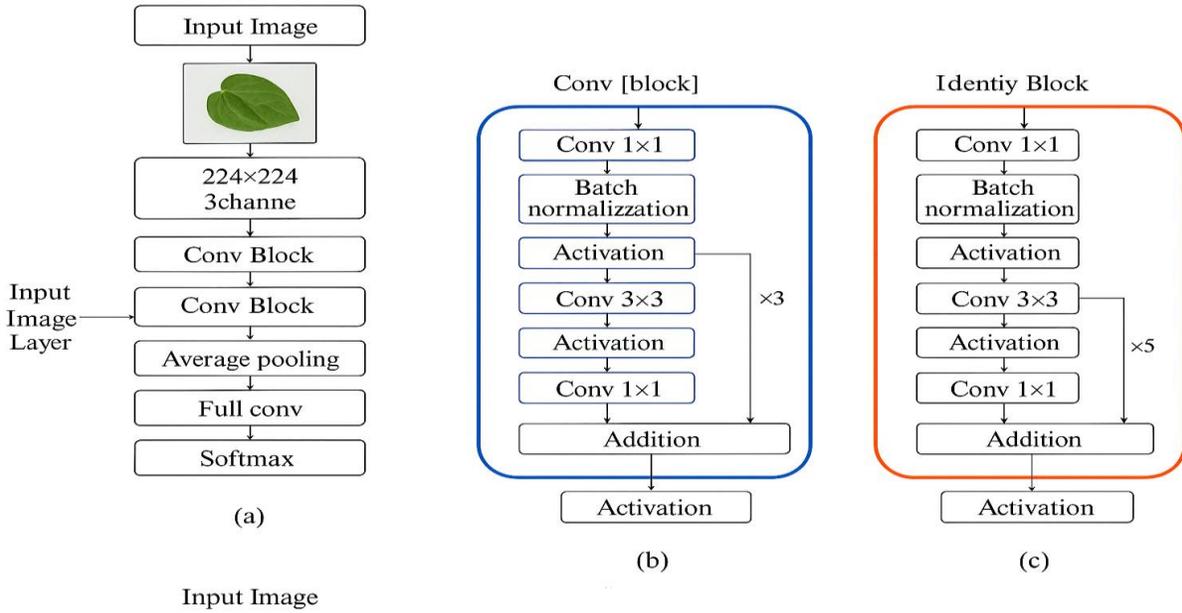


Figure 4. Architecture of the Proposed ResNet50 Model for Cotton Leaf Disease Classification

5.2 Custom Classification Head

A custom classification head was appended to the ResNet50 backbone, comprising a Global Average Pooling layer, Dropout (0.25), and a Dense softmax layer with seven outputs corresponding to the disease classes: Bacterial Blight, Curl Virus, Fusarium Wilt, Healthy, Pest Damage, Powdery Mildew, and Target Spot.

5.3 Fine-Tuning Strategy

Feature specialization was achieved through selective fine-tuning of higher ResNet50 layers, with the learning rate reduced from 1e-3 to 1e-4. This strategy improved disease-specific representation learning while preventing overfitting and resulted in strong generalization performance on the test dataset.

VI. GRAD-CAM ALGORITHM FOR EXPLAINABLE DISEASE CLASSIFICATION

Deep learning models, although highly accurate, often behave as black-box systems with limited interpretability. To address this challenge, the present study integrates Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize and interpret the decision-making process of ResNet50 models. Grad-CAM generates a class-specific heatmap that

highlights the discriminative regions in the leaf image that contributed most to the predicted disease class.

6.1 Conceptual Background

Grad-CAM leverages the gradients flowing into the final convolutional layer of a CNN to produce localization maps. This layer is chosen because convolutional layers retain spatial information essential for visual interpretation, unlike fully connected layers that flatten features.

The heatmap produced by Grad-CAM helps determine whether the model is focusing on biologically relevant symptoms, such as:

- lesion boundaries
- vein discoloration
- fungal powdery patches
- necrotic or chlorotic spots
- curling patterns
- pest bite marks

This confirmatory visualization increases trust and transparency essential for deployment in agricultural settings.

6.2 Mathematical Formulation of Grad-CAM

Let A^k denote the k -th feature map of the selected convolutional layer, and y^c represent the score for class

c (before softmax). The Grad-CAM heatmap is computed as follows:

Step 1: Compute Gradients

Gradients of the class score y^c with respect to the feature map A^k are computed as:

$$\partial y^c / \partial A^k$$

These gradients show how sensitive the class score is to each feature map.

Step 2: Global Average Pooling (Weight Calculation)

The importance weight α^{kc} for each feature map A^k is calculated as:

$$\alpha^{kc} = (1 / Z) \sum_i \sum_j (\partial y^c / \partial A^{k_{ij}})$$

Where

- i, j = spatial indices
- Z = total number of pixels ($H \times W$)

These weights represent how strongly each feature map influences the prediction.

Step 3: Compute the Weighted Feature Map

The Grad-CAM map is computed as:

$$L^c (\text{Grad-CAM}) = \text{ReLU} (\sum_k \alpha^{kc} A^k)$$

ReLU ensures that only positive contributions are included.

Step 4: Heatmap Normalization

The heatmap is normalized as follows:

$$\hat{L} = (L - L(\min)) / (L(\max) - L(\min))$$

The normalized heatmap is then resized to 224×224 and overlaid on the original cotton leaf image.

6.3 Algorithm 2: Grad-CAM for Cotton Leaf Disease Visualization

Input:- Trained CNN model- Input leaf image I- Target class c

Output:

- Grad-CAM heatmap H^c

Algorithm Steps

1. Pre-process image I and feed it into the model.
 2. Identify the last convolutional layer L^* .
 3. Compute prediction scores and obtain y^c .
 4. Compute gradients: $\partial y^c / \partial A^k$
 5. Calculate importance weights α^{kc} using global average pooling.
 6. Compute Grad-CAM map: $L^c = \sum_k \alpha^{kc} A^k$
 7. Apply ReLU.
 8. Normalize and resize the heatmap.
 9. Overlay the heatmap on the original image.
- End.

6.4 Interpretation of Grad-CAM Heatmaps

Grad-CAM heatmaps generated during this study revealed:

- Red/Yellow regions → high model attention
- Blue regions → minimal contribution to prediction

Sample observations include:

- Curl Virus: highlighted leaf edges and twisted regions
- Bacterial Blight: dark necrotic patches
- Powdery Mildew: white fungal spread
- Pest: irregular bite marks
- Target Spot: circular lesion centres
- Fusarium Wilt: vein discoloration and yellowing zones

These confirm that the model focuses on correct biological symptoms rather than background noise.

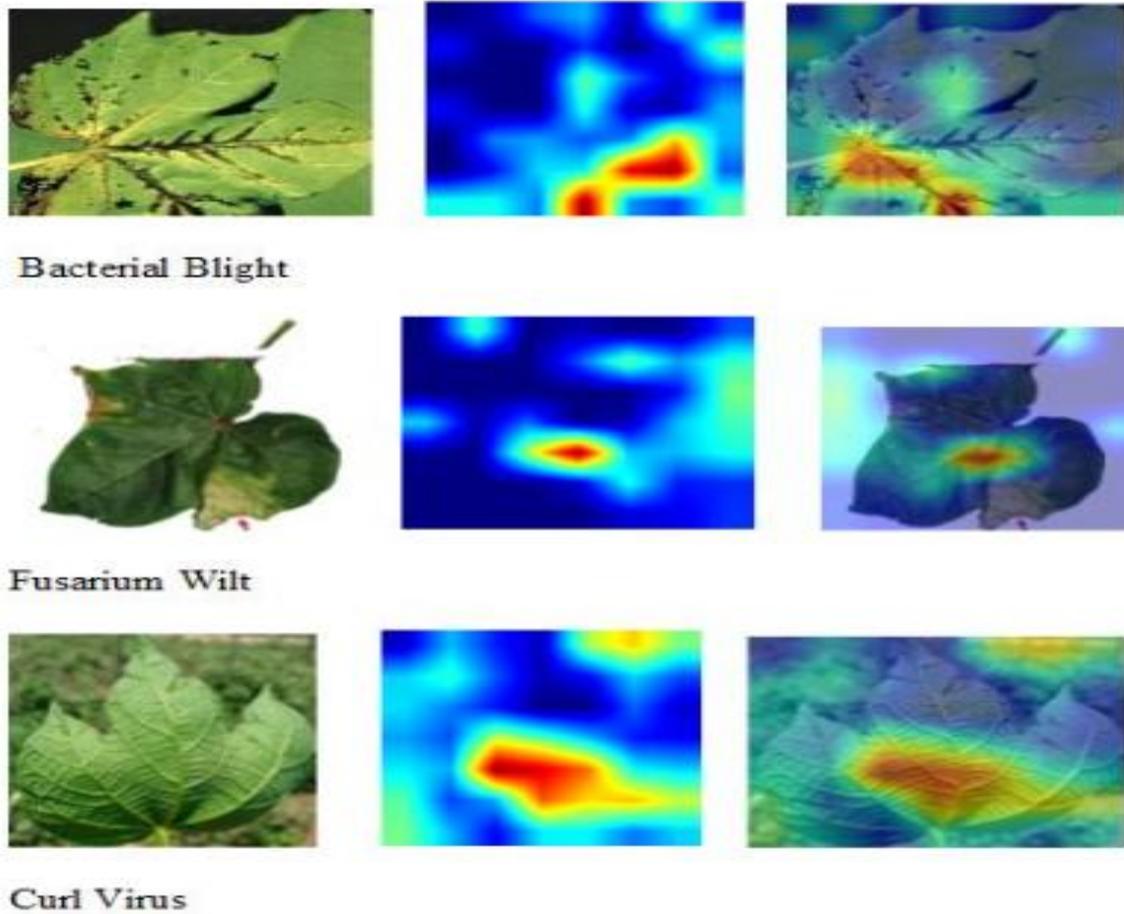


Figure 5. Grad-CAM Heat Map Highlighting Discriminative Regions in Cotton Leaf Images.

6.5 Significance of Grad-CAM in Agriculture

Grad-CAM enhances cotton disease classification by:

- ensuring trust and transparency
- validating correctness of CNN predictions
- helping field experts confirm symptom regions
- avoiding misclassification due to background artifacts
- supporting adoption in real-world farming applications

Thus, integrating explainability transforms the system from a simple classifier into a reliable expert decision-support tool.

VII. RESULTS AND DISCUSSION

This section presents the evaluation results of the proposed deep learning framework, including the performance of ResNet50 models on the test dataset and the interpretability analysis using Grad-CAM. The

results demonstrate that the integration of transfer learning and explainable AI significantly enhances both accuracy and transparency in cotton leaf disease detection.

7.1 Training and Validation Performance

The ResNet50 model was trained using a two-stage transfer learning strategy, as described in Section 4.5. This approach enabled the model to effectively leverage pre-trained ImageNet features while adapting to cotton leaf disease characteristics.

During training, the model demonstrated stable learning behaviour and strong generalization capability. The effectiveness of the training strategy is reflected in the high classification performance achieved on the independent test dataset.

ResNet50 Performance Summary

- Overall Test Accuracy: 99.2%
- Weighted Average Precision: 99.3%

- Weighted Average Recall: 99.3%
- Weighted Average F1-score: 99.3%
- Test set (original) = 562
- Number of Evaluation Samples: 839 (augmented)

Class	Precision	Recall	F1-Score	Support
Bacterial_Blight	1.000	0.980	0.990	151
Curl_Virus	1.000	1.000	1.000	146
Fusarium_Wilt	1.000	1.000	1.000	146
Healthy	0.993	1.000	0.997	147
Pest	1.000	0.964	0.982	83
Powdery_Mildew	1.000	1.000	1.000	83
Target_Spot	0.943	1.000	0.971	83
Accuracy	0.992	0.992	0.993	839
Macro Avg	0.991	0.992	0.991	839
Weighted Avg	0.993	0.993	0.993	839

Table 2. Precision, Recall, and F1-score performance of the ResNet50 model on the test dataset.

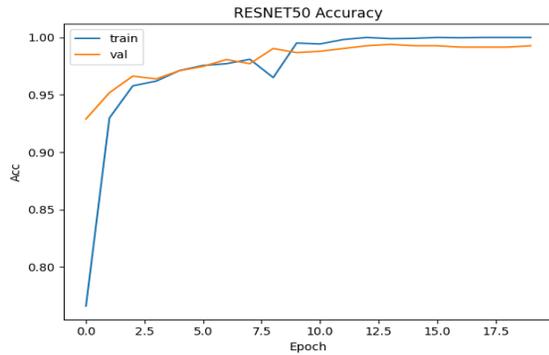


Figure 6. Training accuracy vs validation accuracy across epochs.

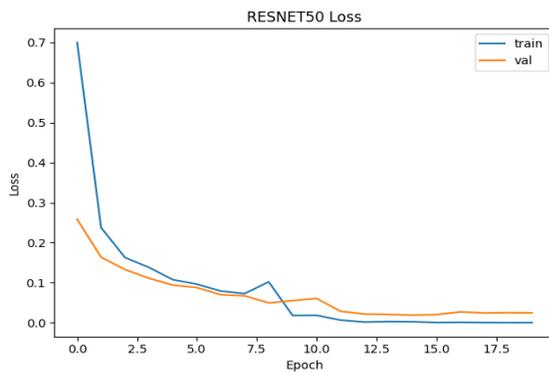


Figure 7. Training loss vs validation loss across epochs.

7.2 Evaluation Metrics

To provide a comprehensive assessment, standard evaluation metrics were calculated for all seven classes:

- Accuracy: Overall predictive performance
- Precision: Correctness of positive predictions
- Recall: Ability to identify all samples of a class
- F1-Score: Balance between precision and recall
- Support: Number of samples per class

ResNet50 achieved high precision and recall for Powdery Mildew, Healthy, and Curl Virus, while slightly lower performance was observed for visually overlapping classes such as Fusarium Wilt and Bacterial Blight.

7.3 Confusion Matrix Analysis

The confusion matrix provides deeper insight into per-class performance. Key findings include:

- High correct classification for Healthy and Powdery Mildew due to distinct texture patterns.
- Moderate confusion between Bacterial Blight and Target Spot, attributed to similarities in brown lesions and dark patches.
- Curl Virus was recognized accurately due to its unique leaf distortion patterns.
- Pest Damage was occasionally confused with Target Spot when bite marks resembled circular lesions.

This analysis confirms the model's strong generalization ability while highlighting biological overlaps that may challenge even human experts.

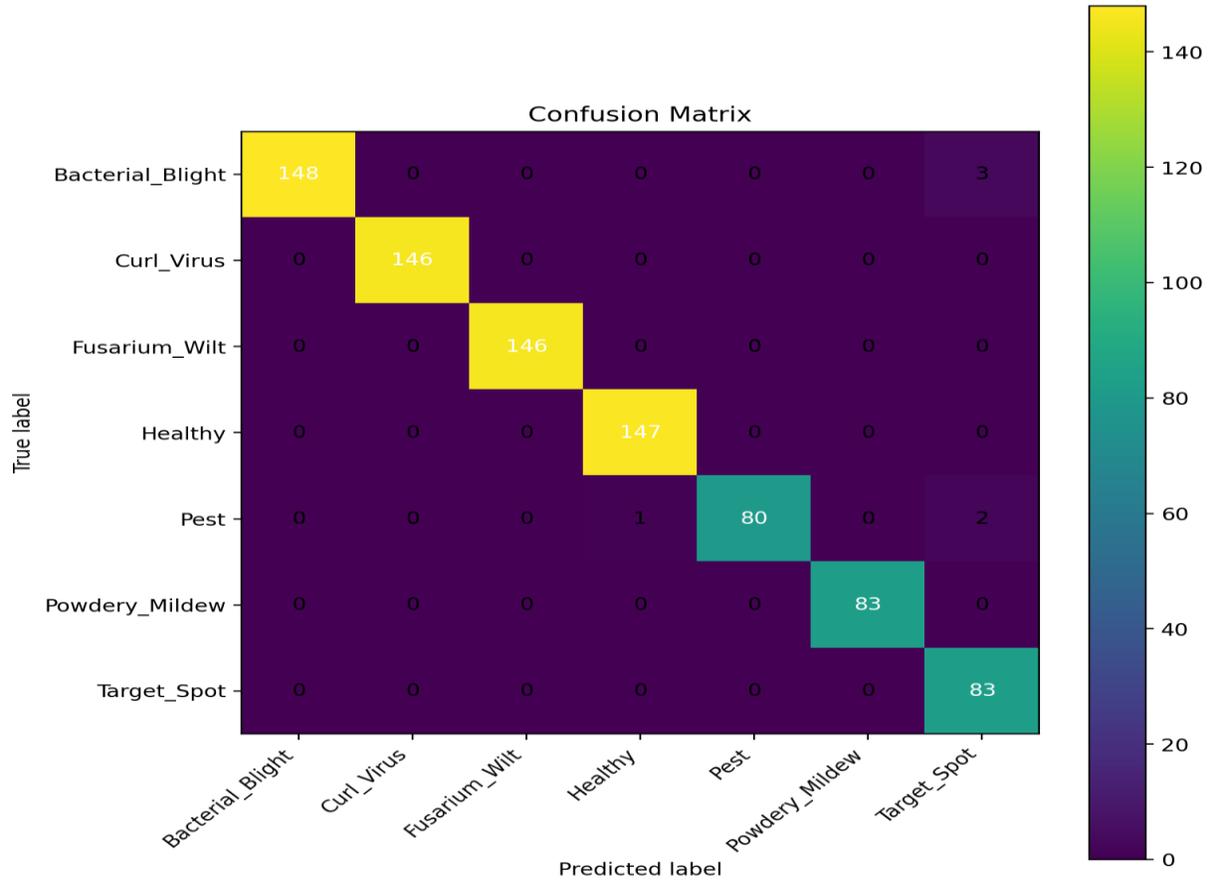


Figure 8 Confusion matrix for seven-class classification on the test dataset

7.4 Grad-CAM Visualization Results

Grad-CAM heatmaps were generated for representative test images of all seven classes. The visualizations revealed that the CNN models consistently focused on biologically relevant regions, confirming the reliability of model predictions.

Key Observations from Grad-CAM Heatmaps

1. Curl Virus

- Heatmaps illuminated curled and twisted leaf edges.
- High attention on vein deformation patterns.

2. Bacterial Blight

- Red zones concentrated on dark necrotic patches.
- Model accurately localized lesion boundaries.

3. Fusarium Wilt

- Focused on vein yellowing and discoloration areas.
- Captured subtle early-stage patterns.

4. Powdery Mildew

- Bright activation over fungal powder-like regions.
- Distinct white patches correctly highlighted.

5. Target Spot

- Activated circular lesion centers and surrounding rings.
- Heatmaps aligned with typical disease morphology.

6. Pest Damage

- Focused on irregular bite marks and torn edges.
- Background leaf areas showed minimal activation.

7. Healthy

- Very low red activation, confirming no disease pattern detected.
- Model attention distributed uniformly across the leaf.

These heatmaps confirm that the model does not base predictions on irrelevant background features such as soil, sky, or hand-held regions—demonstrating strong robustness and reliability.

7.5 Discussion

The integration of Grad-CAM significantly enhances the interpretability of CNN predictions. The ability to visualize disease-specific areas helps validate whether the model is making decisions based on correct biological symptoms rather than background noise.

The integration of Grad-CAM significantly enhances the reliability of deep learning-based cotton disease classification by combining high accuracy with interpretability. The visual explanations validate that predictions are driven by genuine disease symptoms, supporting expert verification and farmer confidence. Overall, the proposed framework demonstrates strong potential for deployment as a transparent, field-ready decision-support system for early cotton disease diagnosis.

Overall, the proposed explainable deep learning framework demonstrates promising potential for developing reliable, field-ready cotton disease diagnosis systems.

VIII. CONCLUSION

This study presented an explainable deep learning framework for cotton leaf disease classification using a fine-tuned ResNet50 model integrated with Grad-CAM visualization. The proposed approach achieved high classification accuracy while providing visual insights into disease-specific regions, improving model transparency and reliability. The results demonstrate the effectiveness of transfer learning for agricultural disease diagnosis and highlight the importance of explainability for real-world deployment. Future work will focus on expanding the dataset with field images and developing a farmer-oriented decision support system.

REFERENCES

- [1] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, 1419.
- [2] Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, 145, 311–318.
- [3] Patil, S. B., Bodhe, S. K., & Pawar, P. S. (2020). Leaf disease severity measurement using image processing. *International Journal of Engineering and Advanced Technology*, 9(3), 217–222.
- [4] Too, E. C., Yujian, L., Njuki, S., & Yingchun, L. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture*, 161, 272–279.
- [5] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (pp. 618–626).
- [6] Barbedo, J. G. A. (2019). Plant disease identification from individual lesions and spots using deep learning. *Biosystems Engineering*, 180, 96–107.
- [7] Chouhan, S. S., Singh, U. P., Jain, S., & Jain, S. (2020). Automated plant disease detection and identification using deep learning. *Journal of Artificial Intelligence and Systems*, 2(1), 1–15.
- [8] Basavaiah, J., & Anthony, A. A., “Tomato leaf disease classification using feature extraction and machine learning techniques,” *International Journal of Engineering Research & Technology*, vol. 10, no. 5, pp. 45–51, 2021.
- [9] Chen, J., Zhang, D., Nanehkaran, Y. A., & Li, D., “Detection of tomato leaf diseases using deep learning and feature fusion techniques,” *Computers and Electronics in Agriculture*, vol. 182, pp. 106046, 2021.
- [10] Agarwal, M., Gupta, S. K., & Biswas, K. K., “Development of efficient CNN model for tomato crop disease identification,” *Sustainable Computing: Informatics and Systems*, vol. 28, pp. 100407, 2020.
- [11] Cengil, E., & Çınar, A., “A hybrid deep learning based approach for plant leaf disease classification,” *Applied Soft Computing*, vol. 108, pp. 107402, 2021.
- [12] Khan, M. A., Akram, T., Sharif, M., Saba, T., & Javed, K., “Automated detection and classification of tomato leaf diseases using

- optimized segmentation and deep learning,” IEEE Access, vol. 7, pp. 128043–128056, 2019.
- [13] Zhao, Y., Zhang, Y., Liu, J., & Xiong, Z., “Improved convolutional neural network with attention mechanism for tomato disease recognition,” *Biosystems Engineering*, vol. 194, pp. 81–94, 2020.
- [14] Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., & Stefanovic, D. (2016). Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience*, 2016, 3289801.
- [15] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- [16] Rangarajan, A. K., Purushothaman, R., & Ramesh, A. (2018). Tomato crop disease classification using pretrained deep learning models. *Procedia Computer Science*, 133, 1040–1047.
- [17] Picon, A., Alvarez-Gila, A., Seitz, M., Ortiz-Barredo, A., Echazarra, J., & Johannes, A. (2019). Crop disease classification using CNNs in the wild. *Computers and Electronics in Agriculture*, 161, 280–290.
- [18] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90.
- [19] Ghosal, S., et al. (2019). Explainable AI for plant stress phenotyping. *Plant Methods*, 15, 123.
- [20] Saleem, M. H., et al. (2022). Plant disease classification using deep learning: A review. *Artificial Intelligence in Agriculture*, 6, 79–94.