International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI: 10.55041/ISJEM05979
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Deep Learning-Based Detection of Plant Diseases

B. Rajasekharam[1], Voruganti Anuradha[2], Madem Satya Nooka Vijay Kumar[3], Saka Mary Jasmine[4], Kashv Yadav[5], Maradana Bhanu Prakash[6]

[1]Assistant Professor, [23456]Students
[123456]Department of Computer Science and Engineering
Visakha Institute of Engineering and Technology
Narava, Visakapatnam
Andhra Pradesh, India

Abstract:
Plant diseases pose a significant threat to agricultural productivity and global food security, making early and reliable detection essential. This paper presents a deep learning-based approach for automated plant disease detection using a fine-tuned ResNet50 model. The system classifies plant leaf images into 38 categories, covering both healthy and diseased conditions, and achieves a validation accuracy of 99.13%, demonstrating strong generalization capability.

To enhance model transparency, Grad-CAM is integrated to visually highlight disease-affected regions, enabling users to understand the reasoning behind predictions. The proposed model is deployed as a web-based application that provides real-time predictions along with confidence scores, visual explanations, condition analysis, and treatment recommendations.

In addition, the system incorporates multilingual support and crop-specific information to improve accessibility and usability. The results indicate that the proposed approach is not only highly accurate but also practical and interpretable, making it suitable for real-world agricultural applications.

Keywords: Deep Learning, Grad-CAM, Precision Agriculture, Transfer Learning, Plant Disease Detection, ResNet50.

1.INTRODUCTION

Agriculture plays a crucial role in ensuring global food security; however, plant diseases remain a major factor that reduces crop yield and quality. Early and accurate detection is essential, but conventional methods based on manual inspection are often time-consuming, subjective, and not suitable for large-scale use.

Recent advances in deep learning have enabled automated and scalable solutions for plant disease detection. Convolutional Neural Networks (CNNs) are capable of learning complex features such as color, texture, and patterns in plant leaves. Among these, ResNet50, with its residual learning capability, provides efficient feature extraction and improved classification performance. Furthermore, transfer learning helps enhance accuracy while reducing training time.

In addition to accuracy, interpretability is important for practical agricultural applications. To address this, the proposed system incorporates Gradient-weighted Class Activation Mapping (Grad-CAM), which highlights the regions of the leaf that influence the model's predictions, thereby improving transparency and user confidence.

The system is developed as a web-based application that allows users to upload plant leaf images and receive real-time predictions along with confidence scores, heatmaps, condition analysis, and treatment recommendations. It also supports multilingual output, enabling users to view results in English, Hindi, Telugu, Spanish, French, and German.

Overall, the proposed system integrates deep learning, explainable artificial intelligence, and real-time deployment to provide an accurate and interpretable solution for plant disease detection.

However, many existing approaches primarily focus on improving prediction accuracy while overlooking interpretability and practical usability. In real-world agricultural environments, users require not only accurate predictions but also clear explanations and actionable insights.

To address these challenges, this work proposes an integrated system that combines high accuracy, explainability, and real-time usability within a unified framework.

2. LITERATURE SURVEY

Deep learning has significantly improved plant disease detection compared to traditional image processing techniques that rely on handcrafted features such as color, texture, and shape, which often fail under varying environmental conditions. Convolutional Neural Networks (CNNs) are particularly effective as they can automatically learn meaningful feature representations from plant leaf images.

Mohanty et al. [1] demonstrated the effectiveness of deep CNN models such as AlexNet and GoogleNet using the PlantVillage dataset, achieving high classification accuracy. Similarly, Sladojevic et al. [2] proposed a CNN-based approach for identifying multiple plant diseases, while Ferentinos [3] reported very high accuracy across different crop species using deep learning techniques.

Transfer learning has further improved performance while reducing training complexity. Too et al. [4] compared several pre-trained models, including ResNet50, VGG16, and DenseNet, and found that ResNet-based architectures perform better due to their residual learning capability. Zhang et al. [5] also showed that fine-tuning pre-trained networks can significantly improve classification accuracy in multi-class disease detection tasks.

Despite these advancements, many existing approaches lack interpretability, which is important for real-world applications. Selvaraju et al. [6] introduced Gradient-weighted Class Activation Mapping (Grad-CAM), a technique that provides visual explanations by highlighting the regions that influence model predictions.

Recent studies have also focused on real-time deployment. Brahimi et al. [7] proposed a CNN-based system for disease detection, while Ramcharan et al. [8] developed a mobile-based solution for field-level diagnosis. However, most of these systems focus mainly on prediction and do not include additional features such as explainability, treatment recommendations, or crop-specific information.

Therefore, there is a need for an integrated system that combines accuracy, interpretability, and real-time usability. The proposed work addresses this gap by incorporating a fine-tuned ResNet50 model with Grad-CAM visualization in a web-based application, along with condition analysis, treatment recommendations, and multilingual support.

3. METHODOLOGY

The proposed system is designed as a complete end-to-end framework for automated plant disease detection, integrating deep learning-based classification, explainable artificial intelligence, and real-time web deployment. The methodology includes several stages such as dataset preparation, preprocessing, model development, training, interpretability, and system integration.

3.1 Dataset Description

The dataset is obtained from the PlantVillage dataset available on Kaggle. It consists of plant leaf images categorized into 38 classes, including both healthy and diseased samples across multiple crop species such as apple, tomato, corn, potato, and grape. The dataset provides sufficient diversity in terms of disease patterns and environmental variations. It is divided into training and validation sets in an 80:20 ratio to ensure reliable and unbiased evaluation.

3.2 Data Preprocessing and Augmentation

All images are resized to 224×224 pixels to match the input requirements of the ResNet50 model. Data augmentation techniques such as random cropping, horizontal flipping, rotation, and color adjustments are applied during training to improve generalization and reduce overfitting. Validation images are resized and center-cropped to maintain consistency. Additionally, all input images are normalized using ImageNet mean and standard deviation values to ensure proper scaling.

3.3 Model Architecture

The system uses a fine-tuned ResNet50 convolutional neural network based on transfer learning. The pre-trained backbone is utilized for feature extraction, while the final layers are modified for multi-class classification. A fully connected layer with 512 neurons

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI: 10.55041/ISJEM05979
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

is added, followed by a dropout layer for regularization, and an output layer with 38 neurons corresponding to the target classes. This architecture ensures efficient feature learning and high classification accuracy.

3.4 Training Strategy

A two-phase training strategy is adopted to optimize model performance. In the initial phase, only the final classification layer is trained using the Adam optimizer with a learning rate of 0.001. In the fine-tuning phase, the entire network is trained with a reduced learning rate of 0.0001 to refine feature representations. The model is trained with a batch size of 64 using categorical cross-entropy loss, ensuring stable convergence and improved generalization.

The model is optimized using categorical cross-entropy loss, defined as:

$$L = - \Sigma (y_i \log(\hat{y}_i))$$

where y_i represents the ground truth label and $\hat{y}_i$ denotes the predicted probability for class i.

3.5 Model Performance

The model achieves a training accuracy of 99.50% and a validation accuracy of 99.13%, indicating strong performance with minimal overfitting. The close alignment between training and validation results demonstrates that the model generalizes well to unseen data.

3.6 Explainability using Grad-CAM

To enhance interpretability, Gradient-weighted Class Activation Mapping (Grad-CAM) is integrated into the system. This method generates heatmaps that highlight the regions of the input image that influence the model's prediction. For diseased samples, the highlighted regions correspond to infected areas, while healthy samples show minimal activation, validating the model's behavior.

3.7 Web-Based System Implementation

The model is deployed as a web-based application using the Flask framework, enabling real-time plant disease detection. Users can upload plant leaf images through a user-friendly interface, after which the images are preprocessed and passed to the trained model for prediction. The system outputs the predicted disease class along with a confidence score.

Grad-CAM is applied to generate visual explanations, allowing users to better understand the basis of the model's prediction. The system also includes a condition analysis module that describes symptoms and possible causes, along with a treatment recommendation module that suggests preventive and corrective measures.

Additional features include a crop information module and a statistical dashboard to support better decision-making. A multilingual translation module enables users to view results in multiple languages, including English, Hindi, Telugu, Spanish, French, and German, through a language selection interface. A disclaimer is also included to inform users about possible variations in prediction accuracy for real-world images.

The system interface displays the uploaded leaf image along with the predicted disease class, confidence score, and Grad-CAM heatmap, providing a clear and informative representation of the results.

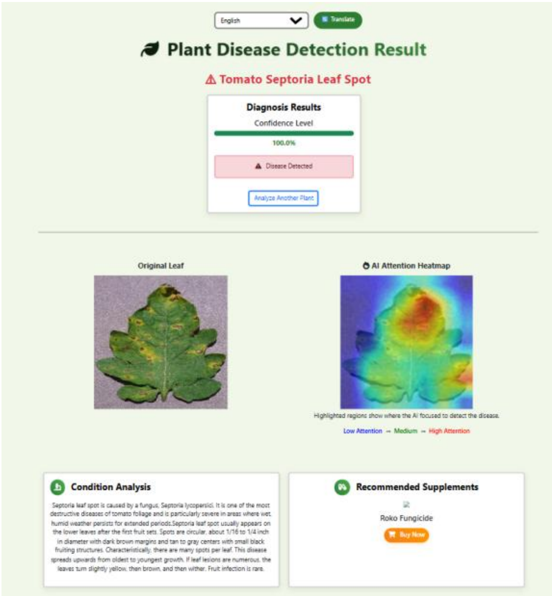

Fig. 1: *Output Interface of the Proposed System showing disease prediction, confidence score, original image, Grad-CAM heatmap, condition analysis, treatment recommendation, and multilingual support.*

3.8 Technology Stack

The system is developed using Python and PyTorch, leveraging a pre-trained ResNet50 model for transfer learning. Data preprocessing and augmentation are implemented using torchvision libraries, and model evaluation is carried out using standard performance metrics. The model training was performed on a GPU-enabled Kaggle environment, which significantly accelerated the training process. The use of hardware acceleration ensured efficient computation and reduced training time.

The web application is built using Flask for backend integration, while the frontend is developed using HTML, CSS, and JavaScript to provide an interactive interface. Additional functionalities, including Grad-CAM visualization, condition analysis, treatment recommendation, and multilingual support, are integrated into the system. Model training is performed on the Kaggle platform.

This technology stack ensures efficient performance, scalability, and suitability for real-world agricultural applications.

4. PROPOSED SYSTEM

The proposed system is designed as an end-to-end framework for automated plant disease detection, integrating deep learning, explainable artificial intelligence, and a web-based interface for real-time usage. The system aims to provide an accurate, interpretable, and practical solution suitable for real-world agricultural environments.

The system consists of multiple modules, namely image input, preprocessing, feature extraction, classification, explainability, and result generation. The interface allows users to upload an image of a plant leaf, preprocessed and sent to a fine-tuned ResNet50 model to be multi-classified into 38 disease classes. This is done so that the identification of plant diseases is efficient and accurate.

Grad-CAM is incorporated into the system to produce heatmaps, which indicate areas of the leaf that affect the model in its prediction and to enhance the interpretability of the model and assist users to learn more about how the model reaches its decision. The system will produce the estimated classification and a confidence rating and visual explanation.

Besides prediction, the system has condition analysis module which describes the detected disease and its features with treatment suggestions in order to help users make the appropriate action. To facilitate a better understanding and decision making a crop information module and a statistical dashboard are also provided.

Overall, the proposed system effectively integrates accurate prediction, interpretability, and user-centric features into a unified framework. This combination enhances its practical applicability and makes it well-suited for real-world deployment in precision agriculture.

5. ARCHITECTURE OF THE PROPOSED SYSTEM

The architecture of the proposed system is shown below, providing a clear representation of the overall workflow. It illustrates how the input image is processed through different stages, including preprocessing, feature extraction, classification, and explainability. Each stage plays an important role in ensuring accurate and reliable disease detection. The structured design also enables efficient data flow between components. This helps improve both the performance and interpretability of the system in real-world applications.

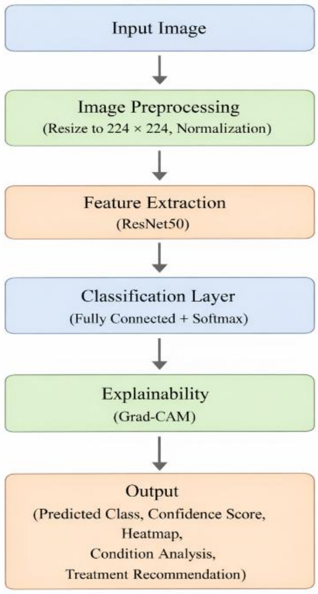

Fig. 2*. The architecture of the proposed system depicting the flow between the prediction and Grad-Cam-based explainability and input image.*

The architecture of the proposed system aims to be structured in the format of the pipeline that incorporates image acquisition, preprocessing, feature extraction, classification, explainability, and result generation as shown in Fig. 2. The design facilitates efficient processing, correct prediction and the interpretability of the detection of the plant diseases is clear.

5.1 Input Layer

The process begins with the input layer, where plant leaf images are provided by the user through the web interface. These images act as the primary input to the system and allow real-time interaction with the model.

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI: 10.55041/ISJEM05979
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

5.2 Image Preprocessing

In the next step, the input images are preprocessed to make them suitable for the deep learning model. This includes resizing all images to 224 × 224 pixels, normalizing them using ImageNet mean and standard deviation values, and converting them into tensor format. These steps help maintain consistency in the data and improve model performance.

5.3 Feature Extraction using ResNet50

After preprocessing, the images are passed through a fine-tuned ResNet50 model for feature extraction. The network learns important patterns such as color variations, textures, and disease-specific characteristics present in plant leaves. The residual connections in ResNet50 help in training deeper networks effectively and enable better feature learning.

5.4 Classification Layer

The extracted features are then passed through fully connected layers, followed by a softmax activation function for classification. In this step, the input image is assigned to one of the 38 predefined classes, along with a confidence score indicating the reliability of the prediction.

5.5 Explainability Module (Grad-CAM)

To improve interpretability, Grad-CAM is integrated into the system. It generates heatmaps that highlight the regions of the input image that influence the model's prediction. This allows users to clearly see which parts of the leaf are affected and understand how the model arrives at its decision. As a result, the model becomes more transparent and easier to trust in practical applications.

5.6 Output Layer

Finally, the system produces the output, which includes the predicted disease class, confidence score, Grad-CAM heatmap, condition analysis, and treatment recommendations. These outputs provide both numerical and visual insights into the prediction, along with useful information to help users understand and manage the detected disease.

Overall, the architecture follows a clear and modular design that ensures smooth data flow from input to output while maintaining accuracy, interpretability, and practical usability.

6. RESULTS AND ANALYSIS

This section presents the performance evaluation of the proposed plant disease detection system using accuracy metrics, training behavior, confusion matrix analysis, and qualitative results. The results help in understanding how well the model performs on both seen and unseen data.

6.1 Dataset Visualization

The following figure presents representative samples from the dataset, illustrating both

healthy and diseased plant leaves.

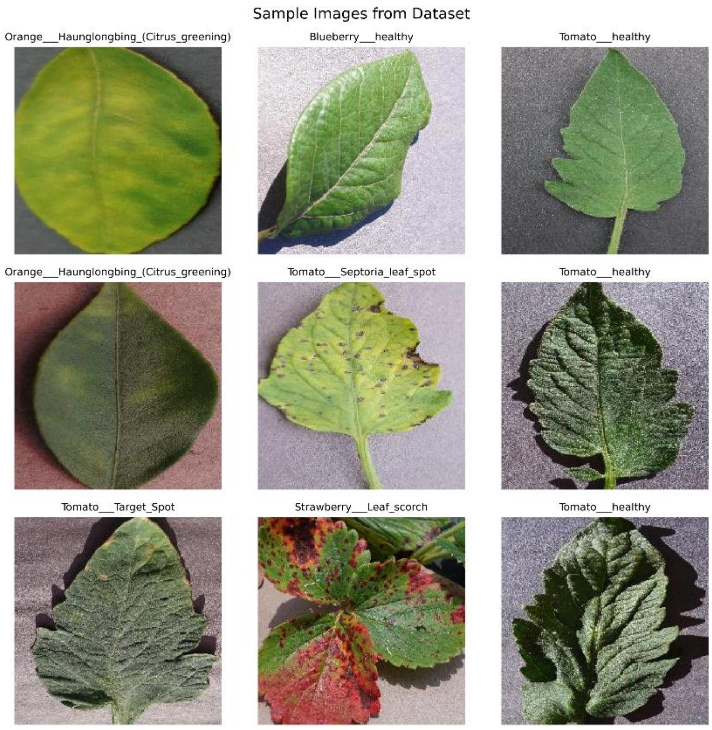

Fig. 3: *Sample Images from the Dataset Representing Healthy and Diseased Plant Leaves.*

The dataset consists of images from 38 different classes, including both healthy and diseased plant leaves across multiple crops. These images show variations in color, texture, and disease patterns, which are important for training the model. Such diversity helps the model learn more robust and generalized features, improving its performance in real-world scenarios.

6.2 Comparative Analysis with Existing Methods

To evaluate the effectiveness of the proposed model, its performance is compared with existing deep learning-based approaches for plant disease detection. Table 1 shows a comparison of classification accuracy with selected state-of-the-art methods.

International Scientific Journal of Engineering and Management (ISJEM)
Volume: 05 Issue: 03 | March – 2026
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata
ISSN: 2583-6129
DOI: 10.55041/ISJEM05979

Table 1: *Comparison of Classification Accuracy with Existing Methods*

Method	**Model**	**Accuracy (%)**
Mohanty et al. (2016)	CNN	99.35
Ferentinos (2018)	CNN	99.53
Too et al. (2019)	Transfer Learning	99.30
Proposed Method	ResNet50	**99.13**

Although the proposed model achieves accuracy comparable to existing approaches, it offers additional advantages in terms of practical usability. Unlike many earlier works, this system integrates explainable AI using Grad-CAM and supports real-time deployment through a web-based interface. These features make the system more useful in real-world agricultural applications.

6.3 Training and Validation Accuracy

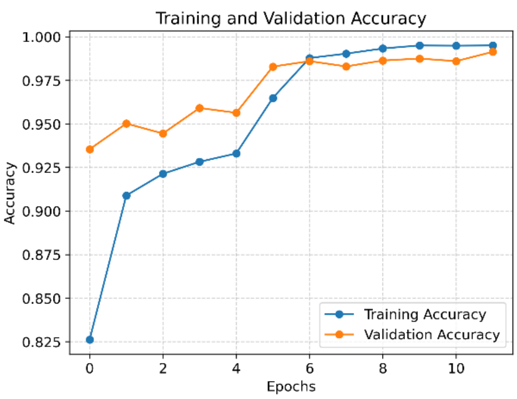

Fig. 4: *Training and Validation Accuracy over Epochs.*

The training accuracy increases steadily from approximately 82% to 99.5%, while the validation accuracy improves to 99.13%. The close alignment between the training and validation curves indicates strong generalization capability and minimal overfitting. This shows that the model is able to learn meaningful features from the training data while maintaining consistent performance on unseen data.

Additionally, the smooth progression of both curves suggests stable learning behavior throughout the training process, without significant fluctuations or instability.

6.4 Training and Validation Loss

These results further indicate that the model is learning effectively and maintaining consistency across training epochs.

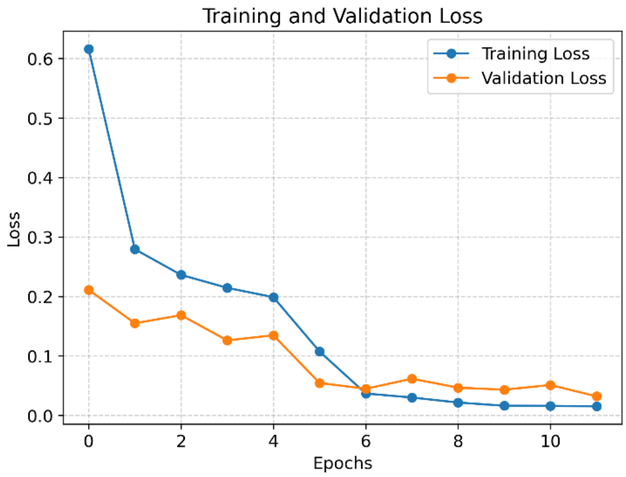

Fig. 5: *Training and Validation Loss over Epochs.*

Both training and validation loss decrease consistently during the training process, with training loss reaching approximately 0.02 and validation loss stabilizing around 0.03–0.05. This indicates stable learning and proper optimization of the model.

The training and validation loss curves provide insight into the optimization process and learning stability of the model. They clearly illustrate how effectively the model minimizes error over successive training epochs. The absence of sharp spikes or divergence in the loss curves further confirms that the model is well-tuned and not suffering from overfitting.

6.5 Confusion Matrix Analysis

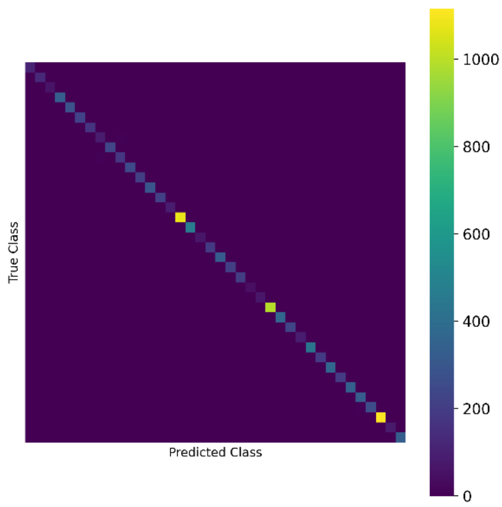

Fig. 6: *Confusion Matrix Illustrating the Performance of the Proposed ResNet50-Based Plant Disease Classification Model across 38 Classes.*

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI: 10.55041/ISJEM05979
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

The confusion matrix shows strong diagonal dominance, indicating that most samples are correctly classified. Minimal off-diagonal values confirm low misclassification and effective distinction between similar disease classes.

This demonstrates that the model has high classification accuracy across different categories and is capable of distinguishing even closely related disease patterns. The overall distribution of values reflects the robustness and reliability of the model in multi-class classification tasks.

6.6 Classification Performance Metrics

The classification results are summarized in terms of precision, recall, and F1-score for selected classes.

Table 2: *Classification Performance Metrics of Selected Classes*

Class Name	Precision	Recall	F1-Score	Support
Apple Scab	1.00	0.97	0.98	115
Apple Black Rot	1.00	1.00	1.00	122
Apple Cedar Rust	1.00	1.00	1.00	58
Apple Healthy	0.99	1.00	1.00	321
Blueberry Healthy	1.00	1.00	1.00	309
Tomato Mosaic Virus	0.96	1.00	0.98	70
Tomato Healthy	1.00	0.99	1.00	324
Overall Accuracy	—	—	99.13 %	—

The classification performance of the proposed model is evaluated using precision, recall, and F1-score metrics. As shown in Table 2, the model achieves consistently high performance across multiple classes, with most classes exhibiting near-perfect scores. The overall accuracy of 99.13% demonstrates the robustness and effectiveness of the proposed ResNet50-based approach.

Furthermore, the balanced precision and recall values indicate that the model performs well in both correctly identifying diseases and minimizing false predictions, which is essential for reliable real-world applications.

6.7 Qualitative Prediction Results

Fig. 7: *Sample Prediction Results Showing Predicted and True Labels.*

The model accurately predicts both healthy and diseased samples on unseen data, confirming its robustness across different plant species and conditions. The close match between predicted and true labels highlights the model's ability to generalize well beyond the training dataset.

These qualitative results further validate that the model is not only accurate in numerical metrics but also reliable in practical scenarios involving diverse leaf images.

6.8 Grad-CAM Visualization

These visual results further confirm that the model focuses on relevant disease-affected regions rather than background features. This improves confidence in the model's predictions and makes the system more reliable for practical use.

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | March – 2026 DOI: 10.55041/ISJEM05979
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

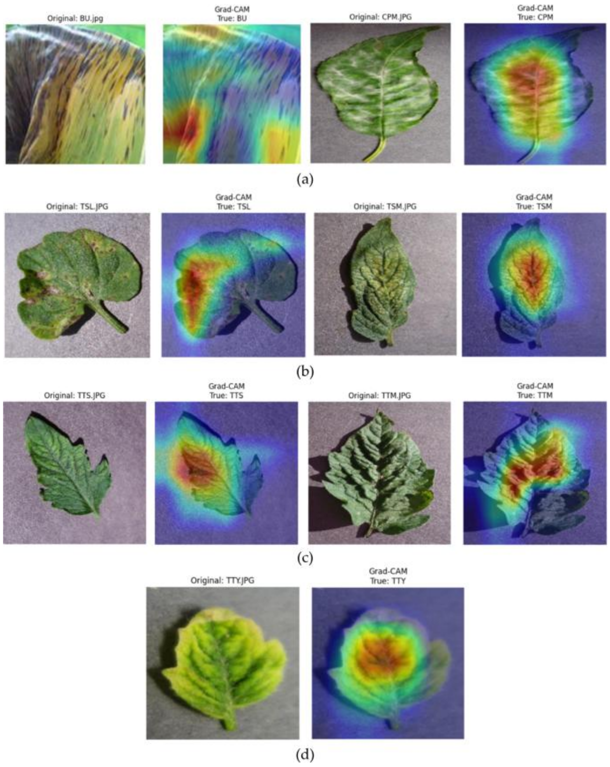

Fig. 8: *Grad-CAM Visualization Highlighting Disease-Affected Regions in Plant Leaf Images.*

To evaluate the interpretability of the proposed model, Gradient-weighted Class Activation Mapping (Grad-CAM) is employed, as shown in Fig. 8. The visualization highlights the regions of the input image that contribute most to the model's prediction. It can be observed that the model focuses on disease-affected areas such as lesions, discoloration, and irregular patterns, which are key indicators for classification.

High-intensity regions in the heatmap correspond to infected portions of the leaf, indicating that the model is effectively learning relevant visual features rather than relying on background information. In contrast, healthy leaf samples exhibit low and uniform activation, reflecting the absence of disease-specific patterns.

These results confirm that the model's predictions are based on meaningful and biologically relevant features. The integration of Grad-CAM therefore improves interpretability, enhances user trust, and supports the practical usability of the system in real-world agricultural applications.

Additionally, this visual explanation capability makes the system more transparent and easier to adopt by users who may not have a technical background in deep learning.

6.9 Model Complexity and Efficiency The model contains approximately 24.5 million parameters, with only about 1.06 million trainable due to transfer learning. This significantly reduces computational cost while maintaining high accuracy, enabling efficient real-time deployment.

This balance between model complexity and performance ensures that the system can be deployed on practical platforms without requiring extremely high computational resources.

6.10 Limitations

Although the model performs well on the dataset, its performance may vary on real-world images due to differences in lighting conditions, background noise, and image quality. These factors can affect prediction accuracy in uncontrolled environments.

To address this, a disclaimer is included in the system to inform users about possible variations and ensure responsible usage. Future improvements may focus on incorporating more diverse real-world datasets to further enhance model robustness.

7. CONCLUSION

This paper presents an end-to-end deep learning-based framework for automated plant disease detection, integrating accurate classification, explainable artificial intelligence, and real-time deployment. The proposed system utilizes a fine-tuned ResNet50 model to classify plant leaf images into 38 categories, achieving a validation accuracy of 99.13%, thereby demonstrating the effectiveness of transfer learning for agricultural image analysis.

To enhance interpretability, Grad-CAM is incorporated to generate visual explanations that highlight disease-affected regions, enabling users to better understand and trust the model's predictions. The system is deployed as a web-based application that provides real-time outputs, including predicted class, confidence score, condition analysis, treatment recommendations, and multilingual support, making it accessible to a diverse group of users.

Experimental evaluation, including accuracy trends, loss convergence, confusion matrix analysis, and qualitative results, confirms the robustness and strong generalization capability of the model across multiple plant species and disease conditions. These results indicate that the proposed approach is reliable and suitable for practical agricultural applications.

ISSN: 2583-6129
DOI: 10.55041/ISJEM05979

Furthermore, the integration of explainable AI with user-oriented features bridges the gap between high-performance models and real-world usability, enhancing the system's potential for adoption in precision agriculture.

Future work will focus on improving robustness using more diverse real-world datasets, enabling mobile-based deployment for field-level usage, and exploring advanced deep learning architectures to further enhance performance and scalability. Overall, the proposed system provides a scalable, interpretable, and efficient solution for intelligent plant disease detection.

REFERENCES

[1] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, 2016.

[2] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks-based recognition of plant diseases," Computational Intelligence and Neuroscience, 2016.

[3] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, 2018.

[4] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," Computers and Electronics in Agriculture, 2019.

[5] S. Zhang, S. Zhang, C. Zhang, and X. Wang, "Cucumber leaf disease identification with global pooling dilated convolutional neural network," Computers and Electronics in Agriculture, 2020.

[6] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2017.

[7] M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep learning for tomato diseases classification," in Proc. IEEE Int. Conf. Image Processing (ICIP), 2017.

[8] A. Ramcharan, K. Baranowski, P. McCloskey, B. Ahmed, J. Legg, and D. P. Hughes, "A mobile-based deep learning approach for cassava disease detection," Frontiers in Plant Science, 2019.

[9] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Proc. Advances in Neural Information Processing Systems (NIPS), 2012.

[10] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.

[11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016.

[12] C. Szegedy, W. Liu, Y. Jia, et al., "Going deeper with convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2015.

[13] F. Chollet, "Xception: Deep learning with depth wise separable convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017.

[14] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017.

[15] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. Int. Conf. Machine Learning (ICML), 2019.

[16] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in Proc. Int. Conf. Learning Representations (ICLR), 2021.

[17] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," Journal of Big Data, 2019.

[18] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[19] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, 2015.

[20] "PlantVillage Dataset," Kaggle. [Online]. Available: https://www.kaggle.com/datasets/abdallahalidev/plantvillage-dataset