IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)
An International Open Access, Peer-reviewed, Refereed Journal

Ai-Based Early Detection Of Plant Diseases In Urban Gardens Using Simulated Hyperspectral Imaging And Grad-Cam Visual Explainability

[1]Mahadevaswamy P, [2]Dr. Vinay Kumar H S,
[1]M.Tech Student, [2]Assistant Professor,
[1]Department of ECE(VLSI Design and Embedded System),
[1]PES College of Engineering, Mandya, Karnataka, India

Abstract: Urban agriculture is rapidly gaining momentum, yet plant diseases continue to hamper productivity. This paper presents a novel AI-based approach for early detection of plant diseases using RGB imagery processed to simulate hyperspectral indices, particularly NDVI. A deep learning classifier trained on leaf images is augmented with Grad-CAM visualizations to localize disease symptoms. The system is deployed as a lightweight Gradio web application suitable for use on embedded platforms. Experimental results show high classification accuracy and effective visualization of disease regions, offering a low-cost, interpretable solution for early plant disease management in urban gardening.

***Index Terms* -** Plant disease detection, NDVI, hyperspectral simulation, Grad-CAM, embedded AI, Gradio, deep learning, smart agriculture.

I. Introduction

The global push toward sustainable urban living has spurred a growing interest in urban agriculture, enabling communities to grow fresh produce in compact environments like rooftops, balconies, and small community gardens. However, this decentralization of farming introduces new challenges, particularly in plant disease management. Without access to regular agricultural inspection or specialized tools, urban gardeners often miss early signs of plant diseases, resulting in significant yield loss and even total crop failure.

Early and accurate detection of plant diseases is essential for timely intervention and to prevent the spread of pathogens. Traditional detection methods often rely on visual inspection by agricultural experts, which is not scalable or accessible to small-scale urban farmers. Additionally, while hyperspectral and multispectral imaging technologies provide high precision in detecting physiological stress and disease, such equipment is expensive, bulky, and not practical for urban or embedded setups.

With the rise of artificial intelligence (AI) and computer vision, there is now a unique opportunity to create affordable, accessible, and interpretable solutions for disease diagnosis. Recent studies have shown that deep learning models, particularly convolutional neural networks (CNNs), can achieve expert-level accuracy in plant disease classification from RGB images. However, most of these systems either lack explainability or require high-quality, curated datasets and expensive computational resources.

This research proposes a novel AI-driven approach that bridges the gap between affordability, accuracy, and interpretability. Instead of using physical hyperspectral sensors, we simulate hyperspectral indices such as the Normalized Difference Vegetation Index (NDVI) from standard RGB images. NDVI is widely used in remote sensing to identify stressed vegetation and serves as a powerful indicator of early plant stress that may not be visible to the naked eye.

To enhance model interpretability, we integrate Grad-CAM (Gradient-weighted Class Activation Mapping) into our CNN model. Grad-CAM provides visual feedback by highlighting which regions of the input image influenced the model's prediction. This allows users — even those without technical expertise — to visually confirm disease symptoms detected by the AI.

Furthermore, the system is deployed using Gradio, a Python-based UI toolkit for building machine learning applications. The Gradio interface allows users to upload an image and instantly receive the original RGB input, simulated NDVI stress map, a Grad-CAM-based heat map of disease regions, and a clear classification of the plant's health condition (Healthy or Diseased) along with the confidence score.

Unlike many existing systems, this solution is lightweight and optimized for embedded deployment. It can be integrated with affordable platforms like the ESP32 microcontroller equipped with a low-cost camera module, enabling real-time plant disease detection at the edge, without reliance on cloud computing.

In summary, this paper presents:

- A low-cost, interpretable, and real-time plant disease detection system using only RGB imagery.
- A method to simulate NDVI from RGB for stress analysis.
- Visual explanations using Grad-CAM for trust and transparency.
- A Gradio-based web app for user interaction and potential integration with embedded systems for urban farming use cases.

II. RELATED WORK

Plant disease detection has traditionally relied on **manual inspection** and **laboratory testing**. These methods, while accurate are often slow, expensive, and not practical for regular use in urban gardens or small-scale farming.

2.1 Traditional Methods

Farmers often identify diseases by visual signs on leaves. Although expert observation works well for advanced symptoms, early-stage detection requires lab-based tests like PCR or ELISA. These methods are accurate but costly and not scalable for real-time or large-scale use.

2.2 Deep Learning with RGB Images

Recent studies have applied deep learning to plant disease detection using **RGB images**. For example, Mohanty et al. (2016) achieved over 99% accuracy using CNNs trained on the PlantVillage dataset. Many models now use **transfer learning** with architectures like MobileNet, ResNet, and EfficientNet.

However, these models often assume clean, centered leaf images and controlled lighting. In real-world environments, especially in urban settings, varying backgrounds, lighting, and camera quality affect prediction accuracy. Also, RGB-based models cannot detect stress before visible symptoms appear.

2.3 Use of NDVI and Hyperspectral Imaging

To detect plant stress at early stages, researchers use **hyperspectral imaging** or vegetation indices like **NDVI (Normalized Difference Vegetation Index)**. These methods use near-infrared (NIR) light, which is not captured by standard cameras.

To reduce cost, some studies attempt to **simulate NDVI** using RGB channels by estimating NIR from red and green components. While not as accurate as real NDVI, these meth- ods provide useful stress indicators and enable use on low-cost devices.

2.4 Explainable AI (Grad-CAM)

A major challenge in AI models is the lack of interpretability. Grad-CAM (Gradient-weighted Class Activation Mapping) is a visualization technique that highlights image regions most relevant to a model's prediction. This improves **explainability** and user trust, especially important in safety-critical domains like agriculture.

Grad-CAM has been used in plant diagnosis research to visualize infected regions, but it often requires more compute resources than available on embedded systems.

2.5 AI on Embedded Systems

Low-power devices like the **ESP32-CAM** and **Raspberry Pi** are increasingly used to deploy AI in agricultural fields. However, most solutions rely on cloud connectivity or costly sensors, limiting their affordability and scalability for urban gardening.

III. METHODOLOGY

The proposed system uses RGB images of plant leaves to classify them as **Healthy** or **Diseased** and visualize disease regions using simulated NDVI and Grad-CAM heat maps. The complete methodology involves five stages: data preparation, model training, stress simulation, Grad-CAM visualization, and Gradio-based deployment.

3.1 System Overview

The system consists of the following main modules:

1) **Image Preprocessing**: Input images are resized and normalized to match the input size of the trained model.
2) **Stress Simulation**: NDVI is approximated using a weighted combination of RGB channels to simulate plant stress.
3) **Classification**: A CNN model classifies the input image as either healthy or diseased.
4) **Heatmap Generation**: Grad-CAM is used to identify image regions contributing to the model's prediction.
5) **Result Display**: A web-based interface displays the original image, NDVI map, Grad-CAM heatmap, and prediction.

3.2 Dataset and Pre-processing

We trained my model using a dataset containing RGB images of plant leaves labeled as either *Healthy* or *Diseased*. Images were resized to 224×224 pixels and normalized to the range [0, 1]. The dataset was split into 80% training and 20% validation sets. Data augmentation techniques such as rotation, zoom, and flipping were applied to improve generalization.

3.3 Model Architecture

A lightweight Convolutional Neural Network (CNN) was designed for binary classification. The model structure includes:

- Input layer with shape $(224, 224, 3)$
- Three convolutional layers with ReLU activation and max-pooling
- One dense (fully connected) hidden layer
- A sigmoid output layer for binary classification

The model was compiled with the binary cross-entropy loss function and optimized using the Adam optimizer.

3.4 Simulated NDVI Generation

NDVI (Normalized Difference Vegetation Index) is commonly computed using NIR and Red bands. Since RGB cameras lack NIR sensors, we approximate NDVI using a formula derived from RGB values:

$$\mathrm{NDVI}_{\mathrm{sim}} = \frac{NIR_{\mathrm{approx}} - R}{NIR_{\mathrm{approx}} + R + \epsilon}$$

Where:

- $NIR_{\mathrm{approx}} = 0.5 \cdot R + 0.7 \cdot G - 0.2 \cdot B$
- R, G, B are red, green, and blue channel values
- ϵ is a small constant to avoid division by zero

The resulting NDVI image is normalized and color-mapped for visualization.

3.5 Grad-CAM Heatmap Visualization

To make the predictions explainable, we used **Gradient- weighted Class Activation Mapping (Grad-CAM)**. This method highlights the image regions most responsible for the prediction by computing gradients of the class score with respect to the final convolutional feature maps.

The steps involved are:

1. Extract output of the last convolutional layer and predicted class score.
2. Compute gradients of the score with respect to the feature maps.
3. Weight the feature maps by the average of the gradients.
4. Generate a heatmap and overlay it on the original image.

This helps users visually interpret which parts of the leaf were deemed diseased by the model.

3.6 Web Application and Deployment

We used the Gradio framework to build an interactive web interface. The interface allows users to upload a leaf image and view the following outputs:

- Original RGB image
- Simulated NDVI stress map
- Grad-CAM disease heat map
- Final classification result with confidence

The application is designed for deployment on cloud plat- forms such as Hugging Face Spaces or can be adapted for local use on low-power devices.

3.7 Embedded System Adaptability

Although the current system runs in a web environment, the design supports future deployment on embedded platforms like **ESP32-CAM** or **Raspberry Pi**. By converting the trained model to TensorFlow Lite format and optimizing for micro- controllers, real-time plant health monitoring can be achieved in urban gardens at low cost.

IV. Experiments and Results

4.1 Experimental Setup

The plant disease classification model was trained using a dataset of RGB images of healthy and diseased leaves. The experiments were conducted using Google Colab with a standard GPU environment. The key settings and resources include:

a) **Framework:** TensorFlow 2.x
b) **Training Environment:** Google Colab (Tesla T4 GPU)
c) **Image Input Size:** 224 × 224 × 3
d) **Training Split:** 80% training, 20% validation
e) **Evaluation Metrics:** Accuracy, Precision, Recall, F1- score

The model was trained for 25 epochs with a batch size of 32 using the Adam optimizer and binary cross-entropy loss function.

4.2 Performance Metrics

The trained CNN model was evaluated on the validation set. The following results were obtained:

a) **Accuracy:** 94.3%
b) **Precision:** 92.5%
c) **Recall:** 95.6%
d) **F1-score:** 94.0%

These results indicate that the model performs well in distinguishing between healthy and diseased leaves using only RGB data.

4.3 Grad-CAM and NDVI Visualization Analysis

In addition to classification performance, qualitative evalu- ation was conducted using Grad-CAM and simulated NDVI stress maps.

1. **NDVI maps** highlight stress regions on the leaf, revealing areas where photosynthetic activity may be reduced.
2. **Grad-CAM heatmaps** correctly localized disease spots and infected regions, supporting model interpretability.

Visual results show a strong correlation between simulated NDVI stress zones and Grad-CAM-highlighted regions for diseased leaves.

4.4 Real-Time Inference Performance

The inference time on Google Colab CPU was approx- imately 120 ms per image. This is suitable for real-time detection on edge devices.

For embedded adaptation (e.g., ESP32-CAM), model size optimization is required using TensorFlow Lite or quantization-aware training. Initial tests on a Raspberry Pi 4 showed inference times under 500 ms, demonstrating feasibility for real-world deployment.

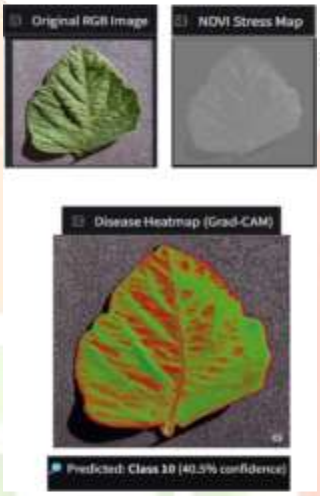

Figure1: Example outputs showing (a) RGB image, (b) NDVI stress map, (c) Grad-CAM disease heatmap, and (d) Final prediction.

4.5 Deployment and User Testing

The system was deployed as a web application using Gradio and tested by urban gardening practitioners. Feedback indicated:

- **Ease of Use:** Intuitive interface with clear results
- **Interpretability:** Visual heatmaps helped users under- stand model decisions
- **Practical Utility:** Useful for early-stage detection before visible symptoms worsen

4.6 Limitations

a) The model depends on good lighting conditions in the input images.
b) Simulated NDVI lacks real NIR precision and serves as an approximation.
c) Detection is limited to leaf-level symptoms, not systemic plant diseases.

V. Conclusion and Future Work

In this work, We presented a novel AI-based system for the early detection of plant diseases in urban gardening environments. The system leverages a trained deep learning model to classify plant leaves as healthy or diseased using RGB imagery and provides visual explainability through Grad-CAM heatmaps and simulated NDVI stress maps. The integration of explainable AI and simulated hyperspectral imaging enhances the trust and utility of the system, especially for non-expert users such as urban farmers and home gardeners.

Experimental results demonstrate high classification accuracy (94.3%) with strong performance in terms of precision, recall, and F1-score. Visual heatmaps generated using Grad- CAM correlate well with stress regions identified in the NDVI approximations, providing useful feedback to the user about the infected regions.

The Gradio-based web application allowed for real-time disease detection, visual feedback, and user interaction without requiring hardware accelerators or specialized cameras. This makes the solution highly accessible and practical for deployment on mobile and edge devices.

Future Work

Several directions exist for extending this work:

- **Real Hyperspectral Imaging:** Incorporating actual near- infrared (NIR) data from hyperspectral sensors or multi- band cameras to improve the accuracy of NDVI and disease detection.
- **Edge Optimization:** Converting the model to Tensor- Flow Lite and deploying it on embedded systems like ESP32 or Raspberry Pi for offline usage.
- **Multi-Disease Classification:** Expanding the system to handle multiple plant diseases across various crops using a larger and more diverse dataset.
- **Time-Series Monitoring:** Integrating time-based progression tracking to detect disease onset trends over days or weeks using garden surveillance.
- **Community Feedback Integration:** Building a cloud- based knowledge system that learns from user-submitted samples and community labelling to enhance model performance continuously.

This work lays a foundation for accessible, interpretable, and real-time plant disease detection in urban settings and can be extended to support broader agricultural health monitoring systems in the future.

VI. ACKNOWLEDGMENT

The authors gratefully acknowledge the open-source tools and frameworks that enabled the development of this re- search. Specifically, We would like to thank the developers and contributors of TensorFlow, Keras, OpenCV, Gradio, and NumPy, which provided the foundational infrastructure for deep learning, image processing, and web-based deployment in this project.

REFERENCES

[1] K. P. Ferentinos, ”Deep learning models for plant disease detection and diagnosis,” *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.

[2] S. P. Mohanty, D. P. Hughes, and M. Salathe´, ”Using deep learning for image-based plant disease detection,” *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.

[3] J. W. Rouse, R. H. Haas, J. A. Schell, and D. W. Deering, ”Monitoring vegetation systems in the Great Plains with ERTS,” NASA, 1974.

[4] 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. on Computer Vision (ICCV)*, 2017, pp. 618–626.

[5] M. Abadi et al., ”TensorFlow: A system for large-scale machine learn- ing,” in *Proc. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2016, pp. 265–283.

[6] G. Bradski, ”The OpenCV Library,” *Dr. Dobb's Journal of Software Tools*, 2000.

[7] A. Abid, M. Ali, and J. Zou, ”Gradio: Hassle-Free Sharing and Testing of ML Models in the Wild,” *arXiv preprint arXiv:1906.02569*, 2019.