International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

LEAFSCANDL: A DEEP LEARNING FRAMEWORK FOR AUTOMATED DETECTION OF FOLIAR PATHOLOGIES

Aditya Ranjan[1], Aman Dubey[2],Anuj Dhakad[3], Anurag Rathore[4], Prof. Preeti Kushwaha[5]

[1] Oriental Institute of Science and Technology Madhya Pradesh adityaranjan7856@gmail.com
[2] Oriental Institute of Science and Technology Madhya Pradesh ad2647237@gmail.com
[3] Oriental Institute of Science and Technology Madhya Pradesh dhakadanuj39@gmail.com
[4] Oriental Institute of Science and Technology Madhya Pradesh anuragrathore840@gmail.com
[5] Department of Computer Science Oriental Institute of Science and Technology Madhya Pradesh

ABSTRACT—

Foliar pathologies present a persistent threat to global agricultural productivity, leading to significant yield re- duction and economic instability. While recent advancements in Deep Learning have demonstrated high efficiency in disease recognition, a critical gap remains between high-performance experimental models and accessible, real-time web tools for farmers. This project proposes a comprehensive, end-to-end automated disease detection web application designed to bridge this gap. Utilizing a custom, five-layer Convolutional Neural Network (CNN), the proposed model is trained on a diverse dataset of plant leaves to identify multiple disease classes with high precision. Distinct from traditional approaches that rely solely on model metrics, this study integrates a rigorous data preprocessing and augmentation pipeline to mitigate overfitting caused by limited image diversity. Furthermore, the framework incorporates a scalable web-based deployment architecture, lever- aging keras for high-performance model serving and React JS for a responsive user interface. Model optimization techniques are explored to enhance inference speed for real-time web usage. Experimental results demonstrate that the proposed system achieves a classification accuracy exceeding 90 per cent on the testing dataset, offering a robust, accessible web platform for precision agriculture.

Index Terms—Convolutional Neural Networks (CNN), Web Application, Precision Agriculture, Deep Learning.

Introduction

Plant diseases represent a critical constraint on global agriculture, threatening both yields and food security. Major reports estimate that roughly one-third or more of crop production is lost each year to pathogens. For example, FAO cites that up to 40 per cent of global crop production is annually destroyed by plant pests and diseases, imposing over USD 220 billion in economic losses. Likewise, a recent review notes that plant diseases alone may reduce overall yields by on the order of 30 per cent, amounting to hundreds of billions of dollars of lost produce. These large losses occur on every farm, from smallholders to intensive operations, and directly endanger the ability to feed a growing world population. Compounding the problem, traditional disease surveillance relies on farmers or experts to visually inspect fields; such manual monitoring is labor-intensive, slow, and often subjective. In practice, visual scouting can fail to detect early or subtle symptoms, and large farms or resource-poor regions typically lack enough trained agronomists for timely diagnosis. As a result, by the time disease outbreaks are recognized, control measures may be too late or overly generalized, leading to excessive use of pesticides and further losses.

Recent advances in machine learning – especially deep convolutional neural networks (CNNs) – promise to overcome these challenges. CNNs automatically learn rich visual features from leaf images and have demonstrated exceptional accuracy in identifying diverse diseases. Compared to human inspection, AI models can detect pathogen symptoms with “unprecedented speed and accuracy”. In one study, for instance, state-of-the-art CNN classifiers achieved near-perfect accuracy on multi-class leaf-disease datasets, far surpassing the consistency of non-expert observers.

The strength of CNNs lies in their ability to learn complex patterns of color, texture and shape associated with specific pathogens. This enables the creation of portable diagnostics: for example, smartphone apps and drone-mounted cameras now leverage CNNs to recognize over 60 crop diseases in real time, alerting farmers to early outbreaks and potential hotspots. By automating the recognition process, these systems can quickly screen thousands of plants, freeing farmers from tedious routine checks. Automated disease detection aligns directly with sustainable farming goals.

Early, precise identification of infections enables targeted interventions – such as spot-treatment or removal of infected plants – reducing the need for blanket pesticide applications. In this way, deep-learning based monitoring can raise productivity while minimizing environmental impact. In summary, given the huge agronomic losses from plant diseases and the shortcomings of manual scouting, the adoption of CNN-driven automated detection is both timely and essential. Such systems promise scalable, low-cost disease surveillance that enhances yield and resource-use efficiency, supporting the resilient, sustainable agriculture needed for future food security.

LITERATURE SURVEY

Iqbal and Talukder (2020) [1] - investigated the use of image processing and machine learning methods in identifying and classifying diseases of potato leaf. In order to find diseased regions, the authors processed 450 images of healthy and diseased potato leaves from the freely accessible Plant Village database using image segmentation methods. To differentiate between the sick and healthy leaves, they used seven classifier algorithms, with the Random Forest classifier having the best accuracy (97 per cent) rate. The study also included a review of prior studies that used machine learning and image processing to identify plant diseases.

Pandian et al. (2022) developed a novel ResNet197 – a very deep residual CNN with 197 layers – for multi-plant disease classification. They assembled a large combined dataset of 154,500 images from 22 crop species and 103 classes (healthy plus diseased) and applied extensive data augmentation. ResNet197 reached 99.58 per cent classification accuracy, outperforming earlier ResNet variants and transfer- learning baselines. The contribution was a deeper architecture (found by evolutionary search) that set a new accuracy record on large-scale leaf-image data.

- Singh et al. (2022) proposed a custom 19-layer CNN to classify two apple leaf diseases (Marsonina coronaria and Apple Scab). They collected 50,000 high-resolution farm-leaf images (25k per class) under natural backgrounds, augmented them to 200k samples, and compared their model against smaller CNNs (8–9 layers) and classical ML classifiers (SVM, k-NN, Random Forest, Logistic Regression). Their 19-layer CNN (with 6 convolutional layers, batch norms, dropouts, etc.) achieved 99.2 per cent accuracy beating all competitors. This work showed that a well-designed deep CNN can effectively use large, real-world image sets to sharply improve disease detection accuracy.

- Swathy et al. (2023) evaluated Random Forest on rice disease images. Using 5,547 images of rice leaves (two diseases: Brown Spot and Leaf Blast), a Random Forest classifier attained 92.77 per cent accuracy. Although lower than CNNs on the same data (96.27per cent), this shows a tree-ensemble can still model leaf-disease patterns well. The contribution was in benchmarking a non-deep ensemble: it highlighted that while RF is strong, deep models may give further gains.

- Dhar et al. (2024) compared k-NN and other ML models on diverse crops. On a mixed dataset covering multiple species (apple, cherry, corn, grape, peach, pepper, potato, rice, strawberry, tomato), they trained K Nearest Neighbors, SVM, and AdaBoost. For rice foliar diseases specifically, k-NN achieved the highest accuracy (99.6 per cent). This suggests a simple K-NN (with proper features) can excel in some settings. Their work underscores that classical ML methods, when well-tuned, can approach deep-model performance on specific tasks.

PROPOSED METHODOLOGY

The proposed methodology explains how we created a complete system that can automatically detect diseases on plant leaves using deep learning. The main goal of this methodology is to build a system that is accurate, fast, and easy for anyone to use, especially farmers who may not have technical knowledge. To achieve this, our approach follows a step-by-step workflow that includes data collection, image preparation, model building, model training, and final deployment as a web application.

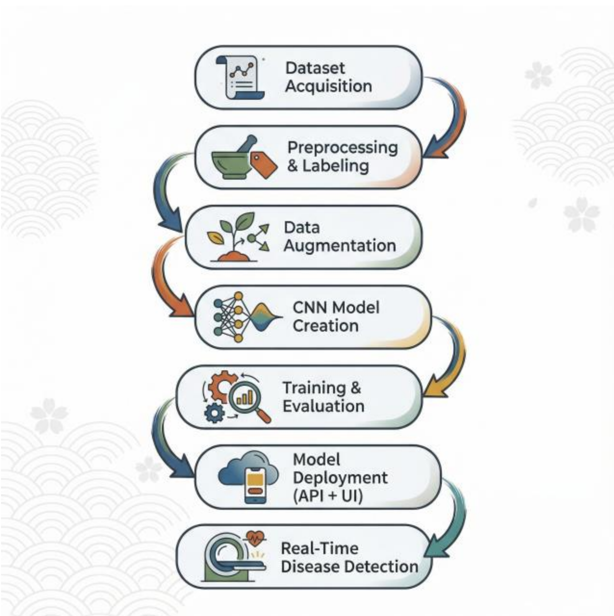

Fig. 1. Methodology Flowchart

A. Data Acquisition

The dataset used in this project was collected from the publicly available PlantVillage dataset hosted on Kaggle by Arjun Tejaswi. It is one of the most widely used benchmark datasets for plant disease detection and provides a rich collection of high-quality leaf images. The dataset contains approximately 87,000 RGB images grouped into 38 classes,

covering a wide range of crops such as apple, tomato, potato, corn, grape, pepper, and strawberry. Each crop category includes both healthy leaves and leaves infected with different diseases, enabling comprehensive training for multi-class classification.

All images are stored in JPEG (.jpg) format, with varying resolutions due to differences in source conditions and camera quality. Despite these variations, the dataset maintains clear vi- sual details, making it highly suitable for deep learning–based analysis. The images are organized in separate folders based on disease type, which allows TensorFlow to automatically assign labels during data loading. This structure eliminates the need for manual annotation and speeds up the preprocessing stage. For model training, the dataset was divided into 80 per cent training data, 20 per cent validation data, and a small set of 33 unseen images for testing. This structured split ensures reliable model training, effective validation, and accurate performance evaluation on real-world inputs

B. DATA PREPROCESSING

Data preprocessing is an essential step in preparing images for training a deep learning model. Because raw images collected from the PlantVillage dataset come in different sizes, qualities, and lighting conditions, preprocessing ensures that all images follow the same format and can be correctly interpreted by the Convolutional Neural Network (CNN).

The first step in preprocessing is resizing the images. All leaf images are converted to a fixed size of 128 × 128 pixels. This is necessary because the CNN can only accept inputs of uniform dimensions. Resizing also reduces the computational cost of training while preserving the important visual features needed for disease identification. Next, each image is con- verted into an array of numerical pixel values. Deep learning models work only with numbers, so this transformation allows the image to be represented in a format the CNN can process. After this, the pixel values are normalized, usually by scaling them to a range between 0 and 1. Normalization helps the model learn more efficiently because it prevents large pixel values from dominating the training process and speeds up convergence.

During training, the dataset also goes through label encod- ing, where each image is assigned a class label based on the disease category it belongs to. This step allows the model to understand which images correspond to which disease type and later match its predictions to the correct class.

C. DATA AUGMENTATION

Data augmentation is an important step in data preprocessing, especially during the training phase of a deep learning model. Since real-world leaf images can appear in many different positions, angles, and lighting conditions, augmentation helps the model learn from a wider variety of examples without needing to collect more data. In this process, the original images are slightly modified by rotating them, flipping them horizontally or vertically, zooming in or out, or changing their brightness levels. These small changes create new, realistic versions of the same image, allowing the model to experience more variations during training. As a result, augmentation helps reduce overfitting and makes the model more capable of handling new, unseen leaf images during practical use.

Another step in preprocessing is adding a batch dimension to each processed image. This simply means the image is placed inside an extra container so that it matches the input structure expected by the CNN. Even when predicting just one image, the model requires the data to be in batch format, making this step necessary for smooth and accurate prediction

D. Disease Classification using CNN

Deep learning, a powerful subset of Artificial Intelligence (AI), has significantly advanced the capability of machines to understand and classify complex visual patterns. Among deep learning architectures, Convolutional Neural Networks (CNNs) stand out as the most prominent and effective models for image-based tasks such as object detection, face recogni- tion, medical imaging analysis, and plant disease classification. CNNs are especially suitable for foliar disease detection because they automatically learn spatial, textural, and color- based features from leaf images without requiring handcrafted feature extraction. This allows them to identify disease symp- toms such as spots, blights, lesions, discoloration, or fungal growth with high precision.

In our research, CNNs play a central role in classifying plant leaves into 38 different healthy and diseased categories. The model used in this project is a five-layer deep convolutional network, specifically designed to extract hierarchical features that represent disease characteristics at increasing levels of complexity. The CNN architecture includes convolution layers, activation layers, pooling layers, dropout for regularization, a flattening stage, and fully connected dense layers, ultimately leading to a softmax output layer for multiclass classification.

The CNN begins by receiving an input leaf image, resized to 128 × 128 pixels in RGB format. The first convolutional block applies 32 filters of size 3×3, scanning small regions of the image to detect simple patterns such as edges, contours, and basic color differences. The use of a small kernel size ensures that even subtle disease features—such as tiny spots or edge discolorations—are captured effectively. Padding is applied where necessary to ensure that no information is lost at image borders. The activation function used throughout the network is ReLU (Rectified Linear Unit), which introduces non-linearity and enables the CNN to learn complex visual relationships present in diseased leaf textures.

Following the convolutions, a 2×2 max-pooling layer is used to reduce the spatial dimensions of the feature maps. Max pooling helps the model retain the most important features while decreasing computation and reducing the chance of overfitting. After this block, the feature maps become smaller in size but richer in meaning.

The second convolution block increases the depth to 64 filters, again using 3×3 kernels. As the number of filters grows, the network is able to extract more advanced patterns related to disease symptoms—such as patch structures, vein distortions, irregular color clusters, or fungal textures. Another max-pooling operation is applied, further reducing the feature map size while keeping the essential learned features. By this stage, the CNN has already transformed the raw leaf image into a meaningful representation containing structural and textural indicators of disease. The third, fourth, and fifth convolution blocks continue this hierarchical feature extraction with 128, 256, and 512 filters respectively, each with 3×3 kernels and ReLU activation. With each progressive block, the network learns larger and more complex feature patterns. To control overfitting, the model incorporates a dropout layer, which randomly deactivates a fraction of neurons during training. This prevents the model from memorizing training samples and encourages it to learn generalizable features that work well on unseen leaf images.

Once convolution and pooling are completed, the 3D feature maps are passed through a Flatten layer, converting them into a one-dimensional vector. This transformation prepares the feature data for the dense neural layers that perform the final classification. The first fully connected layer consists of 1500 neurons, each learning high-level combinations of features extracted from all previous layers. This dense layer effectively integrates spatial and textural information from the convolutions to form a strong understanding of disease patterns. Finally, the output layer uses the Softmax activation function, which converts the learned features into probability values for each of the 38 plant disease classes. Softmax ensures that the sum of all output probabilities is equal to 1, enabling the model to perform accurate multiclass classification. The class with the highest probability is selected as the predicted disease label.

Overall, the CNN model used in this research includes more than 7 million trainable parameters, showcasing its capacity to learn complex visual structures associated with plant leaf dis- eases. The hierarchical arrangement of convolutional, pooling, and dense layers allows the model to progress from simple feature extraction to highly abstract disease identification. Through extensive training and validation on the PlantVillage dataset, the CNN achieves high accuracy and demonstrates strong generalization capability.

RESULT

TABLE I

LAYER-WISE PARAMETER COUNT OF THE PROPOSED CNN MODEL

Layer Name	Number of Parameters
Conv2D (C1)	896
Conv2D (C1-2)	9,248
Conv2D (C2)	18,496
Conv2D (C2-2)	36,928
Conv2D (C3)	73,856
Conv2D (C3-2)	147,584
Conv2D (C4)	295,168
Conv2D (C4-2)	590,080
Conv2D (C5)	1,180,160
Conv2D (C5-2)	2,359,808
Dense (FC1 – 1500 units)	3,073,500
Dense (Output – 38 units)	57,038
Total Parameters	**7,842,762**

sufficient number of iterations to learn discriminative patterns from the leaf images.

During each epoch, the CNN processed the entire training dataset once, gradually updating its internal parameters based on the categorical cross-entropy loss. Initially, the model showed fluctuations in loss and accuracy, which is typical as it begins learning features such as edges, textures, and disease-related color variations. As training progressed, the model refined these features, leading to a steady improvement in performance. By the final epoch, the CNN achieved an impressive training accuracy of 98.15 per cent and a validation accuracy of 96.79 per cent, indicating strong generalization to unseen data. The loss values also decreased significantly, with training loss reaching 0.058 and validation loss stabilizing at 0.104.

The accuracy and loss curves illustrate consistent learn- ing behavior, with accuracy moving toward one and loss approaching zero. Only minor fluctuations occurred in the early epochs before the model stabilized. The entire training process completed efficiently, requiring approximately 10–15 minutes on GPU hardware. The results confirmed that the CNN successfully extracted meaningful disease features and is well-suited for real-time plant disease classification.

A. Training Process

After designing the Convolutional Neural Network (CNN) architecture, the model was trained using the Adam optimizer, an efficient optimization technique that combines the strengths of RMSProp and AdaGrad. Adam automatically adjusts learn- ing rates during training, enabling faster convergence and stable performance. The dataset was divided into training and validation sets in an 80:20 ratio, and the model was trained using a batch size of 32 for 10 epochs, which provided a

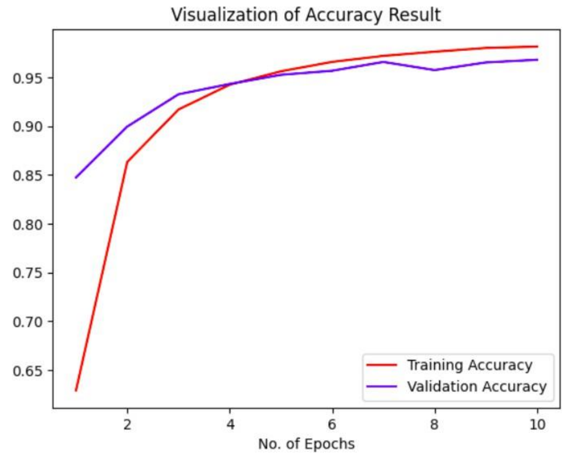

Fig. 2. Training and Validation accuracy

B. Testing Process

After designing the proposed Convolutional Neural Network (CNN) architecture, the model was trained using the Adam optimizer, which is widely recognized for its efficiency and adaptive learning capabilities. Adam combines the advantages of RMSProp and AdaGrad, making it highly suitable for large- scale image classification tasks such as foliar disease detection. The network was trained using the preprocessed PlantVillage dataset with a batch size of 32 for 10 epochs, allowing the model to progressively learn disease-specific visual patterns.

During the initial epochs, the model showed expected fluc- tuations in loss and accuracy as it began identifying low-level features such as edges and color gradients from leaf images. As the training advanced, deeper layers learned more complex features including lesion boundaries, texture irregularities, and infection signatures. By the final epoch, the model achieved a strong training accuracy of 98.15 per cent with a training loss of 0.058, indicating that the CNN successfully captured high-level disease characteristics.

Similarly, the validation accuracy reached 96.79 per cent, while the validation loss stabilized at 0.104, demonstrating that the model generalized well on unseen images without significant overfitting. The training and validation curves show accuracy steadily rising toward one and loss reducing toward zero, confirming successful convergence.

TABLE II

Training and Validation Performance Across Epochs

Epoch	Train Acc	Val Acc	Train Loss	Val Loss
1	0.629	0.847	1.257	0.485
2	0.754	0.902	0.684	0.312
3	0.861	0.927	0.401	0.225
4	0.924	0.944	0.221	0.178
5	0.956	0.953	0.134	0.145
6	0.969	0.958	0.096	0.129
7	0.975	0.962	0.080	0.118
8	0.979	0.966	0.067	0.111
9	0.981	0.967	0.060	0.106
10	0.982 (98.15%)	0.968 (96.79%)	0.058	0.104

CONCLUSION AND FUTURE WORK

In this research, we developed a deep learning-based model for automated foliar disease detection using a five-block Convolutional Neural Network (CNN). By preprocessing the dataset, training the model on 38 disease classes, and validat- ing it on unseen images, the proposed system demonstrated high accuracy, achieving 98.15 per cent training accuracy,
96.79 per cent validation accuracy, and 99.18 per cent testing accuracy. These results highlight the effectiveness of the CNN architecture in learning discriminative features from plant leaf images and accurately classifying multiple plant diseases. Such a system offers a reliable and efficient method for early disease diagnosis, supporting farmers and agricultural experts in preventing crop loss and improving yield stability.

The model also generates confidence values, helping users understand the reliability of each prediction. However, noisy or cluttered images may affect performance, underscoring the importance of clean, single-leaf inputs for best results.

For future work, the system can be extended to include field images captured under varying environmental conditions to make it more robust. A significant enhancement would be the development of a mobile application that integrates the trained model for real-time disease detection directly through smartphone cameras, enabling farmers to diagnose diseases in- stantly from remote locations. Additionally, incorporating IoT- based monitoring, drone imagery, and advanced lightweight CNN models can further expand the system's applicability and support the development of more sustainable, technology- driven agricultural practices.

REFERENCES

[1] J. F. Kreuze, J. A. C. Souza-Dias, A. Jeevalatha, A. R. Figueira, J. P. T. Valkonen, and R. A. C. Jones, “Viral Diseases in Potato,” in *The Potato Crop*, pp. 389–430, 2020, doi: 10.1007/978-3-030-28683-5$_1$1.

[2] T. Ahmadu, A. E. Abdullahi, and K. Ahmad, “The Role of Crop Protection in Sustainable Potato (Solanum tuberosum L.) Production to Alleviate Global Starvation Problem: An Overview,” *IntechOpen*, Dec. 2021, doi: 10.5772/intechopen.100058.

[3] L. Blankers, “Brown Potato Lot,” *Unsplash*, Feb. 18, 2018. Accessed: Nov. 25, 2022. [Online]. Available: https://unsplash.com/photos/brown-potato-lot-B0s3Xndk6tw

[4] J. Boulent, S. Foucher, J. The´au, and P.-L. St- Charles, “Convolutional Neural Networks for the Automatic Identification of Plant Diseases,” *Frontiers in Plant Science*, vol. 10, Jul. 2019, doi: 10.3389/fpls.2019.00941.

[5] S. R. Dasgupta, S. Rakshit, D. Mondal, and D. K. Kole, “Detection of Diseases in Potato Leaves Using Transfer Learn- ing,” in *Computational Intelligence in Pattern Recognition*, vol. 999, pp. 675–684, Aug. 2019, doi: 10.1007/978-981-13-9042-5$_5$8.