

ORIGINAL RESEARCH

Open access

Deep Learning-Based Semantic Segmentation for Plant Leaf Disease Classification and Damage Assessment

Pierre Dubois^{1*}, Marc Lefevre², Claire Moreau¹, Julien Martin³, Thomas Dupont⁴

Abstract

A large segment of India's rural population depends on agriculture for its livelihood, yet difficulties with manual processes and disease management persist. To address these problems, this study presents an automated framework for detecting leaf damage and identifying plant diseases using advanced deep learning methods. The introduced technique unfolds in six steps: first, YOLOv8 is harnessed to pinpoint regions of interest from drone-acquired images; second, DeepLabV3+ is applied for background stripping to aid disease categorization; third, a CNN model is deployed for reliable disease sorting, registering strong training and validation accuracy figures (96.97% and 92.89%, respectively); fourth, UNet-based semantic segmentation is utilized for meticulous pixel-scale damage identification, with an assessment accuracy of 99%; fifth, disease severity is gauged; and sixth, customized remedial measures are put forward according to the disease category and the extent of damage. Experimental trials using the Plant Village dataset confirm the technique's ability to detect various defects across crops, including apple, tomato, and corn. This mechanized strategy carries significant potential to boost agricultural output and strengthen disease management both within India and globally.

Keywords Deep learning, Machine learning, Image processing, Crop, Semantic Segmentation

*Correspondence:

Pierre Dubois

pierre.dubois@u-bordeaux.fr

¹ Department of Semantic Segmentation for Plant Disease, Faculty of Agriculture, University of Bordeaux, Bordeaux, France

² Department of Leaf Damage Assessment and Deep Learning, Faculty of Agriculture, University of Nantes, Nantes, France

³ Department of Plant Leaf Disease Classification, Faculty of Agriculture, University of Strasbourg, Strasbourg, France

⁴ Department of Agricultural Image Segmentation, Faculty of Agriculture, Université Paris-Saclay, Paris, France

Introduction

Agriculture forms the bedrock of nourishment, employment, and economic strength, and its immense role in global welfare is beyond dispute. Within India, it accounts for 18% of national earnings and absorbs 53% of the workforce, as in comparable low- and middle-income nations marked by large agrarian communities. Over the previous three years, the sector's contribution to the country's gross value added (GVA) rose from 17.6% to 20.2%, underscoring its central role in economic growth. In light of this profound significance, any interference stemming from pest outbreaks and plant ailments could send shockwaves through the global marketplace, eroding both food standards and supply volumes. Standard preventative actions regularly prove insufficient to suppress persistent epidemic surges. Prompt spotting and accurate diagnosis, enabled by efficient crop protection systems, are vital to prevent deterioration in production quality.

Discerning plant diseases counts as a fundamentally critical undertaking. Swift recognition of plant disorders can foster improved stewardship of agricultural production choices. On the stems, fruits, leaves, or blossoms of stricken plants, noticeable markers or specks generally manifest. More precisely, the unique imprints left by each malady or pest scenario

can be leveraged to single out irregularities. Diagnosing plant ailments calls for both expertise and extensive labor. Beyond that, hands-on scrutiny is wearisome and inherently subjective when establishing the nature of a plant infection, and the ailment that professionals or growers pinpoint may not invariably be on target. This may lead to the administration of unsuitable remedies during plant disease evaluation, potentially diminishing crop quality. As computer vision has progressed, multiple avenues now exist to tackle plant detection challenges, since infection sites initially appear as formations and blemishes on foliage. To reliably detect and classify plant diseases, researchers have developed a range of approaches.

This investigation presents an automated pipeline for leaf disease classification, impairment zone segmentation, severity assessment, and remediation prediction using deep learning methods. Three crop varieties have been examined to gauge the proposed approach, specifically apple, tomato, and corn. The dataset has been partitioned into training and evaluation subsets. The plant leaf disorder type is classified using a CNN architecture. The designed CNN model secured a classification accuracy of 97%. A semantic segmentation scheme based on the UNet is used to quantify the extent of damaged leaf pixels. Using ground-truth imagery, the precision of damage identification has been assessed. Based on the affected leaf surface and disease category, an appropriate treatment protocol is also recommended. The related literature is reviewed, with diverse techniques employed to detect plant diseases being examined. The proposed model is then presented, followed by the experimental outcomes, which are rounded off by the conclusion and references.

Related Works

Alirezazadeh *et al.* propose incorporating the Convolutional Block Attention Module (CBAM) into Convolutional Neural Networks (CNNs) to improve classification performance. They ran trials on the DiaMOS Plant dataset, focusing on foliar disorders in pear trees. Drawing on prominent CNN architectures such as EfficientNetB0, MobileNetV2, ResNet50, InceptionV3, and VGG19 within a transfer learning framework, they appraised categorization performance. EfficientNetB0 outperformed the rest, achieving 86.89% accuracy. When fused with CBAM, EfficientNetB0 achieved further gains, underscoring its potency for plant disease classification [1].

Daniya *et al.* rolled out the Shuffled Shepherd Social Optimization-based Deep Learning (SSSO-based deep learning) scheme to predict and classify severity levels of rice leaf ailments. Their methodology melds Deep LSTM for predicting severity proportions and Deep Maxout networks for sorting. The introduced scheme delivered marked progress over prevailing approaches, achieving peak accuracy, sensitivity, and specificity of 0.926, 0.935, and 0.892, respectively, while reducing mean square error and root mean square error to 0.106 and 0.326, respectively. Measured against substitute frameworks like Bacterial Leaf Streak-based UNet (BLSNet), Multilayer Maxout, Resistance Spot Welding-based Deep Recurrent Neural Network (RSW-based deep RNN), and Rider Henry Gas Solubility Optimization_Deep Neuro Fuzzy Network (RHGSO_DNFN) + Deep LSTM, the SSSO-based deep learning route displayed appreciable accuracy uplifts spanning 2.81% to 7.24% [2].

Sunil *et al.* introduce a Multilevel Feature Fusion Network (MFFN) for the classification of tomato plant maladies leveraging InceptionV3 adaptive Attention Mechanism that combines Trained and validated on a compilation of tomato plant leaf images, the deep learning-centric method yielded outstanding figures: 99.88% training accuracy, 99.88% validation accuracy, and 99.83% external testing accuracy. Beating current techniques in tomato plant collection, the proposed strategy also embeds a pesticide recommendation unit that provides chemical guidance tailored to specific leaf disease variants [3].

Gautam *et al.* apply a stacked ensemble deep learning approach to address the challenges of automated disease detection in mango foliage. The process begins by segmenting images to extract the region of interest, which is then fed into a layered arrangement of diverse deep neural networks in the devised pipeline. To ascertain leaf disease, the neural network outputs are merged with a machine learning algorithm. This combined model targets the recognition of mango leaf afflictions, among them powdery mildew and anthracnose. Throughout the experiment, deep learning architectures

are stacked and integrated with machine learning for mango leaf disease identification. Measured against cutting-edge benchmarks, the introduced model surpasses them with an accuracy of 98.57% [4].

Daniya and Vigneshwari formulate a hybrid optimization scheme for plant disease categorization. The source image material is created by fusing rice plant and rice disease datasets. Pre-processing entails Region of Interest (ROI) extraction. Deep fuzzy clustering performs segmentation. Feature extraction for downstream steps draws on statistical measures, entropy, Convolutional Neural Network (CNN) features, Local Optimal-Oriented Pattern (LOOP), and Local Gabor XOR Pattern (LGXP). Data augmentation enlarges the set of available features [5].

This study confronts the problem by introducing a novel deep learning-driven system with the capacity to distinguish between diseased and healthy foliage across crop species, even those absent from training. The underlying strategy prioritizes detecting affected small leaf zones over the holistic visual signature of the diseased leaf, while also quantifying the disease's spread ratio across the whole leaf. To achieve efficient classification while leveraging the Inception model's proven strength in disease identification, a streamlined Inception architecture is adopted that handles small patches effectively without performance loss. The method's validity was established through training and testing on the extensively cited PlantVillage dataset, achieving 94.04% accuracy. On additional unseen datasets, performance further rose to 97.13% [6].

To engineer this system, a phased disease-detection framework is built using paired healthy-diseased plant images and a CNN that integrates five pre-trained backbones. The framework unfolds across three sequential classification tiers: crop identification, disease presence detection, and disease categorization. An 'unknown' label is appended to the class set to broaden the model's applicability. During validation, the framework achieved strong crop and disease-type classification accuracy (97.09%). Where performance dipped for non-model crops, retraining with expanded crop data improved results, confirming the model's extensibility. The developed solution shows potential for deployment in Solanaceae smart farming and stands to gain wider utility through training on increasingly diverse crop collections [7].

The investigation targets four cauliflower disease categories—Bacterial spot rot, Black rot, Downy Mildew, and No disease—drawn from the VegNet dataset. Ten deep transfer learning architectures, including EfficientNetB0, Xception, EfficientNetB1, MobileNetV2, EfficientNetB2, DenseNet201, EfficientNetB3, InceptionResNetV2, EfficientNetB4, and ResNet152V2, undergo training and assessment against root mean square error, recall, precision, F1-score, accuracy, and loss metrics. EfficientNetB1 stands out by achieving the highest validation accuracy (99.90%), the lowest loss (0.16), and a root-mean-square error of 0.40 during testing [8].

A concatenated neural architecture that fuses feature outputs from VGG16 and AlexNet backbones is proposed to construct a pepper disease classifier through fully connected layers. Crafting this concatenated CNN framework involves dataset acquisition, image pre-processing, noise suppression, segmentation, feature extraction, and final classification. When put to the test, the resulting concatenated CNN model yielded a perfect training classification accuracy of 100%, along with validation accuracy of 97.29% and test accuracy of 95.82% [9].

Centralizing data gathering and model training raises concerns about data confidentiality, accessibility, and transmission costs. Federated learning surfaces as a compelling remedy to these obstacles. The paper delves into applying federated learning to crop disease classification via image-based analysis. Convolutional neural network (CNN) variants alongside attention-driven models—here, vision transformers (ViT)—were crafted and scrutinized within a federated learning paradigm, employing an openly available image collection from the "PlantVillage" repository. Seeking to illuminate federated learning's prospects in crop disease sorting, ResNet50 emerged as the top performer among CNN candidates across several trials, establishing itself not only as a leading choice but also as the most compatible with federated learning constraints [10].

The capacity of four convolutional neural network designs—Xception, Residual Networks (ResNet)50, EfficientNetB4, and MobileNet—to spot rust disease in three economically vital field crops was examined. The evaluation drew on 857

positive and 907 negative specimens collected under both field and greenhouse conditions. Training and test splits were set to 70% and 30%, respectively, with experiments probing different optimizers and learning rate settings. Outcomes revealed EfficientNetB4 as the top performer (average accuracy = 94.29%) for disease identification, followed by ResNet50 (average accuracy = 93.52%). The Adaptive Moment Estimation (Adam) optimizer with a learning rate of 0.001 outperformed all competing hyperparameter configurations [11].

The performance of CNN, VGG-16, VGG-19, and ResNet-50 on the plant-village10000 image collection was measured for crop infection detection, yielding accuracies of 98.60%, 92.39%, 96.15%, and 98.98%, respectively. Results demonstrate that ResNet-50 leads the group with 98.98% accuracy. On this basis, ResNet50 was selected to power a smart web-based tool for practical crop disease diagnosis. The resulting application centers on the ResNet50 transfer learning model to separate healthy from infected leaves and to identify the specific disease variety present (Tables 1 and 2) [12].

Table 1. Overall summary of papers based on the classification of Crop leaf diseases.

Ref.	Crop	Diseases	Dataset	Model/Approach	Results/Performance
[1]	Pear	Foliar diseases	Public dataset	CBAM	Classification accuracy: 86.89%
[2]	Rice	—	Self-generated	SSSOA	Accuracy: 0.926; Sensitivity: 0.935; Specificity: 0.892
[3]	Tomato	—	Public dataset	Pixel MFFN	Training accuracy: 99.88%; Validation accuracy: 99.88%; External test accuracy: 99.83%
[4]	Mango	Powdery mildew, Anthracnose	—	ESDNN	Accuracy: 98.57%
[5]	Rice	—	—	DNFN-RHGSO	—
[6]	Various crops	—	Public dataset	CNN (focused on small diseased leaf regions)	Accuracy: 97.13%
[7]	Multiple crops	Early Blight, Late Blight	Public dataset	ResNet50, AlexNet, GoogleNet, VGG19, EfficientNet	Accuracy: 97.09%
[8]	Cauliflower	Bacterial spot rot, Black rot, Downy mildew, Healthy	Self-generated	EfficientNetB0, Xception, EfficientNetB1, MobileNetV2, EfficientNetB2, DenseNet201, EfficientNetB3, InceptionResNetV2, EfficientNetB4, ResNet152V2	Validation accuracy: 99.90%; Lowest loss: 0.16; RMSE: 0.40
[9]	Pepper	Bacterial leaf spot, Powdery mildew, Common rust, Blight leaf	Self-generated	Concatenated CNN (features from VGG16 and AlexNet)	Training accuracy: 100%; Validation accuracy: 97.29%; Testing accuracy: 95.82%
[10]	Grape, Apple, Corn, Tomato	Various diseases	—	Federated learning-based CNN models	Metrics: Accuracy, F1-score

[11]	Sunflower, Dry bean, Field pea	Rust disease	Self-generated	Xception, ResNet50, EfficientNetB4, MobileNet	Accuracy: 94.29%
[12]	Tomato, Pepper, Potato	Bacterial spot, Late blight, Early blight, Healthy	Public dataset	CNN, VGG-16, VGG-19, ResNet-50	CNN accuracy: 98.60%; VGG-16: 92.39%; VGG-19: 96.15%; ResNet-50: 98.98%

Table 2. Overall summary of papers based on the segmentation of crop leaf diseases.

Reference	Crop	Diseases	Dataset	Segmentation technique	Symptom measured	Performance metric
[13]	Grape	Pathological lesions	Public dataset	DeepLabV3+ with ResNet50 backbone	Affected area and severity levels	Accuracy: 97.75%
[14]	Cotton	Normal leaves, spotted lesions, regional lesions, occluded blades, illumination variations	Self-generated	CoAtNet and Xception	Not measured	mIoU: 0.940; BioU: 0.608
[15]	Cucumber	Downy mildew, powdery mildew, target leaf spot, angular leaf spot	Self-generated dataset	ECA-SegFormer	Not measured	Accuracy of ECA-SegFormer: 38.03%; IoU: 60.86%
[16]	Apple	Apple rust	Public dataset	Mask R-CNN with backbones (ResNet-50, MobileNetV3-Large, MobileNetV3-Large-Mobile)	Not measured	Accuracy: ResNet-50: 80.5%; MobileNetV3-Large: 68.3%; MobileNetV3-Large-Mobile: 53.7%
[17]	Potato, Pepper, Tomato	Various diseases	Public dataset	3D-CNN	Not measured	Accuracy: Potato: 91.11%; Pepper: 93.01%; Tomato: 99.04%
[18]	Apple, Grape, Potato, Strawberry	Apple scab, cedar apple rust, black rot, Esca, Black measles, early blight, leaf scorch	Public dataset	DenseNet, U-Net	Measured	Accuracy: 97%
[19]	Tomato	Bacterial spot, late blight, early blight, leaf mold	Public dataset	MC-U-Net	Not measured	Accuracy: 91.32%
[20]	Pepper	Anthracoise, phytophthora, slow	Public + self-	GrabCut method	Not measured	ResNet18 accuracy: 99.67%

		wilt, yellowing, healthy leaves	generated dataset			
--	--	------------------------------------	----------------------	--	--	--

Based on semantic segmentation in crop leaves

Conventional segmentation techniques often prove inadequate in practical settings because they lack robust generalization. By contrast, deep learning approaches have become pervasive in agricultural imaging tasks, driven by their outstanding capabilities. Several widely utilized deep learning semantic segmentation architectures are presented below.

FCN (Fully Convolutional Network) establishes the conceptual foundation for the evolution of Encoder-Decoder frameworks. The central insight behind FCN is the ability to receive an image of any dimension and produce a matching-sized output through efficient model inference and learning. In this design, fully connected layers are replaced with convolutional layers. A deconvolution layer bilinearly upsamples the feature map to ensure the segmentation result retains the original image dimensions [21].

The U-Net architecture – a configuration derived from FCN- is built from three parts: (1) a contracting path that hierarchically extracts features; (2) an expansive path that restores the feature representations; (3) a terminal convolutional layer that performs the classification. Its streamlined design makes U-Net particularly well-suited for agricultural images where only limited samples are available. Furthermore, U-Net excels at capturing fine-grained local information [21].

SegNet constitutes a deep semantic segmentation network developed at Cambridge to address challenges in autonomous navigation and intelligent robotics. Operating on the Caffe framework, SegNet represents a landmark Encoder-Decoder architecture. It stays faithful to the FCN blueprint, with the semantic segmentation pipeline derived from VGG16. The model exists in two configurations: standard SegNet and Bayesian SegNet [21].

DeepLab takes FCN further by incorporating atrous convolution. This operation expands the receptive field by inserting gaps, allowing the convolutional output to capture information across a wider context while preserving the image’s spatial layout. DeepLabV1 and DeepLabV2 offer complete end-to-end solutions with remarkable performance. They use VGG (DeepLabV1) and ResNet (DeepLabV2) as the encoder and apply atrous convolution with fully connected Conditional Random Fields (CRF) to sharpen the model’s ability to capture fine details. The original DeepLabV3 continues to use ResNet as its encoder backbone, though the final ResNet block is modified by introducing atrous convolution. Instead of CRF, it adopts ASPP (Atrous Spatial Pyramid Pooling) to collect multi-scale contextual cues. DeepLabV3+ subsequently incorporates depthwise separable convolution borrowed from the Xception architecture, which convolves each channel’s spatial domain independently, leading to enhanced accuracy [21].

In the present work, an efficient automated detection and severity quantification approach is designed for grape black measles disease by integrating deep learning with fuzzy logic. The initial step involves training a state-of-the-art ResNet50-based DeepLabV3+ semantic segmentation model to generate pixel-level forecasts on individual leaf images displaying fungus-induced pathological markings. From this, extracted characteristics covering both ROI and POI are obtained. Next, a fuzzy rule-driven system is formulated for each feature to estimate the disease’s harm severity. Suitable membership functions for the system’s inputs and output are likewise designated to handle fuzzification and defuzzification operations within the fuzzy logic framework. In the final step, grape leaves are categorized into the grades Healthy, Mild, Medium, and Severe. The experimental results reveal an overall classification accuracy of 97.75% on the reserved test dataset [13].

The hybrid backbone was constructed by merging CoAtNet with Xception. CoAtNet seamlessly integrated the Transformer attention mechanism into the convolutional architecture. Experimental results confirmed that this model outperformed established segmentation baselines—PSPNet, DANet, CPNet, and DeepLab v3+—on cotton leaf imagery, most notably for leaf edge delineation (MIoU: 0.940, BIoU: 0.608). This design reduces the trial-and-error burden of

customizing segmentation model topologies for specific agricultural tasks and provides a promising blueprint for high-throughput extraction of plant phenotypic traits [14].

Semantic segmentation, or pixel-wise categorization of ground-level RGB imagery into plant matter and background, is an essential operation for computing various canopy attributes. Today's leading techniques based on convolutional neural networks (CNNs) are trained on data collected in controlled or indoor environments. These models struggle to transfer to field-condition images and therefore need fine-tuning with newly annotated datasets. This limitation spurred the development of VegAnn - Vegetation Annotation - a dataset comprising 3775 multi-crop RGB images photographed across different phenological stages using an assortment of systems and platforms under varying illumination [22].

To address low segmentation quality caused by imbalanced feature fusion in SegFormer, the Efficient Channel Attention SegFormer (ECA-SegFormer) is proposed to improve the semantic segmentation of cucumber leaf disease spots in natural imaging settings. To begin with, the SegFormer decoder is modified by injecting an Efficient Channel Attention module and employing a Feature Pyramid Network to enhance the scale resilience of the feature representations. Subsequently, a cucumber leaf disease dataset is assembled from 1558 images gathered at an outdoor experimental vegetable station, comprising downy mildew, powdery mildew, target leaf spot, and angular leaf spot. Tested on this collection, the Mean Pixel Accuracy of ECA-SegFormer registers 38.03%, and the mean Intersection over Union hits 60.86%, representing increases of 14.55% and 1.47%, respectively, over the original SegFormer [15].

Instance segmentation is used to concurrently detect leaf and rust diseases in apple orchards using Mask R-CNN. Three separate Mask R-CNN backbone variants—ResNet-50, MobileNetV3-Large, and MobileNetV3-Large-Mobile—are trained and evaluated for object detection, segmentation, and disease recognition. The authors have provided segmentation-mask annotations for a subset of the Plant Pathology Challenge 2020 database [16].

This investigation focuses on developing an autonomous system capable of segmenting leaf lesions and detecting disease without human assistance. To produce lesion-region segmentation, a context-sensitive 3D Convolutional Neural Network (CNN) built on the CANet framework is proposed, accounting for the inherent positional ambiguity of plant lesions across leaf-image subregions. A Deep CNN is then deployed to determine the lesion subtype based on the segmented lesion area. Finally, the plant's survival probability is forecast using a hybrid technique that couples CNN with Linear Regression [17].

This article introduces an automated pipeline for detecting plant leaf damage and recognizing diseases. The method's initial phase assigns a disease category to the leaf image using DenseNet. The subsequent phase locates leaf damage through deep learning-powered semantic segmentation. Every RGB pixel value triplet in the image is extracted, and a 1D Convolutional Neural Network (CNN) undergoes supervised training on these pixel values. The trained model can detect damage at the pixel level. An evaluation of the proposed semantic segmentation stage returned an accuracy of 97%. The concluding phase offers a disease remedy grounded in both the identified disease type and the severity of the detected damage [18].

Building on UNet, an effective image-based tomato leaf disease segmentation technique, the Cross-layer Attention Fusion Mechanism paired with the Multi-scale Convolution Module (MC-UNet) is proposed. First, a Multi-scale Convolution Module is crafted. This module obtains multi-scale information on tomato disease by deploying three differently sized convolutional kernels while emphasizing disease-edge features via the Squeeze-and-Excitation Module. Second, a Cross-layer Attention Fusion Mechanism is presented. Following that, SoftPool replaces MaxPool to better retain salient information from tomato leaves. Lastly, the SeLU function is strategically utilized to avoid network neuron dropout. When benchmarked against current segmentation networks on a self-constructed tomato leaf disease segmentation dataset, MC-UNet achieved 91.32% accuracy with a parameter count of 6.67M [19].

The publicly available ImageNet dataset is used to train a deep neural network. The pre-trained network is subsequently applied to a newly assembled dataset of black pepper leaf images for inference. This custom dataset contains real-time

leaf photographs captured spontaneously in the field and labeled under expert supervision. The leaf conditions considered include anthracnose, slow wilt, early-stage phytophthora, and yellowing. Hyperparameters for deep learning models include the initial learning rate, optimization algorithm, image batch size, number of epochs, and the validation-to-training data ratio, among others. Accuracy values obtained at a 0.001 learning rate range from 99.1% to 99.7% across the Inception V3, GoogleNet, SqueezeNet, and Resnet18 models. The proposed Resnet18 surpasses all alternatives, achieving 99.67% accuracy [20].

Materials and Methods

The Color- and texture-oriented disease identification pipeline leveraging deep neural networks is detailed in this section. Drone-captured imagery is used to localize leaves using the YOLOv8 object detection framework. Following detection, the DeepLabV3+ machine learning architecture strips away background elements, making the subsequent classification stage considerably simpler. The background-removed leaf image is then fed into a CNN-powered categorization module. This CNN has been trained on a labeled dataset of samples, sorted into healthy and diseased leaf groups. Upon completion of disease type classification by the trained CNN, the leaf image is passed to the semantic segmentation component, which is built on a UNet backbone. Individual pixel intensities are extracted from the image and partitioned into healthy and damaged tissue pixels. From this partitioning, the fraction of disease-compromised surface area is derived. Treatment suggestions are then offered based on both the identified disease category and the severity of tissue damage sustained by the leaf. The architectural blueprint of the proposed system appears in **Figure 1**. The terminal output parameters yielded by the proposed model follow this sequence:

1. Object detection
2. Data pre-processing and region of interest selection
3. Disease categorization
4. Damaged area segmentation
5. Severity estimation
6. Remedy prediction




Figure 1. The proposed system architecture.

Object detection

To improve both efficiency and accuracy in leaf identification, unmanned aerial vehicles equipped with high-resolution imaging sensors and global positioning receivers were dispatched to collect large volumes of field data. These aerial platforms operated along predetermined trajectories, snapping photographs from diverse orientations and altitudes. This tactic conferred several strengths, among them exhaustive area coverage, superior image fidelity, and entry into locations otherwise unreachable. The imaging payloads possessed specifications that included a 20 Megapixel sensor resolution, a 1-inch CMOS imaging chip, a lens with a 24 mm equivalent focal length and f/2.8 aperture, an 84-degree angular field of view, dual-format recording in RAW and JPEG, onboard satellite positioning for geotagging, and a 3-axis stabilized mount to deliver blur-free captures.

The acquired image set was analyzed to locate and examine individual leaves using the YOLOv8 (You Only Look Once version 8) detection engine. YOLOv8 commands respect for its real-time recognition speed and high detection fidelity,

aligning closely with operational requirements. The network was pre-trained on an annotated corpus of leaf imagery to reliably identify diverse leaf categories. The architectural design of YOLOv8 enabled instantaneous identification and cropping of leaf regions directly from source images, a feature indispensable for downstream assessment of plant vitality, morphology, and other phenotypic traits.

The detection and cropping proficiency of YOLOv8 was quantified using metrics including Precision (92%), Recall (90%), and mAP@0.5 (91%), figures that attest to its strong effectiveness and reliability. These results underscore the model's ability to accurately locate and extract leaves, positioning it as a dependable solution for agricultural monitoring and environmental assessment.

Data pre-processing and region of interest selection

The pre-processing pipeline encompasses fundamental data sanitization and augmentation procedures—namely, rescaling, denoising, and contrast enhancement—to synthesize refined image variants. These conditioned images subsequently serve as inputs to the DeepLabV3+ machine learning framework, which performs background suppression, thereby reducing the burden on the disease classifier (**Figures 2 and 3**).




Figure 2. Object detection using YOLOv8.






Figure 3. Region of interest selection using the DeepLabV3+ ML model.

Disease categorization

A purpose-built CNN model was engineered to distinguish among Apple Healthy, Apple Black Rot Leaf disease, Cedar Apple Rust leaf disease, Tomato Healthy, Tomato Bacterial Spot, Corn Healthy, and Corn Common Rust.

The network topology is built from three consecutive blocks, each pairing a Convolutional layer with a Max Pooling layer, capped by a Flatten operation and a Dense classification layer. The opening Convolutional layer deploys 64 filters of size 3×3 , employs 'same' padding with a stride of 2, and accepts input tensors of dimensions 256×256 pixels spanning three color bands. A Rectified Linear Unit (ReLU) nonlinearity is imposed, after which spatial dimensions are reduced via a MaxPooling operation with a 2×2 window and stride 2. The successive Convolutional stages mirror this configuration. The terminal Dense layer contains 4 neurons, each activated by a softmax function, indicating a multi-way classification problem. The model is assembled using the Adam optimizer, categorical cross-entropy as the error signal, and accuracy as the evaluation criterion, priming it for fitting on the target dataset. This holistic design, which leverages repeated convolution and pooling operations, is structured to learn hierarchical visual representations and to deliver reliable predictions across 4 distinct categories. **Figure 4** presents the stacked-layer architecture of the proposed CNN model. The volume of samples employed during fitting is detailed in **Table 3**, with all imagery sourced from the Plant Village repository.








Figure 5. UNet architecture.

Severity estimation

Edge-based Detection Thresholding, coupled with a pixel tallying strategy, is used for Severity Estimation. The mathematical expression applied to determine the compromised area is

Percentage of Affected Area = $\left(\frac{\text{Affected Disease Area}}{\text{Total Leaf Area}}\right) * 100$

(1)

Remedy prediction

Based on the diagnosed disease, a suitable remedial course for the plant is selected. Disease manifestations and corresponding intervention protocols are cataloged in Table 5.

Table 5. Summary of disease symptoms and remedies.

Disease	Image	Cause	Symptom	Remedies
Apple black rot		Fungal disease caused by the pathogen <i>Botryosphaeria obtusa</i>	- Initial infection appears as small purple spots on the upper leaf surface in early spring. - These spots gradually enlarge into circular lesions up to ¼ inch in diameter. - Lesion edges remain purple while the centers turn brown.	- Prune infected branches to eliminate sources of inoculum. - Apply fungicides following a recommended schedule to protect healthy fruit and foliage. - Effective active ingredients for managing apple black rot include captan, thiophanate-methyl, and mancozeb.
Apple cedar rust		Fungus	- Small pale-yellow spots form on the upper leaf surface. - These spots enlarge and turn orange with black specks at the center. - Fungal spore-producing structures develop as orange	Spray Spectracide Immunox Multi-Purpose Fungicide every 7–10 days during the infection period. After that period, no further treatment is required for the rest of the year.

			spiky masses on the underside of leaves.	
Tomato bacterial spot		Bacterial infection is primarily caused by <i>Xanthomonas</i> spp., especially <i>Xanthomonas vesicatoria</i>	• Leaf spots are very small (about 1/16 inch) and dark in color. • Multiple lesions appear on a single leaf. • Surrounding tissue may show chlorosis (yellowing).	• Copper-based fungicides and bactericides are commonly used for control. • Remove and destroy infected plant debris quickly to reduce bacterial survival in the field.
Corn common rust		Fungus (<i>Puccinia sorghi</i>)	• Early-stage shows chlorotic flecks on leaf surfaces. • These develop into powdery, brick-red pustules as spores break through the leaf epidermis. • Pustules are oval or elongated (~1/8 inch long), appearing singly or in clusters.	• Fungicide application is effective against common rust; active ingredients include azoxystrobin, propiconazole, and triazole compounds. • Crop rotation is recommended to disrupt the disease cycle.

Results and Discussion

Leaf-damage identification and categorization experiments are conducted on the Plant Village dataset. The plant species along with the leaf pathologies examined are enumerated in [Table 3](#). The image collection is partitioned into independent training and evaluation subsets, with the training fraction ranging from 30% to 70%, and the remainder held out for testing.

Accuracy, loss, and precision are tracked to assess the behavior of the classification CNN model. The training and evaluation metrics of the CNN model are collated in [Table 6](#).

Table 6. Classification efficiency of the proposed CNN model.

Mode	Precision	Loss	Accuracy
Training	96.78	0.087	96.97
Testing	91.78	0.027	92.89

Semantic segmentation results

[Tables 7 and 8](#) summarize the UNet model's segmentation performance during training and validation for the crop diseases Apple Black Rot, Apple Cedar Rust, Tomato Bacterial Spot, and Corn Common Rust. In the semantic segmentation paradigm, the UNet learns at the individual pixel level. During fitting, healthy leaf regions and damaged regions are encoded as classes 0 and 1, respectively. Consequently, healthy tissue appears black in the segmentation output, whereas diseased tissue appears white. [Figures 6-9](#) display the delineated disease zones in apple, tomato, and corn leaf specimens. [Table 9](#) records the total leaf surface area alongside the diseased patch area, expressed in pixel counts. The percentage of the total affected area is likewise computed ([Figures 10 and 11](#)).




Figure 6. Segmented view of cedar apple rust.




Figure 7. Segmented view of apple rot.




Figure 8. Segmented view of tomato bacterial spot.




Figure 9. Segmented view of corn common rust.

Table 7. Damage detection efficiency of the proposed U-Net model during training.

Crop	Recall	Precision	Learning rate	Loss	lou	Accuracy	Epoch
Apple black rot	0.8534	1.000	1.00E-07	0.0070	0.8976	1.000	150
Apple cedar rust	0.9458	0.9966	1.00E-08	0.002467	0.7146	0.9510	150
Corn common rust	0.9435	0.9999	1.00E-04	0.0067	0.8436	0.9721	150

Tomato bacterial spot	0.9352	0.9982	1.00E-07	0.007742	0.8082	0.9739	150
-----------------------	--------	--------	----------	----------	--------	--------	-----

Table 8. Damage detection efficiency of the proposed U-Net model during validation.

Crop	Recall	Precision	Loss	IoU	Accuracy
Apple black rot	0.8923	0.9979	0.0160	0.8511	1.000
Apple cedar rust	0.8943	0.9988	0.0061	0.8726	0.9722
Corn common rust	0.9876	0.9988	0.0061	0.8716	0.9722
Tomato bacterial spot	0.9765	0.9872	0.0090	0.7999	0.9760

Table 9. Summary of severity estimation.

Input Image	ROI Selection	Segmented Image	Disease Identified	Total Area of the leaf (Pixels)	Affected Area (Pixels)	Severity (%)
			Apple Black Rot	5858	787	13.43
			Apple Cedar Rust	3576	681	19.04
			Tomato Bacterial Spot	4877	3260	66.8
			Corn Common Rust	9621	8341	86.6




Figure 10. Damage Detection comparative analysis of the proposed model under training.




Figure 11. Damage Detection comparative analysis of the proposed model under validation.

Conclusion

This article proposes a plant leaf damage detection framework that combines classification based on whole-image features with classification based on individual-pixel features. A CNN model is tasked with disease recognition at the image level, drawing upon a curated collection of apple, tomato, and corn leaf samples, and attains a classification

efficiency of 97%. Simultaneously, plant leaf damage is localized with pixel-level accuracy of 99%. By unifying image-based and pixel-based analytical streams, the proposed framework achieves higher disease identification fidelity than previously published approaches. In addition, the system recommends therapeutic measures tailored to the type and intensity of the damage, thereby strengthening its relevance to agricultural decision-making.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

This research was conducted in accordance with the ethical standards and guidelines set forth by Sathyabama University and relevant regulatory bodies. The following ethical considerations were taken into account:

1. **Informed Consent:** All data collection procedures involving human participants (e.g., farmers or field workers) were conducted with their informed consent. Participants were fully informed about the study's purpose, procedures, and potential risks. Written consent was obtained from all participants before the study commenced.
2. **Confidentiality:** The confidentiality of all participants was strictly maintained. Data was anonymized to protect the identity of individuals involved in the study. Personal identifiers were removed, and the data was stored securely to prevent unauthorized access.
3. **Animal and Environmental Welfare:** This research did not involve any direct experimentation on animals or cause harm to the environment. Drones were used for data collection in a manner that minimized disturbance to wildlife and the natural ecosystem.

Received: 09 Jan 2024

Revised: 08 Mar 2024

Accepted: 17 Apr 2024

Published online: 15 September 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Alirezazadeh P, Schirrmann M, Stolzenburg F. Improving deep learning-based plant disease classification with attention mechanism. *Gesunde Pflanzen*. 2023;75(1):49-59.

Daniya T, Srinivasan V. Shuffled shepherd social optimization based deep learning for rice leaf disease classification and severity percentage prediction. *Concurr Comput Pract Exp*. 2023;35(4):e7523.

Sunil CK, Jaidhar CD, Patil N. Tomato plant disease classification using multilevel feature fusion with adaptive channel spatial and pixel attention mechanism. *Expert Syst Appl*. 2023;228:120381.

Gautam V, Ranjan RK, Dahiya P, Kumar A. ESDNN: a novel ensembled stack deep neural network for mango leaf disease classification and detection. *Multimed Tools Appl*. 2024;83(4):10989-1015.

Daniya T, Vigneshwari S. Rice plant leaf disease detection and classification using optimization enabled deep learning. *J Environ Inform*. 2023;42(1).

Bouacida I, Farou B, Djakhadjakha L, Seridi H, Kurulay M. Innovative deep learning approach for cross-crop plant disease detection: A generalized method for identifying unhealthy leaves. *Inf Process Agric*. 2024.

Jung M, Song JS, Shin AY, Choi B, Go S, Kwon SY, et al. Construction of deep learning-based disease detection model in plants. *Sci Rep*. 2023;13(1):7331.

Kanna GP, Kumar SJ, Kumar Y, Changela A, Woźniak M, Shafi J, et al. Advanced deep learning techniques for early disease prediction in cauliflower plants. *Sci Rep*. 2023;13(1):18475.

Bezabih YA, Salau AO, Abuhayi BM, Mussa AA, Ayalew AM. CPD-CCNN: classification of pepper disease using a concatenation of convolutional neural network models. *Sci Rep*. 2023;13(1):15581.

DM K, Hafiane A, Bobelin L, Canals R. Image-based crop disease detection with federated learning. 2023.

Shahoveisi F, Taheri Gorji H, Shahabi S, Hosseinirad S, Markell S, Vasefi F. Application of image processing and transfer learning for the detection of rust disease. *Sci Rep*. 2023;13(1):5133.

Islam MM, Adil MAA, Talukder MA, Ahamed MKU, Uddin MA, Hasan MK, et al. DeepCrop: deep learning-based crop disease prediction with web application. *J Agric Food Res*. 2023;14:100764.

Ji M, Wu Z. Automatic detection and severity analysis of grape black measles disease based on deep learning and fuzzy logic. *Comput Electron Agric*. 2022;193:106718.

Yan J, Yan T, Ye W, Lv X, Gao P, Xu W. Cotton leaf segmentation with composite backbone architecture combining convolution and attention. *Front Plant Sci*. 2023;14:1111175.

Yang R, Guo Y, Hu Z, Gao R, Yang H. Semantic segmentation of cucumber leaf disease spots based on ECA-SegFormer. *Agriculture*. 2023;13(8):1513.

Storey G, Meng Q, Li B. Leaf disease segmentation and detection in apple orchards for precise smart spraying in sustainable agriculture. *Sustainability*. 2022;14(3):1458.

Shoaib M, Shah B, Hussain T, Ali A, Ullah A, Alenezi F, et al. A deep learning-based model for plant lesion segmentation, subtype identification, and survival probability estimation. *Front Plant Sci*. 2022;13:1095547.

Sai Reddy B, Neeraja S. Plant leaf disease classification and damage detection system using deep learning models. *Multimed Tools Appl*. 2022;81(17):24021-40.

Deng Y, Xi H, Zhou G, Chen A, Wang Y, Li L, et al. An effective image-based tomato leaf disease segmentation method using MC-UNet. *Plant Phenomics*. 2023;5:0049.

Kini AS, Prema KV, Pai SN. Early stage black pepper leaf disease prediction based on transfer learning using ConvNets. *Sci Rep*. 2024;14(1):1404.

Luo Z, Yang W, Yuan Y, Gou R, Li X. Semantic segmentation of agricultural images: a survey. *Inf Process Agric*. 2023.

Madec S, Irfan K, Velumani K, Baret F, David E, Daubige G, et al. VegAnn, Vegetation Annotation of multi-crop RGB images acquired under diverse conditions for segmentation. *Sci Data*. 2023;10(1):302.