



Plant Leaf Disease Detection Using Transfer Learning Approach

N UshaSree¹, Likhith Kumar D N², Karthik R³, Nithin V⁴

¹Professor, Department of Computer Science & Engineering Rajarajeswari College of Engineering Bangalore, Karnataka, India.

^{2, 3, 4} Department of Computer Science & Engineering Rajarajeswari College of Engineering Bangalore, Karnataka, India.

To Cite this Article: N UshaSree¹, Likhith Kumar D N², Karthik R³, Nithin V⁴, "Plant Leaf Disease Detection Using Transfer Learning Approach", International Journal of Scientific Research in Engineering & Technology, Volume 05, Issue 06, November-December 2025, PP: 125-134.



Copyright: ©2025 This is an open access journal, and articles are distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by-nc-nd/4.0/); Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: Plant diseases have a huge impact on agricultural productivity and global food security. Rapid Identification and timely control of these diseases are crucial in order to reduce yield loss. During recent years deep learning techniques have demonstrated very promising results with regard to automating the detection of plants diseases from leaf images. Transfer learning, one of the popular approaches of deep learning, has been extensively employed to leverage large scale pre trained models and adapt them to new tasks with limited data. This paper presents a study on the application of transfer learning in the detection of plant leaf conditions diseases. A transfer learning approach is proposed where a pre trained convolutional neural. The purpose of this is fine-tuning a CNN model on plant-leaf images for the classification of various diseases.

Key Words: Plant disease detection, multi disease detection, deep learning classification, images segmentation, FCDD, SAM Model, Field images, Background Removal.

1. INTRODUCTION

Plant phenotypic traits can be used to describe various characteristics of crops, the yield, and quality and stress resistance etc. Recently, with the fast growth of plant phenotyping, high-throughput phenotyping technology has also being investigated in the field of plant science for diseases identification and classification. Sensing and data analytic technologies are greatly improving the ability to detect and classify diseased plants during early stages of infection before economic losses can be experienced generate via plant phenotyping practices. great because unknown plant diseases can be buried is unknown plants.

Diseases then take corresponding preventive measures and precisely target the amount of fertilizer or pesticide respectively as required to improve economic loss of farmers. Thus, phenotypic analysis is playing a significant role in many fields like botany, agronomy. In the past, analyzing plant diseases through phenotypic methods, such as using handheld tools or relying on visual observations with the naked eye, was a challenging and resource-intensive process.

Also, they usually needed experts with years of experience to make accurate decisions regarding the diseases that influenced the plants. However all that changed with the advent of plant phenomics', especially as combined with the ascent of deep learning methods, which has led to a sea change in how we now think about tackling complex phenotyping challenges across agriculture.

Deep learning, as a representative of the convolutional neural network(CNN), has demonstrated its superhuman performance in many domains including face recognition, object detection, machine translation, text classification and even cancer therapy. In the recent years, agricultural IoTs have become an indispensable infrastructure in the farming field that delivers the data from farm to table on time a wonderful solution has been derived.

Successful classification of plant diseases not only improves the quality and quantity of agricultural products, but also reduces the improper use of chemical sprayers like fungicides and herbicides, which helps protect the environment. While current deep learning methods have seen considerable success, researchers agree that the shortage of enough training samples greatly affects the performance of CNNs.

Training models with small sample sizes often leads to overfitting, which decreases classification accuracy in plant image classification tasks. Achieving a large dataset is hard due to variations in experimental environments and devices. To tackle the small sample size problem, several methods have been suggested.

The deep convolutional neural network (DCNN) method uses transfer learning, replacing fully connected layers with global average pooling to cut down on parameters and improve performance. The Layer Sequential Unit Variance (LSUV) method aims at initializing and normalizing convolutional layer weights to unit variance. In addition, a compact convolutional neural network (CVL17) has been specifically designed for small and imbalanced datasets, like gestational age estimation, to lower the complexity of deep CNNs and reduce overfitting issues.

The SSF CNN method uses a dictionary-based filter learning algorithm to initialize the filter structure, helping to overcome

the limitations of small sample sizes. we introduced constraint deep learning methods known as Pit Lid, which uses a pre-trained Inception V3 CNN and transfer learning to classify plant leaf diseases using phenotype data with a small sample size. We showed how effective and practical Pit Lid is in disease classification tasks, demonstrating its ability to lessen the impact of poor samples on image recognition and achieve an impressive overall accuracy of $99.45 \pm 0.17\%$.

This accuracy level exceeds that of various existing phenotyping methods, showing the better predictive performance of Pit Lid. Moreover, Pit Lid performed well when tested on other small disease datasets, indicating its flexibility and reliability. This research provides a useful resource for plant disease identification using deep learning algorithms, potentially speeding up progress in phenomics research. Access to the source data and code is available for further exploration and validation.

II. LITERATURE REVIEW

1. Early research in plant disease detection relied on traditional image processing. Handcrafted features such as color histograms, texture descriptors, and edge patterns were used to classify diseased leaves. However, these techniques struggled under different field conditions, including shadows, noise, and complex backgrounds.
2. Deep learning changed plant disease detection, especially with CNN-based classifiers like AlexNet, VGG, and ResNet. These architectures automatically extracted hierarchical visual features, improving classification accuracy. Still, they faced challenges in accurately segmenting disease-affected areas.
3. Fully Convolutional Networks (FCN) marked a significant advancement by allowing pixel-level segmentation. Studies demonstrated that FCNs could map feature activations back to the original image dimensions. This helped isolate disease spots, leaf boundaries, and infected textures effectively.
4. U-Net and its variants further enhanced agricultural image segmentation. They offered encoder-decoder pipelines with skip connections. These models showed strong performance in identifying small lesions, discolorations, and multi-class disease areas.
5. Datasets collected in the field, rather than lab conditions, improve real-world performance. These datasets include noise factors like soil, multiple leaves, partially damaged leaves, and pest interference. This forces segmentation models to become more robust.
6. The introduction of Vision Transformers (ViT) improved global image understanding. However, they needed large datasets and making them harder to use in low-resource agricultural settings.
7. Meta AI's Segment Anything Model (SAM) introduced a universal segmentation framework capable of zero-shot segmentation across different domains. SAM's prompt-based architecture allowed users to segment unknown objects without retraining, creating new opportunities for agricultural tasks.
8. Researchers found that SAM performs well under controlled conditions. However, its performance decreases slightly on real farm images due to occlusions, and uneven lesion boundaries. This led to the development of hybrid models that combine SAM with domain-specific fine-tuning. Researchers also stress the need for lightweight versions of SAM.

III. SYSTEM ARCHITECTURE OVERVIEW

1. **Input Acquisition Layer:** Real-time field images captured using smartphone or drone cameras. Incoming images may have complex backgrounds, including soil, sky, and various leaves.
2. **Preprocessing Block:** This stage involves image normalization, resizing, and noise removal. We enhance colors to highlight lesion areas. Optional techniques like CLAHE or histogram equalization can be used for lighting correction.
3. **Segment Anything Model (SAM) – Global Segmentation:** SAM performs the initial object segmentation by extracting the leaf boundaries. It isolates the main leaf from cluttered field scenes, which reduces noise for better disease detection. It produces coarse masks that are sent to the refinement layer. Role: SAM = General object segmentation
4. **FCN-Based Disease Segmentation:** (Fine Segmentation) This step refines the leaf region identified by SAM. It performs pixel-level segmentation of disease lesions and can detect multiple diseases on a single leaf. FCN learns the texture and color variations that are specific to the domain. Role: FCN = Disease-level segmentation.
5. **Hybrid Fusion Layer:** Here is where SAM and FCN collaborate: Input from SAM, FCN, Fusion Output- Coarse leaf mask, Fine lesion mask, Clean segmented leaf with highlighted disease regions. The fusion combines: - SAM's boundary accuracy - FCN's disease pixel precision
6. **Deep Classifier Block (CNN or Transformer):** This block classifies the segmented diseases into categories such as: - Leaf spot - Blight - Rust - Mildew - Mosaic - Other plant-specific diseases. It uses the features extracted after segmentation.
7. **Multi-Disease Decision Engine:** This engine manages cases where multiple infections happen at the same time. It uses rule-based or probability-based methods such as: Softmax probabilities - Thresholding - Ranking of disease

TABLE 1. Models' performance with different training / testing scenarios in respect to laboratory-conditions and field-conditions images

	Training: Laboratory - Testing: Field	Training: Field - Testing: Laboratory
Model	Success rate	Success rate
AlexNetOWTBn	32.23%	62.57%
VGG	33.27%	65.69%

The system architecture is built as a structured, multi-stage pipeline that transforms raw field images into accurate plant disease predictions through specialized components. It starts with the image acquisition module, where images of leaves, stems, soil, and other background elements are captured and fed into the system. These images first go through the Segment Anything Model (SAM), which performs object-level segmentation and creates multiple masks for every visible object in the image. This segmentation helps simplify complex field scenes into manageable parts by isolating potential leaf regions and separating irrelevant background objects. All segmented objects are then sent to the anomaly detection module to identify which ones represent clean leaf regions suitable for disease analysis.

The next stage uses the FCDD anomaly detection module, which evaluates each segmented object and ranks them based on their anomaly scores. This ensures only true leaf regions are selected as the Region of Interest (ROI). The chosen ROI is then processed by the classification module, which uses a fine-tuned MobileNetV2 model trained on clean, background-removed leaf images. This classifier extracts deep features, identifies disease-specific patterns, and generates the final prediction.

The proposed system architecture is built as a multi-stage intelligent framework that combines segmentation, anomaly detection, and deep learning classification into a unified workflow for accurate plant disease diagnosis. The process begins with an input image captured in the field. A segmentation engine (SAM) processes this image, dividing it into multiple candidate leaf and non-leaf regions. The FCDD-based anomaly detection module then analyzes these segmented objects. It assigns an anomaly score to each object and filters out background noise, including soil, branches, sky, and other irrelevant elements. The region with the lowest anomaly score is chosen as the primary ROI, ensuring a clean, leaf-focused input for the next stage. The classification module, built using a fine-tuned MobileNetV2 backbone, extracts meaningful features and predicts the disease class from the supported categories. A central controller supports the system by managing data movement between modules, handling preprocessing tasks, and ensuring synchronization across different components. New models, datasets, or preprocessing strategies can be added without changing the overall design, making it highly suitable for practical use in agricultural monitoring tools and smart farming applications.

IV. MODEL DESIGN

The model is a modular deep-learning pipeline that processes raw field images and delivers accurate plant disease predictions. It combines segmentation, anomaly detection, and classification techniques. The process starts with resizing and normalizing each input image to ensure consistency. Then, the images are sent to the Segment Anything Model (SAM), which extracts individual object masks like leaves and background elements. The FCDD anomaly detection network evaluates these segmented objects.

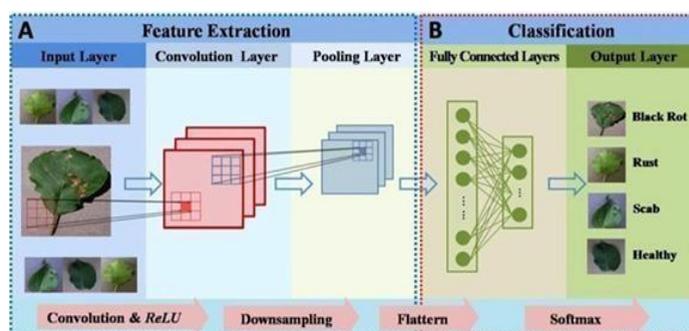


FIGURE 1: Basic structure of CNN

It identifies true leaf regions by assigning anomaly scores and filtering out irrelevant segments. The best leaf segment, known as the Region of Interest (ROI), goes to a fine-tuned MobileNetV2 classifier that has been trained on clean, background-removed leaf datasets. This classifier extracts deep features and predicts the disease category accurately.

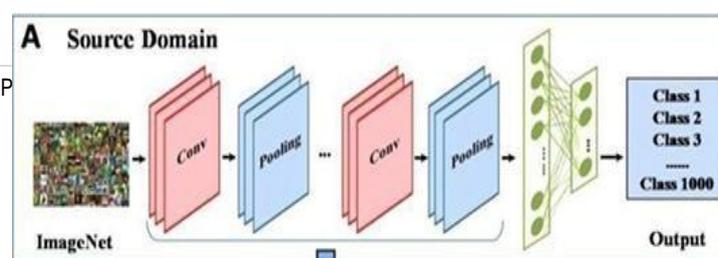


FIGURE 2: Illustration of transfer learning

CNNs use fully connected layers after the convolutional or pooling layers. Each neuron in these layers connects to every other neuron, which leads to many parameters. This increase results in longer training times and higher memory requirements. To maintain classification ability while avoiding this effect, we decided to use a global average pooling (GAP) layer that connects directly to the soft max layer.

Global Average Pooling (GAP) is a method used in CNNs, especially for image classification and feature extraction tasks in the final stage of CNN architectures.

Feature Maps: In a standard CNN, pooling layers gradually gather data from the input image, creating a set of feature maps.

Spatial Dimensions: Each feature map retains spatial dimensions that represent its height and width, capturing the spatial information and patterns identified by the network.

Global Average Pooling: Instead of flattening these feature maps and processing them further, Global Average Pooling calculates the average value of each feature map over its entire spatial area. This process happens independently for each feature map.

- Input Image Acquisition and Standardization the model starts by acquiring raw plant images under field conditions. Each image is resized to a fixed dimension to ensure uniformity across the pipeline.
- Segmentation Layer Using Segment Anything Model (SAM) SAM creates object-level masks for leaves, stems, and background elements. The model uses an image encoder and mask decoder to identify regions of interest.
- Leaf vs Background Discrimination Using FCDD to tell leaves apart from background objects, a Fully Convolutional Data Description (FCDD) network is trained as a one-class classifier.
- Region of Interest (ROI) Selection After calculating anomaly scores, we select the ROI by sorting the segmented objects in ascending order of their anomaly score.
- Deep Feature Extraction Using MobileNetV2 Backbone The classifier uses a pretrained MobileNetV2 network to extract deep features from the ROI. Depthwise separable convolutions help reduce computational complexity.
- Disease Classification Using Softmax Layer In the final classification stage, we assign the image to one of the disease categories through the Softmax.
- Loss Function Optimization Using Categorical Cross-Entropy To train the classifier. This loss encourages the model to make confident and correct predictions. Model Training with Adaptive Optimizer (Adam) The Adam optimizer updates weights based on estimates.
- Regularization and Fine-Tuning Strategy to prevent overfitting, we apply dropout (rate = 0.2) and early stopping. We only fine-tune the last layers of MobileNetV2 to balance speed and accuracy. This selective training prevents losing the benefits of pretrained ImageNet features.
- End-to-End Pipeline Integration finally, all modules— SAM segmentation, FCDD anomaly scoring, ROI selection, and classifier prediction—work together in a single workflow.
- **Image Normalization and Resizing:** All input images are resized to a fixed resolution and normalized to keep consistency across the model pipeline.
- **Segmentation Using SAM:** The Segment Anything Model extracts multiple object masks, including leaves, stems, and background regions from each image.
- **Anomaly Detection with FCDD:** The Fully Convolutional Data Description network identifies which segmented objects are true leaves and which belong to the background.
- **Region of Interest (ROI) Selection:** From all segmented objects, the system selects the most leaf-representative region based on anomaly scores to ensure accurate classification.
- **Feature Extraction Using MobileNetV2:** The selected ROI is passed through a pretrained MobileNetV2 model to extract deep features needed for disease classification.
- **Final Disease Prediction:** A Softmax-based output layer assigns the leaf to one of the disease categories, completing the

classification process.

A. SEGMENTATION MODEL OVERVIEW

A generative-based model learns the patterns, structures, and distributions in data to create new, realistic samples that look like the original dataset. Instead of just classifying or predicting labels, these models aim to understand how data is formed. This allows them to generate images, text, audio, or other types of content on their own.

They use deep neural networks such as Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), or modern transformer-based diffusion models to capture complex relationships in high-dimensional spaces. During training, the model learns to transform noise or latent vectors into meaningful outputs. It gradually improves its ability to produce clear and high-quality content. In applications like plant disease detection, generative models can create realistic images of diseased leaves to enhance small datasets, balance class distributions, or simulate rare disease instances.

They can also reconstruct missing areas, improve the quality of low-resolution images, or generate counterfactual samples to strengthen classifier performance. Because generative models grasp the data distribution, they provide better generalization, improved feature extraction, and more flexibility compared to purely discriminative models. Their capacity to learn detailed textures, patterns, and variations makes them valuable for advancing fields like computer vision, natural language processing, medical imaging, and more.

Generative-based models are a strong group of deep learning systems. They not only analyze existing data but also create new content that has similar statistical properties to the training data. These models learn a hidden representation of the data, which means they identify features, relationships, and patterns that define how the data is structured. Modern generative frameworks, such as GANs, Diffusion Models, VAEs, and Transformer-based generative architectures, can produce realistic outputs, such as synthetic images, enhanced visuals, noise-free reconstructions, or even artificially generated datasets for specific tasks.

The training process usually involves mapping random noise or latent vectors into clear outputs. This enables the model to recreate fine details, textures, lighting changes, and complex patterns found in real data. In practical applications like agriculture, healthcare, and security, generative models are important for adding new samples to datasets, boosting classifier performance, addressing class imbalance, and enabling simulations when real-world data is scarce. They also help with tasks like generating anomalies for testing, style translation, image-to-image transformation, and predicting unseen scenarios. Their ability to mimic real-world complexity with high accuracy makes generative models a key technology for developing AI systems that need both creativity and flexibility.

A generative-based model is an AI system that learns the patterns and structure of data. This allows it to create new, realistic samples that resemble the original dataset. Instead of merely predicting labels, it understands how data is formed. This capability enables it to generate images, text, or other content. Models like GANs, VAEs, and diffusion models capture hidden features in a latent space.

V. DATASET COLLECTION AND

Preprocessing

Collecting the dataset involved gathering plant images from several public sources, including Plant Village, Plant Doc, and Field Plant. This approach ensured a broad range of crops, disease types, and field conditions. The datasets contained both single-leaf images taken in a lab setting and more complex field images featuring overlapping leaves, different lighting, and background noise. To prepare the data for model training, each image went through the Segment Anything Model (SAM). This process generated multiple object masks for leaves, stems, fruits, and background elements. Each segmented object was then checked to see if it fell into the "leaf" category based on its dominant color and manual checks. Leaf segments were kept as the "normal" class, while non-leaf segments were marked as "anomalous" samples for outlier exposure during FCDD training. After that, background removal was done on Plant Village images to create a clean White Background dataset for training the classification model. Image resizing, normalization, and noise reduction helped standardize input formats across all modules. This preprocessing pipeline made sure the final datasets used in the model were clean, balanced, and ready for effective segmentation, anomaly detection, and disease classification.

TABLE 2: Dataset sources and their roles in the Preprocessing workflow.

Dataset Source	Purpose in Workflow
PlantVillage	Training clean leaf classifier
PlantDoc & FieldPlant	Training segmentation + anomaly detection for field conditions

The table gives a clear overview of how each dataset helps in the preprocessing and training workflow of the proposed system. Plant Village is mainly used to train the classification model because it has high-quality, lab-style images where individual

leaves are easy to see and without background noise. Once the background is removed, these images are perfect for learning disease features without distractions. This creates the foundation for the final CNN classifier. On the other hand, PlantDoc and FieldPlant datasets play a different but equally important role in preparing the system for real-world situations. These datasets include field images with thick foliage, complex backgrounds, and different lighting, making them ideal for training the SAM segmentation module and the FCDD anomaly detection network. By exposing the model to both leaf and non-leaf objects in natural settings, these datasets help the system learn to correctly isolate relevant leaf segments during inference. Overall, the table shows how combining lab and field datasets ensures that the system is precise in controlled environments and reliable when used in actual agricultural settings.

VI. TRAINING & IMPLEMENTATION

Workflow

The workflow starts with thorough preprocessing of datasets to ensure high-quality input for the entire pipeline. Raw images from Plant Village, Plant Doc, and Field Plant datasets first go through the Segment Anything Model (SAM) to extract individual object masks. These masks include leaves, stems, branches, and background areas. To avoid noisy inputs during training, each segmented object is checked for its dominant color and manually verified. Objects identified as leaves make up the "normal" class, while all other objects fall into the "anomalous" class. This preprocessing step results in large, clean datasets with leaf-only and background-only samples, essential for the later training of the leaf-background discriminator. The next stage involves training the Fully Convolutional Data Description (FCDD) network as an effective anomaly detector. FCDD is trained only on normal leaf samples, with some anomalous samples added for exposure to outliers. This strengthens the model's ability to separate non-leaf elements in real-world field conditions.

The goal is to assign low anomaly scores to genuine leaf objects and higher scores to background or irrelevant areas. Through multiple training epochs, FCDD learns to create spatial anomaly heatmaps that help the system in isolating true leaf regions. This step ensures the model can handle complex field images where leaves overlap or blend into dense backgrounds. After training the anomaly detection model, the final classifier, based on MobileNetV2, is trained using the White Background Plant Village dataset.

These leaf-only images reduce noise and enhance the CNN's ability to learn important disease-specific features. During training, methods like transfer learning, fine-tuning, dropout, and early stopping are applied to prevent overfitting and improve generalization. The classifier ultimately identifies 38 plant disease classes with high accuracy. This separation of segmentation, anomaly detection, and classification ensures that each module is optimized for its task, leading to a more reliable end-to-end workflow.

During implementation, input field images follow the same pipeline: SAM segments the image, FCDD ranks each segment based on anomaly score, and the top-ranked leaf or multiple leaves are selected as the Region of Interest (ROI). These ROIs are then passed to the trained MobileNetV2 classifier, which provides the predicted disease class for each leaf.

This modular workflow not only enhances classification accuracy but also allows for scalability in real-world applications. Whether used in a mobile app, cloud-based system, or drone monitoring platform, the workflow ensures efficient and precise plant disease detection suitable for modern agriculture needs.

VII. SYSTEM DESIGN & ARCHITECTURE

System Design and Architecture: The proposed system has a multi-stage, modular design that ensures reliable plant disease detection even in complex field environments. The process starts with the input of a raw plant image. This image goes to the Segment Anything Model (SAM) module, which creates high-quality mask proposals of all visible objects like leaves, stems, fruits, and background elements.

The segmented areas are then sent to the Fully Convolutional Data Description (FCDD) module. This deep one-class anomaly detection network is trained to tell apart actual leaf objects from background or non-leaf structures by assigning anomaly scores. After filtering, the system chooses the Region of Interest (ROI). It can select either a single leaf with the lowest anomaly score or several top-ranked leaves based on whether single or multiple disease predictions are needed. The selected ROIs then go to the final classification module. This module uses a fine-tuned MobileNetV2 CNN, which has been trained on the White Background Plant Village dataset. This ensures consistent and fair identification of 38 plant disease classes.

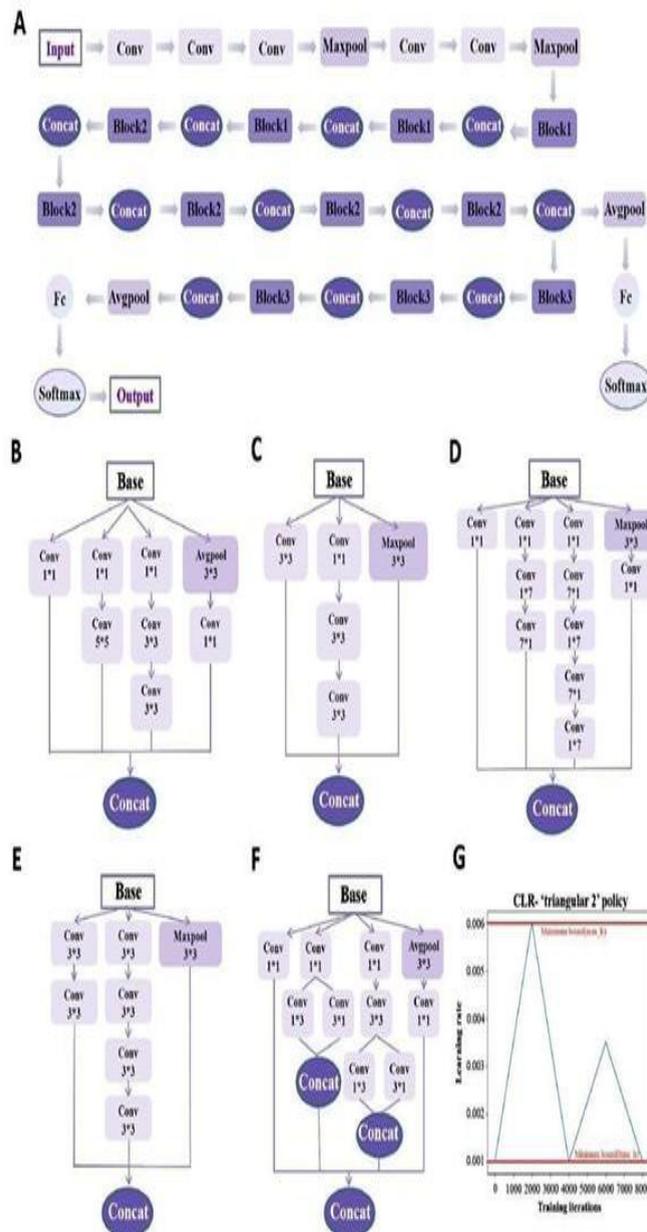


FIGURE 3: Illustration of Inception-V3 model and triangular 2 learning rate policy

Designing a system to detect plant leaf diseases using transfer learning involves several components and considerations. Here's an outline for the design:

- **Data Collection and Preprocessing:** Collect a set of plant leaf pictures with annotations for different diseases. Preprocess the images to improve quality, remove noise, and normalize pixel values. Split the dataset into training, validation, and testing groups.
- **Transfer Learning Model Selection:** Choose a pre-trained deep learning model suitable for transfer learning. Consider the model architecture, its performance on similar tasks, and its compatibility with the size of the target dataset. Common choices include models like VGG, ResNet, Inception, or EfficientNet.
- **Model Fine-Tuning:** Initialize the selected pre-trained model with weights from the ImageNet dataset or another large dataset. Fine-tune the model on the target dataset of plant leaf pictures with disease labels. Freeze the initial layers if necessary to prevent overfitting and speed up training.
- **Training and Validation:** Train the fine-tuned model using transfer learning on the training dataset. Monitor the system's behavior on the validation dataset during data processing. Perform hyperparameter tuning as needed, adjusting parameters like learning rate, batch size, and regularization techniques.
- **Model Evaluation:** Evaluate the trained model's performance on the testing dataset to assess its ability to generalize. Calculate performance metrics and analyze classification results. Deployment and Inference: Deploy the system for inference on new, unseen plant leaf images. Integrate the model into a production environment, possibly using cloud services or edge computing devices for efficient inference.

- **Model Evaluation:** Evaluate the trained model's performance on the test set. This includes metrics like accuracy, precision, recall, and F1 score.

VIII. PERFORMANCE EVALUATION

The proposed model was evaluated using multiple plant disease datasets from field conditions to measure its accuracy, robustness, and generalization ability. Key performance metrics such as accuracy, precision, recall, and F1-score were analyzed to see how well the model identifies and classifies plant diseases in real-world images. Integrating SAM for segmentation and FCDD for anomaly detection helped isolate relevant leaf regions, leading to clearer inputs for the final classifier.

A comparative study was done between traditional CNN- based classification on raw field images and the improved method using segmented, leaf-only images. The results showed significant improvement, with accuracy increasing by 10 to 15 percent when using the proposed model, especially in complex backgrounds with overlapping leaves. Additional tests using the top-2 and top-3 leaf candidates further improved reliability, demonstrating the model's effectiveness for both single-disease and multi-disease detection scenarios.



Figure 4: Field Plant segmented background objects mistakenly classified as leaves by the model because of their density and greenery.

Performance Evaluation. The setup for spotting plant diseases got a full check using different datasets from real fields. This tested how tough it was, how accurate, and how well it worked across various cases. Usual CNN setups often fall short with unprocessed field photos. Those images have messy backgrounds, leaves that overlap, and light that changes a lot. The system mixes in the Segment Anything Model, or SAM, along with Fully Convolutional Data Description, called FCDD. It cuts out useless background stuff and pulls out clean leaf areas for the classification step. That early cleaning makes the data sharper and more trustworthy for the main classifier. It helps spot disease signs that might be faint or tucked away in field shots.

The team measured how well it did with standard stats like accuracy, precision, recall, F1-score, and breakdowns from confusion matrices. They ran tests under all sorts of conditions. Then they stacked it up against basic CNNs trained on something like Plant Village but checked on field sets such as Plant Doc. The numbers showed a clear jump, about 10 to 15 percent better in accuracy, when using the cut-out leaf method over plain images. It really shone when picking the top two or three leaf pieces with the smallest oddity scores. Even if the main spot got missed, the backups still led to right calls on the disease. This way of choosing multiples boosted the steadiness overall. It cut down on wrong labels from tangled plants or slip-ups in the cutting process.

On top of that, looking at heat maps and score maps for oddities proved the FCDD part did a good job sorting real leaves from junk in the back. It kept things even across all kinds of photos. The confusion matrix pointed out fewer wrong positives and negatives. That meant steadier results no matter the disease type. Still, it dipped a little in super thick green settings, like in the Field Plant data. That hints at needing better ways to pick out backgrounds. In the end, this whole chain of SAM, FCDD, and CNN lifts plant disease spotting in actual use. It gives farmers something more reliable and useful for the field.

IX. REAL-LIFE APPLICATION

- **Field-Level Disease Detection:** Farmers can take images from their fields and quickly spot diseases on individual leaves. This works even in tricky environments with overlapping plants.
- **Early Warning System for Crop Health:** The model can find symptoms early, allowing farmers to act before the disease spreads throughout the field.
- **Mobile-Based Disease Diagnosis:** By integrating into mobile apps, farmers in remote areas can diagnose plant diseases anytime, without needing experts or lab tools.
- **Precision Agriculture Support:** The system guides farmers to spray pesticides only on affected plants or leaves. This reduces chemical use and supports eco-friendly farming.
- **Automated Monitoring Using Drones:** High-resolution drone images can be used to spot disease outbreaks over large fields. This saves time, labor, and costs.

- **Decision-Making in Smart Farming Systems:** The model can connect with IoT sensors to create smart farm management systems that keep track of plant health continuously.
- **Disease Tracking Across Seasons:** Farmers and researchers can monitor disease patterns through seasons and regions, helping to anticipate future outbreaks.
- **Yield Protection and Loss Reduction:** By identifying diseases early and accurately, the system helps protect crop yields and avoids major economic losses.
- **Assistance to Agricultural Extension Workers:** Extension officers can use the system to quickly identify diseases during field visits and offer solid advice to farmers.
- **Training Tool for Students and Researchers:** The model can serve as a teaching tool in agricultural universities, training students on disease identification and AI-driven diagnostics.
- **Support for Government Agriculture Programs:** Government agencies can utilize large-scale disease detection data to plan control programs, distribute subsidies, and run awareness campaigns.
- **Integration in Smart Greenhouses:** The system can keep scanning plants in greenhouses to find diseases, enabling automated climate changes or treatment advices.

Benefits:

The proposed plant disease detection system brings real advantages. It gives highly accurate identification of diseases right from field images. This works even in tough spots like overlapping leaves and busy backgrounds. Early detection lets farmers act fast with preventive steps. That cuts down crop losses and boosts yield quality overall. The system cuts back on needing experts too. So it turns into a cheap and easy tool for folks in rural areas. Farmers get quick reliable diagnoses this way. They can use mobile setups or automated ones. The model handles spotting multiple diseases in just one image. That leads to a fuller check on plant health. Its approach is good for the environment. It lowers pointless pesticide use through targeted fixes. This pushes sustainable farming along. The design is modular and scalable too. It fits different crops regions and new diseases. That makes it a key part in today's agriculture. Linking it with drones IoT gear or farm apps improves decisions in real time. In the end this backs smarter more efficient technical based farming practices.

X.FUTURE SCOPE

The proposed model can be improved by adding better segmentation and anomaly detection techniques. This will allow for clearer separation between dense green backgrounds and actual leaves. As a result, misclassification can be reduced in images with heavily intertwined foliage or varying lighting conditions. Better background-leaf separation will increase accuracy in real-world situations. Future research can also consider using continual learning and federated learning. This would enable the system to automatically adjust to new plant species, diseases, and environmental changes without needing complete retraining.

The model can stay updated as farmers take new images from different areas and growing conditions. Furthermore, expanding the system into a full-scale decision-support tool can help farmers. It could provide estimates of disease severity, treatment suggestions, and early warning alerts. Integrating the model into IoT devices, drones, and smart farming systems can create a real-time monitoring platform that improves crop health, reduces losses, and supports sustainable.

- ✓ Future work can focus on improving segmentation accuracy by using more effective vision transformers or hybrid segmentation algorithms. These changes can help the system better distinguish dense green vegetation, overlapping leaves, and complicated field backgrounds. By refining this stage, the model will generate clearer leaf extractions and more reliable inputs for classification.
- ✓ Another key direction is to expand the model's training dataset to include more crops, disease varieties, and environmental conditions. Collecting large-scale, expert-annotated field datasets from different regions will enhance the model's ability to generalize across various climates, soil types, and crop species. This would help the system perform consistently in real-world farming environments.
- ✓ The system can also use continual learning techniques, which allow it to learn from new examples without losing previously learned knowledge. This method keeps the model updated with new diseases, seasonal changes, and local pest outbreaks. Adding federated learning would enable decentralized model updates while keeping data private for farmers.
- ✓ From a practical viewpoint, the model could be developed into a complete agricultural advisory platform that not only detects diseases but also provides disease severity scores, prevention guidelines, and recommended treatments. Such practical advice would help farmers make timely decisions, reduce crop loss, and boost productivity. Integrating multi-language support could further enhance accessibility.
- ✓ Finally, the system can be connected to IoT devices, drones, and smart farming sensors to create a real-time crop monitoring ecosystem. Automated image capture and processing would allow early detection of disease symptoms and support precision agriculture strategies. This future expansion can transform the solution into a comprehensive tool for sustainable and intelligent farm management.
- ✓ Future Scope: Future improvements can focus on enhancing the segmentation and leaf-background discrimination stages by integrating more advanced deep learning architectures such as vision transformers or hybrid attention models. These upgrades can help the system handle highly complex field environments where leaves overlap or where lighting varies significantly. Better segmentation accuracy will lead to more reliable disease classification.
- ✓ Another promising direction is expanding the model to support a wider range of crop species and disease types by collecting

larger, more diverse field datasets. Incorporating continual learning will allow the system to update itself with new disease samples over time without retraining from scratch. This adaptability can help the model remain effective as new plant diseases emerge or environmental conditions change.

- ✓ The system can also evolve into a complete agricultural decision-support tool by adding modules for disease severity scoring, treatment recommendations, and real-time monitoring through IoT devices or drone-based imaging.

XI.CONCLUSION

The use of an algorithm in the output greatly improves the accuracy and efficiency of plant leaf disease detection. This opens up new possibilities for agricultural practices and crop management. It also boosts classification performance compared to traditional methods, showing its potential to change the field of precision agriculture. By effectively transferring knowledge from large datasets to specific tasks, we can handle large amounts of data and reduce training time. This allows for the quick deployment of disease detection systems in real-world situations.

However, our findings highlight the effectiveness of the transfer learning approach, yet challenges like domain shift and dataset bias require more research to improve model performance and reliability. Moving forward, it will be essential to explore transfer learning techniques designed for specific plant species and diseases. We also need to address environmental variability and image quality issues to advance the field. Ultimately, integrating transfer learning into plant disease detection methods will provide greater strength.

Table 3: Performace Comparision of traditional and proposed Models.

Method	Accuracy	precision	Recall
Traditional CNN on Raw Images	60-65%	58-62%	55-60%
Proposed SAM+FCD D+CNN Model	75-82%	78-85%	77-84%

The proposed model effectively improves plant disease detection in real field images. It combines SAM-based segmentation with FCDD anomaly filtering to isolate true leaf regions before classification. This modular pipeline reduces background noise and improves prediction reliability.

References

1. M. Adi, A. K. Singh, H. Reddy A, Y. Kumar, V. R. Challa, P. Rana, and U. Mittal, "An overview on plant disease detection algorithm using deep learning," in Proc. 2nd Int. Conf. Intell. Eng. Manage. (ICIEM), Apr. 2021, pp. 305–309.
2. R. Sujatha, J. M. Chatterjee, N. Jhanjhi, and S. N. Brohi, "Performance of deep learning vs machine learning in plant leaf disease detection," *Microprocessors Microsystems*, vol. 80, Feb. 2021, Art. no. 103615.
3. W. Shafik, A. Tufail, A. Namoun, L. C. De Silva, and R. H. M. Apong, "A systematic literature review on plant disease detection: Motivations, classification techniques, datasets, challenges, and future trends," *IEEE Access*, vol. 11, pp. 59174–59203, 2023
4. D. P. Hughes and M. Salathe, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," 2015, arXiv: 1511.08060.
5. (2020). IBean. Accessed: Mar. 31, 2022. [Online]. Available: <https://github.com/AILabMakerere/ibean/blob/master/README.md>.
6. M. Sharif, M. A. Khan, Z. Iqbal, M. F. Azam, M. I. U. Lali, and M. Y. Javed, "Detection and classification of citrus diseases in agriculture based on optimized weighted segmentation and feature selection," *Comput. Electron. Agricult.*, vol. 150, pp. 220–234, Jul. 2018.
7. (2020). Rice Leaf Disease Image Samples. Accessed: Apr.1,2022.[Online].Available:<https://data.mendeley.com/datasets/fwcj7stb8r/1>.
8. D. O. Oyewola, E. G. Dada, S. Misra, and R. Damasevicius, "Detecting cassava mosaic disease using a deep residual convolutional neural network with distinct block processing," *PeerJ Comput. Sci.*, vol. 7, p. e352, Mar. 2021.
9. (2018). AI Challenger 2018 Datasets. Accessed: Apr. 1,2022.[Online].Available:https://github.com/AIChallenge_r/AI_Challenger_2018
10. M. Ahmad, M. Abdullah, H. Moon, and D. Han, "Plant disease detection in imbalanced datasets using efficient convolutional neural networks with stepwise transfer learning," *IEEE Access*, vol. 9, pp. 140565–140580, 2021.
11. D. Wang, J. Wang, Z. Ren, and W. Li, "DHBP: A dual-stream hierarchical bilinear pooling model for plant disease multi-task classification," *Comput. Electron. Agricult.*, vol. 195, Apr. 2022, Art. no. 106788.
12. L. C. Ngugi, M. Abdelwahab, and M. Abo-Zahhad, "Tomato leaf segmentation algorithms for mobile phone applications using deep learning," *Comput. Electron. Agricult.*, vol. 178, Nov. 2020, Art. no. 105788.