IJRAR.ORG E-ISSN: 2348-1269, P-ISSN: 2349-5138

INTERNATIONAL JOURNAL OF RESEARCH AND ANALYTICAL REVIEWS (IJRAR) | IJRAR.ORG
An International Open Access, Peer-reviewed, Refereed Journal

Plant Leaf Disease Multilevel Classification Using Few Shot Learning

Ashutosh Sabnis
Department of Computer Engineering
PES Modern College of Engineering, Pune

Rhutam Mahajan
Department of Computer Engineering
PES Modern College of Engineering, Pune

Aishwarya Behere
Department of Computer Engineering
PES Modern College of Engineering, Pune

Vaishnavi Salunke
Department of Computer Engineering
PES Modern College of Engineering, Pune

Under the Guidance of:
Prof. Dr. Mrs. S. A. Itkar
Head, Department of Computer Engineering
PES Modern College of Engineering, Pune

Abstract —Plant leaf diseases pose significant threats to global agriculture, yet traditional deep learning methods require extensive labelled data. We present a hierarchical few-shot learning (FSL) framework combining dataset cleaning, structured dataset splitting, and prototypical networks built on an EfficientNet-B0 backbone. We first cleaned the Plant Village dataset using intra-class embedding similarity, retaining 12 robust disease classes (13,135 images). Next, we created 5-way, 5-shot training and validation splits (total 240 images). The model was trained for up to 20 epochs with early stopping, achieving a peak validation combined accuracy of **95.73%** at epoch 6 (plant-level 96.13%, disease-level 95.33%). These results demonstrate reliable classification performance in low-data regimes, paving the way for cost-effective, scalable plant disease diagnostics.

Index Terms—Few-Shot Learning, Plant Disease, Prototypical Network, EfficientNet-B0, Dataset Cleaning, Hierarchical Classification, Early Stopping.

1. Introduction

Plant leaf diseases, caused by bacterial, fungal, and viral pathogens, drastically reduce crop yields and quality. Manual disease inspection across large fields is time-consuming, inconsistent, and often impractical. [1] While convolutional neural networks (CNNs) have improved diagnostic accuracy, they demand thousands of labelled samples per class—an unrealistic requirement for scarce or emerging diseases.[2] Few-shot learning (FSL) mitigates this limitation by enabling generalization from only a handful of examples. In this work, we propose a hierarchical FSL framework that: (1) cleans the dataset to ensure intra-class coherence, (2) constructs balanced k-shot splits, and (3) trains prototypical networks for both coarse- and fine-grained classification. Our contributions include:

1. A **cleaned Plant Village subset**: *12 disease classes filtered via an embedding-based similarity threshold (≥ 0.7), yielding 13,135 images.*
2. **Structured few-shot splits**: 5-way, 5-shot train/validation sets (120 support and 120 query images).
3. A **hierarchical prototypical network** on EfficientNet-B0 achieving 95.73% combined validation accuracy under strict low-data conditions.

2. Related Work

A. Traditional and Deep Learning Approaches

Early plant-disease detection efforts relied heavily on manually engineered features—color histograms, texture descriptors (e.g., GLCM), and shape-based metrics—fed into classical classifiers like SVMs and K-Nearest Neighbors. While these pipelines achieved reasonable accuracy on controlled datasets, they proved brittle in field conditions: varying illumination, complex backgrounds, and leaf occlusions caused a steep drop in performance [3]. With the advent of deep learning, convolutional neural networks (CNNs) revolutionized the field by learning hierarchical feature representations directly from raw pixels [4]. Models such as VGG, ResNet, and later EfficientNet extract multi-scale spatial patterns (vein structure, lesion morphology) in an end-to-end fashion, yielding state-of-the-art accuracy (often >95% on benchmark datasets) [5], [6]. However, this accuracy gain comes at the cost of massive labeled datasets—typical implementations require thousands of annotated images per disease class, which are labor-intensive and costly to acquire in agricultural settings.

B. Transfer Learning and Zero-Shot Learning

To mitigate the data-hungry nature of deep CNNs, researchers turned to transfer learning, repurposing models pretrained on large-scale corpora (e.g., ImageNet) and fine-tuning them on limited plant-disease images.[7] This paradigm reduces annotation requirements and speeds convergence, yet still often demands on the order of tens to hundreds of samples per class to avoid overfitting [8]. Zero-shot learning (ZSL) pushes this boundary further by leveraging semantic embeddings—word vectors, ontology graphs, or attribute annotations—to recognize entirely unseen disease categories at test time [9]. ZSL systems map visual features into a shared semantic space, then infer novel classes via their descriptive vectors. Although promising for emerging or rare diseases, zero-shot approaches hinge on the availability of reliable auxiliary information (e.g., attribute lists, relationships), which is scarce or inconsistent across plant pathology domains.

C. Few-Shot Learning in Agriculture

Few-shot learning (FSL) bridges the gap between transfer learning and ZSL by enabling rapid adaptation to new classes with only a handful of examples. Metric-based FSL methods—prototypical networks, matching networks, and relation networks—construct embedding spaces where samples cluster by class, facilitating classification via simple distance metrics [10], [11]. While these techniques have demonstrated competitive performance on standard benchmarks (Omniglot, mini-ImageNet), their application to plant disease detection remains limited. Existing agricultural FSL studies often neglect two critical factors: (1) **hierarchical disease taxonomy**, where coarse-level (e.g., fungal vs. bacterial) and fine-level (specific pathogen) labels exhibit structured relationships; and (2) **robust dataset cleaning**, ensuring that intra-class visual coherence is preserved despite noisy or mislabeled images [12], [13]. Our work addresses both gaps by implementing a hierarchical prototypical network on an embedding-cleaned subset of PlantVillage, thus providing reliable few-shot classification under realistic, low-data constraints.

3. Methodology

Our pipeline involves (A) dataset cleaning, (B) few-shot split creation, (C) preprocessing and augmentation, and (D) hierarchical prototypical network training.

A. Dataset Cleaning

We loaded 38 classes from the Plant Village train and Val directories, extracted EfficientNet-B0 embeddings normalized to unit length, and computed average intra-class cosine similarity. Classes with mean similarity <0.7 were removed, resulting in **12 classes** (13,135 images)[14].

B. Few-Shot Split Creation

From the cleaned classes, we sampled **10 support** and **10 query** images per class for both training and validation. This yielded 120 support images in train and 120 query images in Val (total 240 images).

C. Preprocessing & Augmentation

Images are resized to 224×224 px, normalized with ImageNet mean/std, and during training augmented with random resized crop, horizontal/vertical flips, color jitter, rotation, and random erasing to enhance generalization[15].

D. Hierarchical Prototypical Networks

We remove the classifier head of EfficientNet-B0 to extract 1280-D features, followed by an embedding head (Linear 1280→768, ReLU, Dropout, Linear 768→384, Batch Norm)[16]. Two prototypical network instances perform:

- **Level 1**: 5-way classification among broad categories (e.g., fungal, bacterial).
- **Level 2**: 5-way classification among specific diseases (e.g., powdery mildew vs. early blight). Class prototypes are computed as the mean of support embeddings, and Euclidean distances to query embeddings produce logits for classification.

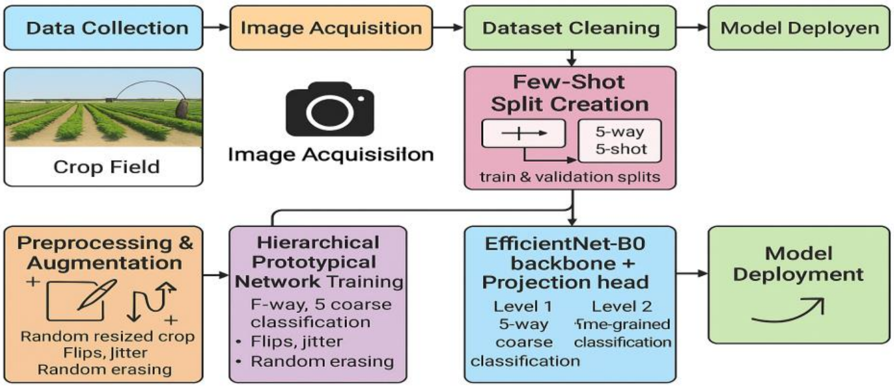

Fig.I. System Architecture

4. ALGORITHM

We employ a hierarchical prototypical network trained episodically for few-shot plant disease classification. In each episode, the model samples 5 classes, with 5 support and 10 query images per class.[17] Support images are embedded using EfficientNet-B0 and a projection head. For each class, a prototype is computed as the mean of its support embeddings. Each query image is embedded, and its squared Euclidean distance to each prototype is calculated.[18] The probability that a query belongs to a class is determined by a SoftMax over the negative distances (optionally scaled by a temperature parameter).[19] The model is optimized using cross-entropy loss with label smoothing. Model parameters are updated via AdamW, with early stopping based on validation accuracy and overfitting criteria. The core of the method is summarized by the following formula for classification probability:

$$p\,(y = c \mid x_q) = \frac{\exp(-d_c(x_q)/T}{\sum_{c'} \exp(-d_c(x_q)/T)}$$

Where $d_c(x_q)$ is the squared Euclidean distance between the query embedding and class prototype, and T is the temperature parameter. This framework enables rapid adaptation to new classes by learning a metric space suitable for distance-based classification[20].

5. Experimental Setup and Results

A. Hardware & Implementation: Experiments ran on an NVIDIA GPU (if available) or CPU, implementation in PyTorch, training for up to 20 epochs with AdamW (lr=5e-5), ReduceLROnPlateau scheduler, gradient scaling when GPU used.

B. Episode Configuration: Each epoch consists of 20 training episodes (5-way, 5-shot, 10 query) and 10 validation episodes (5-way, 5-shot, 15 query). Early stopping triggers after 5 epochs without improvement or detected overfitting (train-Val gap >5%).

C. Quantitative Results:

1) Table I shows epoch-wise metrics, with peak combined validation accuracy at epoch 6 on fewShot dataset.(240 images)

Epoch	Train Loss	Train Combined Acc	Val Plant Acc	Val Disease Acc	Val Combined Acc
1	1.4206	46.50%	89.07%	93.47%	91.27%
2	1.1427	69.40%	94.40%	96.53%	95.47%
3	1.0142	80.30%	93.60%	94.53%	94.07%
4	0.9302	87.10%	91.87%	95.47%	93.67%
5	0.8720	92.00%	91.07%	98.53%	94.80%
6	0.8698	92.60%	**96.13%**	**95.33%**	**95.73%**
7	0.8051	96.20%	92.67%	97.60%	95.13%
8	0.8178	95.50%	91.73%	95.87%	93.80%
9	0.8347	94.70%	94.27%	95.87%	95.07%
10	0.8113	96.00%	94.13%	95.07%	94.60%

Early stopping occurred at epoch 11 with no further improvement.

2) Now that we’ve shown performance on few-shot episodes, we evaluate on the entire cleaned PlantVillage dataset(13,135 images)

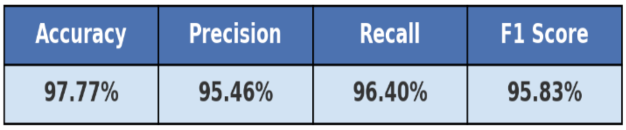

Accuracy	Precision	Recall	F1 Score
97.77%	95.46%	96.40%	95.83%

Fig.2 Evaluation Matrix

D. Dataset Statistics: Post-cleaning: 12 classes, 13,135 images. Few-shot splits: 240 images total (train 120, Val 120).

Fig.3. Accuracy vs Epoch

Results

To test the model in a real-world scenario, we developed a web-based interface using Flask. The interface allows users to upload an image of a plant leaf and receive a prediction of the plant name and its health status.

Figure A shows the upload interface where the user selects an image.

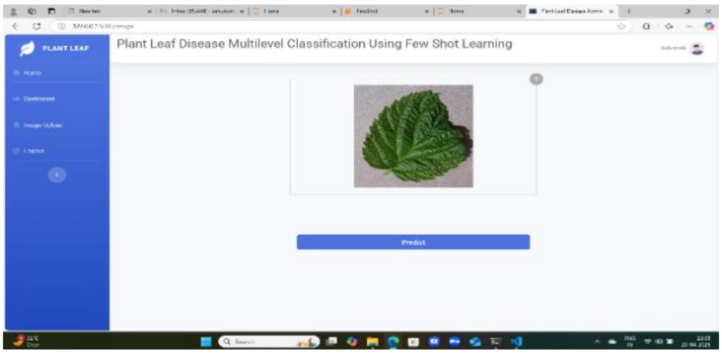

Fig A : Uploaded Image

Figure B displays the result page showing the predicted plant (Raspberry) and the classification (Healthy).

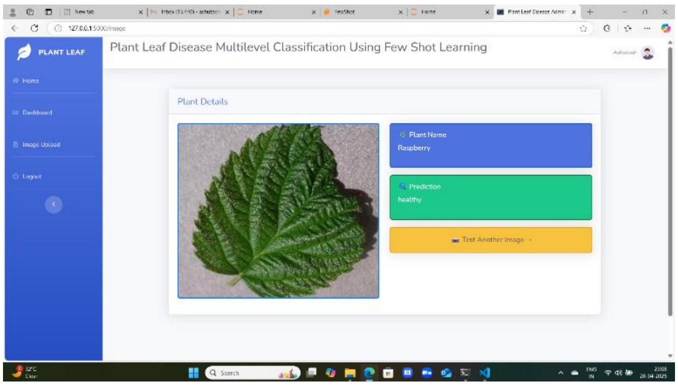
Plant Leaf Disease Multilevel Classification Using Few Shot Learning

Plant Name
Raspberry

Prediction
healthy

Fig B: Prediction Result

6. CONCLUSION AND FUTURE WORK

We demonstrated that a hierarchical prototypical network with EfficientNet-B0 reliably classifies plant leaf diseases using only 10 samples per class for both training and validation, achieving a validation accuracy of 95.73%. This approach underscores the model's capability to learn effectively from limited data, making it suitable for real-world agricultural scenarios where large labelled datasets are often unavailable.

Future work will focus on two key directions:

Zero-Shot Learning: Integrate zero-shot learning techniques to enable the model to recognize novel disease classes without any labelled examples, leveraging semantic embeddings or attribute-based descriptions.

Expanded Taxonomy: Incorporate a broader range of plant species and additional disease types to enhance the model’s generalizability and coverage across diverse agricultural contexts.

7. REFERENCES

1. Bhargava, A., Shukla, A., Goswami, O., Alsharif, M. H., Uthansakul, P., & Uthansakul, M. (2024). *Plant Leaf Disease Detection, Classification and Diagnosis using Computer Vision and Artificial Intelligence: A Review*. IEEE Access.
2. Li, L., Zhang, S., & Wang, B. (2021). *Plant disease detection and classification by deep learning—a review*. IEEE Access, 9, 56683-56698.
3. Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). *Using deep learning for image-based plant disease detection*. Frontiers in Plant Science, 7, 1419.
4. Barbedo, J. G. A. (2016). *A review on the use of machine learning in plant disease detection*. Computers and Electronics in Agriculture, 144, 273-283.
5. Ferentinos, K. P. (2018). *Deep learning models for plant disease detection and diagnosis*. Computers and Electronics in Agriculture, 145, 311-318.
6. Sharma, P., Hans, P., & Gupta, S. C. (2020). *Classification of plant leaf diseases using machine learning and image preprocessing techniques*. IEEE Confluence.
7. Kumar, S., Prasad, K. M. V. V., Srilekha, A., et al. (2020). *Leaf disease detection and classification based on machine learning*. ICSTCEE.
8. Snell, J., Swersky, K., & Zemel, R. (2017). *Prototypical networks for few-shot learning*. NeurIPS.
9. Koch, G., Zemel, R., & Salakhutdinov, R. (2015). *Siamese neural networks for one-shot image recognition*. ICML Deep Learning Workshop.
10. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. (2016). *Matching networks for one shot learning*. NeurIPS.
11. Finn, C., Abbeel, P., & Levine, S. (2017). *Model-agnostic meta-learning for fast adaptation of deep networks*. ICML.
12. Balafas, V., Karantoumanis, E., Louta, M., & Ploskas, N. (2023). *Machine learning and deep learning for plant disease classification and detection*. IEEE Access.
13. Singh, R. S. R., & Sanodiya, R. K. (2023). *Zero-Shot Transfer Learning Framework for Plant Leaf Disease Classification*. IEEE Access.
14. Nitu, V. P., & Srinivasan, R. (2023). *Few-shot plant disease recognition with prototype alignment*. Computers and Electronics in Agriculture, 206, 107656.
15. Jaiswal, A., Gianey, H., & Upadhyay, A. (2022). *Few-shot learning-based plant disease detection: A survey*. Artificial Intelligence in Agriculture.