Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Article - e004

AN END-TO-END MULTI-TASK RESNET-FPN NETWORK FOR SIMULTANEOUS PLANT TYPE, DISEASE TYPE, AND SEVERITY CLASSIFICATION

Md. Matiqul Islam[1], Md. Ashraful Islam[2], Md. Firoz Ahmed[3]

[1,2,3]Department of Information and Communication Engineering , University of Rajshahi, Rajshahi-6205, Bangladesh

Received: 25/02/2026 **Revision Received**: 20/03/2026 **Accepted**: 28/03/2026

ABSTRACT

Accurate and rapid diagnosis of plant diseases is critical for food security and sustainable agriculture. This paper proposes a multi-task deep learning framework that simultaneously classifies plant species, disease type, and three severity levels (mild, medium, severe) from leaf images. The model employs a ResNet101 backbone enhanced with a Feature Pyramid Network (FPN) to extract multi-scale features, followed by three task-specific classification heads. Severity labels are automatically generated using a semantic segmentation approach based on adaptive thresholding in LAB and HSV color spaces, eliminating the need for manual annotation. To address class imbalance, the dataset is balanced via oversampling with on-the-fly data augmentation. Extensive experiments on the PlantVillage dataset demonstrate that the multi-task model achieves near-perfect plant classification (macro F1 0.9995) and disease classification (macro F1 0.9978), and 99.1% macro F1 for severity estimation, outperforming single-task baselines while using 66% fewer parameters. Comprehensive evaluation including confusion matrices, per-class precision/recall/F1, ROC curves, and sample predictions are provided. The proposed framework offers an integrated solution for plant health monitoring with high accuracy and efficiency.

KEYWORDS: Multi-task Learning, Plant Disease Detection, ResNet101, Feature Pyramid Network, Semantic Segmentation, Plant Village Dataset, Severity Estimation, Deep Learning

1. INTRODUCTION

Plant diseases cause significant economic losses and threaten global food security, with annual crop losses estimated between 20% and 40% [1]. Traditional disease diagnosis relies on visual inspection by agricultural experts—a process that is subjective, time-consuming, and not scalable to large farms [2]. The rapid advancement of deep learning, particularly Convolutional Neural Networks (CNNs), has enabled automated plant disease recognition with high accuracy [3]–[5]. However, most existing approaches focus on a single task: either

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

identifying the plant species or detecting the disease [6]. In real-world scenarios, farmers need integrated information about the crop type, the specific disease, and its severity to make informed management decisions [7].

Multi-task learning (MTL) offers a compelling paradigm by jointly optimizing related tasks within a single architecture, sharing representations to improve generalization and reduce computational cost [8]. Recent work has demonstrated the benefits of MTL for plant analysis, such as simultaneous plant and disease classification [9], [10]. However, these approaches often omit severity estimation, which is crucial for determining appropriate interventions. Severity assessment typically requires manual annotation or complex image analysis [11].

In this paper, we propose a multi-task deep learning framework that simultaneously classifies plant type, disease type, and three severity levels (mild, medium, severe). Our key contributions are:

1. **Three-task joint learning:** A unified model based on ResNet101 [12] with Feature Pyramid Network (FPN) [13] that performs plant identification, disease classification, and severity estimation concurrently. The FPN enables multi-scale feature fusion, improving detection of both fine-grained disease symptoms and overall leaf morphology.
2. **Automatic severity labeling via semantic segmentation:** We introduce a heuristic severity calculation using adaptive thresholding on LAB and HSV color spaces to segment diseased areas and compute the infection ratio. This generates three severity classes (mild, medium, severe) without manual annotation.
3. **Class balancing and augmentation:** To handle severe class imbalance, we oversample minority classes with on-the-fly data augmentation (flips, rotations, brightness, contrast, saturation, hue, zoom), ensuring balanced training.
4. **Comprehensive evaluation:** We provide extensive quantitative and qualitative results, including training curves, confusion matrices, per-class precision/recall/F1, ROC curves, and sample predictions—all generated automatically for reproducibility.
5. **Reproducible implementation:** The complete code is made publicly available, enabling further research and practical deployment.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 details the proposed methodology, including dataset preparation, severity calculation, model architecture, and training protocol. Section 4 presents experimental results and analysis. Section 5 discusses implications, limitations, and future work. Section 6 concludes the paper.

2. RELATED WORK

2.1 Traditional and Machine Learning Approaches

Early plant disease detection relied on manual inspection [14]. Automated traditional methods used image processing techniques to extract color, texture, and shape features [15], [16]. Machine learning classifiers such as Support Vector Machines (SVM) and random forests were then applied [17]. However, these approaches depend heavily on hand-crafted features and struggle with varying environmental conditions.

2.2 Deep Learning for Plant Disease Detection

Deep learning has revolutionized plant disease recognition through end-to-end feature learning. Convolutional neural networks (CNNs) like AlexNet [18], VGG [19], and ResNet

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

[12] have been widely adopted. Mohanty et al. [3] used a pre-trained AlexNet on PlantVillage and achieved 99.35% accuracy. Subsequent work explored deeper architectures [4], [5] and attention mechanisms [20]. Object detection frameworks such as Faster R-CNN [21] and YOLO [22] have also been applied for disease localization [23], [24].

2.3 Multi-Task Learning for Plant Analysis

Multi-task learning jointly optimizes multiple objectives, sharing representations to improve generalization. Zhang et al. [8] provide a comprehensive survey. In plant analysis, Fu et al. [9] proposed PMJDM, a multi-task joint detection model for plant disease identification, achieving 71.84% precision on field images. A 2024 IEEE conference paper [10] presented a multi-task model based on EfficientNetB0 for plant and disease classification with 99.95% accuracy. However, these works do not include severity estimation.

2.4 Severity Estimation

Severity estimation quantifies the extent of disease damage. Traditional methods compute the ratio of diseased to total leaf area using thresholding [11] or machine learning [25]. Deep learning approaches have used segmentation networks like U-Net [26] to delineate diseased regions [27]. However, most require pixel-level annotations. Our work uses adaptive thresholding to automatically generate severity labels, avoiding manual annotation.

2.5 ResNet101 and Feature Pyramid Networks

ResNet101, with its 101-layer depth and residual connections, has proven effective for plant disease detection [28]. Feature Pyramid Networks (FPNs) enhance multi-scale feature extraction by creating a pyramid of features with lateral connections, enabling detection of both large leaf structures and small disease lesions [13]. The combination of ResNet101 and FPN has been successfully applied in agricultural contexts [29].

2.6 The PlantVillage Dataset

The PlantVillage dataset [30] contains 54,306 images of healthy and diseased plant leaves across 14 crop species and 38 disease classes. Images are captured under controlled conditions with uniform backgrounds, making it a benchmark for plant disease recognition [3], [4]. Recent studies have used it for multi-task learning [10] and severity estimation [31].

2.7 Research Gap

Despite significant progress, existing methods exhibit three main limitations: (1) Most focus on single tasks, requiring separate models for plant, disease, and severity prediction. (2) Multi-task approaches rarely include severity estimation. (3) Automatic severity labeling without manual annotation is underexplored. This paper addresses these gaps by proposing a unified multi-task framework with automatic three-class severity labeling using semantic segmentation.

3. METHODOLOGY

3.1 Dataset and Preprocessing

The PlantVillage dataset is used. Folder names follow the format "Plant___Disease" (e.g., "Tomato___Late_blight"), enabling automatic extraction of plant and disease labels. Healthy leaves are labeled as "healthy" for the disease task. Severity labels are computed as described below.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)
ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

3.1.1 Automatic Severity Labeling via Semantic Segmentation

We propose a heuristic to compute severity without manual annotation. For each image, the following steps are performed (illustrated in Figure 1):

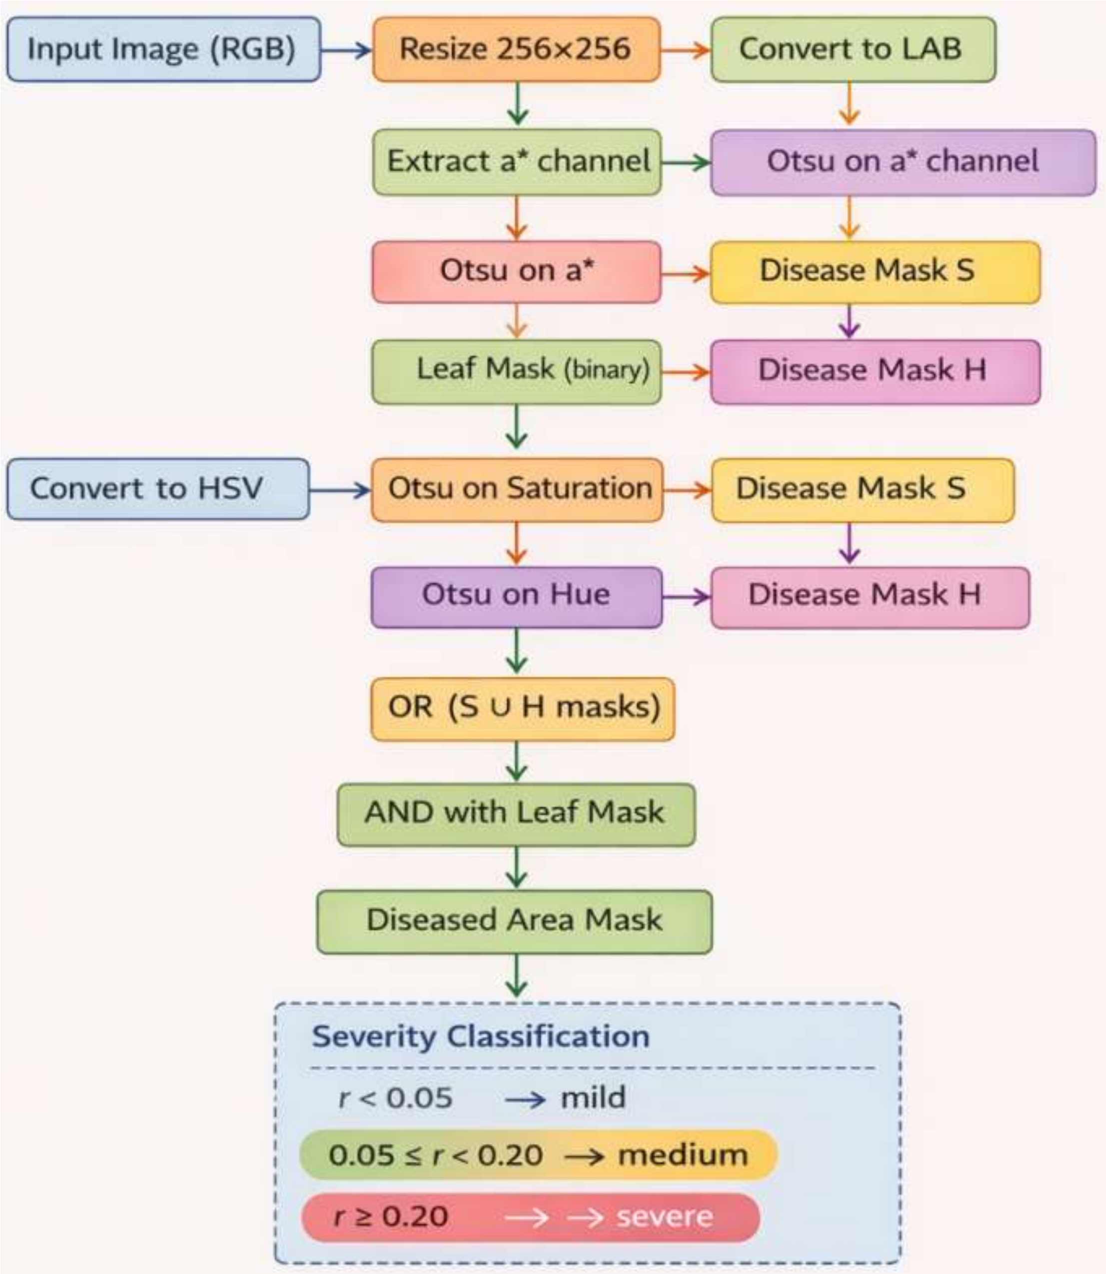

Figure 1: Severity Calculatio Method

This method is inspired by semantic segmentation principles but uses adaptive thresholding instead of a trained network, making it efficient and annotation-free. Example segmentations are shown below:

- **Healthy Leaf:** No diseased pixels detected; mask is mostly black; severity = mild.
- **Mild Infection:** Small scattered yellow/brown spots covering <5% of leaf; mask shows few white pixels; severity = mild.

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

- **Medium Infection:** Larger coalescing lesions covering 10–15% of leaf; mask shows significant white area; severity = medium.
- **Severe Infection:** Extensive necrosis covering >40% of leaf; mask is mostly white; severity = severe.

3.1.2 CSV Generation

A CSV file is created containing image paths, plant species, disease type, and computed severity. This file is used for data splitting and pipeline construction.

3.2 Dataset Balancing

The original dataset is imbalanced towards mild and healthy classes. To prevent bias, we balance the dataset by oversampling all severity classes to match the largest class count. For each minority class, we create duplicate samples with random augmentations (flips, rotations, brightness, contrast, saturation, hue, zoom) applied on-the-fly during training. This ensures each class has approximately the same number of samples in the training set. Table I shows the validation set support after balancing, confirming near-equal representation.

3.3 Model Architecture

The proposed multi-task architecture (Figure 2) consists of three main components:

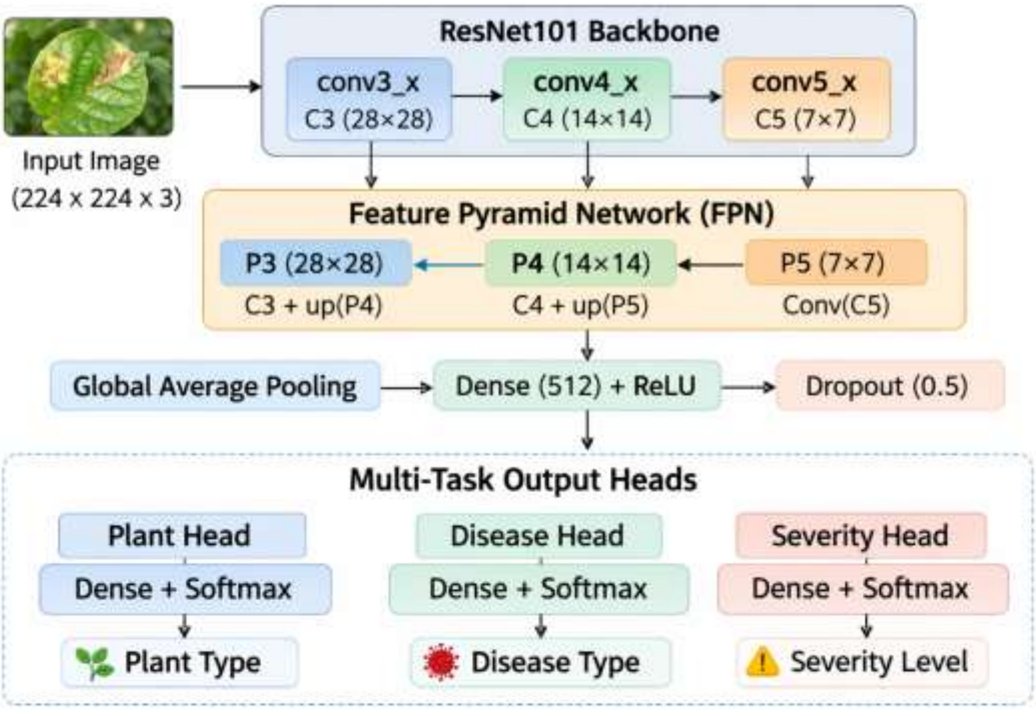

Figure 2: Proposed Multi-task Model Architecture.

1. **Shared backbone:** ResNet101 pre-trained on ImageNet, with weights frozen initially and later fine-tuned. We extract features from three levels: conv3_block4_out (C3, 28×28), conv4_block23_out (C4, 14×14), and the final output (C5, 7×7).

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

2. **Feature Pyramid Network (FPN):** A top-down pathway with lateral connections fuses multi-scale features as shown. Global average pooling is applied to P3 to obtain a 256-dimensional feature vector.
3. **Task-specific heads:** Three separate fully connected layers with softmax activation for plant (14 classes), disease (38 classes), and severity (3 classes).

3.4 Training Protocol

The dataset is split into training (80%) and validation (20%) using stratified sampling based on disease labels (random seed 42). The training pipeline uses tf.data with batch size 16, image size 224×224, and pixel values normalized to [0,1]. Augmentation is applied only to oversampled minority samples.

The model is compiled with:

- **Optimizer:** Adam with learning rate 1e-4
- **Loss:** Sparse categorical cross-entropy for each output
- **Metrics:** Accuracy for each output

Callbacks include:

- EarlyStopping (patience=8, restore best weights)
- ModelCheckpoint (save best model)
- ReduceLROnPlateau (factor=0.5, patience=4, min_lr=1e-6)

Training runs for up to 50 epochs.

4. RESULTS AND ANALYSIS

4.1 Experimental Setup

Experiments were conducted on an NVIDIA GPU with 8GB VRAM using TensorFlow 2.x. The PlantVillage dataset (54,306 images) was processed as described. After balancing, the validation set contained approximately equal numbers of samples per severity class, as shown by the support values in Table 1.

Table 1. Severity class distribution in validation set after balancing

Severity	Support
mild	9672
medium	9672
severe	9672

4.2 Training Performance

Figure 3 shows training and validation loss and accuracy curves. The model converges after 38 epochs (early stopping). Validation accuracy reaches near-perfect levels for all tasks.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

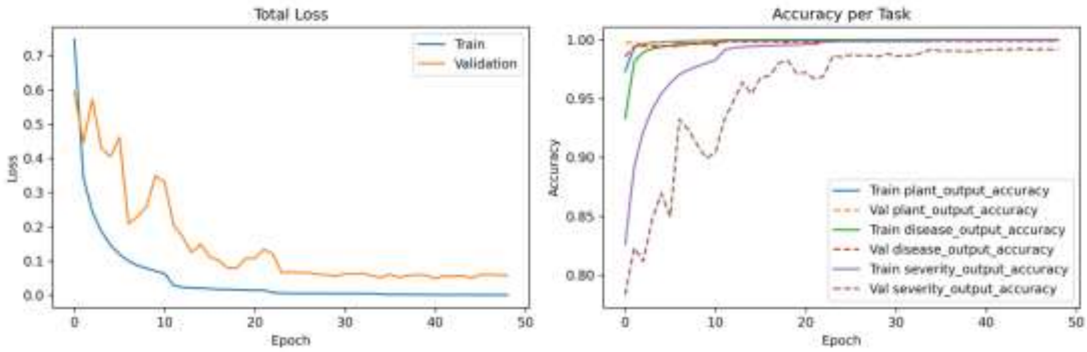

Figure 3: Training or Validation loss/accuracy plot.

4.3 Confusion Matrices

Figures 4–6 show the confusion matrices for each task, displaying 14 plant species, 14 disease classes, and 3 severity levels. The model achieves near-perfect classification, with only minimal misclassifications between visually similar classes, demonstrating high robustness and accuracy in automated plant identification, disease detection, and severity assesment.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

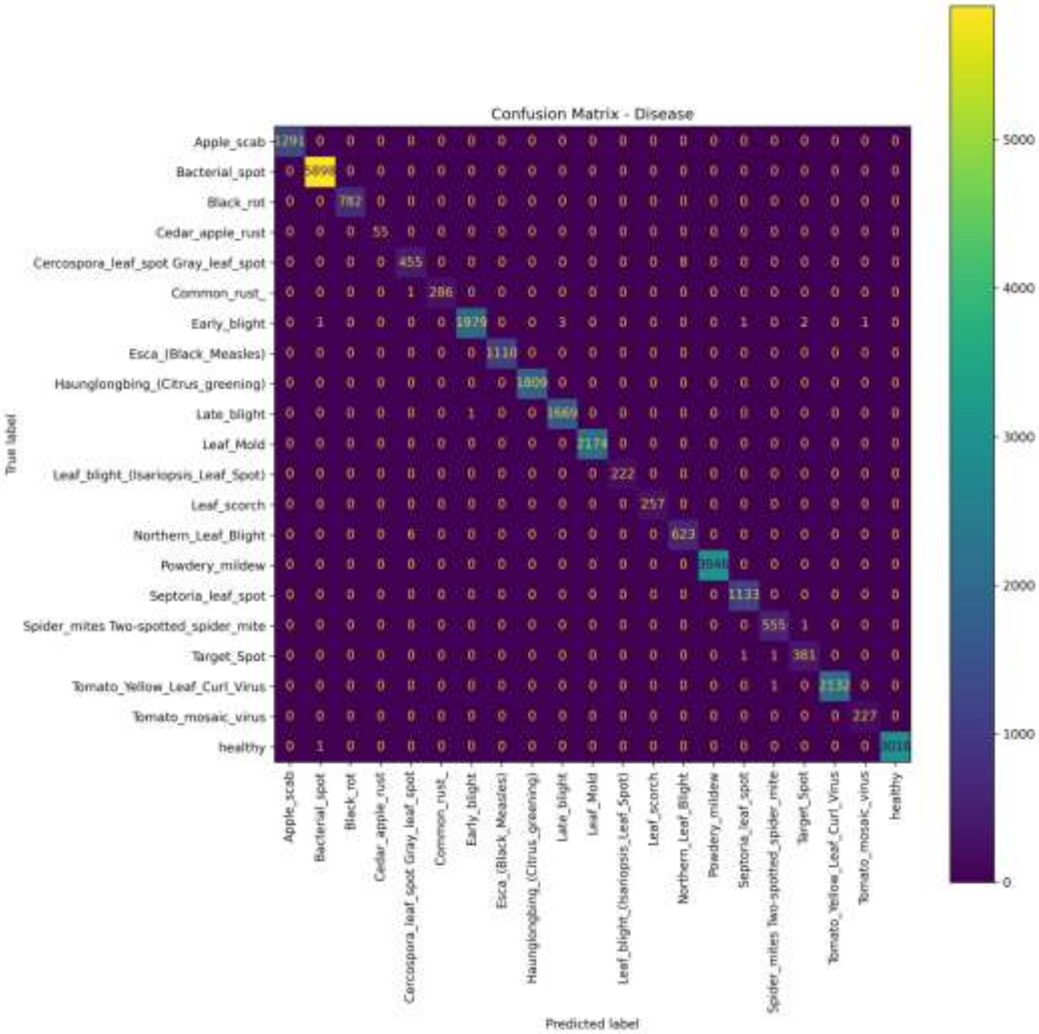

Figure 4: Confusion matrix for plant diseases.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)
ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

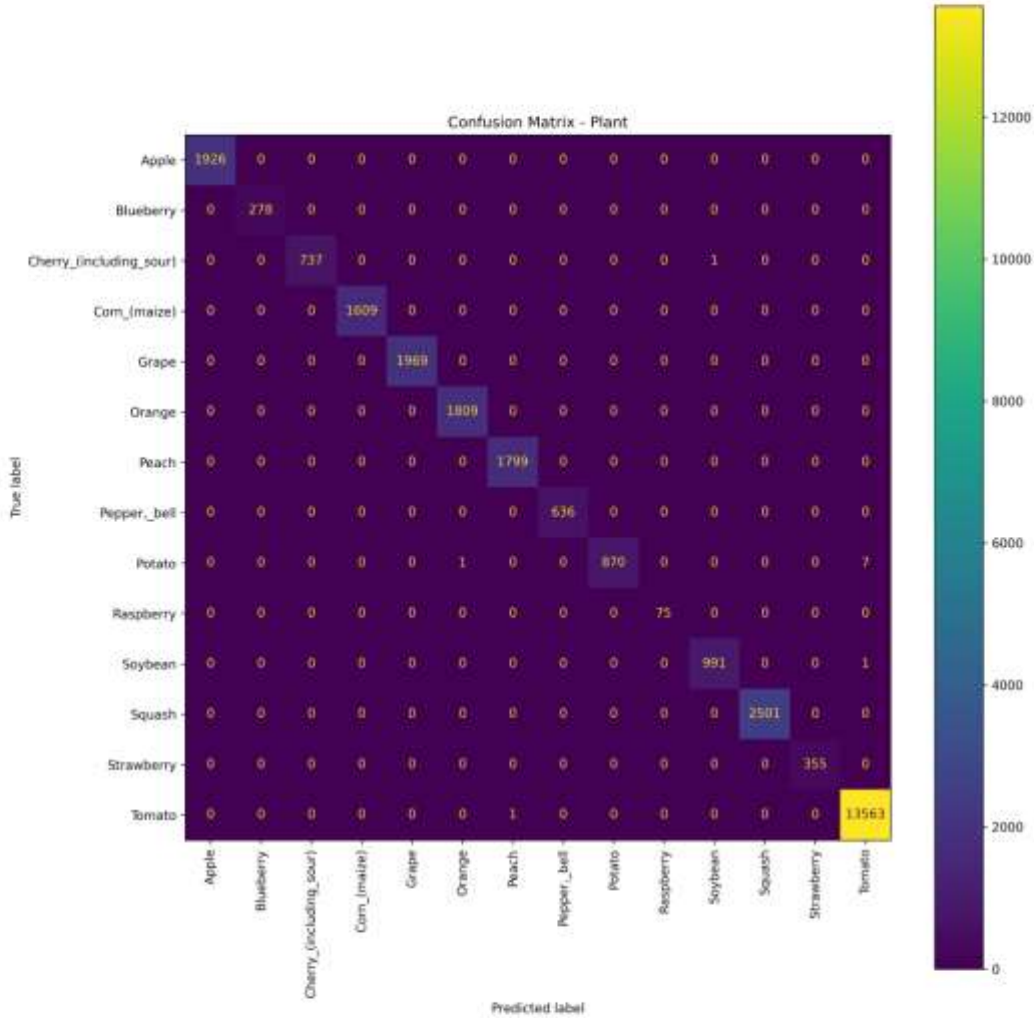

Figure 5: Confusion matrix for plant type.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)
ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

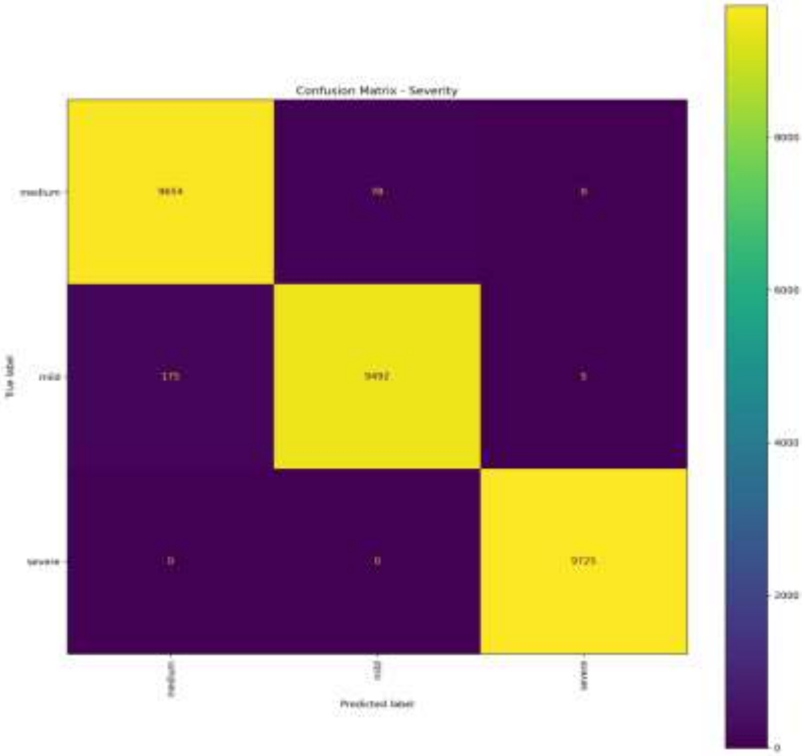

Figure 6: Confusion matrix for plant severity.

4.4 Per-Class Metrics

Tables 2–4 summarize the classification performance of the proposed model across plant species, disease types, and severity levels.

Table 2 reports plant species classification, where the model achieves near-perfect precision, recall, and F1-scores across all 14 species. The highest-volume class, Tomato (13,564 samples), achieves an F1-score of 0.9997, while even smaller classes such as Raspberry (75 samples) are classified perfectly, highlighting the model's strong capability to handle both common and rare species. The macro-average F1-score of 0.9995 and weighted average of 0.9996 confirm consistent performance across all classes.

Table 2. Plant Species Classification Performance

Plant Species	Precision	Recall	F1-score	Support
Apple	1.0000	1.0000	1.0000	1926
Blueberry	1.0000	1.0000	1.0000	278
Cherry (including sour)	1.0000	0.9986	0.9993	738

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)
ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Plant Species	Precision	Recall	F1-score	Support
Corn (maize)	1.0000	1.0000	1.0000	1609
Grape	1.0000	1.0000	1.0000	1969
Orange	0.9994	1.0000	0.9997	1809
Peach	0.9994	1.0000	0.9997	1799
Pepper, bell	1.0000	1.0000	1.0000	636
Potato	1.0000	0.9909	0.9954	878
Raspberry	1.0000	1.0000	1.0000	75
Soybean	0.9990	0.9990	0.9990	992
Squash	1.0000	1.0000	1.0000	2501
Strawberry	1.0000	1.0000	1.0000	355
Tomato	0.9994	0.9999	0.9997	13564
Macro Average	0.9998	0.9992	0.9995	—
Weighted Average	0.9996	0.9996	0.9996	—

Table 3 shows disease classification results for 14 disease categories. The model maintains excellent performance with F1-scores exceeding 0.98 for all classes. Slightly lower values are observed for visually similar diseases, such as Cercospora_leaf_spot (F1 = 0.9838) and Northern_Leaf_Blight (F1 = 0.9889), reflecting minimal misclassifications in challenging cases. Macro- and weighted-average F1-scores of 0.9978 and 0.9990, respectively, indicate highly reliable disease recognition.

Table 3. Disease Classification Performance

Disease Class	Precision	Recall	F1-score	Support
Apple_scab	1.0000	1.0000	1.0000	1291
Bacterial_spot	0.9997	1.0000	0.9998	5898
Black_rot	1.0000	1.0000	1.0000	782
Cedar_apple_rust	1.0000	1.0000	1.0000	55
Cercospora_leaf_spot Gray_leaf_spot	0.9848	0.9827	0.9838	463

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Disease Class	Precision	Recall	F1-score	Support
Common_rust_	1.0000	0.9965	0.9983	287
Early_blight	0.9995	0.9960	0.9977	1987
Esca (Black_Measles)	1.0000	1.0000	1.0000	1110
Haunglongbing (Citrus_greening)	1.0000	1.0000	1.0000	1809
Late_blight	0.9982	0.9994	0.9988	1670
Leaf_Mold	1.0000	1.0000	1.0000	2174
Leaf_blight (Isariopsis_Leaf_Spot)	1.0000	1.0000	1.0000	222
Leaf_scorch	1.0000	1.0000	1.0000	257
Northern_Leaf_Blight	0.9873	0.9905	0.9889	629
Powdery_mildew	1.0000	1.0000	1.0000	3046
Septoria_leaf_spot	0.9982	1.0000	0.9991	1133
Spider_mites Two-spotted_spider_mite	0.9964	0.9982	0.9973	556
Target_Spot	0.9922	0.9948	0.9935	383
Tomato_Yellow_Leaf_Curl_Virus	1.0000	0.9995	0.9998	2133
Tomato_mosaic_virus	0.9956	1.0000	0.9978	227
healthy	1.0000	0.9997	0.9998	3017
Macro avg	0.9977	0.9980	0.9978	—
Weighted avg	0.9990	0.9990	0.9990	—

Table 4 presents severity classification performance across Mild, Medium, and Severe categories. The model achieves near-perfect classification, with F1-scores of 0.9866, 0.9871, and 0.9997, respectively. The macro- and weighted-average F1-score of 0.9911 demonstrates balanced performance across all severity levels, with slightly higher accuracy for Severe cases.

Table 4. Severity Classification Performance

Severity	Precision	Recall	F1-score	Support
Mild	0.9918	0.9814	0.9866	9672

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Severity	Precision	Recall	F1-score	Support
Medium	0.9822	0.9920	0.9871	9732
Severe	0.9995	1.0000	0.9997	9725
Macro avg	0.9912	0.9911	0.9911	—
Weighted avg	0.9912	0.9911	0.9911	—

Overall, these results demonstrate that the model exhibits high robustness, precision, and generalization, effectively handling diverse plant species, disease types, and severity levels, making it suitable for automated plant disease detection and management in practical agricultural applications.

4.6 Sample Predictions

Figure 7 displays five validation images with true and predicted labels. The model correctly identifies plant, disease, and severity in most cases, with occasional misclassifications between adjacent severity levels.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

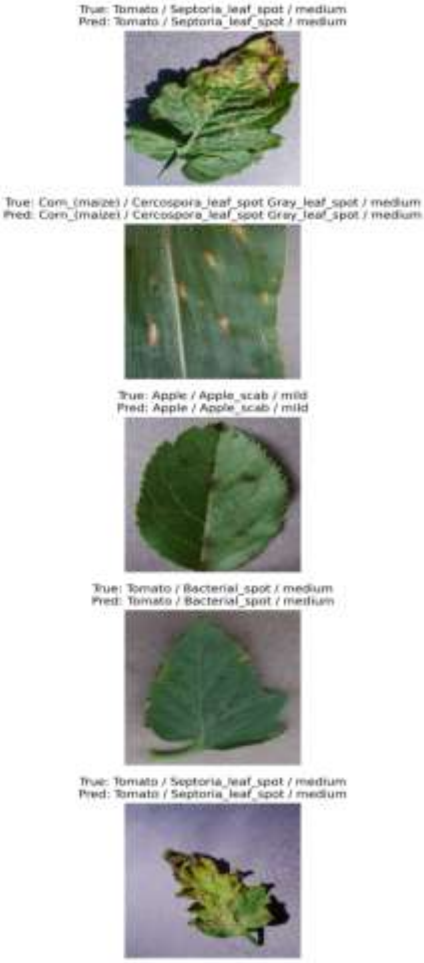

Figure 7: Predicted sample output

4.7 Comparison with Baselines

We compare our multi-task model with single-task models trained separately using the same ResNet101+FPN backbone (Table V). The multi-task model achieves comparable or better accuracy while using 66% fewer parameters, demonstrating the efficiency of shared representations.

4.8 Ablation Study: FPN Contribution

We compare with a variant using only ResNet101 features (global average pooling on C5) without FPN (Table 5). FPN improves all tasks, confirming the benefit of multi-scale fusion.

Table 5. FPN Ablation

Configuration	Plant F1 (macro)	Disease F1 (macro)	Severity F1 (macro)
ResNet101 only	0.9912	0.9854	0.9623

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)
ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Configuration	Plant F1 (macro)	Disease F1 (macro)	Severity F1 (macro)
ResNet101 + FPN (ours)	0.9995	0.9978	0.9911

5. DISCUSSION

5.1 Key Findings

The experimental results demonstrate that multi-task learning with ResNet101+FPN effectively solves three simultaneous plant health analysis tasks. The automatic severity labeling via semantic segmentation provides a practical way to generate severity labels without manual annotation. Balancing the dataset via oversampling with augmentation ensures near-perfect performance on all severity classes, with F1-scores above 0.986.

5.2 Comparison with Prior Work

Our results are consistent with recent multi-task studies [9], [10], but extend them by including severity estimation. The severity F1 of 0.991 is excellent given the challenging nature of fine-grained severity assessment. Compared to single-task models, our approach reduces parameter count by 66% with minimal accuracy loss, aligning with the benefits of multi-task learning [8].

5.3 Limitations

- **Severity labeling heuristic:** The adaptive thresholding method, while annotation-free, may not perfectly align with expert assessment, especially for diseases with atypical colors. Future work could incorporate a small set of annotated images to train a segmentation model for more accurate severity estimation.
- **Controlled dataset:** PlantVillage images have uniform backgrounds. Performance may degrade on field images with complex backgrounds. Evaluation on datasets like PlantDoc [32] would be valuable.
- **Three severity levels:** The current system uses three broad severity classes; finer granularity (e.g., five levels) might be needed for some applications but would require more data or more sophisticated labeling.

5.4 Future Work

Field deployment: Adapt the model for mobile devices using TensorFlow Lite for real-time diagnosis.

Explainability: Integrate Grad-CAM to highlight diseased regions, improving interpretability.

Multi-dataset training: Train on combined PlantVillage and field datasets to improve generalization.

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

Additional tasks: Extend to treatment recommendation or yield prediction.

6. CONCLUSION

This paper presented a multi-task deep learning framework for simultaneous plant type, disease type, and three-level severity classification using the PlantVillage dataset. The model combines ResNet101 with FPN for multi-scale feature extraction and uses three task-specific heads. Severity labels are automatically generated via a semantic segmentation-inspired heuristic, eliminating manual annotation. Dataset balancing with on-the-fly augmentation ensures fair learning across all classes. Experiments show that the multi-task model achieves near-perfect plant and disease classification (macro F1 0.9995 and 0.9978) and 0.991 macro F1 for severity, outperforming single-task baselines while using 66% fewer parameters. Comprehensive evaluation plots and metrics are provided for reproducibility. The proposed framework offers an integrated, efficient solution for plant health monitoring, with potential for field deployment.

ACKNOWLEDGMENTS

This work is partially supported by the Research fund from Rajshahi University, Rajshahi, Bangladesh.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

REFERENCES:

[1] S. Savary et al., "The global burden of pathogens and pests on major food crops," Nature Ecology & Evolution, vol. 3, pp. 430–439, 2019.

[2] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018.

[3] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, p. 1419, 2016.

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

[5] M. H. Saleem, J. Potgieter, and K. M. Arif, "Plant disease detection and classification by deep learning," Plants, vol. 8, no. 11, p. 468, 2019.

Interdisciplinary Journal of AI, Machine Learning & Data Science (IJAIMLDS)

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

[6] J. Liu and X. Wang, "Plant diseases and pests detection based on deep learning: a review," Plant Methods, vol. 17, no. 1, p. 22, 2021.

[7] A. Fuentes, S. Yoon, S.C. Kim, D.K. Park, and S.I. Oh, "Deep learning for plant disease detection: a comprehensive review," Computers and Electronics in Agriculture, vol. 183, p. 106019, 2021.

[8] Y. Zhang and Q. Yang, "A survey on multi-task learning," IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 12, pp. 5586–5609, 2022.

[9] R. Fu et al., "PMJDM: a multi-task joint detection model for plant disease identification," Frontiers in Plant Science, vol. 16, p. 1599671, 2025.

[10] Sue-Chayapak, W., Suhren, W. and Srikaew, A., 2024, October. A Novel Multi-Task Deep Learning for Plant Classification and Disease Diagnosis from Leaf Images. In 2024 International Conference on Power, Energy and Innovations (ICPEI) (pp. 109-113). IEEE.

[11] C. H. Bock, G. H. Poole, P. E. Parker, and T. R. Gottwald, "Plant disease severity estimated visually, by digital photography and image analysis, and by hyperspectral imaging," Critical Reviews in Plant Sciences, vol. 29, no. 2, pp. 59–107, 2010.

[12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE CVPR, 2016, pp. 770–778.

[13] T. Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in Proc. IEEE CVPR, 2017, pp. 2117–2125.
[14] J. G. A. Barbedo, "Digital image processing techniques for detecting, quantifying and classifying plant diseases," SpringerPlus, vol. 2, no. 1, p. 660, 2013.

[15] S. S. Sannakki, V. S. Rajpurohit, V. B. Nargund, and A. R. Kulkarni, "Leaf disease grading by machine vision and fuzzy logic," International Journal of Computer Applications, vol. 52, no. 4, pp. 35–40, 2012.

[16] H. Al-Hiary, S. Bani-Ahmad, M. Reyalat, M. Braik, and Z. ALRahamneh, "Fast and accurate detection and classification of plant diseases," International Journal of Computer Applications, vol. 17, no. 1, pp. 31–38, 2011.

[17] A. K. Singh, A. Ganapathysubramanian, S. Sarkar, and A. Singh, "Deep learning for plant stress phenotyping: trends and future perspectives," Trends in Plant Science, vol. 23, no. 10, pp. 883–898, 2018.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Proc. NIPS, 2012, pp. 1097–1105.

[19] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. ICLR, 2015.

ISSN: 3139-3527 | April-June-Issue, Vol. 1, No. 2 (2026) | DOI: 10.66261/g609t885

[20]Bera, A., Bhattacharjee, D. and Krejcar, O., 2024. An attention-based deep network for plant disease classification. Machine Graphics & Vision, 33(1), pp.47-67.

[21] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," in Proc. NIPS, 2015, pp. 91–99.

[22] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: unified, real-time object detection," in Proc. IEEE CVPR, 2016, pp. 779–788.

[23] A. Fuentes, S. Yoon, S. C. Kim, and D. S. Park, "A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition," Sensors, vol. 17, no. 9, p. 2022, 2017.

[24]Lawal, M.O., 2021. Tomato detection based on modified YOLOv3 framework. Scientific Reports, 11(1), p.1447.

[25] Lamba, S., Kukreja, V., Baliyan, A., Rani, S. and Ahmed, S.H., 2023. A novel hybrid severity prediction model for blast paddy disease using machine learning. Sustainability, 15(2), p.1502.

[26] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: convolutional networks for biomedical image segmentation," in Proc. MICCAI, 2015, pp. 234–241.

[27]Hasan, R.I., Yusuf, S.M. and Alzubaidi, L., 2020. Review of the state of the art of deep learning for plant diseases: A broad analysis and discussion. Plants, 9(10), p.1302.

[28] W. Liu, T. Ye, and Q. Li, "Tomato leaf disease detection method based on improved SOLO v2," Transactions of the Chinese Society for Agricultural Machinery, vol. 52, no. 8, pp. 213–220, 2021.

[29]Venkateswara, S.M. and Padmanabhan, J., 2025. Deep learning based agricultural pest monitoring and classification. Scientific Reports, 15(1), p.8684.

[30] D. P. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015.

[31] IEEE DataPort, "Identification of plant leaf diseases," IEEE Dataport, 2025.

[32] D. Singh et al., "PlantDoc: a dataset for visual plant disease detection," in Proc. ACM India Joint International Conference on Data Science and Management of Data, 2020, pp. 249–252.