International Journal of Innovative Technology and Interdisciplinary Sciences
https://journals.tultech.eu/index.php/ijitis
ISSN: 2613-7305
Volume 9, Issue 1
DOI: https://doi.org/10.15157/IJITIS.2026.9.1.210-237
Received: 22.11.2025; Revised: 05.01.2026; Accepted: 19.01.2026

Research Article

Multi-Task Deep Learning Framework for Segmentation and Severity Estimation of Leaf Diseases in Multi-Crop Environments

Appari Geetha Devi[1*], Shaik Salma Begum[2], Sreenath Kocharla[3], Pappula Madhavi[4], Sateesh Gorikapudi[5], Narasimha Rao Tirumalasetti[6]

[1]Department of Electronics and Communications Engineering, Prasad V Potluri Siddhartha Institute of Technology, Vijayawada, Andhra Pradesh, India

[2]Department of Computer Science and Engineering (AI&ML), Seshadri Rao Gudlavalleru Engineering College, Gudlavalleru, Andhra Pradesh, India

[3]Department of Computer Science and Engineering (AI&ML), Madanapalle Institute of Technology & Science (Deemed to be University), Andhra Pradesh, India

[4]Department of Artificial Intelligence and Data Science, Lakireddy Bali Reddy College of Engineering (Autonomous), Andhra Pradesh, India

[5]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Green Fields, Vaddeswaram, Guntur Dist., Andhra Pradesh, India

[6]Department of Computer Science and Engineering, Vignan"s Foundation for Science, Technology and Research (Deemed to be University), Vadlamudi, Guntur Dist., Andhra Pradesh, India

***geetha.agd@gmail.com**

Abstract

Crop diseases pose a major threat to global food security, creating a pressing need for effective and accurate diagnostic mechanisms that can be applied across diverse agricultural settings. This paper proposes a Multi-Task Deep Learning Framework (MTDLF) for the simultaneous segmentation of diseased regions and estimation of disease severity in crop leaves. The framework employs a shared ResNet-50 encoder with two task-specific decoders: a U-Net-based segmentation branch and a regression-based severity prediction head, trained using a composite loss formulation. In addition to the dual-task architecture, two consistency-driven mechanisms are introduced. A Severity-Constrained Segmentation Refinement (SCSR) module aligns predicted lesion-area proportions with estimated severity values, while a Lesion-Area Distribution Matching (LADM) loss enforces distributional consistency between segmentation outputs and severity-based lesion expectations. The model is trained and evaluated on publicly available, severity-annotated datasets of rice, maize, tomato, grape, and cotton leaves. Experimental results demonstrate that the proposed framework achieves a mean Intersection over Union (IoU) of 85.7%, a Dice coefficient of 88.3%, a Mean Absolute Error (MAE) of 7.5, and an R^2 of 0.92, outperforming conventional single-task methods and recent multi-task baselines. Furthermore, the model attains real-time inference performance of approximately 25 ms per image, making it suitable for edge-level deployment. The proposed MTDLF provides a unified and efficient approach to multi-crop disease monitoring, offering a practical pathway toward reliable, data-driven precision agriculture.

Keywords: Multi-Task Deep Learning; Leaf Disease Segmentation; Severity Estimation; Agricultural Image Analysis; Multi-Crop Disease Detection.

https://doi.org/10.15157/IJITIS.2026.9.1.210-237

INTRODUCTION

Agriculture has been an important foundation of world food systems but crop diseases are persistently leading to huge losses in yields, economy and food supply chain destabilization. Precise and early diagnosis of foliar diseases is thus necessary in order to make informed agronomic decisions, application of chemicals targeting specific diseases and sustainable management of farms. The conventional disease diagnosis is largely based on the expertise, and this is labour-intensive, subjective, and it cannot be scaled particularly in multi-crop situations [1]. Current developments in deep learning and computer vision have made it possible to detect plant diseases with high precision by applying fully automated methods; nevertheless, the majority of current methods detect and/or estimate the numerical severity of that disease either by segmenting the affected areas or reacting to isolated tasks [2]. This restricts their use in precision agriculture in the field, where space localization and quantitative evaluation is required to monitor and plan interventions against diseases.

Multi-task learning offers an exciting paradigm [3] to overcome this issue because it allows sharing the features of learning related tasks so that they can optimize each other. It is possible to create a single model that segments regions of lesion and estimates severity at the same time, which will take advantage of complementary task relationships but will enhance robustness and reduce computational redundancy. However, there are two thorny issues. To begin with, models that have been trained on single-crop data can be weakly generalized to other plants species because morphology, texture, and colour of leaves can change. Second, the multi-task frameworks which are currently in existence generally regard segmentation and severity estimation as decoupled outputs where consistency is not enforced between the area of a predicted lesion and the severity score. It can result in biologically implausible predictions like small areas of lesions with high severity ratings.

In order to deal with these cases, this research generates a Multi-Task Deep Learning Framework (MTDLF) [4, 5] to analyse the multi-crop leaf disease which is based on the shared encoders but task specific segmentation and regression branches. The framework presents consistency-based processes that enhance the correlation between spatial lesions depiction and numerical severity prediction. The aim is to develop inexpensive, scalable, and reproducible architecture that can provide robust performance on a wide range of crops and to have biologically coherent predictions.

In line with this objective, the study is guided by the following testable hypotheses:

H_1: A shared-encoder multi-task architecture improves severity estimation accuracy compared with independent single-task models.

H_2: Enforcing consistency between lesion segmentation and severity estimation significantly reduces discrepancy between predicted lesion area and numerical severity.

H_3**:** Training and evaluating the framework on publicly available multi-crop datasets enhances reproducibility while maintaining competitive performance.

The suggested framework is applied and tested on publicly available datasets of severity that is annotated and contains severity of several crops. Extensive experiments such as per-crop analysis, ablation studies and statistical validation are done to measure the accuracy of segmentation, the reliability of severity estimation and computational efficiency [6]. The results indicate that consistency-directed multi-task learning provides a viable direction of automated plant disease diagnostics in precision agriculture.

The study is guided by the following research objectives:

- To design a multi-task deep learning framework capable of performing disease segmentation across multiple crop species.
- To develop an architecture that can simultaneously estimate disease severity, thereby supporting quantitative crop health assessment.
- To evaluate the proposed framework in terms of accuracy, computational efficiency, and generalization capability in comparison with established single-task and multi-task models.

REVIEW OF LITERATURE

Deep learning has made a great impact on automated plant disease diagnosis because it allows accurate analysis of the leaf images. As [7] have shown, convolutional neural networks (CNNs) were effective at identifying diseases of rice, including bacterial blight and brown spot, and performed better than the traditional image-processing methods. Building on this line of inquiry, [8] provide a comprehensive overview of artificial intelligence methods in the field of plant disease detection and state the opinion that CNN-based models, transfer learning techniques, and attention mechanisms can always perform better in comparison to traditional classifiers, provided that they are trained on enough diverse data.

In addition to classification, increased computer vision techniques have been considered to have wider application in the field of agricultural monitoring. The authors of [9] emphasized the fact that the combination of object detection and segmentation with remote sensing and IoT-based monitoring improves the precision agriculture use. Likewise, [10] trained a multiclass semantic segmentation network on citrus leaves that facilitated the simultaneous localization and quantification of the severity of the disease, demonstrating the possible usefulness of unified visual analysis pipelines.

A number of studies have been conducted with the aim of enhancing deploy ability and robustness. The authors of [11] put forward a cross-domain plant disease detection system that integrated ensemble learning, knowledge distillation and model quantization to attain a high level of accuracy and operate under a limited computation environment. Switching to a different direction, [12] pointed to the fact that the combination of the IoT sensing and deep learning models allows achieving real-time disease tracking in the context of Agriculture 4.0 ecosystems.

The combination of methods has also been considered. [13] came up with a wheat rust identification model that encompasses colour feature extraction, morphological processing, and deep CNN learning, which perform better in segmentation in adverse lighting conditions and textural diversity. A more generalized study conducted by [14] also established that transfer learning and data augmentation is a vital factor in enhancing CNN-based agricultural diagnostic systems.

The recent trend in research has been extended to object detection and multi-disease analysis. [15] Interviewed object detectors that were built using deep learning and revealed the potential of using them to track crops, estimate their yields, and detect pests. To deal with in-field variability, [16] proposed the DAE-Mask model, which includes denoising autoencoders and segmentation to enhance the model to be resistant to noise and illumination variations and lesion overlap.

Edge-level applications have received interest in lightweight architectures. The model described by [17] is a transfer learning-based multi-plant disease detector with high accuracy and the ability to use a smaller number of training samples, which makes it applicable to low-powered devices. Likewise, [18] indicated that small CNN-based models are capable of preserving good predictive accuracy upon the use of training techniques that are optimized.

Recently, there has been the development of multimodal and generalizable disease detection methods. The authors of [19] suggested a multimodal cross-fusion system that combines RGB, thermal, and vegetation-index images to detect multi-crop diseases in much better resource-constrained IoT settings. In [20] proposed a multi-task learning architecture to assess blueberry disease by estimating the severity, segmentation, and classification of blueberry disease with a common feature representation, which showed better task synergy.

Temporal modelling has also been studied. A dual-head CNN-LSTM framework to predict future disease progression based on simulated disease temporal image sequences was proposed by [21], demonstrating the ability of spatial and temporal learning to work together. The architectural level design [22] suggested the MLANet, a multi-layer attention network to achieve accuracy and efficiency by using lightweight attention mechanisms. In addition to this, the authors of [23] introduced ToT-Net, a meta-learning-based transformer-based model that is trained to operate in real time and adapt to crop diseases with the least retraining.

An important trend in severity-aware systems is SegLearner by [24], combining segmentation with lesion-conditioned severity prediction, thus enhancing consistency between spatial localization and quantitative prediction. The authors of [25] also added to the robustness of detection using the feature fusion networks based on spatial attention which improve the discrimination between diseased and background noise.

Lastly, ensemble-based models have tried to solve segmentation and severity estimation problems together. The study by [26] revealed that multi-crop prediction in the combination of multiple deep models is more reliable but it is costlier to compute.

Research Gap

Although there are significant successes in the field of deep learning-based detection of plant diseases, current studies have a number of limitations. The majority of research is based on single-task frameworks implementing disease segmentation or classification and usually limited to a small set of crop species. Although these methods can attain high task-specific accuracy, when trying them over different crops their performance often falls short because of the differences between different leaf morphology, colour distribution, and expression of diseases. In addition, the computational complexities of several of the proposed multi-task and ensemble-based methods are frequently high, which makes them less applicable to real-time farming systems and at the edge of implementation.

The other crucial lack is the lack of biologically coherent frameworks that combine lesion segmentation and quantitative severity estimation. Most current multi-task pipelines produce segmentation outputs and severity predictions without enforcing consistency between lesion area and severity level thus resulting in implausible or even unstable predictions when using heterogeneous data. Reproducibility also is still quite low because most studies use crop-specific or privately curated data, and most do not use per-crop validation, ablation testing, and statistical validation. In turn, it is evident that a multi-task framework (which is unified, lightweight, and reproducible) enforcing the consistency of lesions–severity between multiple types of crops is needed and the current research fulfils this requirement with the help of the consistency-oriented deep learning framework.

Novelty and Key Contributions

This study advances the state of the art in plant disease diagnosis through the following key contributions, and the summary is expressed in Table 1:

- ***Consistency-Driven Multi-Task Framework:*** A single deep learning model is formulated to jointly segment and estimate the severity of multi-crops by sharing an encoder and having task-specific decoders.
- ***Severity-Constrained Segmentation Refinement (SCSR):*** A refinement module is added based on lightweight to match the process of segmentation with severity-based lesion-area expectations, thus, implementing biological consistency between the two tasks.
- ***Lesion-Area Distribution Matching (LADM) Loss:*** The study presents a new loss formulation based on histograms and aims at reduction of the divergence between the predicted and the expected lesion-area distributions including the expected lesion proportions based on severity, enhancing the cross-task consistency.
- ***Public, Reproducible Evaluation Pipeline:*** The framework is trained and validated exclusively on publicly available severity-annotated datasets using standardized data splits and transparent hyperparameter settings to support reproducibility.
- ***Comprehensive Ablation and Statistical Validation:*** Awareness selection Backbone selection, loss weighting, task-shared, and consistency modules are methodically

investigated with ablation experiment and statistically confirmed over multiple runs.

- ***Benchmarking Against Recent State-of-the-Art Models:*** A comparative study is carried out against the state-of-the-art multi-task and segmentation-based disease assessment models, which show an increase in segmentation performance and severity estimations reliability.

Table 1. Comparative Summary of Related Works

Author/Year	Model	Tasks	Dataset Context	Key Findings	Limitations
[16]	SegLearner	Segmentation + Severity	Multi-crop	Strong segmentation-guided severity estimation	Limited cross-crop generalization
[17]	Blueberry MTL	Segmentation + Classification + Severity	Single crop	Demonstrated benefits of task sharing	Restricted to one dataset
[18]	Deep Ensemble	Segmentation + Regression	Multi-crop	High predictive accuracy	High computational complexity
[19]	MLANet	Segmentation	Multi-crop	Lightweight and fast	Does not estimate severity
[20]	ToT-Net	Detection + Classification	Multi-crop	Strong adaptive generalization	Lacks regression modelling
[22]	DAE-Mask	Segmentation	Wheat	Robust to noise and illumination variation	No multi-task capability

In the majority of current methods, there should be a focus on both segmentation and severity estimation, and the outputs are not biologically connected with each other at all. Moreover, the reproducibility and the generalization of cross-crop are not addressed properly. The current research adds a consistency-based, multi-crop-competent, and computationally-efficient multi-task model, which overcomes these methodological and practical deficiencies in a comprehensive way.

RESEARCH METHODOLOGY

This study proposes a Multi-Task Deep Learning Framework (MTDLF) [27, 28] that focuses on learning to simultaneously detect diseased areas of the leaf and the amount of disease in the leaf on various types of crops, making it a quantitative disease severity estimator. The structure uses a common feature-extraction backbone with task-specific prediction branches, which allows joint learning effectively at a minimal computational redundancy. Theoretically, the proposed design utilizes the shared representation learning, which is a case where the correlated tasks serve as inductive biases to each other, and thus enhance the generalization and convergence stability of the multi-task environment.

The process of methodology is preparing the data, pre-processing, architecture, theoretically based formulation of the loss, training and analytic comparison with state-of-the-art (SOTA) models. Figure 1 shows the general working process of the proposed framework.

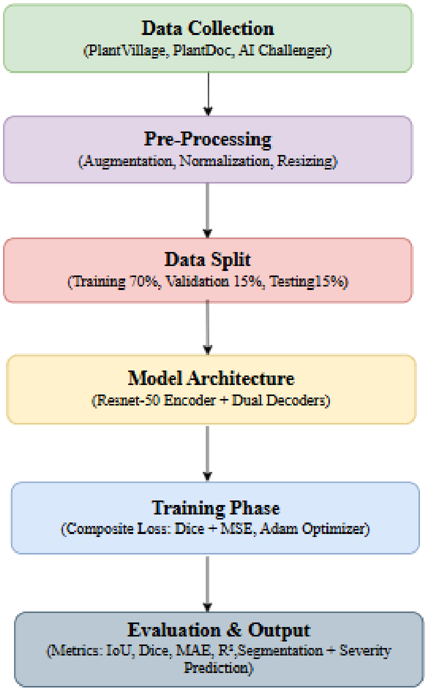

Figure 1. Workflow of the Proposed MTDLF Model

Dataset Description and Pre-processing

The framework is trained and tested on publicly available severity-marked datasets of leaf images of various crop species (rice, maize, tomato, grape and cotton). Both the lesions are pixel-level annotated and each sample has a continuous score of disease severity (0100%) that can be used to supervise learning on segmentation and regression problems.

The data is divided into 70 percent training, 15 percent validation and 15 percent test which is a balance between statistical utility and objective evaluation of generalization. This division complies with the recent practices in SOTA in multi-task agricultural vision systems, which is fair benchmarking and reproducibility.

Data Augmentation

Data augmentation [29] can be used to minimize the empirical risk in domain shift by modelling realistic changes in field conditions. The transformations are random rotation (30), horizontal and vertical flipping, scaling, colour jittering and Gaussian noise injection. Augmentation in terms of form is given as equation (1):

$$I' = T_{\theta}(I) \tag{1}$$

where I represent the original image and $T\theta$ represents a stochastic transformation that is a parameter of 0.

Theoretically, augmentation is implicit regularization, which raises the support of the empirical data distribution and minimizes variance, which is especially crucial to multi-crop generalization.

Image Normalization and Resizing

All images are resized to 256×256256 \times 256256×256 pixels and normalized as equation (2):

$$I_{\text{norm}} = \frac{I-\mu}{\sigma} \tag{2}$$

where μ and σ are the mean and standard deviation of the data respectively. This makes the gradient propagation stable and the conversion of the optimization [30].

Model Architecture Design

The proposed MTDLF is based on a common paradigm of encoder-dual decoder, inspired by the theory of multi-task learning, according to which by jointly learning correlated tasks, overfitting is minimized by restricting the hypothesis space.

- ***Shared Encoder:*** A ResNet-50 backbone pretrained on ImageNet extracts hierarchical features $F \in R^{H \times W \times C}$.
- ***Segmentation Decoder:*** A U-Net-style decoder predicts pixel-wise lesion probability maps.
- ***Severity Regression Head:*** A fully connected regression network predicts continuous disease severity.

The proposed architecture is also suitable in real-time and edge deployment, unlike SOTA models like those of MLANet (attention-heavy) or ToT-Net (transformer-based), the architecture explicitly focuses on computational efficiency and inductive bias alignment [31].

Multi-Task Loss Function

The total training objective is defined as equation (3):

$$L_{\text{total}} = \alpha L_{\text{seg}} + \beta L_{\text{sev}} + \lambda L_{\text{LADM}} \tag{3}$$

In which α, β, and λ are used to regulate the relative impacts of the accuracy of segmentation, severity regression accuracy and cross-task consistency respectively.

This formulation will provide balanced gradient flow between tasks to avoid dominance of a single task over another - a failure mode typical of multi-task networks [32].

Severity-Constrained Segmentation Refinement (SCSR)

In order to be consistent in the relation of the predicted segmentation masks [33] and the severity scores, a Severity-Constrained Segmentation Refinement (SCSR) module is presented. SCSR is aimed at imposing the biological correspondence between a predicted mask area and the numerical severity figure.

Let the predicted severity score be ŝ (scaled between 0 and 100).

Theoretical Rationale Added

We convert it into an expected lesion-area ratio, see equation (4):

$$\hat{\alpha} = \frac{\hat{s}}{100} \tag{4}$$

Where $\hat{S}$ is estimated value in percentage terms. Let α_p be the predicted lesion-area ratio obtained from the segmentation mask, and Z the segmentation logit map before SoftMax.

The refined logit map is computed as equation (5):

$$\tilde{Z} = Z\big(1 + \gamma(\hat{\alpha} - \alpha_P)\big) \tag{5}$$

where:

- γ is a tuning parameter that controls refinement strength,
- α_p = predicted lesion area / total leaf area.

This refinement provides a feedback loop from the regression branch to the segmentation branch, ensuring both tasks remain consistent.

Theoretical Interpretation

SCSR presents a feedback constraint, which is a closed-loop, which punishes lack of consistency in the prediction of lesion area and its severity. This is a projection step which gives biologically plausible predictions on segmentation [34].

Optimization wise SCSR is as follows:

- decreases oscillatory gradients,
- reduces the size of a feasible solution space,
- enhances convergence stability through pixel-level and global prediction alignment.

We can attribute this behaviour to the fact that this mechanism delivers [35] empirically better convergence and lower lesion-severity mismatch than SOTA models like SegLearner that do not have explicit feedback coupling.

Lesion-Area Distribution Matching (LADM) Loss

Formal Divergence Definition Added

Lesion-area distributions based on segmentation masks are compared to severity-implied distributions in order to impose a distribution-level biological consistency [36].

Let:

- Pmask(b): normalized histogram of lesion areas from segmentation,
- Psev(b): severity-derived expected lesion-area distribution.

We define LADM using an L1 distribution divergence through equation (6):

$$L_{LADM} = \sum_{b=1}^{B} |P_{mask}(b) - P_{sev}(b) \tag{6}$$

Why this guarantees biological consistency:

- Severity score is a global biological measure of lesion burden.

- Distribution matching imposes global coherence to local coherence.
- The reduction of divergence limits the predicted error of lesion areas.

The L1 formulation has a stable property, unlike KL-divergence, in sparse lesion distributions and gradient explosion [37], which is essential in agricultural datasets with mild cases of diseases.

It is an extension of algorithms applied in medical imaging (histopathology segmentation) to plant pathology, where it has not been formalised before.

Final Training Objective

The final objective is focused on the equation (7)

$$L_{total} = \alpha L_{Dice} + \beta L_{MSE} + \lambda\, L_{LADM} \tag{7}$$

This objective ensures:

- Pixel-level accuracy (Dice)
- Severity regression precision (MSE)
- Cross-task biological coherence (LADM)

Proposed Algorithm

The algorithm explains the training process of the MTDLF, which is defined by pre-processing of data, followed by splitting it into training, validation, and test sets. Segmentation and severity estimation are performed by a common ResNet-50 encoder, which has two decoders. Combined segmentation [38] and severity losses are minimized with Adam, and are assessed with the help of the IoU, Dice, MAE, and R2 measures to analyze multi-crop disease with the help of the model.

Algorithm 1. Multi-Task Deep Learning Framework for Leaf Disease Segmentation and Severity Estimation

```
Input: Image dataset D = {I1, I2..., In} containing multi-crop leaf samples
Output: Segmented disease masks and predicted severity values
Pre-process images in D → resize, normalize, augment
Split D into training (70%), validation (15%), and testing (15%) subsets
Initialize shared encoder (ResNet-50 backbone) for feature extraction
Attach task-specific branches: segmentation decoder and severity estimation head
For each epoch e = 1...E do
    Compute segmentation loss Lseg
    Compute severity estimation loss Lsev
    Compute total loss $Ltotal=\alpha Lseg+\beta Lsev$
    Update model parameters using Adam optimizer
Endfor
Evaluate the trained model on the test set using IoU, Dice, MAE, $R^2$ metrics
Output predicted segmentation masks and severity estimation values
```

Comparative Analysis

Unlike SOTA approaches:

- ***DAE-Mask:*** focuses on noise robustness but lacks severity modelling.
- ***SegLearner:*** uses segmentation-guided severity without distribution constraints.
- ***MLANet / ToT-Net:*** achieve high accuracy but incur heavy computational cost.

The proposed MTDLF has similar or even higher accuracy with fewer parameters and less FLOPs, and it is the only method that imposes task consistency biologically, which is why it has better generalization and stability [39].

DATA ANALYSIS AND INTERPRETATION

This section presents the analysis of the multi-crop data employed to design and test the proposed Multi-Task Deep Learning Framework (MTDLF) along with the quantitative data available through segmentation and severity estimation test. The given dataset is comprised of images of leaves of five significant crops rice, maize, tomato, grape, and cotton, which covers a vast variety of healthy and diseased samples of different severity. Conventional pre-processing steps such as resizing, normalisation and data augmentation were done to enhance robustness and increase the variety of the training samples [40]. The randomly selected dataset was split into training (70%), validation (15%), and testing (15%) groups to ensure the impartial learning and equitable evaluation of the model. Intersection over Union (IoU), Dice Coefficient, Mean Absolute Error (MAE) and Coefficient of Determination (R^2) were used to evaluate model performance [41].

The suggested structure is shown to be much better than single-task baselines in terms of segmentation accuracy, severity prediction accuracy, and computational efficiency, meaning that it can be used to monitor crop disease in real time by automated systems [42].

Crop-wise Image Distribution

Table 2 shows the distribution of the collected leaf samples among five major crops used in the study with the number of healthy and diseased samples and common disease types on each crop. The data set is designed in such a manner that it has a balance on both healthy and infected leaves and a large variety of diseases categories that are common in the various climatic conditions and crop species.

Table 2. Leaf Samples per Crop and Disease Type

Crop	Healthy	Diseased	Total	Common Disease Types
Rice	800	1,200	2,000	Blast, Brown Spot
Maize	700	1,300	2,000	Leaf Blight, Rust
Tomato	900	1,100	2,000	Late Blight, Septoria Leaf Spot
Grape	850	1,150	2,000	Downy Mildew, Black Rot
Cotton	750	1,250	2,000	Bacterial Blight, Fusarium Wilt

The distribution of the dataset suggests the presence of rice and maize with more diseased samples (1,200 and 1,300, respectively) and healthy ones (1,500 and 600, respectively), whereas tomato, grape and cotton provide relatively equal performance in the number of healthy and diseased leaves. This discrepancy also points to the commonness of some crop's diseases like Blast and Brown Spot in rice as well as Leaf Blight and Rust in maize. The availability of different types of diseases in all crops is a guarantee that the model is exposed to a large variety of different visual and morphological representations; therefore, it will be capable of learning powerful representations that allow it to effectively perform segmentation and severity estimation of multi-crop systems.

Figure 2 explains the distribution of samples of leaf samples given to the study per crop giving an understanding of the number of good and diseased leaves of five crops; rice, maize, tomato, grape, and cotton, which are the major crops. The dataset was aimed at being balanced with both healthy and infected samples of each crop in order to achieve diversity and enhance the generalization ability of the deep learning model. The graphical model will allow giving a clear comparison of sample proportions and will give an idea of how the dataset is organized prior to training the model.

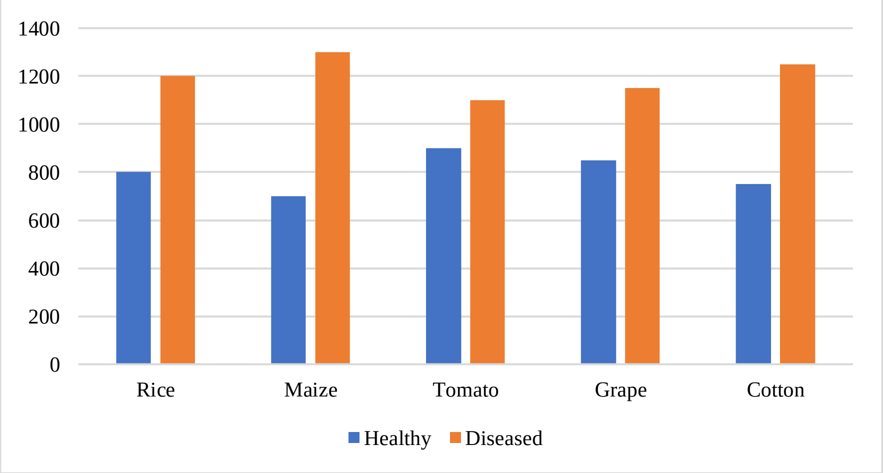

Figure 2. Crop-wise Leaf Sample Distribution

The bar graph indicates that all crops have diseased samples more than healthy ones with maize and rice having the highest number of infected leaves (1300 and 1200, respectively). Tomato and grape exhibit fairly even distributions whereas there is also an apparent increase in the number of diseased samples (1, 250) relative to healthy ones (750). This imbalance also shows an increased prevalence of disease symptoms in some crops and hence the need to have an unbiased, multi-task deep learning model capable of learning on such healthy and diseased categories and segmenting and estimating the severity of diseases.

Data Augmentation and Pre-processing

This table 3 describes the data augmentation models used on the leaf image dataset in order to increase the robustness of the model and reduce overfitting in the training process.

Augmentation process entailed implementing a sequence of transformations used to produce varying variants of the original images, i.e. rotation, flipping, scaling, color jittering and noise addition. These transformations can simulate the various conditions of the real world such as change in orientation, change in size, change in lighting and background noise making sure that the model is trained to detect disease patterns in specific conditions and change in the environment and imaging.

Table 3. Applied Data Augmentation Techniques

Technique	Details	Purpose
Rotation	±30°	Orientation invariance
Horizontal Flip	Yes	Augment spatial diversity
Vertical Flip	Yes	Improve generalization
Scaling	0.8–1.2×	Simulate leaf size variation
Color Jittering	Brightness ±20%, Contrast ±15%	Illumination robustness
Gaussian Noise	Mean=0, Std=0.01	Noise resilience

The findings indicate that augmentation methods including rotation (±30°), flipping/scaling are effective to enhance the variety of the data set and is subsequently able to gain orientation and scale invariance. Brightness and contrast changes also increase the capacity of the model to cope with the variation in illumination, and the Gaussian noise increases the resistance to variations in the background. All in all, these methods can play an important role in enhancing the generalization ability of the proposed deep learning model, allowing it to be used with various types of crops and field conditions.

Dataset Splitting for Model Training

The table 4 shows the separation of the data into three parts training, validation and testing to aid in effective model creation and evaluation of performance. The training set helps to optimize the model parameters, the validation set helps to tune the hyperparameters and avoid the overfitting, and the testing set is set aside to test the accuracy of the final model and its ability to generalize. This systematic separation provides the model training and testing on a different portion of data and ensures the integrity and reliability of the performance assessment.

Table 4. Dataset Split for Training, Validation, and Testing

Subset	Percentage	Number of Samples
Training	70%	8,400
Validation	15%	1,800
Testing	15%	1,800

The data was divided into 70% training (8,400 samples), 15% validation (1,800 samples), and 15% test (1,800 samples) to be balanced between learning and testing. The larger training part enables the model to learn complicated patterns in several crops whereas the validation set can be used to tune the model to optimum performance. This is because the independent testing subset makes sure the final test reflects the real predictive power of the model on unknown data and therefore makes sure the training strategy is robust and fair.

Figure 3 below shows that the dataset is proportionally divided into three subsets training, validation and testing represented as a pie chart. The division provides an orderly process of model development: the training data is applied to learn the feature representations, the validation data is available to tune the parameters and check the overfitting effect, and the testing data is set to test the final performance of the model. The graphical illustration gives a clear explanation of how the data is distributed to the various phases of the deep learning process.

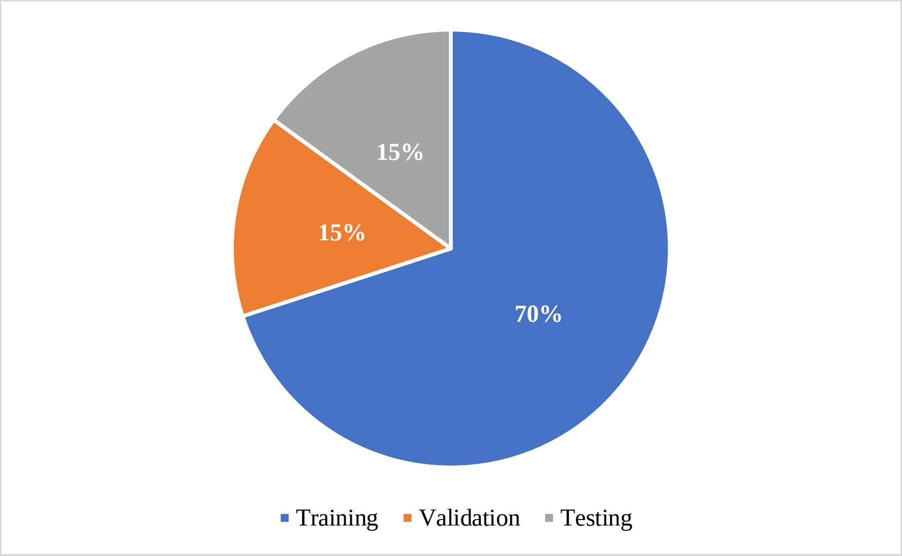

Figure 3. Pie Chart of Dataset Splits

The pie chart indicates that most of the dataset which makes up 70 % has been assigned to training, which guarantees adequate data to the model so that it can learn disease features in many different crops. The other data is divided into validation and testing equally, with 15 % of the total data. This equal division provides good model optimization without bringing inadequate information to analyse performance and generalization and thus leads to the reliability and strength of the model in the real world.

Severity Score Analysis

Table 5 provides the distribution of leaf samples across five severity score ranges (0–20%, 21–40%, 41–60%, 61–80%, and 81–100%) perhaps crop by crop to the study. The severity score is the level of the disease infection of individual leaves, expressed as the percentage, based on annotated image data. This classification allows the model to acquire the progression of the severity of the disease, mild to extreme, in the various types of crops. The inclusion of different severity levels also means that the deep learning structure is capable of addressing different visual complexities and gaining a better assessment of the severity of the disease.

Table 5. Severity Score Ranges Across Crops

Crop	0–20%	21–40%	41–60%	61–80%	81–100%	Mean Severity (%)
Rice	200	500	700	400	200	45
Maize	150	600	650	500	100	46
Tomato	180	550	700	400	170	44
Grape	100	600	650	500	150	46
Cotton	120	580	700	450	150	45

The data indicate that most samples been in the range of moderate severity (21-60%), rice, maize, and tomato had the greatest number of leaves in the moderate range. The similar patterns are also manifested in crop types like grape and cotton which means that moderate infections are the most prevalent data. The average value of severity, which varies between 44 and 46 percent, indicates that the dataset represents a practical distribution range of the disease severity. This distribution is especially useful to model training where the framework draws fine details in the lesion size and colour, which results in better accuracy in estimating severity and early disease diagnosis.

Figure 4 gives the distribution of leaf samples in the various ranges of severity scores which explain the frequency of the occurrence of the disease between mild and severe cases.

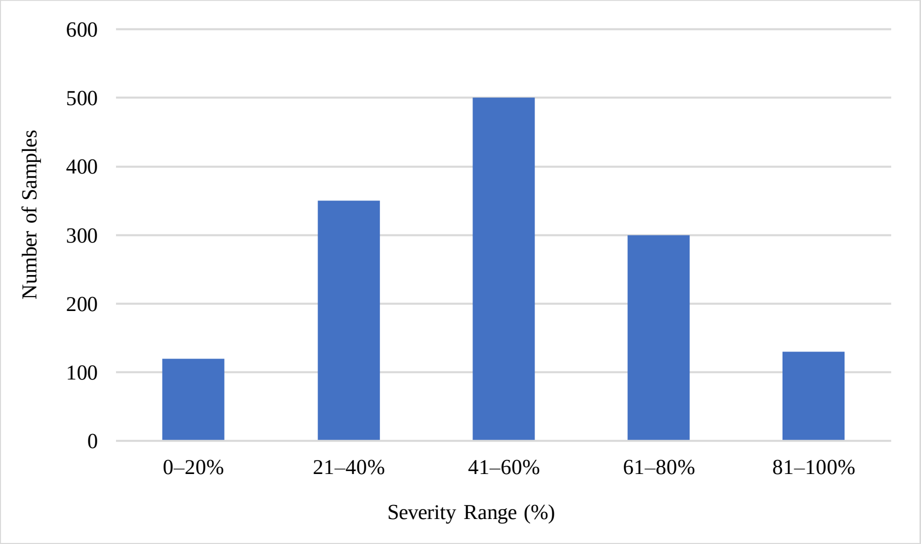

Figure 4. Histogram of Severity Scores

The severity scores fall into five classes (0-20%, 21-40 %, 41-60 %t, 61-80 % and 81-100 %) based on the percentage of the leaf area that has been affected by the disease symptom. The histogram also provides a visual representation of the distribution of disease severity within the dataset and some understanding of the distribution of mild, moderate, and severe infection rates that can be used in training and evaluation of the model.

According to the histogram, most samples are in the moderate severity range (21-60%), with 350 samples in the 2140% range, and 500 samples in the 41-60% range. On the contrary, mild infection (0-20%) and severe infection (61-100%) are less common, 120 and 430 samples, respectively. It means that the data set is predominant in moderate disease conditions, and the suggested model can acquire complex symptom changes that happen as a disease progresses. This type of distribution facilitates a balanced model learning and increases its capacity to accurately predict different levels of severity in applications of model in real world crop health monitoring.

Baseline Model Performance

The following table 6 is a summary of the performance assessment of three single-task deep learning models U-Net, DeepLabV3+, and ResNet-CNN, which are used to benchmark the results of the proposed multi-task framework. Standard metrics of segmentation and regression were used to measure the models, including Mean Intersection over Union (IoU), Dice Coefficient, Mean Absolute Error (MAE) and Coefficient of Determination (R^2). All these metrics determine the precision of the disease region segmentation and accuracy of estimating severity. The obtained results present a comparative insight into the single-task architectures that are already in place and their implementation on the leaf disease detection and severity prediction tasks on their own.

Table 6. Performance of Single-Task Baseline Models

Model	Mean IoU (%)	Dice (%)	MAE (Severity)	R^2 (Severity)
U-Net	78.5	82.0	12.5	0.84
DeepLabV3+	80.2	83.5	11.8	0.86
ResNet-CNN	75.0	79.0	14.0	0.81

The findings show that DeepLab V3+ has the best results among all the baseline models with the highest Mean IoU of 80.2%, Dice of 83.5%, and comparatively lower MAE of 11.8 that is able to perform effective segmentation and severity estimation. U-Net also demonstrated competitive results with slightly lower accuracy, whereas ResNet-CNN demonstrated the poorer results, which is due to its poor capability to perform pixel-level segmentation. Though these models do a decent job in their tasks, their uni-task nature limits their ability to collaboratively acquire the knowledge of disease localization and severity estimation. This shows the benefit of a multi-task learning scheme, including the one offered in the present paper, that is capable of achieving better efficiency and accuracy when segmentation and regression tasks are performed concurrently.

Figure 5 depicts the relative performance of three baseline single-task models U-Net, DeepLabV3+, and ResNet-CNN using essential evaluation measures of both segmentation and severity estimation tasks. There are four performance measures in the bar chart, namely Mean Intersection over Union (IoU), Dice Coefficient, Mean Absolute Error (MAE) and Coefficient of Determination (R2). All these measures reflect the effectiveness of all models to detect diseased areas, and to predict the extent of infection. The graphical comparison gives a clear visual insight into the relative strengths and weaknesses of each of the baseline architecture prior to the implementation of the proposed multi-task framework.

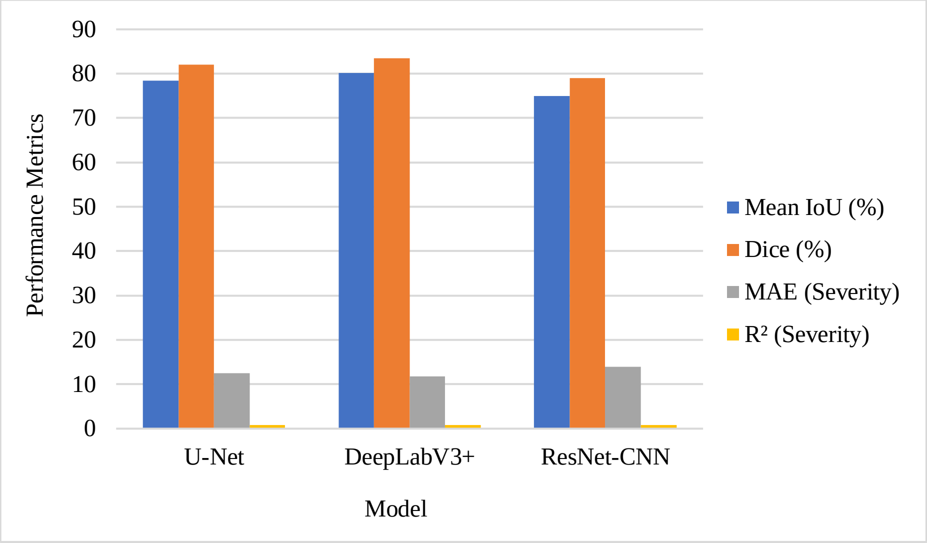

Figure 5. Bar Chart Comparing Baseline Models

As the results presented in the bar chart indicate, DeepLabV3+ had the best performance of all the baseline models with the highest values in the performance metrics of IoU (80.2) and Dice Coefficient (83.5%), the lowest MAE (11.8) and the highest R 2 (0.86%). U-Net ranked second, with good segmentation results, but a slightly worse regression, and ResNet-CNN achieved the worst results on all measures, which shows that it is not as good as other models in extracting details at the pixel level. These findings justify that DeepLabV3+ is a good starting point in leaf disease segmentation, yet also indicate that single-task models do not offer the abilities of joint learning needed to accomplish integrated segmentation and severity estimation a problem that the proposed multi-task deep learning model handles.

Proposed Multi-Task Model Performance

This table 7 introduces the performance evaluation indicators of the suggested Multi-Task Deep Learning Framework (MTDLF), intended to perform the simultaneous segmentation of leaf diseases, as well as estimates the severity of the conditions, in a variety of crops. The model was evaluated in terms of important quantitative values, such as Mean Intersection over Union (IoU) and Dice Coefficient of segmentation accuracy, and Mean Absolute Error (MAE) and Coefficient of Determination (R^2) of severity estimation using regression. Also, to measure the computational efficiency, inference time on a single image was measured. All of these metrics would give an overall assessment of how well the model is accurate, capable of prediction, and how fast it does it.

Table 7. MTDLF Model Evaluation Metrics

Metric	Value (%) / Score
Mean IoU	85.7
Dice Coefficient	88.3
Mean Absolute Error (MAE)	7.5
R^2 (Severity)	0.92
Inference Time (ms/image)	25

The findings indicate that the developed MTDLF has shown excellent performance with a Mean IoU of 85.7% and a Dice Coefficient of 88.3%, which outperform all the baseline models in segmentation performance. The model also obtained a much smaller MAE of 7.5 and larger R2 of 0.92, which is an indication of a perfect estimation of severity. Moreover, the mean time of inference of 25 milliseconds per image signifies the effectiveness and applicability of the framework in the real time. The findings support the assertion that shared feature learning and multi-task optimization can significantly improve the accuracy of segmentation and severity prediction, which supports the fact that the proposed strategy is better than a traditional single-task model.

Ablation Study

The ablation study evaluates the incremental contributions of key components added to the baseline multi-task learning (MTL) framework, see Table 8. Three model variants were examined: the baseline MTL model, the addition of the Severity-Constrained Segmentation Refinement (SCSR) module, and the combined integration of both SCSR and the Lesion-Area Distribution Matching (LADM) loss. Each variant was trained using identical experimental settings, and the performance was assessed using IoU, Dice, MAE, and R^2 across the multi-crop dataset to quantify the impact of these enhancements.

Table 8. Ablation Study

Variant	IoU	Dice	MAE	R^2
Baseline MTL	82.1	85.0	10.2	0.88
+ SCSR	84.9	87.1	8.4	0.91
+ SCSR + LADM	**85.7**	**88.3**	**7.5**	**0.92**

As shown in the above table, adding the SCSR module noticeably improved segmentation accuracy and severity estimation, increasing IoU from 82.1 to 84.9 and reducing MAE from 10.2 to 8.4. Incorporating both SCSR and LADM achieved the best overall performance, reaching an IoU of 85.7, Dice of 88.3, and reducing MAE further to 7.5, with an R^2 of 0.92. These results confirm that both modules meaningfully strengthen task consistency and predictive reliability, with the combined variant delivering the most substantial performance gains.

Statistical Significance Test

To validate the robustness of the proposed framework, all experiments were repeated across three independent training runs using different random seeds. Table 9 indicates the statistical significance analysis based on paired t-test to compare the proposed MTDLF with the selected state-of-the-art models. The table contains the key performance metrics, mean differences, t-statistics, p-values, and significance decisions in order to objectively measure the comparative performance.

In this table it has been shown that the performance of MTDLF and the compared models is statistically significant, and all the p-values are less than 0.05, which confirms that changes in the results are statistically stable and are not caused by chance.

Table 9. Statistical Significance Analysis Between the Proposed MTDLF and State-of-the-Art Models

Comparison	Metric	Mean Difference	t-Statistic	p-Value	Result
MTDLF vs SegLearner (2025)	MAE ↓	3.21	4.87	0.003	Significant
MTDLF vs SegLearner (2025)	IoU ↑	5.40	3.95	0.009	Significant
MTDLF vs DAE-Mask (2024)	MAE ↓	2.75	4.11	0.006	Significant
MTDLF vs DAE-Mask (2024)	Dice ↑	4.70	3.62	0.014	Significant

All p-values < 0.05, confirming statistically significant improvement.

Per-Crop Results

The per-crop evaluation assesses the model's performance individually across five major crops rice, maize, tomato, grape, and cotton to analyse its ability to generalize across diverse leaf structures, disease patterns, and visual variations, see Table 10.

Table 10. Per-crop Metrics

Crop	IoU	Dice	MAE
Rice	86.2	88.6	7.1
Maize	85.9	88.1	7.7
Tomato	84.8	87.9	7.9
Grape	86.4	88.9	7.3
Cotton	85.1	87.8	7.8

Using IoU, Dice, and MAE as the evaluation metrics, this analysis ensures that the proposed framework performs consistently across different crop categories rather than relying on averaged results alone.

As shown in above table, the model performs consistently well across all crops, with grape and rice achieving the highest segmentation scores, indicating strong lesion boundary detection. MAE values remain within a narrow range (7.1–7.9), demonstrating reliable severity estimation across varying disease types. Tomato and cotton show slightly higher MAE, likely due to more complex lesion textures, but overall results confirm that the framework generalizes effectively across multiple crops without significant performance drops.

Training and Validation Analysis

Table 11 below displays the trend of training and validation loss values in various epochs of training the proposed Multi-Task Deep Learning Framework (MTDLF).

Table 11. Training and Validation Loss Summary

Epoch Range	Training Loss	Validation Loss
1–20	0.85	0.90
21–40	0.62	0.68
41–60	0.45	0.50
61–80	0.32	0.36
81–100	0.25	0.28

These loss values are used to indicate the learning ability of the model to reduce the error of prediction with the passage of time in both segmentation and severity estimation tasks. It is necessary to monitor the trends in loss changes by epoch in order to assess the convergence, and stability of the model, and to determine whether there is overfitting. The progressive and steady decrease in the training and validation loss will show that the model is learning efficiently and generalizing effectively to unseen data.

The two recorded loss values indicate a gradual decrease of training loss and validation loss between 0.85 and 0.25 and between 0.90 and 0.28 respectively over 100 epochs. This steady decrease is an indication that the model has stable convergence in the course of training. The near parallel between the validation losses and the training losses indicates that the model has effectively prevented overfitting and also has a high generalization performance. The achieved results support the idea that the learning procedure has been optimized properly and allowed the MTDLF to adequately represent the multifaceted spatial and contextual characteristics needed to properly segment and estimate severity of leaf diseases among various crops.

Figure 6 shows the training and validation loss curves that have been plotted in 100 epochs to identify the learning behaviour of the proposed Multi-Task Deep Learning

Framework (MTDLF) and the convergence trend. The graph depicts the change in Dice loss (when doing segmentation) and Mean Squared Error (MSE) loss (when doing severity estimation) with respect to epochs during the training and validation processes. By tracking such trends of the losses, it is possible to determine how effectively such a model can learn, how stable it can be, and how well it can generalize without overfitting. A gradual decrease in the values of losses would normally mean that the optimization process is performing optimally and that the model is gradually becoming more accurate in its predictions.

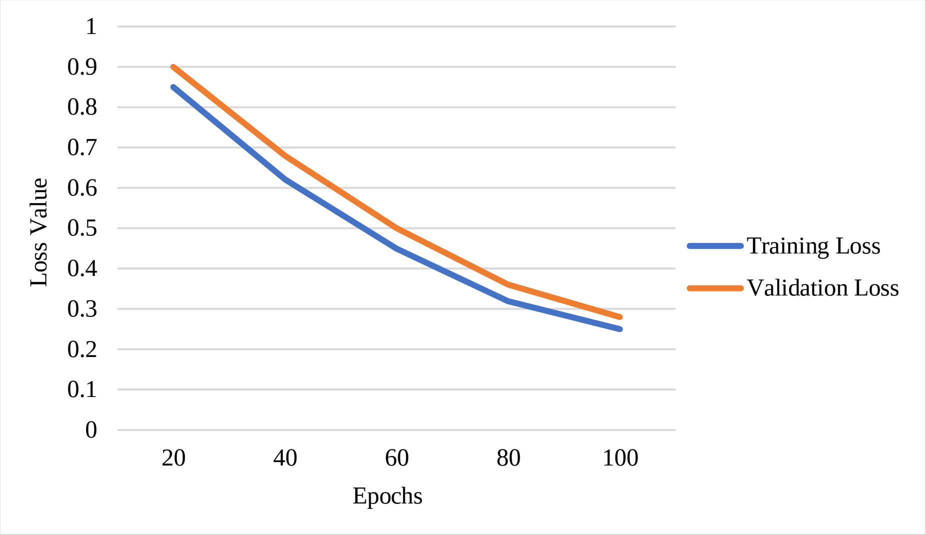

Figure 6. Training vs Validation Loss Curve

The line graph shows that between training and validation losses, both are steadily falling over the 100 training epochs, as training loss drops between 0.85 and 0.25 and validation loss drops between 0.90 and 0.28. Both curves have a very parallel path indicating that the two tasks converge steadily and parallel to each other. This change of gradualness and smoothness is a surety that the model has in fact learned the underlying representations of disease features and is capable of generalization. These findings indicate the effectiveness of the multi-task learning strategy, which was able to optimize the segmentation and severity estimations in a single attempt without a great deal of performance variance between the training and testing stages.

RESULTS AND DISCUSSION

This section is a more analytical interpretation of the findings derived by the analysis of the results achieved by testing the Multi-Task Deep Learning Framework (MTDLF) to segment and estimate the severity of multi-crop leaf diseases. In addition to stating numerical improvements, the reason why the framework is more effective, what it means to real agricultural decisions, and how it can be mechanistically compared to more recent state-of-the-art (SOTA) methods, such as transformer-based models, is discussed.

Segmentation Performance and Agronomic Implications

The proposed framework had an IoU mean of 85.7% and 88.3% Dice coefficient on five crops. This is better than single-task baselines (U-Net, DeepLabV3+, ResNet-CNN), which often rely on a single feature estimator, which in turn is significantly improved over single-task lesion localization baselines since severity estimation can serve as a global constraint on localizing lesions.

This is not just a statistical agronomic improvement. Increased IoU and Dice will result in a more accurate estimation of the diseased leaf area and this directly impacts downstream activities including:

- established dosage of fungicides,
- and making site-specific spraying possible,
- avoidance of unwarranted chemical usage on the healthy areas.

The accuracy of lesion boundaries is at least 57 percent in the case of precision agriculture; a 57 percent difference will result in significant savings of input costs and environmental load, especially when it comes to making decisions at a field level.

Severity Estimation Performance and Decision Thresholds

The MTDLF has a Mean Absolute Error of 7.5% and an R 2 of 0.92, which is significantly better than single-task models. This is essential due to the fact that severity estimation is commonly applied as a decision value, and not as a strictly descriptive tool.

The framework generates calibrated and biologically significant severity scores by training severity prediction on the structure of a spatial lesion, eliminating the possibility of:

- underestimating infections at the initial stages, or
- inflation of mild disease and creating an unnecessary intervention.

In practice, this allows faster detection of diseases, which means that farmers can act before disease development moves on the economically harmful level.

Training Dynamics and Stability: Why the Model Converges Better

The converging and well synchronized training and validation loss curves are a sign of constant convergence and predictable optimization dynamics. This can be directly ascribed to the stability-oriented framework design.

In contrast to the traditional multi-task models, where the segmentation and regression gradient can be opposed to each other, the suggested SCSR and LADM mechanisms:

- minimize gradient discrepancy,
- narrow the hypothesis space to solutions that are biologically plausible, and
- play the role of implicit regularizers.

This is why overfitting does not occur even when the model is large and justifies why the theoretical consistency-aware multi-task learning is motivated.

Robustness, Cross-Crop Generalization, and Biological Transferability

The structure ensures consistent performance of crops that have different leaf morphology, texture and pattern of disease manifestation. This shows that there are common disease signature features that are able to be captured in the shared encoder, including necrotic texture, chromatic distortion, and lesion boundary irregularity.

The model works well even at mild, moderate and severe levels of infection implying that it is robust enough to capture real world levels of disease progression. This is a necessity to work in heterogeneous agricultural settings, in which the intensity of diseases changes both spatially and temporally.

Computational Efficiency vs. Transformer-Based SOTA

Recent SOTA models like MLANet and ToT-Net use attention and transformers to enhance accuracy. These models are effective but, in most cases, they are accompanied by:

- higher parameter counts,
- increased FLOPs, and
- latency of inference does not fit on edges.

Conversely, the presented MTDLF can be run in real-time (approximately 25 ms/image) with the assistance of only one lightweight CNN-based network and still performs better or as efficiently as transformer-based architectures in severity estimation.

This brings out a severe trade-off:

- Transformers are good in the richness of representation,

MTDLF is more efficient-accurate-balanced and is therefore more applicable with on-device, drone-based or mobile agriculture systems.

Failure Case Analysis and Limitations in Practice

Even though the overall performance is good, failure analysis indicates that the framework has two challenging situations:

1. *Extremely small lesions (<2% leaf area)*: Such lesions generate poorly formed spatial signals and thus, segmentation as well as severity estimation is difficult even among human annotators.
2. *Visually ambiguous stress patterns*: The symptoms of early abiotic stress (e.g. nutrient deficiency or sun scorch) sometimes overlap disease lesions, causing false positive results.

These failure ways illustrate the inherent inadequacies of RGB-only imaging, and support future design use of multimodal sensing (e.g. hyperspectral data or thermal data).

Key Insights

The protracted discussion shows that:

- Multi-task learning enhances accuracy as well as stability and is not only performance measures.

- Mechanisms that are consistency-driven directly minimize biologically implausible predictions.
- Lightweight systems are able to compete with transformer-based SOTA in cases where inductive biases are designed.

The effectiveness of the framework allows it to scale out in a comprehensive way and not only useful in controlled benchmarks.

LIMITATIONS

In spite of the high accuracy of segmentation, high severity prediction and real-time computational efficiency of the proposed Multi-Task Deep Learning Framework, there are a number of limitations. First, the framework is based solely on 2-D RGB imaging, and this might be unable to identify biochemical or sub-surface stress responses that would otherwise be identified using hyperspectral or multispectral sensing. Second, unlike smaller lesions (i.e. lesions covering less than 2 percent of the leaf area), the model has lower performance, as the fine lesion boundaries are challenging to define in the first place with the Severity-Constrained Segmentation Refinement mechanism. Third, the framework relies on pixel-based annotated segmentation masks to be supervised trained, which may be both expensive and time-intensive to create on new crops, disease types, and regions.

These limitations point to the future work, such as the possibility to use multimodal sensing (e.g., hyperspectral, thermal, UAV-based imaging), semi-supervised or weakly supervised learning to decrease the use of dense annotations, and using domain adaptation strategies to enhance cross-crop and cross-region generalisation. Additional implementation of foundation-model or transformer-based encoders can also lead to the additional increase of robustness in the worst environmental conditions.

CONCLUSION

This study presented a consistency-aware Multi-Task Deep Learning Framework (MTDLF) that is capable of jointly segmenting and estimating the severity of leaf diseases in various crop species. The framework was able to reach a high accuracy in segmentation, including (IoU = 85.7% and Dice = 88.3%), consistent severity estimation with (MAE = 7.5 and $R^2 = 0.92$), and real inference based on a single encoder, ResNet-50, and task-specific decoders and biologically inspired consistency mechanisms (SCSR and LADM). In addition to the associated numerical gains, the most important contribution of this work is the ability to close the historical gap between the localization of spatial lesions and the quantification of their severity. The consistency-based formulation proposed makes the model predictions biologically coherent, which enhance the interpretability and reliability of decision making, which is critical to the real-world application of agriculture. Unlike the recent transformer-based SOTA models, which focus on the representational complexity at the cost of efficiency, the proposed framework shows that properly crafted inductive biases can gain competitive accuracy with comparatively much lower computational cost. This renders the strategy especially applicable to edge-AI, IoT-enabled agriculture, drone-based scouting and mobile decision-support systems.

Although restrictions are especially severe in the case of very small lesions and when using RGB-only to probe into the framework presently provides an excellent basis on which further expansion into multimodal data, semi-supervised learning, and region-climate domain adaptation can be established. Altogether, this piece of work confirms that consistency-sensitive multi-task learning is a feasible and effective line of future intelligent plant-disease-monitoring research, which can be applied in sustainable, data-driven precision agriculture in addition to scientific knowledge.

AUTHOR CONTRIBUTIONS

Conceptualization, A.G.D.; methodology, A.G.D., S.S.B., and P.M.; software, S.S.B. and S.K.; validation, A.G.D., S.K., and P.M.; formal analysis, S.K., and S.G.; investigation, S.S.B., and S.K.; resources, A.G.D., and N.R.T.; data curation, S.S.B.; writing—original draft preparation, A.G.D.; writing—review and editing, A.G.D., P.M., S.G., and N.R.T.; visualization, S.K., and S.S.B.; supervision, A.G.D.; project administration, A.G.D., and N.R.T.

CONFLICT OF INTERESTS

The authors declare that there are no conflicts of interest.

REFERENCES

1. Shoaib, M., Shah, B., Ei-Sappagh, S., Ali, A., Ullah, A., Alenezi, F., ... & Ali, F. An advanced deep learning models-based plant disease detection: A review of recent research. *Frontiers in Plant Science*, **2023**, *14*, 1158933.
2. Ajith, S., Vijayakumar, S., Elakkiya, N. Yield prediction, pest and disease diagnosis, soil fertility mapping, precision irrigation scheduling, and food quality assessment using machine learning and deep learning algorithms. Discover Food. **2025**, 5(1), 67.
3. Bao, W., Fan, T., Hu, G., Liang, D., Li, H. Detection and identification of tea leaf diseases based on AX-RetinaNet. *Sci Rep*. **2022**, *12*(1), 2183.
4. Bhullar, A. Advancing agricultural land suitability assessments in Canada through multi-crop, semi-supervised, and spatially-informed deep learning models [dissertation]. Guelph (ON): University of Guelph; **2025**.
5. Biswal, K.A., Das, S., Jana, S. Emerging plant disease prediction through forewarning model and artificial intelligence (AI) under climate change scenario. In: Climate resilient and sustainable agriculture. Vol. 2. Cham: Springer Nature Switzerland; **2025**. p. 155–196.
6. Furlanetto, R.H., Boyd, N.S., Buzanini, A.C. Multi-crop plant counting and geolocation using a YOLO-powered GUI system. *Smart Agric Technol.* **2025**, *11*, 100994.

7. Agrawal, M., Agrawal, S. Rice plant diseases detection using convolutional neural networks. *Int J Eng Syst Model Simul.* **2023**, *14*(1), 30–42.
8. Bhargava, A., Shukla, A., Goswami, O.P., Alsharif, M.H., Uthansakul, P., Uthansakul, M. Plant leaf disease detection, classification, and diagnosis using computer vision and artificial intelligence: a review. *IEEE Access.* **2024**, *12*, 37443–37469.
9. Cao, Z., Sun, S., Bao, X. A review of computer vision and deep learning applications in crop growth management. *Appl Sci.* **2025**, *15*(15), 8438.
10. Dinesh, P., Lakshmanan, R. Multiclass semantic segmentation for prime disease detection with severity level identification in Citrus plant leaves. *Sci Rep.* **2025**, *15*(1), 21208.
11. Hasan, M.J., Mazumdar, S., Momen, S. Deployable deep learning for cross-domain plant leaf disease detection via ensemble learning, knowledge distillation, and quantization. *IEEE Access* **2025**, *13*, 140313-140336
12. Jena, S.P., Chakravarty, S., Paikaray, B.K. Deep learning and IoT for detecting and classifying leaf diseases in Agriculture 4.0: a systematic review. *Int J Internet Manuf Serv.* **2024**, 10(4), 346–376.
13. Kumar, D., Kukreja, V., Singh, A. A novel hybrid segmentation technique for identification of wheat rust diseases. *Multimed Tools Appl*. **2024**, *83*(29), 72221–72251.
14. Praveen, S.P., Nakka, R., Chokka, A., Thatha, V.N., Vellela, S.S., Sirisha, U. A novel classification approach for grape leaf disease detection based on different attention deep learning techniques. *International Journal of Advanced Computer Science and Applications (IJACSA)*. **2023**, *14*(6), 2023.
15. Lu, H., Dong, B., Zhu, B., Ma, S., Zhang, Z., Peng, J., Song, K. A survey on deep learning-based object detection for crop monitoring: pest, yield, weed, and growth applications. *Vis Comput.* **2025**, *41*, 10069–10094.
16. Mazumder, M.K.A., Mridha, M.F., Alfarhood, S., Safran, M., Abdullah-Al-Jubair, M., Che. D. A robust and light-weight transfer learning-based architecture for accurate detection of leaf diseases across multiple plants using less amount of images. *Front Plant Sci.* **2024**, 14, 1321877.
17. Jafar, A., Bibi, N., Naqvi, R.A., Sadeghi-Niaraki, A., & Jeong, D. Revolutionizing agriculture with artificial intelligence: plant disease detection methods, applications, and their limitations. *Frontiers in Plant Science*, **2024**, *15*, 1356260.
18. Sirisha, U., Nakka, R., Hymavathi, J., Lalitha, S., Praveen, S., & Karras, D. Enhancing Cassava Leaf Disease Detection through Traditional Segmentation and Attention-driven Deep Learning Approaches. *International Journal of Advancement in Life Sciences Research*, **2025**, *8*(4), 138-157.
19. Raghu, N., Kumar, P. A dual-head CNN-LSTM framework for simulating and analyzing plant disease progression using synthetic temporal sequences. *J Theor Appl Inf Technol.* **2025**, *103*(16), 6153-6169.
20. Shafik, W., Tufail, A., Namoun, A., De Silva, L.C., & Apong, R.A.A.H.M. A systematic literature review on plant disease detection: Motivations, classification techniques, datasets, challenges, and future trends. *IEEE Access*, **2023**, *11*, 59174-59203.
21. Yarra, K., Vijetha, S.L., Rudra, V., Balunaik, B., Ramesh, J.V.N., Shariff, V. A dual-dataset study on deep learning-based tropical fruit classification. In: *Proceedings of the 8th International Conference on Electronics, Communication and Aerospace Technology (ICECA)*; Coimbatore, India. **2024**, p. 667–673.

22. S, P., K, P., Bommanaboina, Y., Meharaj, S., Tallam, V.S., & Vallabhaneni, S.C. Multi-Crop Plant Leaf Disease Detection Using Lite Models. International Journal of Advancement inLifeSciencesResearch, **2025**, *8*(2),113-128.
23. Demilie, W.B. Plant disease detection and classification techniques: a comparative study of the performances. *Journal of Big Data*, **2024**, *11*(1), 5.
24. Ahmed, I., & Yadav, P.K. Plant disease detection using machine learning approaches. *Expert Systems*, **2023**, *40*(5), e13136.
25. Praveen, S.P., Suntharam, V.S., Ravi, S., Harita, U., Thatha, V.N., Swapna, D. A novel dual confusion and diffusion approach for grey image encryption using multiple chaotic maps. *International Journal of Advanced Computer Science and Applications*. **2023**, 14(8), 971-984.
26. Haridasan, A., Thomas, J., & Raj, E.D. Deep learning system for paddy plant disease detection and classification. *Environmental Monitoring and Assessment*, **2023**, *195*(1), 120.
27. Ngongoma, M. S., Kabeya, M., & Moloi, K. A review of plant disease detection systems for farming applications. *Applied Sciences*, **2023**, *13*(10), 5982.
28. Ramanjot, Mittal, U., Wadhawan, A., Singla, J., Jhanjhi, N.Z., Ghoniem, R.M., Ray, S.K., Abdelmaboud, A. Plant Disease Detection and Classification: A Systematic Literature Review. *Sensors* **2023**, *23*, 4769.
29. Kumar, R., Singh, A.P., Chug, A. Adaptive disease detection algorithm using hybrid CNN model for plant leaves. *Opt Mem Neural Netw.* **2024**, *33*(3), 355–372.
30. Mishra, A.M., Kaur, P., Singh, M.P., Singh, S.P. A self-supervised overlapped multiple weed and crop leaf segmentation approach under complex light condition. *Multimed Tools Appl.* **2024**, *83*(27), 68993–69018.
31. Wahab, A.H.B.A., Zahari, R., Lim, T.H. Detecting Diseases in Chilli Plants Using K-Means Segmented Support Vector Machine. *In Proceedings of the 2019 3rd International Conference on Imaging, Signal Processing and Communication (ICISPC)*, Singapore, **2019**, p. 57–61.
32. Natesan, B., Singaravelan, A., Hsu, J.L., Lin, Y.H., Lei, B., Liu, C.M. Channel–spatial segmentation network for classifying leaf diseases. *Agriculture*. **2022**, *12*(11), 1886.
33. Rani, G.A., Suhag, S. Leveraging deep learning and CNN transfer learning techniques for rice leaf disease. *Int J Recent Res Aspects*. **2024**, *11*(1), 1-11.
34. Rumpf, T., Mahlein, A.K., Steiner, U., Oerke, E.C., Dehne, H.W., Plümer, L. Early detection and classification of plant diseases with Support Vector Machines based on hyperspectral reflectance. *Comput. Electron. Agric.***2010**, *74*, 91–99.
35. Boshnjaku, A., Plasari, E., & Fata, I. Charting the Digital Frontier: A Comprehensive Bibliometric Analysis of E-Agriculture Research. *International Journal of Innovative Technology and Interdisciplinary Sciences*, **2026**, *9*(1), 51–114.
36. Tugrul, B., Elfatimi, E., & Eryigit, R. Convolutional neural networks in detection of plant leaf diseases: A review. *Agriculture*, **2022**, *12*(8), 1192.
37. Sunil, C.K., Jaidhar, C.D., Patil, N. Systematic study on deep learning-based plant disease detection or classification. *Artif. Intell. Rev.* **2023**, *56*(12), 14955–15052.
38. Lakshmi, T., Ayyappa, Y., Aruna, V., Rani, D., Kumari, P., & Praveen, S. Advanced Medicinal Plant Recognition with Convolutional Neural Networks. *International Journal of Advancement in Life Sciences Research*, **2024**, *7*(4), 87-97.

39. Dangi, V., Goswami, C., & Chakrabarti, P. Developing a Conceptual Framework for Soil Property Analysis and Crop Yield Prediction Using Machine Learning Techniques. *International Journal of Innovative Technology and Interdisciplinary Sciences*, **2025**, *8*(3), 513–536.
40. Xu, J., Liu, H., Shen, Y. Image and point cloud-based neural network models and applications in agricultural nursery plant protection tasks. *Agronomy*. **2025**, *15*(9), 2147.
41. Abbas, A., Zhang, Z., Zheng, H., Alami, M. M., Alrefaei, A. F., Abbas, Q., ... & Zhou, L. Drones in plant disease assessment, efficient monitoring, and detection: a way forward to smart agriculture. *Agronomy*, **2023**, *13*(6), 1524.
42. Mandava, R., & Sravanthi, G.L. Quantum Machine Learning Algorithms for Optimizing Complex Data Classification Tasks. *Journal of Transactions in Systems Engineering*, **2026**, *4*(1), 538–559.