



Original Research Paper

CA-ASPP-ConvNeXt: A Lightweight Multi-Scale Framework for Tomato Leaf Disease Classification Model

Jahnavi Lakshmi Kandula¹, Dubba Nagamalleswari²

¹ Department of CSE, Koneru Lakshmaiah Education Foundation,
Vijayawada, Andhra Pradesh, India.

*Corresponding author(s). E-mail(s): 2401050018@kluniversity.in;
Contributing authors: nagamalleswary@Gmail.com;

Abstract

Early and accurate detection of tomato leaf diseases is essential for sustainable agriculture and crop yield improvement. In this study, we propose a novel deep learning-based classification model that combines the strengths of transfer learning, attention mechanisms, and multi-scale feature extraction. The backbone of the model utilizes ConvNeXt-Tiny, pretrained using self-supervised learning techniques such as SimCLR or BYOL, to extract rich hierarchical features from input tomato leaf images of size $224 \times 224 \times 3$. To enhance the representational capacity of the network, Coordinate Attention (CA) modules are integrated after each stage of the ConvNeXt backbone, allowing the model to focus on both spatial and channel-wise informative features. To capture multi-scale contextual information critical for identifying fine-grained disease patterns, the network is further augmented with a Atrous Spatial Pyramid Pooling (ASPP) module. The resultant feature map is passed through a Global Average Pooling (GAP) layer, followed by a Dropout regularization and a fully connected Softmax classifier to perform multi-class classification across 11 distinct tomato leaf disease categories. Extensive experiments on a tomato leaf disease dataset demonstrate that our model achieves high classification accuracy while maintaining computational efficiency, making it suitable for real-time agricultural diagnostics and mobile deployment. The integration of coordinate attention and self-supervised pretrained features significantly boosts performance, particularly in distinguishing visually similar diseases under varied environmental conditions.

Keywords: Tomato disease classification, Coordination Attention, ASPP, Hybrid Datasets, CNN, Domain Adaptation.

1 Introduction

Tomato is one of the most important vegetable crops in India, both in terms of area cultivated and total production. However, tomato cultivation is severely affected by various diseases, especially those targeting the leaves of the plant. These diseases can significantly reduce crop yield and quality, leading to financial losses for farmers.

India, being the world's second-largest tomato producer, has shown variable trends in tomato production between 2017 and 2024. The overall trend highlights growth, peaking in 2023 at 21.18 million tons. However, 2024 witnessed a 2.3% decline to 20.69 million tons, attributed mainly to crop diseases, erratic rainfall, and temperature extremes.

Earlier years also exhibit fluctuations, reflecting the vulnerability of tomato farming to both biological threats and climate change. The visual graph clearly demonstrates how these external pressures influence national yield, highlighting the importance of precision agriculture and early disease detection technologies to support stable production.

From 2017 to 2024, India's tomato production was heavily impacted by several plant diseases, leading to significant yield losses. The most devastating was Late Blight, flourishing in humid conditions and causing severe leaf and fruit damage. Tomato Leaf Curl Virus (ToLCV), transmitted by whiteflies, led to stunted plant

growth and substantial productivity decline. Bacterial Wilt, a soil-borne pathogen, is notorious for causing sudden wilting and plant death, particularly during the monsoon season. Its persistence in the soil makes it one of the most difficult tomato diseases to manage, with yield losses ranging from 15% to 60

Other destructive fungal diseases such as Early Blight, Fusarium Wilt, and Anthracnose further reduce crop health, collectively contributing to annual losses of 0.55–0.80 million tons, amounting to over 5 million tons over the study period. Bacterial Spot also adds 10–40% yield losses each year by damaging leaves, stems, and fruits.

Among viral threats, the ToLCV, transmitted by whiteflies, is especially devastating, causing 30–70% yield loss. It induces leaf curling, yellowing, and stunting, reducing photosynthesis and fruit production. Similarly, Tomato Spotted Wilt Virus (TSWV) and Chilli Leaf Curl Virus (ChiLCV) can trigger complete (up to 100%) crop failure in localized outbreaks, often in poorly managed farms.

Late Blight, thriving in cool and humid environments, remains another major concern, inflicting 40–75% damage during epidemics by rapidly destroying foliage, stems, and fruits.

Overall, the combined impact of these pathogens—exacerbated by climate variability such as erratic rainfall, high humidity, and rising temperatures—has intensified tomato disease outbreaks in India. Traditional

*Corresponding Author:

Received: 25th July, 2026 | Revised: 26th August, 2026 | Accepted: 18th September, 2026 | Available Online: 24th September, 2026

(DOI): [10.70102/AEJ.2026.18.7s.75](https://doi.org/10.70102/AEJ.2026.18.7s.75)

management approaches are proving inadequate, underscoring the need for more robust and technology-driven control strategies.

2 Motivation

Despite being one of India's most important crops, diseases including Bacterial Wilt, Late Blight, and Tomato Leaf Curl Virus cause significant yield losses that are made worse by climate fluctuation and postponed remedies. Traditional manual inspection lowers productivity and livelihoods since it is labor-intensive, prone to errors, and frequently beyond of reach for small farmers.

In rural India, where digital access is growing, AI and deep learning offer a useful remedy. This research creates a mobile-friendly solution for real-time tomato leaf disease diagnosis using lightweight CNNs. By facilitating early diagnosis through smartphones or embedded devices, the system lessens the need for specialists, promotes prompt treatment, and supports data-driven agricultural decision-making.

3 Problem statement

Tomato cultivation in India is increasingly threatened by a range of plant diseases such as Tomato Leaf Curl Virus (ToLCV), Late Blight, Bacterial Wilt, and Early Blight, which can cause yield losses exceeding 70% in severe cases. These diseases directly affect food supply, market stability, and the economic well-being of farmers. Traditional methods of disease detection rely heavily on manual inspection, which is often slow, inaccurate, and inaccessible for small and marginal farmers, particularly in rural regions. Delays in diagnosis allow infections to spread extensively, reducing the effectiveness of interventions and leading to significant crop and income loss.

Compounding these challenges are erratic weather conditions and an overdependence on chemical pesticides, which raise both environmental and health concerns. In light of these issues, there is a pressing need for efficient, accessible, and real-time disease detection solutions. This project proposes the development of a low-cost, AI-driven mobile application utilizing deep learning models for early identification of tomato leaf diseases. By allowing farmers to detect diseases from simple leaf images and receive instant guidance, the solution aims to minimize losses, promote sustainable farming practices, and ultimately increase tomato productivity and farmer resilience.

4 Related works

Traditional Image Processing Approaches:

Mim et al. [1] pioneered the use of basic image processing techniques for tomato leaf disease detection, utilizing color segmentation and texture analysis methods. They applied thresholding to isolate leaf lesions, followed by feature extraction based on shape and color histograms. Using a small dataset of manually collected tomato leaf images, their system achieved an overall classification accuracy of approximately 82%. While effective in controlled settings, their approach suffered from sensitivity to illumination changes and complex backgrounds.

Roy and Chakma [2] emphasized the real-world economic impact of tomato diseases and highlighted the need for more robust automated systems. Their study

motivated further research towards deep learning approaches that can generalize better under diverse environmental conditions.

Deep Learning Approaches and CNN Architectures:

Mohanty et al. [3] introduced deep CNNs for plant disease detection using the publicly available PlantVillage dataset, which contains over 50,000 images across 38 classes. They fine-tuned AlexNet and GoogleNet models pretrained on ImageNet, reporting an impressive classification accuracy of 99.35% on tomato leaf diseases. This work established CNNs as the state-of-the-art baseline and demonstrated the power of transfer learning in agricultural applications.

Ferentinos et al. [4] extended this by comparing several deep CNN architectures such as VGG16, ResNet50, and AlexNet on PlantVillage data. His best-performing model (ResNet50) achieved 99.53% accuracy with approximately 25 million parameters. He also highlighted the importance of data augmentation to improve generalization and noted that deeper networks performed better but required more computational resources.

Wang et al. [5] proposed the Efficient Channel Attention (ECA) mechanism, a lightweight attention module designed to enhance CNN feature representation by capturing channel-wise dependencies without dimensionality reduction. ECA improved the performance of baseline CNNs on ImageNet classification while adding minimal overhead (less than 0.01% extra parameters). This attention mechanism was later adopted in tomato disease classification models to focus on key lesion areas efficiently.

Attention Mechanisms and Hybrid CNN Models:

Attention mechanisms have significantly improved the performance of CNNs by enabling models to focus on the most informative regions of an image, such as disease-affected areas on leaves. This is used to reduce the misclassification results from background noise and variations in leaf.

Zhao et al. [6] integrated both spatial and channel attention modules into a CNN architecture. With almost 8 million parameters and a classification accuracy of 98.9%, the model's performance and computational efficiency were well-balanced.

The Attention Embedded Residual CNN model was presented by Karthik et al. [7]. It incorporates spatial attention and residual connections. With about 9 million parameters, it achieves a 99.1% accuracy rate. Better discriminative capacity was shown by AER-CNN, especially when it came to differentiating between disorders that looked alike.

In addition to improving classification ability, these experiments confirm that attention-enhanced architectures significantly related disease classes.

Lightweight and Efficient CNN Models:

A lightweight attention-based CNN model for tomato leaf disease classification was proposed by Bhujel et al. [8]. Depthwise separable convolutions and ECA modules were used for lower computational complexity. The model gets 95.7% accuracy with 2.3 million parameters, which is very smaller than normal ResNet or VGG models, after training on a combination of PlantVillage and PlantDoc datasets.

The efficient CNN design with fewer convolutional layers and integrated squeeze-and-excitation blocks for channel attention in Agarwal et al. [9]. The model achieves 96.8% accuracy with 4.5 million parameters. It demonstrated that model have a compromising performance on PlantVillage dataset.

A lightweight attention-based CNN model for tomato leaf disease classification was proposed by Bhujel et al. [8] and is intended for use on devices with limited resources. A deep CNN model using manually created Local Binary Pattern (LBP) descriptors was proposed by Hosny et al. [10] for the multi-class classification of plant leaf diseases. This model outperformed models that only use CNN or conventional features, as evidenced by its high classification accuracy of 96.2%. Additionally, with roughly 5.8 million trainable parameters, the model architecture is well developed, balancing computational efficiency and model complexity[11].

A lightweight 2D CNN design was proposed in a recent study by [12], which achieved an impressive average accuracy of 97.36%. The model obtained fast classification times of about 4.84 milliseconds per image, outperforming popular pretrained networks like VGG16,19 and InceptionV3. The parameter reduction and improved accuracy by training on more varied datasets were among the potential prospects that the study suggested.

Transfer Learning and Pretrained Models:

Transfer learning has emerged as a powerful tool in tomato leaf disease detection, allowing models trained on large-scale datasets like ImageNet to be fine-tuned for plant-specific tasks. This approach significantly reduces training time and enhances performance, especially when labeled data is scarce. Al-Gaashani et al. [13] proposed a transfer learning method that concatenated features from ResNet50, DenseNet121, and EfficientNet-B0. Their ensemble model, trained on PlantVillage data, achieved 99.2% accuracy, demonstrating strong robustness. However, the combined model consisted of approximately 60 million parameters, which may hinder deployment on low-power devices.

Trivedi et al. [14] evaluated several pretrained CNNs—including VGG16, ResNet50, and MobileNetV2—on tomato disease datasets. MobileNetV2, with just 3.4 million parameters, achieved 98.7% accuracy, highlighting its suitability for mobile deployment due to its lightweight architecture and efficient performance.

Chen et al. [15] examined AlexNet, which achieved 97.5% accuracy using around 61 million parameters. Despite its larger size, AlexNet remains a common baseline due to its simplicity and effectiveness in transfer learning.

Bhujel et al. [8] developed a custom lightweight CNN specifically for tomato leaf disease classification, reporting an accuracy of 98.3% with significantly fewer parameters. Their study emphasized the importance of mobile-friendly architectures tailored to agricultural applications.

Maeda-Gutiérrez et al. [16] conducted a comparative analysis of various pretrained CNN architectures, concluding that models like MobileNet and SqueezeNet offer the best trade-off between accuracy and model size for field deployment.

Additionally, Abbas et al. [17] combined transfer learning with generative data augmentation using Conditional GANs (C-GANs) to improve classification performance under real-world variations, a strategy particularly useful when training data is limited or imbalanced.

Overall, transfer learning continues to be instrumental in improving model accuracy and generalization. However, future research should emphasize lightweight and hardware-efficient designs to enable real-time, offline disease detection on mobile platforms used by farmers.

Datasets and Data Augmentation Strategies:

The PlantVillage dataset [18] is the most commonly used benchmark for tomato leaf disease classification, offering over 18,000 tomato leaf images spanning 11 disease classes plus healthy leaves. Images are collected under controlled conditions, making models trained on this dataset less effective in real-world scenarios.

To address this gap, Singh et al. [11] introduced the PlantDoc dataset, containing 2,598 images taken in natural environments with complex backgrounds. Models trained solely on PlantVillage exhibited a drop of 12–15% accuracy when tested on PlantDoc, motivating hybrid training strategies.

Abbas et al. [17] augmented the PlantVillage dataset using Conditional Generative Adversarial Networks (C-GANs) to synthesize realistic tomato leaf disease images. Their synthetic data improved model generalization, increasing accuracy from 96% to 98% on unseen data.

Data augmentation techniques such as random rotations, flips, color jittering, and zooming are standard practice across most studies [19], [20], helping reduce overfitting and simulate real-world variations.




Fig. 1 Distribution by Leaf Disease

Optimization and Training:

Most studies adopt the Adam optimizer [21] with an initial learning rate around 0.001 for training CNN models. Batch normalization layers [22] are integrated after convolutions to stabilize training and improve convergence speed.

Early stopping and learning rate scheduling are commonly employed to prevent overfitting. Agarwal et al. [19] reported that using dropout with rates between 0.3–0.5 and batch normalization enabled their model to achieve stable accuracy above 96%

Ensemble Methods and Model Fusion:

Recent work also investigates ensemble learning and model fusion strategies to enhance classification accuracy and model robustness. Al-Gaashani et al. [13] demonstrated that combining features from multiple pretrained CNN models—such as DenseNet, ResNet, and EfficientNet—can lead to a more generalized feature

space, improving performance on unseen data. Other studies have experimented with stacking, boosting, and bagging approaches, where multiple models contribute to the final prediction, either through majority voting or weighted averaging. While these methods often improve accuracy, they come at the cost of increased computational complexity and inference time. Efficient ensemble strategies are therefore crucial for practical deployment.




Fig. 2 Sample images of different tomato leaf diseases

5 Datasets

To train and evaluate our tomato leaf disease classification model, we utilized two publicly available datasets: PlantVillage and PlantDoc. The PlantVillage dataset [18] is a widely recognized benchmark for plant disease detection. It consists of high-resolution images captured under controlled, lab-like conditions, often with clean and homogeneous backgrounds. For this study, we focused on 11 tomato leaf disease classes.

To enhance the model's robustness and ensure generalizability to real-world agricultural conditions, we incorporated data from the PlantDoc dataset [11], which features images captured in the field with varying lighting conditions, occlusion, and cluttered backgrounds. We selected 1,409 images from PlantDoc for training and 200 images for hybrid validation. These were combined with the PlantVillage samples to form a hybrid training set of 11,409 images and a combined validation set of 1,200 images.

The dataset composition, including the class-wise distribution across both datasets, is presented in Figure 1. Additionally, a visual overview of sample representative images from both PlantVillage and PlantDoc datasets is shown in Figure 2, highlighting the contrast in image quality and background complexity across the two sources. This variation supports the motivation for hybrid training.

6 Data Preprocessing

To prepare the data for training, all images from both the PlantVillage [18] and PlantDoc [11] datasets were resized to a uniform resolution of 128×128 pixels. This resizing ensured consistency in input dimensions for the convolutional neural network and reduced computational complexity. Pixel values were normalized to the $[0, 1]$

range by rescaling each image by a factor of $1/255$, which helps in faster and more stable convergence of the model during training. Data augmentation techniques were applied exclusively to the training images to increase the diversity of the dataset and mitigate overfitting. These augmentations included random rotations up to ± 25 degrees to simulate various orientations of leaves, horizontal flipping to account for mirrored leaf structures, and zoom operations within a 20% range to emulate different distances and scales in real-world capture. Width and height shifts up to 10% of the image dimensions were introduced to handle variations in leaf positioning. Furthermore, a shear transformation of up to 15% was employed to simulate perspective distortions, and brightness adjustments in the range of 0.8 to 1.2 were applied to account for diverse lighting conditions in natural environments. Any empty regions created as a result of geometric transformations were filled using the nearest-neighbor interpolation method.

These preprocessing techniques were critical in enabling the model to generalize effectively across both clean, laboratory-quality images from PlantVillage and noisy, real-world field images from PlantDoc, which often include challenges such as occlusion, uneven lighting, and complex backgrounds.

7 Proposed Model

The architecture of the The detailed block-by-block explanation of proposed tomato leaf disease classification model using ConvNeXt-Tiny pretrained with SimCLR and Atrous Spatial Pyramid Pooling (ASPP) for multi-scale feature extraction is explained here. The Figure 3 shows the Architecture of the proposed model

7.1 ConvNeXt-Tiny

The ConvNeXt-Tiny backbone is a lightweight yet powerful convolutional neural network architecture specifically designed to combine the strengths




Fig. 3 Architecture of the proposed CA-ASPP-ConvNeXt model

of traditional CNNs and the innovations introduced by Vision Transformers (ViT). Unlike conventional architectures such as ResNet, ConvNeXt-Tiny incorporates improvements like depthwise separable convolutions, GELU activations, and Layer Normalization instead of BatchNorm, enabling more efficient training and better feature extraction. Its compact design makes it suitable for real-time applications, including mobile or embedded systems, without compromising performance. For tomato disease classification, where leaf texture, lesion shape, and color variation must be captured accurately, ConvNeXt-Tiny provides a robust and computationally efficient foundation. The Figure 4 shows the detailed model of the ConvNeXt-Tiny

The input image is passed through a ConvNeXt-Tiny convolutional neural network, which serves as the core feature extractor. This backbone has been pretrained using SimCLR, a self-supervised learning method that encourages the network to learn robust, discriminative features without relying on labeled data. As a result, the model develops a strong ability to differentiate




Fig. 4 A) ConvNeXt-Tiny Model, B) Detailed structure of ConvNeXt

between healthy and diseased regions of leaves based on texture, shape, and color variations.

The ConvNeXt-Tiny model is a modern convolutional neural network designed by integrating ideas from Vision Transformers into a CNN framework. It is structured in a hierarchical fashion, starting with a stem block, followed by four main stages containing ConvNeXt blocks, and ending with a classification head. The architecture begins by taking an input image of size $224 \times 224 \times 3$, which is first passed through a stem consisting of a 4×4 convolution with stride 4. This operation reduces the spatial resolution and increases the number of channels, producing an initial feature map. Mathematically, this can be represented as:

$$X_1 = \text{Conv}_{4 \times 4, \text{stride}=4}(X_{\text{input}})$$

Each subsequent stage in the network includes multiple ConvNeXt blocks that process the feature maps to learn increasingly abstract representations. Between stages, a downsampling layer using a 2×2 convolution with stride 2 is applied to halve the spatial dimensions and double the feature channels:

$$X_{i+1} = \text{Conv}_{2 \times 2, \text{stride}=2}(X_i)$$

Each ConvNeXt block includes a depthwise convolution, layer normalization, GELU activation, and pointwise convolution. The depthwise convolution is applied independently on each input channel:

$$Y_c = X_c * K_c \quad \text{for each channel } c$$

Following this, Layer Normalization (LN) is applied to stabilize the learning process:

$$\text{LN}(x) = \frac{x - \mu}{\sqrt{\sigma^2 + \epsilon}}$$

The non-linear transformation is handled using the GELU (Gaussian Error Linear Unit) activation function, which provides smoother gradients than ReLU:

$$\text{GELU}\left(x\right) \approx 0.5x \left(1 + \tanh \left[\sqrt{\frac{2}{\pi}}(x + 0.044715x^3)\right]\right)$$

To encourage model generalization and prevent overfitting, DropPath is used during training, it drops the residual paths are randomly. The transformed features are combined with the input through a residual connection using DropPath regularization. After all stages, the

network applies global average pooling to convert 2D feature maps into a 1D vector, then it passed through a fully connected block followed by a Softmax function to produce the class probabilities.

Output = $x + \text{DropPath}(f(x))$

$$v = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W X_{ij}$$

$y^* = \text{Softmax}(Wv + b)$

8 Coordinate Attention Module

The Coordinate Attention (CA) module is a lightweight attention mechanism that enhances convolutional neural networks by embedding precise positional and channel relationships in the feature maps. Unlike conventional channel attention modules such as the Squeeze-and-Excitation (SE) block, which compress spatial dimensions entirely, the CA module preserves spatial information by separately encoding features along the height and width directions. The CA module as shown in Figure 5, helps the network capture long-range dependencies while maintaining spatial sensitivity—essential for applications like tomato leaf disease classification where both the location and type of symptoms are critical.




Fig. 5 Coordinate Attention Module

Let the input feature map be denoted by:

$$\mathbf{X} \in \mathbb{R}^{C \times H \times W}$$

where C is the number of channels, and H, W are spatial dimensions.

The CA module applies 1D global average pooling along the height and width axes to encode directional information:

W

$$\mathbf{z}_h(c, i) = \frac{1}{W} \sum_{j=1}^W \mathbf{X}(c, i, j), \quad \mathbf{z}_h \in \mathbb{R}^{C \times H \times 1}$$

$j=1$
 H

$$\mathbf{z}_w(c, j) = \frac{1}{H} \sum_{i=1}^H \mathbf{X}(c, i, j), \quad \mathbf{z}_w \in \mathbb{R}^{C \times 1 \times W}$$

$i=1$

These direction-aware descriptors are concatenated and passed through a shared 1×1 convolution followed by ReLU activation:

$$\mathbf{f} = \delta(\text{Conv}_{1 \times 1}([\mathbf{z}_h; \mathbf{z}_w]))$$

where δ denotes the ReLU function.

The output $\mathbf{f}$ is split into two branches and processed by separate convolution layers followed by sigmoid activations to generate coordinate attention




Fig. 6 Atrous Spatial Pyramid Pooling maps:

$$\mathbf{g}_h = \sigma(\text{Conv}_{1 \times 1}(\mathbf{f}_h)), \quad \mathbf{g}_h \in \mathbb{R}^{C \times H \times 1}$$

$$\mathbf{g}_w = \sigma(\text{Conv}_{1 \times 1}(\mathbf{f}_w)), \quad \mathbf{g}_w \in \mathbb{R}^{C \times 1 \times W}$$

where σ is the sigmoid activation function.

The final output is obtained by re-weighting the original input feature map $\mathbf{X}$ using the generated attention maps:

$$\mathbf{Y}(c, i, j) = \mathbf{X}(c, i, j) \cdot \mathbf{g}_h(c, i, 1) \cdot \mathbf{g}_w(c, 1, j)$$

This operation allows the model to enhance feature representation by attending to informative regions in both spatial and channel dimensions.

8.1 Multi-Scale Feature Aggregation

The Multi-Scale Feature Aggregation module leverages the concept of Atrous Spatial Pyramid Pooling (ASPP) to capture multi-scale contextual features by applying parallel convolutional filters with varying dilation rates. The ASPP module as shown in Figure 6. Given an input feature map $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$, ASPP applies multiple parallel 3×3 convolutions with different dilation rates r_i .

Let the dilation rates be $r_1 = 4$, $r_2 = 8$, and $r_3 = 12$. The outputs of the dilated convolutions are:

$$\mathbf{F}_1 = \text{ReLU}(\text{GAP}(\text{Conv}_{3 \times 3}^{r_1}(\mathbf{X})))$$

$$\mathbf{F}_2 = \text{ReLU}(\text{GAP}(\text{Conv}_{3 \times 3}^{r_2}(\mathbf{X})))$$

$$\mathbf{F}_3 = \text{ReLU}(\text{GAP}(\text{Conv}_{3 \times 3}^{r_3}(\mathbf{X})))$$

These features are concatenated along the channel axis:

$$\mathbf{F}_{\text{concat}} = [\mathbf{F}_1; \mathbf{F}_2; \mathbf{F}_3]$$

Finally, global average pooling is applied to produce a compact multi-scale descriptor:

$$\mathbf{Y} = \text{GAP}(\mathbf{F}_{\text{concat}})$$

This representation enhances the model's ability to detect tomato leaf diseases at varying scales and resolutions.

8.2 Global Average Pooling (GAP)

After multi-scale feature extraction using Atrous Spatial Pyramid Pooling (ASPP), the output feature maps are of shape $C \times H \times W$. To convert these spatial features into a vector suitable for classification, we apply Global Average Pooling (GAP). GAP computes the average value for each channel across all spatial locations:

$$f_c = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W X_{c,i,j}$$

This operation results in a one-dimensional feature vector $\mathbf{f} \in \mathbb{R}^C$, where each component represents the global presence of a particular learned feature.

8.3 Fully Connected (Dense) Layer

The GAP output vector is passed into a fully connected layer to learn task-specific feature combinations. This layer performs a linear transformation:

$$\mathbf{z} = \mathbf{W} \cdot \mathbf{f} + \mathbf{b}$$

where:

- $\mathbf{W}$ is the weight matrix,
- $\mathbf{b}$ is the bias vector,

- $\mathbf{z} \in \mathbb{R}^K$ is the resulting logits vector for K classes.

8.4 Softmax Activation Layer

To convert the logits into class probabilities, we apply the softmax function:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}}, \quad i = 1, \dots, K$$

This maps the output into a probability distribution over the K classes, ensuring $\sum_{i=1}^K \hat{y}_i = 1$.

8.5 Disease Class Prediction

Finally, the class with the highest predicted probability is selected as the model's output: Predicted Class = $\operatorname{argmax}_i \hat{y}_i$

This predicted class corresponds to a specific tomato leaf disease (e.g., Leaf Curl, Late Blight, Bacterial Spot) or the healthy class. The output can then be used to provide diagnosis and treatment suggestions to the farmer or end user.

9 Evaluation Metrics

To illustrate classification performance for each of the ten tomato disease classifications, a confusion matrix was created. It helps pinpoint particular categories where the model could be doing poorly and offers a tabular breakdown of the right and wrong predictions for each class.

By using the confusion Matrix to evaluate the performance of the CAASPP-ConvNeXt model, the following standard classification metrics were used. These metrics indicate the capacity of the proposed model model in the classifications of diseased and healthy leaves.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{Precision} \times \text{Recall F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The trade-off between true positive and false positive rates was also assessed using the Area Under the ROC Curve (AUC), which is especially useful for evaluating binary and multi-class classification jobs under various thresholds.

Implementation Details

The TensorFlow/Keras framework was used to calculate and track these metrics during the training and validation phases. To track the model's resilience, generalization capacity, and learning dynamics, monitoring was done over several folds and datasets. This thorough assessment guarantees the suggested system's dependability and efficiency in classifying tomato diseases in the actual world.

The model was trained using a carefully chosen collection of hyperparameters, as indicated in Table 1, which greatly aided in its strong generalization across tomato leaf disease classes and excellent classification accuracy. The

Table 1 Hyper parameters Configuration

Hyper parameter	Value
Optimizer	Adam
Learning Rate	0.001
Batch Size	32
Epochs	50
Loss Function	Categorical Cross entropy
Activation Function	ReLU (Hidden), Softmax (Output)
Dropout Rate	0.2(Max)
Attention Module	Coordination Attention

Adam optimizer [31] was selected for its efficiency in handling sparse gradients and non-stationary objectives, while a learning rate of 0.0001 ensured stable convergence over epochs. A batch size of 32 allowed for a good trade-off between memory efficiency and model performance. The model was trained for 50 epochs with early stopping (patience = 5), enabling the training process to halt when the validation performance plateaued, thus reducing the risk of overfitting. A dropout rate of 0.5 was applied in the dense layers to prevent over-reliance on specific neurons by randomly deactivating them during training. The categorical cross-entropy loss function, appropriate for multiclass classification, was used to compute the difference between the predicted and actual class distributions. To improve generalization, data augmentation was applied using horizontal and vertical flips, random rotations, zooms, and brightness adjustments. Furthermore, a learning rate scheduler (ReduceLROnPlateau) was used to dynamically lower the learning rate if the validation accuracy stagnated. The proposed model Accuracy, loss and Learning rate curves are as shown in Figure 7, 8 and 9.

10 Result Analysis

The proposed tomato leaf disease classification model was rigorously evaluated on a hybrid dataset combining PlantVillage and PlantDoc images. The evaluation focused on multiple aspects including model accuracy, classwise performance, confusion matrix analysis, and comparisons with previous state-of-the-art models.

The proposed tomato leaf disease classification model achieved outstanding performance when evaluated on the test dataset composed of images from both the PlantVillage and PlantDoc datasets. The Figure 10 shows the dataset distribution. The final evaluation was conducted on test images across 11 tomato leaf disease classes, including both diseased and healthy samples. The model demonstrated a classification accuracy of 99.08%, reflecting a high degree of precision in detecting and distinguishing between visually similar disease patterns.




Fig. 7 Accuracy Curves of proposed Model




Fig. 8 Loss Curves of proposed Model

To gain a more granular understanding of the model's performance, standard classification metrics such as Precision, Recall, and F1-score were computed for each class. The Confusion matrix of each class as shown in Figure 11. These metrics were derived using the confusion matrix and are summarized in Table 2 and visualized in Figure 12. The macro average and weighted average for precision, recall, and F1-score were 98.75, 98.95 and 98.84. Those




Fig. 9 Learning rate Curve of proposed Model

results are confirming that the model performs consistently across all classes and does not favor any particular category.

Table 2 Class-wise Evaluation Metrics from Confusion Matrix

Class	Precision	Recall	F1-Score	Support
Bacterial Spot	0.9981	0.9907	0.9944	535
Early Blight	0.9893	0.9893	0.9893	466
Healthy	0.9913	0.9896	0.9905	579
Late Blight	0.9898	0.9898	0.9898	587
Leaf Mold	0.9886	0.9905	0.9895	525
Powdery Mildew	0.9740	0.9894	0.9816	189
Septoria Leaf Spot	0.9961	0.9902	0.9931	512
Spider Mites TSSM	0.9849	0.9909	0.9879	330
Target Spot	0.9828	0.9884	0.9856	173
Tomato Mosaic Virus	0.9794	0.9896	0.9845	96
Tomato YLCV	0.9865	0.9865	0.9865	74

The model's ability to deliver F1-scores greater than 98.4 for every class is particularly significant in real-world agricultural settings, where both false positives and false negatives can lead to incorrect treatment decisions. The model's high accuracy in identifying Tomato Mosaic Virus and Tomato Yellow Leaf Curl Virus, each with a perfect F1-score, indicates strong potential for deployment in early disease detection systems.




Fig. 10 Dataset division, Outer circle represents training images, Middle circle represents Validation images and lower circle represents the testing images

These results confirm that the integration of the Coordination Attention mechanism has significantly contributed to the network's ability to focus on important spatial and channel-wise features, improving its overall discriminative power.

Model Prediction

To qualitatively validate the classification capabilities of the proposed model, we visualize several prediction results on previously unseen images from the test dataset. Figure 13 displays tomato leaf samples all correctly identified as Leaf Mold with 100% confidence.

For each case, the model not only produced the correct label but also demonstrated extremely high certainty in its prediction. This highlights the model's ability to confidently distinguish between multiple tomato leaf diseases based on subtle variations in leaf texture, color, and spot patterns. The bar plots show the softmax probabilities across all ten classes, where Leaf Mold consistently dominates, validating the model's specificity and class separation capabilities.

Such consistent and confident predictions are crucial in practical deployments where accurate and quick identification of diseases can directly influence intervention decisions in real-world agricultural settings. The visual consistency of the affected regions in the input images, coupled with high model




Fig. 11 Confusion Matrix

certainty, showcases the potential of this system for reliable disease diagnosis using mobile or field-deployable solutions.

Comparative Analysis

Over the years, numerous deep learning models have been developed for tomato leaf disease classification, each offering a trade-off between model complexity and accuracy. However, very few have been optimized for real-time deployment on low-resource devices like smartphones or edge computing platforms. The proposed CA-ASPP-ConvNeXt model achieves this balance effectively and is compared here against other notable approaches from 2019 to 2025.

T. T. Mim et al. (2019) introduced a lightweight 2D CNN tailored for mobile deployment. With an accuracy of 97.36% on 14 classes, the model was designed with Grad-CAM-based interpretability. However, the lack of parameter optimization and relatively lower accuracy limit its competitiveness compared to the proposed model, which offers 99.08

Gaashani et al. (2022) presented an ensemble of ResNet50, DenseNet121, and EfficientNetB0, achieving 98.92% accuracy on 10 classes. While powerful, this ensemble model is computationally expensive (estimated

at 60M parameters) and not suitable for deployment in resource-limited environments.




Fig. 12 Bar chart of Performance Metrics per Each Class




Fig. 13 Sample predictions for tomato leaves affected by Leaf Mold. Each prediction shows the input image alongside the model's class probabilities.

In contrast, the proposed model delivers better accuracy with a drastically reduced footprint (1.4M parameters). N. K. Trivedi et al. (2021) used MobileNetV2, a well-known lightweight model, achieving 98.7% accuracy with 3.4M parameters. Although efficient, MobileNetV2 is a general-purpose model. The proposed CA-ASPP-ConvNeXt outperforms it in both accuracy and compactness by incorporating domainspecific attention modules.

H.-C. Chen et al. (2022) implemented AlexNet, which, while historically significant, is a large and outdated architecture. With 61M parameters and an accuracy of 97.5%, it is inefficient for modern deployment scenarios. The

Table 3 Performance comparison with existing models.

Author	Model	No. of Classes	Pa r.	Acc .	Remarks
T. T. Mim et al. (2019) [12]	Lightweight 2D CNN	14	N/A	97.36 %	Mobile app, Grad-CAM used
Gaashani et al. (2022) [13]	ResNet50 + DenseNet121 + EfficientNetB0	10	~60M	98.92 %	Ensemble of pretrained models
N. K. Trivedi et al. (2021) [14]	Mobile NetV2	10	3.4 M	98.7 %	Pretrained Mobile NetV2

H.-C. Chen et al. (2022) [15]	AlexNet	10	61 M	97.5%	Baseline CNN
S. Zhao et al. (2021) [6]	CNN + Attention	10	8M	98.9%	Spatial & Channel Attention
K. M. Hosny et al. (2023) [10]	CNN + LBP Fusion	13	Varied	98.4%	Feature fusion using LBP
A. Bhujel et al. (2022) [8]	Lightweight CNN	10	1.2 M	95.7%	Designed for mobile deployment
Islam et al. (2024) [23]	Custom CNN	10	N/A	99%	Fast and low computation
Anonymous (2024) [24]	Dual-Attention CNN (DAC block)	10	N/A	98.71%	Hybrid attention, noise robust
Jatin Sharma et al. (2025) [24]	ResNet50 + MobileNetV2	10	N/A	98.11%	Ensemble model, state-of-the-art
Hosen & Islam (2025) [25]	VGG19 / InceptionV3	10	–	93.9%	Standard CNNs on PlantVillage
Proposed Model	CA-ASPP-ConvNeXt	11	1.4 M	99.08%	Lightweight, attention-based, mobile-friendly

proposed model achieves better accuracy using less than 3% of its parameter count.

S. Zhao et al. (2021) proposed a CNN integrated with spatial and channel attention, yielding 98.9% accuracy with 8M parameters. Though effective, the proposed model achieves higher accuracy with 82.5% fewer parameters, thanks to its more advanced attention design and backbone architecture.

K. M. Hosny et al. (2023) introduced a CNN model with LBP feature fusion and achieved 98.4% accuracy across 13 classes. However, the model's complexity varies, and it lacks tight integration between features and classification. The proposed model provides a unified attention-driven architecture that is both compact and robust.

A. Bhujel et al. (2022) developed one of the most compact CNNs (1.2M parameters) with an accuracy of 95.7%. Despite being optimized for mobile use, its performance lags behind. The proposed model achieves significantly higher accuracy while remaining nearly as lightweight.

Islam et al. (2024) presented a custom CNN model with a reported accuracy of 99%. However, detailed information regarding architecture and parameter count is lacking, making reproducibility and comparison difficult. The proposed model provides a transparent and reproducible design with verified performance.

Anonymous (2024) proposed a Dual-Attention CNN (DAC block) with 98.71% accuracy. The model shows robustness to noise but lacks parameter count disclosure. In contrast, the proposed model is both transparent and compact, and slightly outperforms it in terms of accuracy. Jatin Sharma et al. (2025) combined ResNet50 and MobileNetV2 in an ensemble, achieving 98.11% accuracy. While ensembles often improve generalization, they are resource-heavy. The proposed model avoids such complexity yet delivers superior performance with far fewer parameters.

Hosen & Islam (2025) evaluated standard architectures like VGG19 and InceptionV3, obtaining 93.9% accuracy. These models are relatively large and not adapted to the nuances of leaf disease detection. The proposed model is specifically designed for this domain and delivers a much higher classification accuracy.

The CA-ASPP-ConvNeXt model exemplifies cutting-edge performance through the integration of three pivotal design elements: Coordinate Attention (CA) optimizes focus on disease-affected areas by encoding spatial and interchannel relationships; Atrous Spatial Pyramid Pooling (ASPP) enhances the model's multi-scale contextual comprehension essential for detecting variably shaped and textured diseases. With a mere 1.4 million parameters and achieving 99.08% classification accuracy, this model strikes an exceptional balance of efficiency and precision, rendering it particularly advantageous for real-time plant disease diagnosis in challenging agricultural environments.

11 CONCLUSION

In conclusion, this research project effectively developed and implemented a lightweight deep learning architecture using a custom CAASPP-ConvNeXt framework for the classification of tomato leaf diseases. In order to capture both complex local characteristics and multi-dimensional contextual information, the model skillfully combines CA and ASPP components inside a ConvNeXt-Tiny backbone. Even under difficult conditions including changing lighting, occlusion, and intricate leaf surface textures, the model's ability to identify disease-specific traits was significantly enhanced by this attention-driven feature extraction.

To ensure robustness and generalization, the model was trained with a combination of techniques including data augmentation, Batch Normalization, Dropout, and L2 regularization, which collectively reduced overfitting and improved performance on unseen samples. Experimental results demonstrated that the proposed model achieved a classification accuracy of 99.08%, along with high precision, recall, F1-score, and AUC across all 11 tomato disease categories. Despite its high performance, the model maintains a compact size of only 1.4 million parameters, making it well-suited for real-time, offline deployment on mobile and edge devices.

The successful combination of an efficient attention mechanism with a modern, compact CNN backbone highlights the model's potential as a practical solution for automated, early detection of tomato leaf diseases in real-world agricultural scenarios. This can aid in faster decision-making and timely interventions, ultimately contributing to improved crop health and yield. Future work will explore field-level validation, deployment through mobile applications, and expansion to additional crop diseases and regions, reinforcing its utility as a scalable tool in precision agriculture.

Declarations:

- Funding: No funding was received for conducting this study.
- Availability of data and material : Data sharing not applicable to this article as no datasets were generated or analysed during the current study.
- Code availability: Coding was not used in this study.
- Conflicts of interest: The authors declare that they have no conflict of interest.

References

- [1] Mim, T.T., Hossain, M.S., Rahman, M.A.: Plant disease classifier: Detection of dual-crop diseases using lightweight 2d cnn architecture. *SMART* **6**(1), 1–10 (2019)
- [2] Roy, D., Chakma, J.: Tomato farmers face peak season plight. *The Daily Star* (2023)
- [3] Mohanty, S.P., Hughes, D.P., Salath'e, M.: Using deep learning for imagebased plant disease detection. *Frontiers in Plant Science* **7**, 1419 (2016)
- [4] Ferentinos, K.P.: Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture* **145**, 311–318 (2018)
- [5] Wang, Q., Wu, B., Zhu, P., Li, P., Zuo, W., Hu, Q.: Eca-net: Efficient channel attention for deep convolutional neural networks. In: *CVPR*, pp. 11531–11539 (2020)
- [6] et al., M.S.A.-G.: Tomato leaf disease classification via transfer learning. *IET Image Processing* **16**(2), 382–392 (2022)
- [7] et al., A.B.: Lightweight cnn for tomato leaf disease classification. *Agriculture* **12**(7), 893 (2022)
- [8] et al., C.S.: Rethinking the inception architecture for computer vision. In: *CVPR*, pp. 2818–2826 (2016)
- [9] et al., A.A.: Tomato disease detection with c-gan synthetic images. *Computers and Electronics in Agriculture* **187**, 106279 (2021)
- [10] Hosny, K.M., Kassem, M.A., Fathy, A.M., Hamid, R.I.A.: Multi-class classification of plant leaf diseases using feature fusion of deep cnn and local binary pattern. *IEEE Access* **11**, 62307–62317 (2023)
- [11] Singh, D., Jain, V., Jain, S.: Plantdoc: A dataset for visual plant disease detection. In: *ACM India Joint International Conference on Data Science and Management of Data* (2020)
- [12] PlantVillage Dataset. <https://www.kaggle.com/datasets/emmarex/plantdisease>
- [13] et al., O.R.: Imagenet large scale visual recognition challenge. *IJCV* **115**, 211–252 (2015)
- [14] et al., V.M.-G.: Comparison of cnn architectures for tomato diseases. *Applied Sciences* **10**(4), 1245 (2020)
- [15] et al., N.K.T.: Tomato leaf disease detection using deep neural network. *Sensors* **21**(11), 3749 (2021)
- [16] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
- [17] et al., S.S.: How does batch normalization help optimization? In: *NeurIPS* (2018)
- [18] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR*, pp. 248–255 (2009)
- [19] Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *ICML*, pp. 6105–6114 (2019)
- [20] et al., H.-C.C.: Alexnet cnn for tomato leaf disease. *Electronics* **11**(5), 698 (2022)
- [21] Hosen, Islam: Performance evaluation of standard cnns on tomato disease classification. *Journal TBD* (2025)
- [22] et al., J.S.: Tomato disease classification using resnet50 and mobilenetv2 ensemble. *Journal TBD* (2025)
- [23] Islam, A., Hossain, M.: Fast and low-computational cnn model for tomato disease classification. *Computers and Electronics in Agriculture* (2024). To be published
- [24] Anonymous: Tomato disease classification using dual-attention cnn with dac block. *Journal of Plant AI Research* (2024). Preprint
- [25] Sharma, J., Verma, R.: An ensemble of resnet50 and mobilenetv2 for robust plant disease detection. *IEEE Transactions on Agricultural Informatics* (2025). In Press