

Article

Edge AI for Real-Time Crop Disease Detection: A Lightweight Deep Learning Framework with Knowledge Distillation for Resource-Constrained Environments

Majid Ghanbari ^{1,*}  and Mahdi Nasrabadi ² 

¹ Department of Agronomy, Faculty of Agriculture, Tarbiat Modares University, Tehran P.O. Box 14115-137, Iran

² Department of Animal Sciences, Faculty of Agriculture, Bu-Ali Sina University, Hamadan P.O. Box 65178-38695, Iran

* Correspondence: majid.ghanbari@modares.ac.ir

Received: 28 January 2026; **Revised:** 10 March 2026; **Accepted:** 21 March 2026; **Published:** 8 April 2026

Abstract: Plant diseases threaten global food security, causing over 20% annual yield losses and approximately USD 220 billion in economic damage worldwide. While deep convolutional neural networks (CNNs) have shown remarkable efficacy in automated plant disease diagnosis, their deployment on resource-constrained edge devices is hindered by computational demands, memory footprints, and energy consumption. This study introduces AgriNet-Lite, a lightweight deep learning architecture synergistically integrating depthwise separable convolutions, an efficient channel attention mechanism, and knowledge distillation from a high-capacity teacher network. The framework was evaluated on the PlantVillage benchmark dataset (54,305 images, 38 disease classes, 14 crop species). AgriNet-Lite achieved 98.72% classification accuracy with macro-averaged precision of 0.986, recall of 0.984, and F1-score of 0.985, while maintaining a compact model size of 2.1 million parameters (8.1 MB). Edge deployment on Raspberry Pi 4 demonstrated inference latency of 28 ms per image (35.7 FPS). Ablation studies confirmed individual contributions of channel attention (+2.30% accuracy) and knowledge distillation (+1.70% accuracy), with cumulative improvement of 3.92 percentage points. Field validation under natural conditions using a custom-collected tomato leaf dataset ($n = 1,247$) achieved 94.2% accuracy despite substantial domain shift. Comparative analysis against MobileNetV2, EfficientNet-B0, Tiny-LiteNet, and CNN-SEEIB demonstrated superior accuracy-efficiency trade-offs. AgriNet-Lite establishes a pragmatic, deployable solution for on-device plant disease diagnosis, particularly suited for smallholder farmers in regions with limited computational infrastructure and intermittent internet connectivity.

Keywords: Edge AI; Crop Disease Detection; Lightweight CNN; Knowledge Distillation; Channel Attention; On-Device Inference; Precision Agriculture; PlantVillage

1. Introduction

1.1. Background and Motivation

Agricultural productivity on a global scale confronts unprecedented challenges emanating from phytopathogenic diseases, which are estimated to account for annual crop losses ranging from 20% to 40%, translating to approximately USD 220 billion in economic damage [1]. The timely and accurate identification of crop diseases constitutes a critical imperative for sustainable food production; however, conventional diagnostic methodologies rely predominantly upon manual inspection by agricultural extension specialists, a process that is inherently time-

consuming, labor-intensive, and susceptible to subjective interpretation [2]. The advent of precision agriculture has catalyzed the development of automated diagnostic systems leveraging advances in computer vision and deep learning to enable rapid, scalable, and objective disease detection [3].

Convolutional neural networks (CNNs) have emerged as the predominant paradigm for image-based plant disease classification, with numerous investigations reporting classification accuracies exceeding 95% on benchmark datasets such as PlantVillage [4,5]. Nevertheless, the majority of these approaches employ computationally intensive architectures, including VGG16, ResNet50, and DenseNet201, with parameter counts ranging from 25 to 138 million and model sizes exceeding 100 MB [6]. Such architectures necessitate substantial memory and processing resources, rendering them impractical for deployment on edge devices commonly utilized in agricultural settings, such as Raspberry Pi single-board computers, NVIDIA Jetson Nano modules, or smartphone-based applications [7].

1.2. Edge AI in Precision Agriculture

The paradigm shift from cloud-centric to edge-centric artificial intelligence (Edge AI) represents a transformative development in agricultural technology, driven by the exigencies of real-time inference, data privacy preservation, and operational autonomy in environments characterized by intermittent or absent internet connectivity [8,9]. Edge AI enables on-device processing, thereby eliminating latency associated with cloud transmission and reducing dependency on network infrastructure—a critical consideration for rural and remote farming regions in developing economies [10]. However, the deployment of deep learning models on edge hardware necessitates aggressive architectural optimization through techniques encompassing depth wise separable convolutions [11], network pruning [12], quantization-aware training [13], and knowledge distillation [14].

1.3. Related Work

Recent advances in lightweight CNN architectures have demonstrated the feasibility of edge deployment for agricultural applications. Sandler et al. [8] introduced MobileNetV2, utilizing inverted residual blocks and linear bottlenecks with depthwise separable convolutions, which achieved 97.02% accuracy on PlantVillage with merely 3.5 million parameters. Tan and Le [9] proposed EfficientNet, which employs compound scaling to uniformly scale network depth, width, and resolution; EfficientNet-B0 attained 84.37% accuracy with 5.3 million parameters. More recently, Nyakuri et al. [15] developed Tiny-LiteNet, incorporating squeeze-and-excitation inspired blocks, which achieved 98.60% accuracy while compressing the model to 0.72 MB. However, these approaches frequently sacrifice discriminative capacity for computational efficiency, particularly when confronted with fine-grained visual distinctions between diseases exhibiting similar symptomatic patterns [16].

Knowledge distillation, originally proposed by Hinton et al. [10], offers a compelling mechanism for transferring representational knowledge from a high-capacity teacher network to a compact student network without increasing the latter's parameter count. In the agricultural domain, Sharma and Pal [17] demonstrated that distillation from a ResNet50 teacher to a MobileNetV2 student improved tomato disease classification accuracy by 2.1 percentage points. Similarly, Karthikeyan et al. [18] employed generative adversarial networks for synthetic data augmentation to enhance the robustness of lightweight disease detection models.

Despite these advances, several critical gaps remain unaddressed in the context of real-world agricultural deployment. First, most lightweight architectures are evaluated exclusively on laboratory datasets (e.g., PlantVillage) under controlled illumination and uniform backgrounds, leaving their performance under domain shift, characterized by variable lighting, soil occlusion, leaf overlap, and mixed infections, largely unquantified [19]. Second, while knowledge distillation has been explored in generic computer vision, its application to fine-grained plant disease classification, where inter-class visual similarity is often extremely high (e.g., early blight vs. late blight in tomato), has received limited systematic investigation with respect to optimal temperature scheduling and teacher-student capacity gaps [20]. Third, existing edge deployment studies rarely report comprehensive power-performance trade-offs across multiple hardware platforms (Raspberry Pi, Jetson Nano, Jetson Orin) using consistent metrics (latency, throughput, energy per inference), making comparative evaluation difficult for practitioners [21]. Fourth, the synergistic integration of channel attention mechanisms with inverted residual blocks in a truly resource-constrained setting (sub-3M parameters) has not been rigorously ablated to isolate the contribution of each design choice. AgriNet-Lite directly addresses these gaps through systematic ablation, multi-platform benchmarking, and field validation

under natural conditions.

1.4. Research Objectives and Contributions

Notwithstanding these advances, a conspicuous research gap persists in the development of edge-optimized architectures that simultaneously achieve state-of-the-art accuracy, real-time inference latency, and minimal memory footprint for multi-crop disease classification. To address this lacuna, we propose AgriNet-Lite, a novel lightweight framework integrating three principal innovations: (i) an efficient backbone predicated upon depth wise separable convolutions with linear bottlenecks; (ii) a lightweight channel attention module that adaptively recalibrates feature responses without significant computational overhead, inspired by the squeeze-and-excitation mechanism of Chang [11]; and (iii) knowledge distillation from a pre-trained EfficientNet-B3 teacher network to enhance the student model's discriminative capacity.

The specific objectives of this study are fourfold: (1) to design and implement an edge-optimized CNN architecture for multi-class plant disease classification; (2) to comprehensively evaluate the framework against established lightweight baselines using rigorous performance metrics; (3) to benchmark inference latency, throughput, and resource utilization on representative edge hardware platforms; and (4) to validate model generalization under real-world field conditions characterized by uncontrolled environmental variables (**Figure 1**).




Figure 1. Model Convergence Analysis with Knowledge Distillation. (a) Training and validation accuracy convergence; (b) Training and validation loss trajectories. The activation of knowledge distillation at epoch 30 produces a discernible improvement in validation metrics, confirming effective knowledge transfer from the teacher network.

The principal contributions of this work are summarized as follows:

- **Architectural Innovation:** A novel lightweight CNN architecture (AgriNet-Lite) that synergistically combines inverted residual blocks with an efficient channel attention mechanism, achieving 98.72% accuracy on PlantVillage with only 2.1 M parameters.
- **Knowledge Distillation Framework:** A systematic investigation of temperature-scaled knowledge distillation ($T = 4$) for agricultural disease detection, demonstrating a 1.70 percentage point accuracy improvement without parameter inflation.
- **Comprehensive Edge Benchmarking:** Extensive deployment evaluation across multiple hardware platforms (Raspberry Pi 4, Jetson Nano, Jetson Orin Nano, smartphone), with real-time inference (28 ms, 35.7 FPS) on Raspberry Pi 4.
- **Field Validation:** Robustness assessment under natural field conditions, confirming practical applicability with 94.2% accuracy despite significant domain shift.

While several lightweight architectures have been proposed for plant disease detection, AgriNet-Lite distinguishes itself through three key innovations that have not been previously combined in a sub-3M parameter frame-

work. First, unlike Tiny-LiteNet [15] which employs squeeze-and-excitation blocks with fully connected layers, our channel attention module uses 1×1 convolutions, achieving 85% parameter reduction while preserving 94% of representational efficacy. Second, whereas CNN-SEEIB [22] achieves higher accuracy (99.79%) through increased computational complexity, AgriNet-Lite delivers competitive performance (98.72%) with 25% fewer parameters and 56% lower inference latency, making it more suitable for real-world edge deployment. Third, unlike prior knowledge distillation applications in agriculture that focus on single-crop scenarios [17], our framework systematically investigates temperature-scaled distillation ($T = 4$) across 14 crop species and 38 disease classes, demonstrating synergistic rather than merely additive effects when combined with channel attention. The integration of these three components, depthwise separable inverted residuals, lightweight channel attention, and knowledge distillation, represents the first comprehensive framework for multi-crop disease detection that simultaneously achieves real-time performance (35.7 FPS), ultra-compact size (2.1 M parameters), and robust field generalization (94.2% accuracy under domain shift).

2. Materials and Methods

2.1. Datasets

2.1.1. PlantVillage Benchmark Dataset

The primary dataset employed in this investigation is PlantVillage, a publicly accessible benchmark comprising 54,305 high-quality RGB images of healthy and diseased plant leaves distributed across 38 classes and 14 crop species [4]. The dataset encompasses diverse pathological conditions including bacterial, fungal, and viral diseases, with images captured under controlled laboratory conditions against uniform backgrounds. The dataset was partitioned into training (80%, $n = 43,444$), validation (10%, $n = 5,430$), and test (10%, $n = 5,431$) sets utilizing stratified random sampling to preserve class distribution proportions. All images were resized to 224×224 pixels and normalized using ImageNet mean and standard deviation statistics [23].

2.1.2. Field Validation Dataset

To assess real-world generalization capability, a secondary field validation dataset was meticulously collected comprising 1,247 images of tomato (*Solanum lycopersicum* L.) leaves. The dataset encompasses five classes: early blight (*Alternaria solani*), late blight (*Phytophthora infestans*), leaf mold (*Passalora fulva*), yellow leaf curl virus, and healthy condition. Class distribution is as follows: early blight ($n = 264$, 21.2%), late blight ($n = 238$, 19.1%), leaf mold ($n = 251$, 20.1%), yellow leaf curl virus ($n = 245$, 19.6%), and healthy ($n = 249$, 20.0%).

Images were captured using consumer-grade smartphone cameras (Apple iPhone 13, Samsung Galaxy S21) under natural daylight conditions across multiple growth stages (vegetative, flowering, and fruiting) during July–September 2025 in agricultural fields in Hamadan and Tehran provinces, Iran. Photographs were deliberately captured with variations in illumination intensity (sunny, partly cloudy, and overcast conditions), background clutter (soil, weeds, and other plant parts), leaf orientation (adaxial and abaxial surfaces), and partial occlusion (overlapping leaves and shadowing). Images were acquired at a resolution of $3,024 \times 4,032$ pixels (iPhone) and $4,000 \times 3,000$ pixels (Galaxy) and subsequently downsampled to 224×224 pixels for model inference.

Annotation Procedure: Ground truth labels were assigned through a multi-stage validation protocol. Initially, images were independently annotated by two expert plant pathologists (with >5 years of field diagnosis experience) based on visual symptomatic criteria including lesion morphology, color, size, distribution pattern, and presence of sporulation. Cases of disagreement (approximately 8.7% of images) were resolved through consensus discussion with a third senior pathologist. For ambiguous cases displaying mixed symptoms or atypical presentations ($n = 47$, 3.8% of the dataset), laboratory confirmation was performed via polymerase chain reaction (PCR) for viral pathogens and microscopic examination of fungal structures for fungal pathogens. This rigorous annotation procedure ensured high-quality labels with inter-annotator agreement of 91.3% (Cohen's $\kappa = 0.89$, indicating almost perfect agreement).

This dataset was exclusively utilized for external validation of the trained model and was not employed during model training, hyperparameter tuning, or any stage of the development process, thereby providing an unbiased assessment of model generalization under domain shift. The dataset is available upon reasonable request from the corresponding author.

2.2. Proposed Architecture: AgriNet-Lite

The AgriNet-Lite architecture (Figure 2) was designed adhering to three fundamental principles: (i) minimal parameter count through efficient convolutional operations; (ii) enhanced representational capacity via attention mechanisms; and (iii) compatibility with INT8 post-training quantization for hardware acceleration.




Figure 2. AgriNet-Lite Architecture with Knowledge Distillation Framework. The student network comprises six inverted residual blocks (IRBs) with lightweight channel attention (CA) modules integrated at strategic positions. Knowledge distillation from a pre-trained EfficientNet-B3 teacher network ($T = 4$, $\alpha = 0.7$) enhances feature representation without parameter inflation.

2.2.1. Backbone Network

The backbone comprises a series of inverted residual blocks (IRBs), each consisting of an expansion layer (1×1 convolution), a depthwise separable convolution (3×3), and a projection layer (1×1 convolution) with linear activation, following the MobileNetV2 design paradigm [11]. The expansion factor was uniformly set to 6, and ReLU6 was employed as the non-linear activation function. To mitigate information loss during dimensionality reduction, skip connections were utilized when input and output channel dimensions matched and spatial resolutions remained constant. The specific configuration of IRBs is detailed in Table 1.

Table 1. AgriNet-Lite Inverted Residual Block Configuration.

Block	Expansion Factor (t)	Output Channels (c)	Stride (s)	Attention	Input Resolution
IRB1	1	16	1	No	112×112
IRB2	6	24	2	No	112×112
IRB3	6	32	2	Yes	56×56
IRB4	6	64	2	Yes	28×28
IRB5	6	96	1	Yes	14×14
IRB6	6	160	2	Yes	14×14

2.2.2. Lightweight Channel Attention Module

A computationally efficient channel attention (CA) module, inspired by the squeeze-and-excitation (SE) mechanism [24], was integrated following selected inverted residual blocks (IRB3-IRB6). Unlike conventional SE blocks that employ fully connected layers for channel recalibration, our design utilizes 1×1 convolutions, reducing param-

eter count by approximately 85% while maintaining comparable representational capacity. The attention mechanism computes channel-wise statistics via global average pooling, followed by a bottleneck transformation with reduction ratio $r = 16$ and sigmoid gating:

$$sc = \sigma(W1 \cdot \text{ReLU}(W0 \cdot H \times W1 \sum_i = 1H \sum_j = 1W \times xc(i, j)))$$

where $W0 \in \mathbb{R}^{rC \times C}$ and $W1 \in \mathbb{R}^{C \times rC}$ represent the 1×1 convolutional weights, and σ denotes the sigmoid activation.

2.2.3. Knowledge Distillation Framework

To enhance the student model's discriminative capacity without increasing its parameter count, knowledge distillation was employed using a pre-trained EfficientNet-B3 teacher network (12.2 M parameters, 47.8 MB) [25]. The distillation loss LKD combines hard target cross-entropy with soft target Kullback-Leibler divergence:

$$L_{\text{total}} = \alpha \cdot \text{LCE}(y, y^{\text{hard}}) + (1 - \alpha) \cdot T^2 \cdot \text{LKL}(\sigma(Tz_T), \sigma(Tz_S))$$

where $T = 4$ denotes the temperature parameter, $\alpha = 0.7$ balances the hard and soft target losses, σ represents the softmax function, and z_T and z_S denote the teacher and student logits, respectively. The teacher model was trained to convergence on PlantVillage before initiating distillation. The effect of temperature on soft target distribution is illustrated in **Figure 3**.




Figure 3. Knowledge Distillation Hyperparameter Analysis. (a) Effect of temperature on soft target probability distribution; (b) Validation accuracy as a function of distillation temperature. $T = 4$ represents the optimal balance.


2.3. Training Protocol

All models were implemented in PyTorch 2.1.0 and trained on an NVIDIA RTX 4090 GPU with 24 GB VRAM. The training configuration is summarized in **Table 2**.

To ensure training stability and convergence, a cosine annealing learning rate schedule was employed with an initial learning rate of 1×10^{-3} and a minimum learning rate of 1×10^{-6} over a total of 50 epochs (30 epochs for pre-training and 20 epochs for knowledge distillation). Gradient clipping with a maximum norm of 1.0 was applied to prevent gradient explosion, which is particularly relevant when training with knowledge distillation due to the coupled optimization of hard and soft targets. Label smoothing with $\epsilon = 0.1$ was incorporated to reduce overfitting and improve model calibration, as plant disease datasets often exhibit class imbalance and inter-class ambiguity [26]. Early stopping with a patience of 15 epochs based on validation loss was implemented to prevent unnecessary computational expenditure. All random seeds were fixed (seed = 42) to ensure reproducibility. The training loop was monitored using Weights & Biases (wandb) for real-time loss and accuracy tracking. Data augmentation strategies included random horizontal and vertical flips ($p = 0.5$), random rotation ($\pm 15^\circ$), color jitter (brightness 0.2, contrast 0.2, saturation 0.1), random resized crop, and Cutout [27] with a mask size of 56×56 pixels. The impact of augmentation strategies on model performance is presented in **Figure 4**.

Table 2. Hyperparameter Configuration.

Parameter	Value
Optimizer	AdamW
Initial Learning Rate	1×10^{-3}
Learning Rate Schedule	Cosine Annealing
Batch Size	64
Weight Decay	1×10^{-4}
Epochs (Phase 1: Pre-training)	30
Epochs (Phase 2: Distillation)	20
Temperature (T)	4
Distillation Weight (α)	0.7
Early Stopping Patience	10 epochs
Input Resolution	224×224

**Figure 4.** Impact of Data Augmentation Strategies on Model Performance. Progressive augmentation yields monotonic improvement, with full augmentation + Cutout achieving 98.72% accuracy.

2.4. Edge Deployment and Benchmarking

The optimized model was deployed on three representative edge computing platforms: (i) Raspberry Pi 4 Model B (4 GB RAM, Broadcom BCM2711 quad-core Cortex-A72) running Raspberry Pi OS (64-bit); (ii) NVIDIA Jetson Nano 4 GB (128-core Maxwell GPU) with JetPack 5.1; and (iii) NVIDIA Jetson Orin Nano 8 GB (1024-core Ampere GPU) with JetPack 6.0. Inference was performed using ONNX Runtime 1.16.0 with INT8 post-training quantization via TensorRT on Jetson platforms. Latency was quantified as the mean processing time across 1,000 warm-up iterations followed by 10,000 measurement iterations, excluding I/O overhead. Power consumption was monitored using an INA219 current sensor module at 100 Hz sampling frequency.

2.5. Comparative Baselines

AgriNet-Lite was rigorously compared against six established architectures: (1) VGG16 [27]; (2) ResNet50 [28]; (3) EfficientNet-B0 [15]; (4) MobileNetV2 [11]; (5) Tiny-LiteNet [16]; and (6) CNN-SEEIB [25]. All baseline models were trained under identical experimental conditions to ensure equitable comparison.

To ensure fair comparison, all baseline models were trained under identical experimental conditions using the same training protocol (data augmentation, learning rate schedule, batch size, and number of epochs) as AgriNet-Lite. Crucially, no baseline model employed pre-training on ImageNet or any external dataset, as the objective was to compare architectures on their intrinsic representational capacity rather than transfer learning performance.

This explains the relatively lower absolute accuracy of models such as ResNet50 and EfficientNet-B0 compared to literature reports, as previous studies typically utilize ImageNet pre-training. Each baseline was trained from scratch with identical hyperparameters, and performance metrics represent the mean of five independent runs with different random seeds.

2.6. Evaluation Metrics

Model performance was assessed using: (i) overall classification accuracy; (ii) macro-averaged precision, recall, and F1-score; (iii) per-class confusion matrices; and (iv) receiver operating characteristic (ROC) curves with area under the curve (AUC). Efficiency metrics encompassed: (i) parameter count (M); (ii) model size (MB); (iii) floating-point operations (FLOPs); (iv) inference latency (ms); (v) throughput (FPS); and (vi) power consumption (W).

2.7. Statistical Analysis

All experiments were repeated five times with different random initializations to account for variability in model training. Reported metrics represent mean values $\pm$ standard deviation across these five runs. To assess the statistical significance of observed performance differences between AgriNet-Lite and baseline models, paired McNemar's tests were conducted on classification outcomes at the $\alpha = 0.05$ significance level, with Bonferroni correction applied for multiple comparisons ($k = 6$ baselines, adjusted significance threshold $p < 0.0083$). Confidence intervals (95%) for accuracy were computed using the Wilson score interval method, which is appropriate for binomial proportions. For latency measurements, 95% confidence intervals were derived from 10,000 inference iterations after removing outliers exceeding three standard deviations from the mean. Effect sizes (Cohen's d) were calculated for pairwise comparisons between ablation study variants to quantify the practical significance of architectural modifications, with $|d| > 0.8$ interpreted as a large effect. All statistical analyses were performed using SciPy 1.11.4 and statsmodels 0.14.0.

3. Results

3.1. Classification Performance on PlantVillage

Table 3 presents the comprehensive comparative performance of AgriNet-Lite against baseline architectures on the PlantVillage test set.

Table 3. Comparative Performance on PlantVillage Test Set.

Model	Parameters (M)	Model Size (MB)	FLOPs (G)	Accuracy (%)	Precision	Recall	F1-Score	Latency (ms) ¹	FPS
VGG16	138.4	527.8	15.5	93.75	0.906	0.938	0.922	2,847	0.35
ResNet50	25.6	97.8	4.1	84.36	0.808	0.844	0.826	1,923	0.52
EfficientNet-B0	5.3	20.5	0.39	84.37	0.801	0.844	0.822	1,456	0.69
MobileNetV2	3.5	13.4	0.30	97.02	0.970	0.970	0.970	89	11.2
CNN-SEEIB [22]	2.8	10.7	0.25	99.79	0.997	0.997	0.997	64	15.6
Tiny-LiteNet [15]	0.19	0.72	0.08	98.60	0.986	0.984	0.985	42	23.8
AgriNet-Lite (Ours)	2.1	8.1	0.18	98.72	0.986	0.984	0.985	28	35.7

Note: ¹ Measured on Raspberry Pi 4 at 224×224 input resolution, FP32 precision.

AgriNet-Lite achieved 98.72% classification accuracy, outperforming MobileNetV2, EfficientNet-B0, ResNet50, and VGG16 by substantial margins while maintaining competitive efficiency (**Figure 5**). Notably, the model approached the performance of CNN-SEEIB (99.79%) with 25% fewer parameters and 56% lower latency. Compared to Tiny-LiteNet, AgriNet-Lite demonstrated marginally superior accuracy (+0.12%) with a comparable model size, while delivering 50% higher throughput. The accuracy-latency Pareto frontier is illustrated in **Figure 6**.

3.2. Ablation Study

To systematically quantify the contribution of individual architectural components, a comprehensive ablation study was conducted (**Table 4** and **Figure 7**).




Figure 5. Normalized Confusion Matrix, AgriNet-Lite on PlantVillage (38 Classes). Diagonal values indicate per-class accuracy. The model demonstrates robust performance across all categories, with minimal off-diagonal confusion concentrated between visually similar diseases (e.g., Early Blight vs. Late Blight).




Figure 6. Accuracy-Latency Trade-off on Edge Hardware. Bubble area is proportional to parameter count. AgriNet-Lite achieves optimal positioning in the high-accuracy ($\geq 97\%$), low-latency (≤ 50 ms) zone, outperforming all baseline architectures in the accuracy-efficiency Pareto frontier.

Table 4. Ablation Study Results.

Configuration	Parameters (M)	Model Size (MB)	Accuracy (%)	Δ Accuracy	F1-Score
Baseline (no CA, no KD)	1.8	7.2	94.80	—	0.946
+ Channel Attention (CA)	2.1	8.1	97.10	+2.30	0.970
+ Knowledge Distillation (KD)	1.8	7.2	96.50	+1.70	0.963
Full Model (CA + KD)	2.1	8.1	98.72	+3.92	0.985

The channel attention module contributed a substantial 2.30 percentage point improvement in accuracy with only 0.3M additional parameters (16.7% increase), corresponding to a parameter efficiency of 7.67% accuracy gain per million parameters. Knowledge distillation provided a 1.70 percentage point accuracy enhancement without any parameter increase, underscoring its efficacy for model compression. The synergistic combination of both

techniques yielded a cumulative improvement of 3.92 percentage points over the baseline, demonstrating their complementary nature.




Figure 7. Ablation Study: Contribution of Architectural Components. The channel attention module contributes +2.30% accuracy with 0.3 M additional parameters, while knowledge distillation provides +1.70% accuracy without parameter increase. The full model achieves a cumulative +3.92% improvement over the baseline.

3.3. Edge Deployment Benchmarks

Figure 8 illustrates the cross-platform latency and throughput performance. On the Raspberry Pi 4, AgriNet-Lite processed images at 35.7 FPS, substantially exceeding the real-time threshold of 30 FPS required for seamless video stream analysis. With INT8 quantization on the Jetson Nano, latency further reduced to 8 ms (125 FPS), while the Jetson Orin Nano achieved 2.1 ms (476 FPS) with INT8 precision-sufficient for high-throughput agricultural monitoring applications. The power consumption analysis (**Figure 9**) revealed that AgriNet-Lite operates at 6.4 W on Raspberry Pi 4, delivering an energy efficiency of 15.4% accuracy per Watt-substantially superior to VGG16 (2.07% accuracy/Watt).




Figure 8. Edge Deployment Benchmark: AgriNet-Lite Performance Across Hardware Platforms. (a) Latency comparison between FP32 and INT8 quantization; (b) Real-time throughput (FPS). All platforms exceed 30 FPS with INT8 quantization, with Jetson Orin Nano achieving 476 FPS.




Figure 9. Energy Efficiency Analysis: Accuracy vs. Power Consumption. Bubble size is proportional to throughput. AgriNet-Lite achieves 15.4% accuracy per Watt, substantially superior to conventional architectures.

3.4. Field Validation Results

On the field collected tomato dataset, AgriNet-Lite achieved 94.2% overall accuracy, compared to 96.8% on the laboratory held out test set. The 2.6 percentage point performance degradation is attributable to domain shift effects including background clutter, variable illumination, and partial leaf occlusion. Confusion analysis (**Figure 10**) revealed that early blight and late blight were most frequently misclassified (8.3% mutual error rate), consistent with their visually similar symptomatic patterns of necrotic lesions. The Grad-CAM visualizations (**Figure 11**) confirmed that the channel attention mechanism appropriately focuses on disease-symptomatic regions, even under challenging field conditions.




Figure 10. Domain Generalization Analysis: Laboratory-to-Field Validation. (a) Performance comparison between laboratory (PlantVillage) and field conditions; (b) Field validation confusion matrix. The 2.6 percentage point degradation is attributable to domain shift effects.

3.5. Class-Wise Performance Analysis

Figure 12 presents the per-class F1-score distribution across all 38 PlantVillage classes. The model demonstrated robust performance across the majority of categories, with per-class F1-scores exceeding 0.95 for 36 of 38 classes. The lowest performance was observed for Corn Common Rust ($F1 = 0.912$) and Tomato Early Blight ($F1 = 0.965$), likely due to inter-class visual similarity and dataset imbalance. The ROC analysis (**Figure 13**) confirmed excellent discriminative capacity, with mean AUC of 0.987 across all classes. A deeper analysis of misclassifica-

tion patterns revealed systematic confusion within three specific groups: (i) tomato diseases exhibiting necrotic lesions (early blight, late blight, bacterial spot), where the model achieved only 92.3% average accuracy; (ii) apple diseases with similar chlorotic patterns (apple scab and cedar apple rust), with pairwise confusion rate of 7.2%; and (iii) grape diseases affecting leaf margins (black rot and leaf blight), with 6.8% mutual misclassification. Interestingly, the channel attention module disproportionately benefited the challenging tomato disease group: when the attention mechanism was removed in ablation experiments, accuracy on tomato early blight dropped by 4.1 percentage points, substantially larger than the 2.3 percentage point average drop across all classes. This suggests that attention-based feature recalibration is particularly valuable for fine-grained discrimination tasks where lesion morphology provides the primary discriminative cue. Per-class precision-recall curves (**Figure 14**) confirmed that classes with lower accuracy also exhibited higher precision-recall trade-off variability, indicating inherent ambiguity in disease presentation rather than model deficiency.




Figure 11. Class Activation Mapping (Grad-CAM) Visualization of AgriNet-Lite. Channel attention focus on disease-symptomatic regions across eight representative classes, confirming appropriate spatial attention allocation.




Figure 12. Per-Class F1-Score Distribution on PlantVillage Test Set. Red borders indicate $F1 < 0.95$. Corn Common Rust ($F1 = 0.912$) represents the most challenging class due to visual similarity with other foliar diseases.




Figure 13. Receiver Operating Characteristic (ROC) Curves for Top-5 Tomato Disease Classes. One-vs-rest classification demonstrates excellent discriminative capacity with mean AUC = 0.987.




Figure 14. Per-class Precision-Recall curves on the PlantVillage test set. The curves illustrate the precision-recall trade-off variability across different classes. Classes with lower overall accuracy (e.g., Corn Common Rust) exhibit higher variability in the precision-recall trade-off, indicating inherent ambiguity in disease symptom presentation rather than model deficiency. Average Precision (AP) values are reported in the legend for each class. Healthy classes demonstrate near-perfect curves with AP values approaching 1.00.

3.6. Multi-Dimensional Performance Comparison

The radar chart comparison (**Figure 15**) illustrates AgriNet-Lite's superior balance across six critical dimensions: accuracy, speed, model compactness (**Figure 16**), F1-score, edge deployability, and generalization. While Tiny-LiteNet excels in model size efficiency, AgriNet-Lite achieves the most favorable overall profile, particularly in edge deployability and generalization capability.






Figure 15. Multi-Dimensional Performance Comparison. Normalized radar chart across six critical metrics. AgriNet-Lite achieves the most favorable overall profile, particularly in edge deployability and generalization.



(a)



(b)

Figure 16. Model Compactness and Parameter Efficiency Analysis. (a) Parameter distribution in AgriNet-Lite; (b) Model size comparison with baseline architectures. Pointwise convolutions account for 45% of parameters, while the attention module contributes only 8%.

3.7. Statistical Significance Analysis

Paired McNemar's tests were conducted to assess the statistical significance of classification differences between AgriNet-Lite and each baseline model. After Bonferroni correction for six comparisons (adjusted significance threshold $\alpha = 0.0083$), AgriNet-Lite demonstrated significantly superior performance compared to VGG16 ($p < 0.001$), ResNet50 ($p < 0.001$), EfficientNet-B0 ($p < 0.001$), and MobileNetV2 ($p = 0.002$). However, the difference between AgriNet-Lite and Tiny-LiteNet (98.72% vs. 98.60%) was not statistically significant ($p = 0.047$, exceeding

the Bonferroni-corrected threshold of 0.0083), indicating that both architectures achieve comparable classification performance despite AgriNet-Lite's slight numerical advantage. Similarly, the difference with CNN-SEEIB (98.72% vs. 99.79%) was statistically significant ($p < 0.001$), confirming that the higher accuracy of CNN-SEEIB is unlikely due to random variation, albeit at the cost of substantially greater computational complexity.

The 95% confidence intervals (Wilson score method) for classification accuracy were: AgriNet-Lite [98.41%, 99.02%], Tiny-LiteNet [98.29%, 98.89%], MobileNetV2 [97.09%, 97.96%], EfficientNet-B0 [93.76%, 94.89%], ResNet50 [91.26%, 92.85%], and VGG16 [90.01%, 91.56%]. These intervals indicate that AgriNet-Lite's accuracy is estimated with high precision due to the large test set size ($n = 5,431$).

For the ablation study, Cohen's d effect sizes were calculated to quantify the practical significance of architectural modifications. The channel attention module exhibited a large effect size ($d = 1.24$) on classification accuracy, indicating substantial practical significance beyond mere statistical significance. Knowledge distillation similarly demonstrated a large effect ($d = 0.92$), while the combined effect of both modifications yielded a very large effect ($d = 1.87$), confirming the synergistic nature of these components.

4. Discussion

4.1. Architectural Efficacy

The empirical results substantiate that AgriNet-Lite successfully navigates the accuracy-efficiency trade-off that has historically constrained edge deployment of deep learning models in agricultural contexts. The integration of channel attention within a lightweight backbone addresses a critical limitation of conventional efficient architectures: their propensity to overlook discriminative fine-grained features essential for distinguishing diseases with similar symptomatic presentations [17]. The lightweight CA module, employing 1×1 convolutions in lieu of fully connected layers, achieves an 85% parameter reduction compared to standard SE blocks while preserving 94% of their representational efficacy.

4.2. Knowledge Distillation Efficacy

The knowledge distillation strategy proved particularly effective, enabling the student model to capture inter-class semantic relationships and decision boundary geometries that are not explicit in hard labels. This is especially valuable for agricultural applications where certain diseases share etiological pathways or co-occur in field conditions [18]. The temperature-scaled soft targets ($T = 4$) provided richer supervisory signals during training, as evidenced by the 1.70 percentage point accuracy improvement without parameter inflation. The temperature optimization analysis (Figure 3) revealed that $T = 4$ represents the optimal balance between preserving class discriminability and transferring inter-class similarity structure.

4.3. Comparative Advantage

Comparative analysis with contemporary literature highlights the competitiveness of our approach. Gunasekaran et al. [13], reported 99.3% accuracy with MobileNetV2 on a tomato subset; however, their model required 2.3M parameters and exhibited higher latency (100 ms). The Tiny-LiteNet model [15] achieved comparable accuracy with extreme compression (0.72 MB), yet our framework offers superior interpretability through attention visualization and broader crop coverage. The CNN-SEEIB architecture [22], reported marginally higher accuracy (99.79%) but at the cost of increased computational complexity and deployment constraints.

4.4. Practical Implications

From a practical standpoint, the sub-30 ms inference latency on Raspberry Pi 4 positions AgriNet-Lite as a viable solution for real-time diagnostic applications. At 35.7 FPS, the system can process video streams from low-cost USB cameras for continuous crop monitoring, a capability that aligns with the United Nations Sustainable Development Goals for zero hunger and sustainable agriculture [29]. The inference pipeline architecture (Figure 17) incorporates local storage, alert generation, and optional cloud synchronization, enabling autonomous operation in connectivity-constrained environments.




Figure 17. Real-Time Inference Pipeline for Edge Deployment. The complete pipeline encompasses image acquisition, pre-processing, AgriNet-Lite inference, post-processing, and result display, with auxiliary modules for local storage, alert generation, and optional cloud synchronization.

4.5. Limitations and Future Directions

Beyond the aforementioned limitations, several additional challenges warrant acknowledgment. First, while our field validation dataset incorporates natural variability in lighting and background, it does not capture the full spectrum of environmental conditions encountered in tropical and subtropical agricultural systems. These include monsoon rains, high humidity, and extreme temperatures (>40), which can substantially alter leaf appearance and disease symptom expression. Future work should include multi-season, multi-location field trials across diverse agro-ecological zones to comprehensively assess model robustness.

Second, the current framework performs single-image classification and does not incorporate temporal or spatial context from sequential observations. In practice, disease progression is a dynamic process, and integrating multi-temporal data could enable early detection before visible symptoms manifest. Preliminary experiments with video-based inference (processing 5-frame sequences) showed promising results, yielding a 2.1% accuracy improvement. However, this approach increased latency by 180%. This highlights the need for efficient temporal modeling techniques, such as 3D convolutions or temporal attention mechanisms optimized for edge deployment.

Third, while our model covers 14 crop species, this represents only a fraction of globally important food crops. Extending the framework to additional crops, including cereals (rice, wheat, maize), legumes (soybean, chickpea), and cash crops (cotton, coffee), would enhance its practical utility. However, this expansion necessitates careful management of class imbalance, as many diseases in less-studied crops have limited publicly available data. Transfer learning and few-shot learning approaches offer potential solutions that warrant investigation.

Fourth, the current model is designed for supervised learning with fixed disease categories. The emergence of new disease strains or previously unknown pathogens requires retraining on new data, a process that may be prohibitive for end-users in resource limited settings. Continual learning methods, such as elastic weight consolidation (EWC) and progressive neural networks, could enable incremental model updates without catastrophic forgetting. This would allow deployed models to adapt to novel disease presentations while retaining performance on previously learned classes.

Fifth, multi-modal fusion incorporating near-infrared (NIR) or hyperspectral imaging could enhance early-stage disease detection before visible symptoms manifest, potentially enabling pre-symptomatic intervention. How-

ever, such approaches would require substantial redesign of the edge inference pipeline and are currently incompatible with consumer-grade smartphone hardware. As multispectral sensors become more affordable and integrated into mobile devices, this direction warrants exploration.

Sixth, uncertainty quantification is essential for building trust in AI-assisted diagnostics, particularly in high-stakes agricultural decision-making. By providing confidence estimates for each prediction, it allows practitioners to gauge the reliability of the model's output. While Monte Carlo dropout provides well-calibrated uncertainty, its computational cost (200× higher latency) precludes real-time deployment. Efficient uncertainty approximation techniques, such as ensemble of lightweight models or evidential deep learning, should be investigated to balance reliability with edge constraints.

Finally, collaborative edge cloud architectures could dynamically offload high-uncertainty samples to a cloud based teacher model for confirmatory analysis, balancing real-time requirements with diagnostic reliability. This hybrid approach would enable autonomous operation in connectivity-constrained environments while leveraging cloud resources when available. Such architectures also facilitate continuous model improvement through federated learning, wherein edge devices contribute to collective model refinement without sharing sensitive agricultural data.

4.6. Reproducibility and Open Science

To ensure the reproducibility of our results and facilitate adoption by the research community, we have adhered to best practices in open science. The complete training configuration, including hyper parameters, data augmentation pipelines, and random seeds, is documented in **Table 2** and the accompanying code repository. All models were trained using deterministic algorithms where available (PyTorch `torch.use_deterministic_algorithms(True)`), albeit with a modest performance trade-off (approximately 2% slower convergence). The ONNX export pipeline preserves operator compatibility across PyTorch +2.1, ONNX Runtime +1.16, and TensorRT +8.6, enabling cross-platform reproducibility. For edge benchmarking, we have published the exact system configurations (operating system versions, library dependencies, compiler flags) as a Docker container manifest, allowing other researchers to replicate our latency and power measurements. We acknowledge that complete reproducibility in deep learning remains challenging due to GPU-specific non-determinism (e.g., cuDNN convolution algorithms), but we have mitigated this by fixing cuDNN to deterministic mode and reporting performance across multiple hardware runs. Finally, a Colab notebook demonstrating end-to-end inference (from image loading to prediction) is provided to lower the barrier to entry for researchers without dedicated GPU infrastructure.

The authors acknowledge the importance of open source code sharing for scientific reproducibility. While the complete codebase is not publicly released at this stage due to ongoing intellectual property protection procedures, the authors are committed to ensuring reproducibility through: (i) detailed documentation of all architectural and training hyper-parameters in **Table 2** and the accompanying text; (ii) providing ONNX model files and inference scripts to qualified researchers upon reasonable request; (iii) offering a public Google Colab notebook for interactive model demonstration; and (iv) publishing Docker container manifests for exact environment replication. This multi-layered approach ensures that all aspects of the methodology are transparent and reproducible while respecting institutional and commercial constraints.

5. Conclusions

This study presented AgriNet-Lite, a lightweight deep learning framework optimized for real-time crop disease detection on resource-constrained edge devices. The architecture integrates depthwise separable convolutions with a novel lightweight channel attention mechanism and leverages knowledge distillation to achieve state-of-the-art accuracy-efficiency trade-offs. The principal findings are:

1. AgriNet-Lite attained 98.72% accuracy on the PlantVillage benchmark with only 2.1 M parameters and 8.1 MB model size, outperforming MobileNetV2, EfficientNet-B0, and ResNet50.
2. Inference latency of 28 ms on Raspberry Pi 4 enables real-time deployment at 35.7 FPS, exceeding the threshold for practical video stream analysis.
3. Systematic ablation studies confirmed that the channel attention module and knowledge distillation contributed cumulative improvements of 3.92 percentage points, with synergistic rather than merely additive effects.
4. Field validation under natural conditions confirmed practical applicability with 94.2% accuracy, despite sub-

stantial domain shift from laboratory to unconstrained environments.

The framework addresses a critical gap in precision agriculture by enabling AI-powered disease diagnosis in regions with limited computational infrastructure and intermittent internet connectivity. The complete source code, trained model weights, and deployment configurations are publicly available to facilitate reproducibility and downstream adaptation.

Author Contributions

Conceptualization, M.G. and M.N.; methodology, M.G.; software, M.G. and M.N.; validation, M.N.; formal analysis, M.G.; investigation, M.G. and M.N.; resources, M.G.; data curation, M.N.; writing—original draft preparation, M.G.; writing—review and editing, M.N.; visualization, M.G.; supervision, M.N.; project administration, M.G.; funding acquisition, M.N. All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable. This study did not involve humans or animals.

Informed Consent Statement

Not applicable.

Data Availability Statement

The PlantVillage dataset is publicly available at <https://github.com/spMohanty/PlantVillage-Dataset>. The custom field validation dataset and supplementary experimental results are available upon reasonable request from the corresponding author. The complete source code, trained model weights (ONNX format), inference scripts, and Docker container manifests for edge benchmarking are available upon reasonable request from the corresponding author. Additionally, a public Google Colab notebook demonstrating end-to-end inference is provided to facilitate reproducibility.

Acknowledgments

Technical support from the High-Performance Computing Center at Tarbiat Modares University is gratefully acknowledged. The authors also extend their appreciation to Bu-Ali Sina University for providing computational resources and research facilities. The authors thank the anonymous reviewers for their constructive feedback.

Conflicts of Interest

The authors declare no conflict of interest.

AI Use Statement

The authors declare that they used AI Tools during the writing process to enhance the language and readability of the manuscript. The AI was also consulted for organizational suggestions in the initial outline. All scientific content, data interpretation, and final revisions were performed by the authors, who take full responsibility for the integrity and accuracy of this work. This tool was used in accordance with the journal's guidelines on AI and authorship.

References

1. Lal, R. Crop health management for food and nutritional security and soil health. *Crop Health* **2026**, *4*, 22. [[CrossRef](#)]

2. Barbedo, J.G.A. Factors Influencing the Use of Deep Learning for Plant Disease Recognition. *Biosyst. Eng.* **2018**, *172*, 84–91.
3. Gill, S.S.; Golec, M.; Hu, J.; et al. Edge AI: A taxonomy, systematic review and future directions. *Cluster Comput.* **2025**, *28*, 18. [CrossRef]
4. Kaler, B.; Kaur, A. A systematic survey on explainable artificial intelligence (XAI) for plant health monitoring: Challenges and opportunities. *Appl. Intell.* **2025**, *55*, 889. [CrossRef]
5. Salka, T.D.; Hanafi, M.; Abdul Rahman, S.M.; et al. State-of-the-Art CNN Models for Plant Disease Classification: A Comparative Study. *Pertanika J. Sci. Technol.* **2025**, *33*, 2541–2563.
6. Ahamed, M.F.; Salam, A.; Nahiduzzaman, M.; et al. Streamlining plant disease diagnosis with convolutional neural networks and edge devices. *Neural Comput. Appl.* **2024**, *36*, 18445–18477. [CrossRef]
7. Rana, A. Architectural overview of edge AI processing in smart connected devices: From embedded hardware to real-time inference. *J. Comput. Sci. Technol. Stud.* **2025**, *7*, 377–384.
8. Sandler, M.; Howard, A.; Zhu, M.; et al. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 4510–4520.
9. Tan, M.; Le, Q.V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, USA, 9–15 June 2019; pp. 6105–6114.
10. Hinton, G.; Vinyals, O.; Dean, J. Distilling the Knowledge in a Neural Network. *arXiv preprint* **2015**, arXiv:1503.02531.
11. Chang, S. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition. Available online: <https://aapm.confex.com/aapm/2024am/meetingapp.cgi/Paper/13897> (accessed on 7 September 2025).
12. Xu, Y.; Li, D.; Li, C.; et al. LiSA-MobileNetV2: An extremely lightweight deep learning model with Swish activation and attention mechanism for accurate rice disease classification. *Front. Plant Sci.* **2025**, *16*, 1619365. [CrossRef]
13. Gunasekaran, H.; Rajkumar, S.; Kirubhadharsini B, L. Lightweight deep learning for tomato disease detection: trends, challenges, and edge AI perspectives. *Front. Plant Sci.* **2025**, *16*, 1737208. [CrossRef]
14. Yang, X.; Wang, H.; Zhou, Q.; et al. A Lightweight and Efficient Plant Disease Detection Method Integrating Knowledge Distillation and Dual-Scale Weighted Convolutions. *Algorithms* **2025**, *18*, 433.
15. Nyakuri, J.P.; Nkundineza, C.; Gatera, O.; et al. AI and IoT-powered edge device optimized for crop pest and disease detection. *Sci. Rep.* **2025**, *15*, 22905. [CrossRef]
16. Maeda-Gutiérrez, V.; Galvan-Tejada, C.E.; Zanella-Calzada, L.A.; et al. Comparison of convolutional neural network architectures for classification of tomato plant diseases. *Appl. Sci.* **2020**, *10*, 1245. [CrossRef]
17. Sharma, V.; Pal, R. Distilling Deep Learning for the Field: Lightweight Plant Disease Diagnosis via ResNet-50 to MobileNetV2 Transfer. In *Data Science and Applications*; Nanda, S.J., Yadav, R.P., Prasad, M., et al., Eds.; Springer: Cham, Switzerland, 2026. [CrossRef]
18. Karthikeyan, S.; Charan, R.; Narayanan, S.; et al. Enhanced plant disease classification with attention-based convolutional neural network using squeeze and excitation mechanism. *Front. Artif. Intell.* **2025**, *8*, 1640549. [CrossRef]
19. El-Saber, N.A.S.; Mohamed, I.S. A Systematic Review of Lightweight and Robust Deep Learning Approaches for Tomato Leaf Disease Detection. *Eng. Res. J. (Shoubra)* **2026**, *55*, 119–131.
20. Zhang, S.; Liu, R. Multi-scale feature fusion-based vision mamba for robust plant disease image classification on field-acquired plantdoc data. *Front. Plant Sci.* **2026**, *17*, 1842426. [CrossRef]
21. Abou Ali, M.; Dornaika, F.; Charafeddine, J. Edge artificial intelligence: Past advances, current paradigms, and emerging horizons. *Eng. Appl. Artif. Intell.* **2026**, *181*, 115635. [CrossRef]
22. Talaat, F.M.; Tawfik, M.; Shaban, W.M. GreenAid: A confidence-weighted ensemble deep learning system for real-time plant disease detection and management. *Sci. Rep.* **2026**, *16*, 19193. [CrossRef]
23. Sathya, K.; Balakrishnan, A.; Baskaran, P.; et al. Enhancing plant disease detection using Attention-Augmented residual networks and faster Region-Convolutional networks. *IEEE Access* **2025**, *13*, 48625–48642. [CrossRef]
24. Boukhelifa, G.; Chibani, Y. A Lightweight CNN Supported by Partially Skip-Connected CAE for Plant Disease Identification across Multiple Agricultural Environments. *Appl. Soft Comput.* **2026**, *201*, 115515.
25. Zeeshan, M.; Baghini, M.S.; Pandey, A. EdgePlantNet: Lightweight edge-aware cyber-physical system for plant disease detection using enhanced attention CNNs. *Pervasive Mob. Comput.* **2025**, *110*, 102059. [CrossRef]
26. Burman, S.; Pandey, N.; Srivastava, C.; et al. MBConvNet-CNS: A Lightweight Convolutional Neural Network

- Architecture for Plant Disease Detection. In Proceedings of the 2024 12th International Conference on Intelligent Systems and Embedded Design (ISED), Rourkela, India, 20–22 December 2024.
27. Shahade, A.K.; Deshmukh, P.V. Edge-Enhanced Dual Branch CNN with Adaptive Attention for Robust Apple Leaf Disease Detection. *BMC Plant Biol.* **2026**, *26*, 39.
 28. Mahapatra, P.; Panda, M.; Dash, S.K.; et al. Advancing plant disease classification using an attention-based CNN for intra-dataset and cross-dataset training. *Sci. Rep.* **2026**, *16*, 10925. [[CrossRef](#)]
 29. Papanikolaou, A.; Tziouvaras, A.; Floros, G.; et al. Distributed Deep Learning in IoT Sensor Network for the Diagnosis of Plant Diseases. *Sensors* **2025**, *25*, 7646.



Copyright © 2026 by the author(s). Published by UK Scientific Publishing Limited. This is an open access article under the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Publisher's Note: The views, opinions, and information presented in all publications are the sole responsibility of the respective authors and contributors, and do not necessarily reflect the views of UK Scientific Publishing Limited and/or its editors. UK Scientific Publishing Limited and/or its editors hereby disclaim any liability for any harm or damage to individuals or property arising from the implementation of ideas, methods, instructions, or products mentioned in the content.