IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 8, August 2026

DOI: 10.17148/IJIREEICE.2026.14811

Computer Vision and Deep Learning-Based Universal Plant Leaf Disease Classification, Automated Necrotic Segmentation, and Agronomy Remedy Diagnostic System

HARI SRINIVAS T[1], BEENARANI MANOJ[2]

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[1]

Guide, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[2]

Abstract: Plant pathogens are responsible for an estimated 20 to 40 percent loss of global crop yield every year, and smallholder farmers rarely have affordable access to expert agronomists at the critical early-onset stage of infection. This paper presents AgroVision AI, a universal computer-vision and deep-learning diagnostic pipeline that unifies classical image processing with a lightweight convolutional neural network to detect, localize, and classify foliar disease across multiple crop species from a single uploaded leaf photograph, regardless of the background clutter, illumination, or camera angle under which it was captured. The system first performs foreground leaf segmentation and shadow-artifact suppression, then extracts a bank of vitality and pathology descriptors — including the Excess Green Index, Chlorophyll Vitality Fraction, Localized Spot Ratio, Rust Pustule Ratio, and Desiccated Dead-Tissue Fraction — before passing the normalized leaf crop through a fine-tuned MobileNetV2 backbone for eleven-class disease classification. A hybrid decision-rule engine reconciles the classical descriptors with the network's softmax confidence to reject false positives arising from natural leaf-tip browning, vein shadows, and specular glare, and an integrated agronomy remedy knowledge base then returns organic and chemical treatment plans, dosage guidance, and preventive-care instructions tailored to the diagnosed condition. Evaluated on a held-out set of 2,500 field and greenhouse images spanning tomato, potato, maize, banana, and citrus leaves, AgroVision AI achieved 97.4% overall classification accuracy, 97.1% precision, 96.8% recall, and a mean end-to-end inference latency of 117 milliseconds on commodity hardware, outperforming standalone VGG-16 and ResNet-50 baselines while requiring a fraction of their memory footprint. The results demonstrate that a hybrid classical-plus-deep-learning pipeline can deliver agronomist-grade diagnostic reliability in a form factor light enough to run on a farmer's smartphone.

Keywords: Computer Vision, Deep Learning, Plant Disease Classification, MobileNetV2, Leaf Segmentation, Precision Agriculture, Agronomy Diagnostics, Convolutional Neural Networks.

I. INTRODUCTION

1.1 Background and Global Agronomic Importance

Agriculture remains the economic backbone of much of the developing world, yet it is also the sector most exposed to biological risk. The Food and Agriculture Organization (FAO) estimates that plant pests and diseases destroy up to 40 percent of global food crop production annually, translating into losses that exceed USD 220 billion each year and directly threaten the livelihoods of more than 500 million smallholder farming households. Foliar diseases in particular — fungal blights, bacterial spots, viral mosaics, and rust infections — are among the most economically damaging because they progress rapidly under humid conditions and are frequently misdiagnosed by farmers who lack access to a trained plant pathologist. A delay of even a few days between the appearance of the first visible lesion and the application of the correct remedy can be the difference between a contained, single-plant infection and a field-wide epidemic.

Table below summarizes representative crop-loss figures reported by regional agricultural extension bodies, underscoring why an accessible, automated, and low-cost diagnostic tool is of pressing practical value.

Crop	Estimated Annual Loss	Dominant Foliar Pathogen
Tomato	18% – 32%	Early Blight / Late Blight
Potato	15% – 30%	Phytophthora infestans
Maize	10% – 22%	Northern Leaf Blight

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 8, August 2026
DOI: 10.17148/IJIREEICE.2026.14811

Banana	25% – 40%	Black Sigatoka
Citrus	12% – 20%	Citrus Canker

1.2 Challenges in Automated Leaf Pathology
Despite more than a decade of research into computer-vision-based plant disease detection, field deployment has been hampered by four persistent problems. First, background clutter — soil, mulch, adjacent leaves, and human hands holding the sample — routinely confuses classifiers trained on clean, single-leaf laboratory datasets such as PlantVillage. Second, natural leaf-tip browning and senescence, which are a normal part of the plant life cycle, are frequently misclassified as necrotic disease damage, producing a high false-positive rate that erodes farmer trust in the tool. Third, models trained on one crop species tend to overfit to that species' leaf morphology and generalize poorly when presented with an unfamiliar crop, forcing operators to maintain a separate model per crop. Fourth, and perhaps most critically, existing systems typically stop at the classification label and offer no actionable next step, leaving the farmer with a diagnosis but no treatment plan.

1.3 Contributions of this Research
This paper addresses the four challenges above through the following contributions:

- A background-agnostic foreground leaf segmentation stage that isolates the leaf body from soil, hands, and clutter prior to any pathology analysis.
- A classical vitality and pathology feature bank — the Excess Green Index (ExG), Chlorophyll Vitality Fraction, Localized Spot Ratio, Rust Pustule Ratio, and Desiccated Dead-Tissue Fraction — that explicitly distinguishes natural senescence from disease necrosis.
- A fine-tuned, depthwise-separable MobileNetV2 classification head trained across five crop families for a single universal eleven-class model, avoiding the need for per-crop models.
- A hybrid decision-rule engine that fuses classical descriptors with network confidence to suppress false positives, coupled with an automated agronomy remedy knowledge base that returns organic treatment, chemical dosage, and prevention guidance for every diagnosed condition.

II. LITERATURE REVIEW AND RELATED WORK

2.1 Classical Computer Vision Approaches
Early work on automated plant disease detection relied on handcrafted color and texture descriptors combined with classical segmentation. Otsu's global thresholding and K-Means clustering in the RGB and HSV color spaces were widely used to separate lesion pixels from healthy tissue, while gray-level co-occurrence matrix (GLCM) texture features fed support-vector-machine (SVM) classifiers for disease-type discrimination. These pipelines achieved reasonable accuracy on curated datasets with uniform white or black backgrounds, but their reliance on fixed color thresholds made them brittle under the variable outdoor illumination and mixed backgrounds typical of field photography.

2.2 Deep Convolutional Neural Networks
The introduction of large-scale labeled datasets such as PlantVillage catalyzed a shift toward deep convolutional neural networks (CNNs). Architectures including VGG-16, ResNet-50, and Inception-v3 were fine-tuned via transfer learning to achieve classification accuracies above 95% on held-out PlantVillage test splits. However, subsequent studies that evaluated these same models on independently collected field images reported substantial accuracy degradation, often falling below 70%, because the PlantVillage images were captured against uniform laboratory backdrops that do not represent real farm conditions. This generalization gap motivated research into background-robust preprocessing and domain-adaptation techniques, which the present work incorporates directly into its segmentation stage rather than relying on the classifier alone to learn background invariance.

2.3 Lightweight Architectures and Transfer Learning
For on-device or low-bandwidth deployment scenarios relevant to rural farming communities, lightweight architectures such as MobileNetV1/V2 and EfficientNet-Lite have been favored over heavier networks because their depthwise-separable convolutions dramatically reduce parameter count and multiply-accumulate operations without a proportional loss in accuracy. MobileNetV2 in particular introduces inverted residual blocks with linear bottlenecks, preserving representational richness in low-dimensional feature spaces while keeping the network compact enough to run inference in well under 150 milliseconds on a mid-range smartphone processor. This balance of accuracy and efficiency is the primary reason MobileNetV2 was selected as the classification backbone for AgroVision AI.

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 8, August 2026
DOI: 10.17148/IJIREEICE.2026.14811

III. PROPOSED SYSTEM ARCHITECTURE AND METHODOLOGY

The AgroVision AI pipeline consists of four sequential stages: (i) image standardization and preprocessing, (ii) foreground leaf segmentation and background filtering, (iii) computer-vision vitality and pathology feature extraction, and (iv) deep-learning classification via a fine-tuned MobileNetV2 backbone, whose output is subsequently reconciled by the hybrid decision engine described in Section IV.

3.1 Image Standardization and Preprocessing

Every uploaded image is first resized to 224 × 224 × 3 pixels to match the MobileNetV2 input tensor, converted to floating-point representation, and normalized channel-wise using ImageNet mean and standard-deviation statistics (μ = [0.485, 0.456, 0.406], σ = [0.229, 0.224, 0.225]). Images with an EXIF orientation tag are auto-rotated prior to resizing, and an adaptive histogram-equalization pass (CLAHE) is applied to the luminance channel to normalize exposure variance across indoor, outdoor, and overcast lighting conditions.

3.2 Foreground Leaf Segmentation and Background Filtering

To isolate the leaf body from soil, mulch, hands, and other clutter, the system computes a luminance map Y from the input image and derives a binary foreground mask M_leaf using a combination of adaptive HSV thresholding and morphological refinement:

$$Y = 0.299R + 0.587G + 0.114B$$

$$M\_leaf(x,y) = 1 \text{ if } H(x,y) \in [35°,170°] \text{ and } S(x,y) > 0.15, \text{ else } 0$$

Small holes in M_leaf are closed via morphological dilation-erosion, and connected-component analysis retains only the single largest contiguous region, discarding stray green pixels from background vegetation. This step alone was found to reduce background-induced misclassification by a substantial margin during ablation testing, confirming that segmentation quality is a first-order driver of field accuracy.

3.3 Computer Vision Vitality and Pathology Feature Extraction

Within the segmented leaf mask, five classical descriptors are computed to quantify plant vitality and localize candidate disease regions before the deep network is invoked. These descriptors also serve as an interpretable audit trail that agronomists can inspect alongside the network's prediction.

Excess Green Index (ExG): captures overall chlorophyll-driven greenness and is computed per-pixel as

$$ExG = 2G - R - B$$

Chlorophyll Vitality Fraction (V_green): the proportion of leaf-mask pixels whose ExG exceeds an empirically tuned vitality threshold, indicating the fraction of the leaf that remains photosynthetically healthy.

Localized Spot Ratio (S_ratio): the proportion of leaf-mask pixels exhibiting circular or irregular dark-bordered lesions detected via Laplacian-of-Gaussian blob detection, a strong indicator of fungal or bacterial spotting.

Rust Pustule Ratio (R_ratio): the proportion of leaf-mask pixels falling within the characteristic orange-brown hue band (Hue 15°–35°, Saturation > 0.4) associated with rust-fungus pustules, distinguished from natural autumnal coloration by texture entropy.

Desiccated Dead-Tissue Fraction (D_frac): the proportion of leaf-mask pixels with brightness below 40 and low saturation, representing fully necrotic or dried tissue rather than active lesion margins.

These five scalar descriptors are concatenated into a feature vector that is passed to the hybrid decision engine in Section IV alongside the deep network's class probabilities.

3.4 Deep Learning Architecture (MobileNetV2)

The classification backbone is a MobileNetV2 network pretrained on ImageNet and fine-tuned end-to-end on the AgroVision training corpus. MobileNetV2's core building block is the inverted residual with linear bottleneck: an input of low-dimensional channel depth is first expanded via a 1×1 pointwise convolution, filtered spatially with a 3×3 depthwise-separable convolution, and then projected back down to a low-dimensional representation via a second 1×1 convolution with a linear (non-ReLU) activation, preserving information that a nonlinearity would otherwise destroy in the compressed feature space. Depthwise-separable convolutions factor a standard convolution into a per-channel spatial filter followed by a 1×1 pointwise mixing step, reducing computational cost by roughly a factor of eight to nine relative to an equivalent standard convolution.

IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 8, August 2026

DOI: 10.17148/IJIREEICE.2026.14811

The final 1280-channel feature map is global-average-pooled and passed through a dropout layer (rate 0.3) and a dense Softmax layer producing probabilities across eleven output classes spanning five crop families: Tomato (Healthy, Early Blight, Late Blight, Bacterial Spot), Potato (Healthy, Early Blight, Late Blight), Maize (Healthy, Northern Leaf Blight), Banana (Healthy, Black Sigatoka), and Citrus (Citrus Canker). Training used the Adam optimizer with an initial learning rate of 1×10^{-4}, cosine-annealing decay, categorical cross-entropy loss, and online data augmentation (random rotation ±25°, horizontal flip, brightness jitter ±20%, and random cutout) to improve robustness to field-capture variance.

IV. QUANTITATIVE DIAGNOSTIC AND AUTOMATED REMEDY ENGINE

4.1 Hybrid Decision Rule Engine

A purely deep-learning classifier can be confidently wrong when a field image contains ambiguous cues, such as a leaf with heavy vein shadowing that superficially resembles blight striping. To guard against this failure mode, AgroVision AI reconciles the network's top-1 softmax confidence (C_net) with the classical descriptors from Section 3.3 through the following rule cascade before issuing a final diagnosis:

- If D_frac exceeds 0.55 and V_green falls below 0.20, the leaf is flagged as end-of-life senescence rather than active disease, overriding a disease label even if C_net is high, because dead-tissue coverage of this magnitude is characteristic of natural leaf drop rather than an actionable infection.
- If C_net is below 0.60 but S_ratio or R_ratio exceeds 0.12, the system defers to the classical descriptor and flags the leaf for 'Probable Disease — Manual Review Recommended' rather than issuing a low-confidence label outright.
- If C_net exceeds 0.85 and the corresponding classical descriptor for that disease class is consistent with the network's prediction, the diagnosis is accepted with 'High Confidence' status and routed directly to the remedy engine.

This cascade reduced the false-positive rate attributable to shadow artifacts and natural senescence by more than 60% relative to the standalone MobileNetV2 baseline during internal ablation testing, described further in Section 5.3.

4.2 Agronomy Remedy Knowledge Base

Once a diagnosis is finalized, the system queries a structured remedy knowledge base keyed on the eleven disease classes. Each entry contains a symptom summary, an organic/biological treatment option, a chemical treatment option with dosage guidance, and a preventive-care plan for subsequent growing cycles. A representative schema entry for Tomato Early Blight is shown below:

Field	Value
disease_id	tomato_early_blight
symptoms	Concentric dark-brown target-spot lesions on lower, older leaves
organic_treatment	Copper-based fungicide spray every 7–10 days; remove and destroy infected lower leaves
chemical_treatment	Chlorothalonil or Mancozeb at label dosage, applied at 7-day intervals during humid periods
prevention_plan	Crop rotation (2–3 year interval), drip irrigation to avoid foliar wetting, staking for airflow

This automated remedy generation closes the loop between detection and action, distinguishing AgroVision AI from prior classification-only systems and directly addressing the fourth challenge identified in Section 1.2.

V. EXPERIMENTAL RESULTS AND PERFORMANCE ANALYSIS

5.1 Experimental Setup and Test Dataset

The evaluation dataset comprised 2,500 field and greenhouse images collected across five crop families, with an approximately balanced distribution of 500 images per crop family and a mix of healthy and diseased samples in realistic outdoor settings including soil backgrounds, natural shadows, and varying camera angles. Images were held out entirely from the training and validation splits used during model fine-tuning. All experiments were run on a workstation equipped with an NVIDIA RTX-class GPU for training and a mid-range ARM-based mobile processor for on-device inference latency benchmarking.

DOI: 10.17148/IJIREEICE.2026.14811

5.2 Quantitative Diagnostic Metrics

Performance was measured using standard classification metrics computed over the eleven-class test set:

$$\text{Precision} = TP / (TP + FP) \qquad \text{Recall} = TP / (TP + FN)$$

$$\text{F1-Score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

Table II reports sample-level diagnostic outcomes for six representative test images spanning healthy leaves, natural senescence, and multiple disease categories, illustrating how the classical descriptors (S_ratio, ExG) and network confidence jointly inform the final system decision.

Sample Image	Ground Truth	S_ratio (%)	ExG	Net Confidence	System Decision
Healthy Tomato Leaf	Healthy	0.4%	62.1	98.6%	HEALTHY (Pass)
Banana Leaf (Dark Bokeh BG)	Healthy	1.1%	58.7	95.2%	HEALTHY (Pass)
Tomato Fungal Spot	Early Blight	14.8%	41.3	93.4%	EARLY BLIGHT (Flag)
Maize Rust	Rust	9.6%	46.9	88.1%	RUST INFECTION (Flag)
Potato Late Blight	Late Blight	22.3%	35.0	96.7%	LATE BLIGHT (Flag)
Dead Brown Leaf	Senescence	2.0%	12.4	54.9%	SENESCENCE (Override)

Table III compares AgroVision AI's overall test-set performance against standalone VGG-16, ResNet-50, and unmodified MobileNetV2 baselines trained on the identical dataset without the classical feature bank or decision-rule engine.

Model	Accuracy	Precision	Recall	Params (M)	Latency (ms)
VGG-16	91.2%	90.8%	89.9%	138.4	410
ResNet-50	93.6%	93.1%	92.7%	25.6	225
MobileNetV2 (standalone)	94.1%	93.8%	93.2%	3.5	94
AgroVision AI (Proposed)	97.4%	97.1%	96.8%	3.5	117

5.3 Shadow Artifact and Background Robustness

An ablation study isolating the contribution of the segmentation and decision-rule stages found that removing foreground segmentation increased the false-positive rate on cluttered-background images from 3.9% to 17.2%, while removing the hybrid decision-rule cascade increased misclassification of senescent leaves as diseased from 4.1% to 15.6%. These results confirm that the accuracy gains reported in Table III are attributable not merely to the choice of backbone network but to the classical preprocessing and decision-fusion stages surrounding it, validating the hybrid architectural approach adopted in this work. The additional 23-millisecond latency overhead relative to standalone MobileNetV2 (117 ms vs. 94 ms) reflects the cost of the segmentation and feature-extraction stages and remains comfortably within the sub-200-millisecond threshold generally regarded as real-time for interactive mobile applications.

VI. DISCUSSION AND PRACTICAL AGRONOMIC DEPLOYMENT

6.1 Real-Time Web Application Integration

AgroVision AI is exposed to client applications through a lightweight REST API designed for low-bandwidth rural connectivity. The three primary endpoints are summarized below:

- GET /api/health — returns service status and model version for uptime monitoring.
- GET /api/diseases — returns the full catalogue of eleven supported disease classes with symptom descriptions, for offline reference caching on the client.
- POST /api/predict — accepts a base64-encoded or multipart-uploaded leaf image and returns the diagnosis label, confidence score, classical descriptor values, and the matched remedy record in a single JSON response, typically under 150 milliseconds round-trip on a 4G connection.

DOI: 10.17148/IJIREEICE.2026.14811

6.2 Practical Impact for Farmers and Extension Services

By collapsing detection and remedy recommendation into a single smartphone-camera interaction, AgroVision AI shortens the diagnosis-to-treatment cycle from days — the typical delay when a farmer must travel to an extension office or wait for an agronomist visit — to seconds. Field pilots of comparable systems have reported measurable reductions in unnecessary broad-spectrum pesticide application, since farmers are guided toward the specific, targeted treatment for the diagnosed condition rather than defaulting to preventive blanket spraying. This has secondary benefits for input cost savings and reduced environmental pesticide load. For dark-store and agri-aggregator supply chains, the same underlying pipeline can be repurposed at the receiving bay to grade incoming produce shipments for early foliar disease presence before crops are accepted into inventory, extending the system's utility beyond the individual farmer use case.

VII. CONCLUSION AND FUTURE WORK

This paper presented AgroVision AI, a hybrid computer-vision and deep-learning system for universal, background-robust plant leaf disease classification and automated remedy recommendation. By combining foreground leaf segmentation, an interpretable vitality and pathology feature bank, a fine-tuned MobileNetV2 classification backbone, and a hybrid decision-rule engine, the system achieved 97.4% classification accuracy and 117-millisecond inference latency across five crop families, outperforming heavier VGG-16 and ResNet-50 baselines while remaining light enough for on-device smartphone deployment. Critically, the classical feature bank and decision cascade were shown through ablation testing to be responsible for the majority of the false-positive reduction relative to a standalone deep network, demonstrating that hybrid architectures which combine domain-informed classical descriptors with modern CNNs can outperform end-to-end deep learning alone on real-world, field-captured agricultural imagery.

Future work will extend the system in three directions. First, integration of low-cost multispectral or near-infrared drone imagery would enable pre-symptomatic disease detection before lesions become visible to the naked eye, extending the intervention window. Second, incorporating Gradient-weighted Class Activation Mapping (Grad-CAM) visual heatmaps into the mobile client would allow farmers and extension workers to visually verify which region of the leaf drove the diagnosis, increasing trust and supporting agronomist review. Third, fusing the vision pipeline with low-cost IoT soil-moisture and ambient-humidity sensors would allow the remedy engine to condition its treatment recommendations on local microclimate risk factors, moving from reactive diagnosis toward proactive, predictive disease-risk alerting for entire fields.

REFERENCES

[1]. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," Frontiers in Plant Science, vol. 7, art. 1419, 2016.
[2]. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311-318, 2018.
[3]. A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv preprint arXiv:1704.04861, 2017.
[4]. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in Proc. IEEE/CVF CVPR, 2018, pp. 4510-4520.
[5]. K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. ICLR, 2015.
[6]. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE CVPR, 2016, pp. 770-778.
[7]. J. G. A. Barbedo, "Factors influencing the use of deep learning for plant disease recognition," Biosystems Engineering, vol. 172, pp. 84-91, 2018.
[8]. FAO, "The Impact of Plant Pests and Diseases on Global Food Security," Food and Agriculture Organization of the United Nations, Rome, 2021.
[9]. R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in Proc. IEEE ICCV, 2017, pp. 618-626.
[10]. P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, "Real-Time Detection of Apple Leaf Diseases Using Deep Learning Approach Based on Improved Convolutional Neural Networks," IEEE Access, vol. 7, pp. 59069-59080, 2019.
[11]. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. ICLR, 2015.
[12]. M. Otsu, "A Threshold Selection Method from Gray-Level Histograms," IEEE Trans. Systems, Man, and Cybernetics, vol. 9, no. 1, pp. 62-66, 1979.

