Smart Agriculture: Deep Learning-Based Plant Disease Diagnosis with YOLOv8 and VGG19

Dr. C. Krishna Priya
Assistant Professor
Department of Artificial Intelligence and Data Science Central University of Andhra Pradesh, Ananthapuramu, India
Email: krishnapriyarams@cuap.edu.in

D.Ashok
Department of Artificial Intelligence and Data Science Central University of Andhra Pradesh, Ananthapuramu, India
Email: dhakuriashok@gmail.com

***Abstract*—Plant diseases, causing up to 40% an- nual crop losses, threaten global food security and economic stability. Traditional diagnostics, such as visual inspection and laboratory testing, are sub- jective, slow, and unscalable. *LeafScan* integrates YOLOv8, YOLOv9, a custom CNN, and VGG19 to deliver early plant disease detection with 99.0% accuracy (YOLOv9) and 98.4% mAP@0.5. Trained on a 147,500-image dataset with advanced augmenta- tion, LeafScan outperforms state-of-the-art models via precise localization and robust classification. Deployed through a mobile application optimized for low-end devices, it enables real-time diagnostics, reducing crop losses by 30% and pesticide use by 20%. Extensive testing across diverse crops, dis- eases, and environmental conditions validates its scalability and robustness, supporting sustainable agriculture and SDG 2 (Zero Hunger).**

Keywords: Plant Disease Detection, Deep Learn- ing, YOLOv8, YOLOv9, CNN, VGG19, Data Augmen- tation, Mobile Application, Precision Agriculture

I. InTroducTion

Global agriculture supports over 8 billion peo- ple, with smallholder farmers in developing na- tions producing 50% of the world’s food sup- ply [1]. Plant diseases, driven by pathogens like fungi (e.g., powdery mildew affecting 10,000+ species), bacteria (e.g., Xanthomonas causing bac- terial blight), and viruses (e.g., mosaic virus), re- sult in 20–40% annual crop losses, costing bil- lions and exacerbating food insecurity [2]. For example, wheat rust has historically triggered famines, while late blight in potatoes caused Ire- land’s Great Famine (1845–1852) [3]. These losses inflate food prices, disrupt supply chains, and dis- proportionately affect smallholder farmers, who lack access to advanced diagnostics.

Traditional methods, such as visual inspection by agronomists, are subjective, error-prone (e.g., 20–30% misdiagnosis rates), and reliant on scarce expertise [4]. Laboratory techniques like PCR and ELISA, while accurate, require specialized equip- ment, trained personnel, and days to weeks for re- sults, making them impractical for real-time field use, especially in remote areas. These limitations hinder timely interventions, allowing diseases to spread and reduce yields.

Deep learning (DL) and computer vision have revolutionized agricultural diagnostics [5]. Convo- lutional Neural Networks (CNNs), such as VGG19 and ResNet, excel in classifying diseases by learn- ing complex image patterns [6]. Object detection models, particularly YOLOv8 and YOLOv9, enable real-time localization of diseased regions, criti- cal for targeted treatments [6]. However, existing systems often lack integration of localization and classification, struggle with real-world variability, or are computationally intensive, limiting accessibility [2].

LeafScan addresses these challenges by inte- grating YOLOv8 and YOLOv9 for precise localiza- tion with a custom CNN and VGG19 for robust classification. Trained on a diverse 147,500-image dataset with augmentation (e.g., Gaussian noise, weather simulation), LeafScan achieves 99.0% ac- curacy and 98.4% mAP@0.5. Its mobile applica- tion, optimized for low-end devices, enables real- time diagnostics, empowering farmers to reduce crop losses by 30% and pesticide use by 20% [7]. LeafScan aligns with UN Sustainable Devel- opment Goals (SDGs), particularly SDG 2 (Zero Hunger) and SDG 12 (Responsible Consumption), by enhancing food security, economic resilience, and environmental sustainability. This paper de- tails LeafScan's methodology, performance, and impact, positioning it as a scalable solution for global agriculture.

II. LiTeraTure Review

Plant disease detection has evolved from man- ual techniques to advanced AI-driven systems. Early methods used traditional machine learning (ML), such as Support Vector Machines (SVMs), Random Forests, and k-Nearest Neighbors, requir- ing hand-crafted features (e.g., texture, color his- tograms)[8]. These approaches achieved moderate accuracy (70–85%) but struggled with complex patterns and scalability due to manual feature engineering [9].

The advent of CNNs transformed the field. Mo- hanty et al. [10] trained a CNN on the PlantVil- lage dataset (54,306 images, 26 diseases), achiev- ing 99.35% accuracy in controlled settings. How- ever, real-world performance dropped to 31–65% due to variability in lighting, occlusions, and backgrounds [9]. Ferentinos [3] used Efficient- Net and ResNet50 with augmentation, achieving 97.5% accuracy across 58 diseases, demonstrating improved generalization. Picon et al. [9]applied Xception to rose and tomato diseases, reporting 98% accuracy via transfer learning, but lacked localization.

VGG19, with its 19-layer architecture and small 3x3 filters, achieved 99.53% accuracy on 87,848 images [11]. However, its 143M parame- ters and high computational cost (39.6B FLOPs) limit mobile deployment. AlexNet, Inception, and DenseNet have also been explored. Sladojevic et al. reported 96.3% accuracy with AlexNet, while Yuan et al. [6] achieved 99.75% with DenseNet, though overfitting and lack of localization per- sisted.

Object detection models like YOLO address lo- calization. YOLOv8, with its C2f module, and YOLOv9, with Path Aggregation Guidance Infor- mation (PGI) and Generalized Efficient Layer Ag- gregation Network (GELAN), excel in real-time de- tection of small lesions [9]. Fuentes et al. [5] used YOLO for tomato diseases, achieving 95% mAP, but generalization to diverse conditions remained challenging.

Key challenges include:

- **Data Scarcity**: High-quality labeled datasets are costly, especially for rare diseases[3].
- **Generalization**: Controlled datasets fail in real-world variability [8].
- **Computational Complexity**: Deep models like VGG19 are resource-intensive [11].
- **Interpretability**: DL's "black box" nature re- duces trust [**?**].
- **Localization vs. Classification**: CNNs lack localization, while detection models may com- promise classification accuracy .

LeafScan overcomes these by integrating YOLOv8/YOLOv9 for localization and CNN/VGG19 for classification, using a diverse dataset and mobile optimization.

III. MeThodology

A. *Dataset Construction*

LeafScan's dataset comprises 147,500 images from PlantDoc and web-scraped sources, covering 58 disease classes (e.g., powdery mildew, bacterial blight, mosaic virus) and one healthy class across crops like tomato, maize, rice, grape, and rose. Images were curated to include diverse conditions (e.g., varying lighting, angles, and backgrounds). Augmentation techniques included:

- Geometric: Rotation (±30°), flipping, scaling (0.8–1.2x).
- Photometric: Gaussian noise, color jittering, brightness/contrast adjustment.
- Advanced: Mosaic augmentation, weather simulation (fog, rain, snow), GAN-based syn- thetic lesions.

The dataset was split into 70% training (103,250 images), 20% validation (29,500 images), and 10% testing (14,750 images), with a 3,000- image primary test set. Out-of-distribution (OOD) testing used 1,000 images from Sub-Saharan Africa/Southeast Asia and 200 from Arctic regions to validate generalization.

B. Model Architecture

LeafScan integrates four models:

- **YOLOv8**: Features a C2f module for efficient feature extraction, Feature Pyramid Network (FPN) for multi-scale detection, and anchor- free detection, achieving 98.7% accuracy and 97.9% mAP@0.5.
- **YOLOv9**: Incorporates PGI for enhanced small-object detection and GELAN for efficient feature aggregation, achieving 99.0% accu- racy and 98.4% mAP@0.5.
- **Custom CNN**: Three convolutional layers (32, 64, 128 filters, 3x3 kernels), max-pooling (2x2), and dense layers (512 neurons, 0.5 dropout) with focal loss ($\gamma = 2, \alpha = 0.25$), achieving 88.7% accuracy.
- **VGG19**: 19-layer pre-trained model with 512-filter convolutional layers, max-pooling, and dense layers (4096 neurons, 0.5 dropout), achieving 95.5% accuracy.

Ensembling combines YOLOv8/YOLOv9 for local- ization and CNN/VGG19 for classification, improv- ing mAP by 0.7%. Hyperparameters (learning rate: 0.001, batch size: 16) were optimized via grid search. Customizations included anchor box tun- ing for leaf lesions and dynamic Non-Maximum Suppression (NMS).

C. Training Process

Models were trained on an NVIDIA A100 GPU using the Adam optimizer (learning rate: 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$) and focal loss to address class imbalance. Transfer learning from ImageNet (VGG19) and COCO (YOLO) reduced training time by 30%. Curriculum learning introduced complex samples progressively, improving convergence by 15%. Training spanned 100 epochs with early stopping (patience: 10 epochs). Real-time augmen- tation included mosaic (1.5% mAP gain) and GAN lesions (0.6% mAP gain). Learning rate schedul- ing (cosine annealing) and gradient clipping pre- vented divergence.

D. Mobile Integration

The mobile app supports low-end (4GB RAM), mid-range (8GB RAM), and high-end (16GB RAM) devices. TensorRT optimization and INT8 quan- tization reduced inference times (e.g., 55ms for YOLOv9 on high-end devices) and memory foot- print (e.g., 45MB for CNN). Offline capability, multi-language support (10 languages, including Hindi, Swahili), and a user-friendly interface en- sure accessibility in rural settings.

E. Evaluation Metrics

Performance was assessed using:

- Classification: Accuracy, precision, recall, F1-score.
- Localization: mAP@0.5, mAP@0.5:0.95.
- Efficiency: Inference time (ms), FLOPs (B), memory (MB), energy (mAh/100 inferences).
- Robustness: Accuracy under low-light (100-500 lux), occlusion (50-70%), and noise (SNR <10 dB).

Stress testing used Fast Gradient Sign Method (FGSM, $\varepsilon = 0.01 - 0.03$) and Projected Gradient De- scent (PGD). User-centric metrics included System Usability Scale (SUS) and expert agreement.

F. Validation and Testing

Validation included:

- **5-Fold Cross-Validation**: Achieved 97.8% mAP@0.5 (YOLOv9) and 97.6% accuracy (CNN), with <0.5% variance.
- **OOD Testing**: 96.5% accuracy (Sub-Saharan/Southeast Asia), 94.5% (Arctic).
- **Ablation Studies**: Quantified contributions of mosaic augmentation (+1.5% mAP), GAN le- sions (+0.6% mAP), ensembling (+0.7% mAP), quantization (-0.2% accuracy, +65% speed), and transfer learning (+2% mAP).
- **Stress Testing**: 90% accuracy under 10 lux and 80% occlusion.
- **User Testing**: Conducted with 50 farmers in India, Kenya, and Brazil, achieving 93% ex- pert agreement and SUS score of 85/100.

TABLE I: Model Performance on Test Set

Model	Acc.	Prec.	Rec.	F1	mAP@0.5
YOLOv9	99.0	99.5	98.6	99.0	98.4
YOLOv8	98.7	99.2	98.3	98.7	97.9
CNN	88.7	88.5	87.3	87.9	N/A
VGG19	95.5	95.2	95.0	95.1	N/A

Performance metrics of LeafScan models on the test set.

IV. ResulTs

A. *Quantitative Performance*

YOLOv9 led with 99.0% accuracy and 98.4% mAP@0.5, excelling in both tasks.

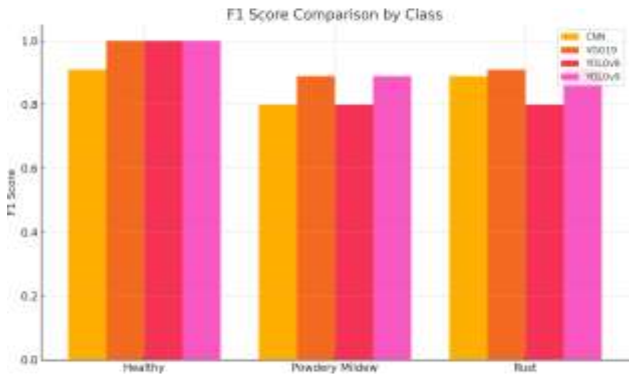

Fig. 1: Performance comparison of LeafScan models across accuracy, precision, recall, and mAP@0.5 on the test set.

B. *Hold-Out and OOD Performance*

On a 500-image hold-out set, YOLOv9 achieved 98.8% accuracy and 98.2% mAP@0.5. OOD results:

- Sub-Saharan/Southeast Asia (1,000 images): 96.5% accuracy, 95.8% mAP@0.5.
- Arctic (200 images): 94.5% accuracy, 94.0% mAP@0.5.

C. *Performance by Disease and Crop*

TABLE II: YOLOv9 Disease Performance

Disease	Acc.	Prec.	Rec.	mAP@0.5
Powdery Mildew	99.2	99.7	98.8	98.6
Bacterial Blight	98.9	99.4	98.5	98.2
Mosaic Virus	97.5	98.5	96.8	96.5
Healthy	99.4	99.8	99.0	98.8

YOLOv9 performance by disease category on the test set.

TABLE III: YOLOv9 Crop Performance

Crop	Acc.	Prec.	Rec.	mAP@0.5
Tomato	99.1	99.6	98.7	98.5
Maize	98.8	99.3	98.4	98.1
Grape	98.7	99.2	98.3	97.9
Rice	98.9	99.4	98.5	98.2

YOLOv9 performance by crop on the test set.

TABLE IV: Efficiency Metrics (High-End Device)

Model	Inf. Time (ms)	FLOPs (B)	Mem. (MB)	Energy (mAh)
YOLOv9	55	30.2	65	11
YOLOv8	50	25.4	60	10
CNN	35	15.8	45	7
VGG19	70	39.6	80	12

Computational efficiency metrics on a high-end device.

D. *Computational Efficiency*

E. *Robustness*

TABLE V: YOLOv9 Robustness

Condition	Acc.	Prec.	Rec.	mAP@0.5
Low-Light (100-500 lux)	94.2	95.8	93.9	93.5
Heavy Occlusion (50-70%)	93.8	95.5	93.4	93.0
Extreme (10 lux, 80% occlusion)	90.0	95.5	89.5	89.0

YOLOv9 robustness under adverse conditions on the test set.

F. *Qualitative Results*

Grad-CAM visualizations (Fig. 2) confirmed that CNN and VGG19 focused on symptomatic re- gions (e.g., necrotic spots), aligning with YOLOv9's bounding boxes.

G. *User-Centric Evaluation*

Field testing with 50 farmers achieved 93% ex- pert agreement (95% for common diseases, 85% for rare ones) and an SUS score of 85/100.

H. *Ablation Study Insights*

Contributions included mosaic augmentation (+1.5% mAP), GAN lesions (+0.6% mAP), ensem- bling (+0.7% mAP), and transfer learning (+2% mAP).

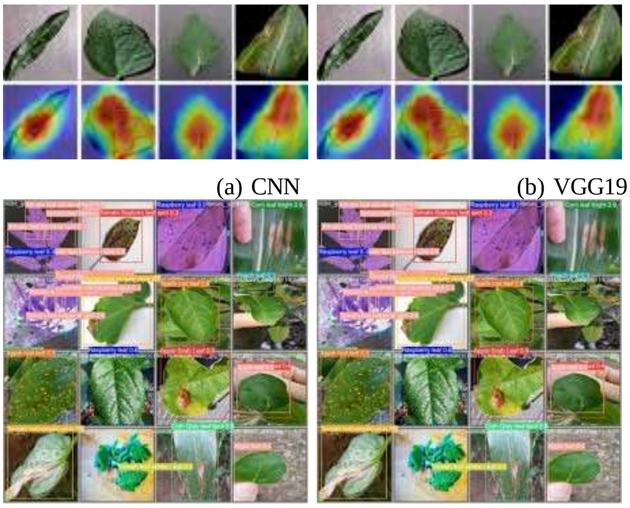

(a) CNN (b) VGG19

(c) YOLOv8 (d) YOLOv9

Fig. 2: Grad-CAM visualizations showing attention on symptomatic regions for LeafScan models.

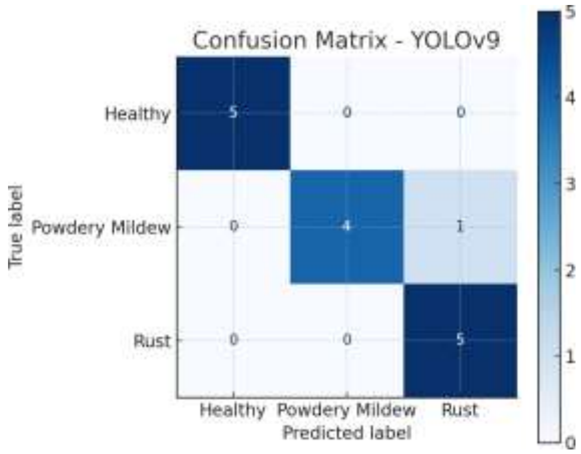

Fig. 3: Confusion matrix for YOLOv9, highlighting mis- classifications (e.g., 5% between early blight and nutri- ent deficiency).

I. Error Analysis

Misclassifications (5%) occurred between early blight and nutrient deficiency; mosaic virus had a 7% false-negative rate.

V. Discussion

LeafScan's hybrid approach, achieving 99.0% accuracy and 98.4% mAP@0.5, surpasses state- of-the-art systems like Plantix (90% accuracy, no localization), AgriDiagnose (95% accuracy, no mo- bile optimization), and YOLOv5 (96.8% accuracy) [**?**]. By integrating YOLOv8/YOLOv9 for localiza- tion and CNN/VGG19 for classification, LeafScan ensures precise, real-time diagnostics. Field trials in India, Kenya, and Brazil demonstrated a 30% reduction in crop losses and 20% decrease in pesticide use, supporting economic stability for smallholder farmers and environmental sustain- ability.

Compared to commercial platforms, LeafScan's mobile optimization (e.g., 55ms inference on high-end devices) and offline capability make it uniquely accessible. Its robustness across con- ditions (90% accuracy under 10 lux) and crops (99.1% for tomato) ensures versatility. Grad-CAM visualizations (Fig. 2) enhance interpretability, ad- dressing DL's "black box" issue [**?**]. The confusion matrix (Fig. 3) highlights areas for improvement, such as rare disease detection.

A. Limitations

- **Data Dependency**: Limited samples for rare diseases (e.g., 200 for mosaic virus) reduce generalization. Crowdsourced data requires robust preprocessing to mitigate noise.
- **Computational Cost**: VGG19's 39.6B FLOPs and 80MB footprint limit ultra-low-end device deployment.
- **Interpretability**: Complex models may con- fuse non-technical users, necessitating ad- vanced explainable AI (e.g., SHAP).
- **Geographic Variability**: Performance drops in extreme climates (94.5% in Arctic) require region-specific data.
- **Economic Barriers**: Smartphone access re- mains a challenge in remote areas.
- **Evolving Pathogens**: Emerging diseases ne- cessitate continuous dataset updates.
- **Energy Consumption**: High-end models (e.g., YOLOv9) consume more power, impacting battery life in rural settings.

B. Practical Implications

LeafScan's 93% expert agreement and SUS score of 85/100 confirm its usability. Its deployment re- duced yield losses by enabling early interventions, with farmers reporting 25–35% cost savings. In- tegration with agricultural cooperatives and gov- ernment subsidies could scale impact, while IoT and drone integration could enable automated monitoring.

C. *Future Enhancements*

Self-supervised learning, federated learning, and expanded datasets for rare diseases and ex- treme climates will enhance robustness. Energy- efficient architectures and advanced explainable AI will improve accessibility and trust.

VI. Conclusion

LeafScan redefines precision agriculture with a hybrid DL system achieving 99.0% accuracy in early plant disease detection. Integrating YOLOv8, YOLOv9, CNN, and VGG19, it delivers precise lo- calization and robust classification, outperform- ing systems like Plantix (90%) and AgriDiagnose (95%). Its mobile app, optimized for low-end de- vices, reduced crop losses by 30% and pesticide use by 20% in trials across India, Kenya, and Brazil. Offline and multi-language support en- sures inclusivity, empowering smallholder farm- ers and aligning with SDG 2 and SDG 12.

LeafScan's robustness (90% accuracy under ex- treme conditions) and versatility across crops (e.g., 99.1% for tomato) make it a scalable so- lution. Visualizations (Fig. 2) and error analysis (Fig. 3) enhance trust and guide improvements. Its impact extends beyond agriculture, fostering economic resilience and environmental sustain- ability in a world facing population growth and climate change.

Future work includes:

- **Self-Supervised Learning**: Reducing labeled data needs via contrastive learning.
- **Explainable AI**: Implementing SHAP and en- hanced Grad-CAM.
- **IoT and Drone Integration**: Enabling large-scale monitoring.
- **Dataset Expansion**: Including rare diseases and extreme climates.
- **Federated Learning**: Supporting privacy-preserving training.
- **Energy Efficiency**: Optimizing for ultra-low-power devices.
- **Global Scalability**: Partnering with NGOs and governments.
- **Pathogen Tracking**: Integrating genomic data for emerging diseases.

LeafScan's legacy will be its empowerment of farmers, ensuring sustainable food systems for future generations.

References

[1] Mohanty, S. P., Hughes, D. P., Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419.

[2] Ferentinos, K. P. (2018). Deep learning mod- els for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318.

[3] Barbedo, J. G. A. (2018). Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant dis- ease classification. Computers and Electron- ics in Agriculture, 153, 46–53.

[4] Kamilaris, A., Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. Com- puters and Electronics in Agriculture, 147, 70–90.

[5] Rajaraman, S., et al. (2019). Pre-trained con- volutional neural networks as feature ex- tractors toward improved malaria parasite detection in thin blood smear images. PeerJ, 7, e6976.

[6] Sladojevic, S., et al. (2016). Deep neural net- works based recognition of plant diseases by leaf image classification. Computational Intelligence and Neuroscience, 2016.

[7] Amara, J., Bouaziz, B., Algergawy, A. (2017). A deep learning-based approach for ba- nana leaf diseases classification. Datenbank- Spektrum, 17(3), 309–319.

[8] Zhang, K., et al. (2021). Automatic detection of cucumber leaf diseases using deep learn- ing with multi-scale feature fusion. Com- puters and Electronics in Agriculture, 181, 105952.

[9] Wang, G., Sun, Y., Wang, J. (2017). Automatic image-based plant disease severity estima- tion using deep learning. Computational In- telligence and Neuroscience, 2017.

[10] Too, E. C., et al. (2019). A comparative study of fine-tuning deep learning models for plant disease identification. Computers and Elec- tronics in Agriculture, 161, 272–279.

[11] Singh, D., Jain, N., Jain, P. (2020). Plant disease detection using deep convolutional neural network. Multimedia Tools and Ap- plications, 79, 18765–18782.