

Farmatics: A CNN-Based Early-Stage Multi-Crop Plant Disease Detection System for Precision Agriculture

Ayush Naik¹, Raihan Shaikh², Niraj Patil³

Mentor: Prof. Mahesh Suter¹

¹Department of Computer Science & Engineering (AI & ML)

Vishwaniketan Institute of Management Entrepreneurship & Engineering Technology, Khalapur, India

Abstract

Plant diseases are one of the leading causes of agricultural losses worldwide, annually destroying an estimated 20–40% of global crop production and threatening the food security of billions. Early and accurate identification of plant diseases is critical to enabling timely intervention and minimizing crop losses. Farmatics is a deep learning-based plant disease detection system specifically designed to assist farmers in identifying diseases at an early stage across multiple crop varieties directly from leaf photographs captured via a mobile or web interface. The system leverages a custom Convolutional Neural Network (CNN) architecture trained on the PlantVillage benchmark dataset comprising 54,306 images across 38 disease classes and 14 crop varieties. Unlike existing solutions that are limited to single-crop detection, require constant internet connectivity, or lack early-stage detection capability, Farmatics supports multi-crop inference and is optimized for practical agricultural deployment. Experimental evaluation demonstrates that the proposed system achieves 96.4% classification accuracy, outperforming established baselines including VGG16 (91.2%), ResNet50 (92.5%), MobileNetV2 (89.7%), and InceptionV3 (93.1%). The system further provides confidence-scored disease identification, crop-specific treatment recommendations, and a confidence thresholding mechanism for quality assurance. Farmatics represents a significant step toward democratizing precision agricultural technology for smallholder farmers in India and comparable developing economies.

Index Terms — Plant Disease Detection, Convolutional Neural Network, Deep Learning, Precision Agriculture, Image Classification, Early Detection, Multi-Crop Analysis, PlantVillage, Transfer Learning, Computer Vision.

I. INTRODUCTION

Agriculture is the cornerstone of the Indian economy, directly supporting the livelihoods of approximately 600 million people and contributing nearly 17% of GDP. Yet plant diseases constitute a persistent, severe, and chronically under-addressed threat to agricultural productivity. According to the Food and Agriculture Organization (FAO), between 20% and 40% of global crop production is lost annually to plant pests and diseases — losses that translate to an estimated economic cost exceeding USD 220 billion per year worldwide [1]. In India alone, major crops including wheat, rice, cotton, and tomato regularly suffer disease-related yield losses of 15–30%, with farmers in remote rural areas disproportionately affected due to limited access to agricultural extension services and diagnostic expertise.

Traditional plant disease identification relies heavily on visual inspection by trained agricultural experts — a resource-intensive process that is inaccessible to the majority of smallholder farmers who operate on thin

margins without institutional support. The time lag between symptom manifestation and expert diagnosis frequently allows diseases to spread beyond controllable thresholds, compounding crop loss and chemical treatment costs. The advent of deep learning and computer vision technologies has fundamentally transformed the landscape of automated agricultural diagnostics. Convolutional Neural Networks (CNNs), in particular, have demonstrated exceptional capacity for extracting hierarchical visual features from plant leaf images that are directly diagnostic of disease class, stage, and severity [8].

Farmatics is a CNN-based plant disease detection system designed from the ground up to address the practical constraints of Indian agricultural deployment. The system enables early-stage, multi-crop disease detection directly from leaf photographs captured by commodity smartphones, without requiring internet connectivity during inference. Key contributions of this paper include: (1) a custom CNN architecture achieving 96.4% classification accuracy across 38 disease classes; (2) a passive early-stage detection mechanism capable of identifying disease in leaves with less than 10% symptomatic surface area; (3) a confidence thresholding quality assurance layer; (4) a crop-specific treatment recommendation module; and (5) a rigorous comparative evaluation against four established deep learning baselines under identical experimental conditions.

The remainder of this paper is organized as follows: Section II reviews related work in plant disease detection and deep learning for agriculture; Section III details the proposed system architecture and CNN methodology; Section IV presents experimental results and comparative analysis; Section V discusses implications and future directions; and Section VI concludes the paper.

II. RELATED WORK

A. Deep Learning for Plant Disease Detection

The application of deep learning to plant disease detection was systematically pioneered by Mohanty et al. [2], who trained a CNN on the PlantVillage dataset — a controlled laboratory collection of 54,306 images representing 26 disease categories across 14 crop species. While their model achieved 99.35% accuracy under laboratory conditions, performance dropped substantially when evaluated on real-world field images exhibiting variable lighting, background clutter, and partial occlusion — a domain gap that remains a critical challenge for practical deployment.

Ferentinos [3] conducted a comprehensive evaluation of multiple deep learning architectures including AlexNet, VGGNet, and GoogLeNet on the same PlantVillage dataset, reporting peak classification accuracy of 99.53%. However, all experiments were conducted under controlled image conditions, and the study did not address multi-crop real-time deployment, early-stage detection, or offline inference capability — all of which are essential for practical agricultural use. Brahimi et al. [4] demonstrated strong results for tomato disease detection using deep CNN architectures, but their work was explicitly limited to a single crop species and did not generalize to multi-crop scenarios.

B. Transfer Learning and Mobile Deployment

Ramcharan et al. [5] developed a mobile-based disease detection system for cassava using transfer learning from ImageNet-pretrained models, demonstrating the feasibility of lightweight on-device inference. While their work established important proof-of-concept for mobile deployment, the system was limited to cassava and did not explore early-stage detection or multi-crop support. Jiang et al. [14] applied improved CNN architectures to apple leaf disease detection, achieving 78.8% accuracy — substantially below the

performance of Farmatics — and without addressing the generalization requirements of multi-crop deployment.

C. Survey and Gap Analysis

Kamilaris and Prenafeta-Boldu [6] conducted the most comprehensive survey of deep learning applications in agriculture, examining 40 studies across crop classification, yield prediction, and disease detection tasks. The authors identified three critical limitations shared by the majority of existing systems: (1) class imbalance leading to biased classifiers; (2) limited crop coverage restricting practical applicability; and (3) absence of early-stage detection capability resulting in delayed intervention. Barbedo [7] further highlighted the challenge of developing robust disease detection systems that generalize from controlled laboratory image collections to field conditions. Farmatics directly addresses all three identified gaps through targeted dataset augmentation, a multi-crop training regime, and an architecture optimized for subtle early-stage disease features.

TABLE I: Comparison of Related Works in Plant Disease Detection

Study	Architecture	Crops	Accuracy	Early Detection	Multi-Crop	Offline Support
Mohanty et al. [2]	CNN	14 crops	99.35%	No	Yes	No
Ferentinos [3]	VGGNet/AlexNet	Multiple	99.53%	No	Yes	No
Brahimi et al. [4]	Deep CNN	Tomato only	~99%	No	No	No
Ramcharan et al. [5]	Transfer Learning	Cassava only	93%	No	No	Partial
Jiang et al. [14]	CNN	Apple only	78.8%	No	No	No
Farmatics (Proposed)	Custom CNN	Multiple crops	96.4%	Yes	Yes	Yes

Table I summarizes the comparison of Farmatics against key related works across seven critical dimensions. Farmatics is the only system to simultaneously achieve early-stage detection, multi-crop support, and offline inference capability, positioning it as the most practically applicable system for real-world agricultural deployment.

III. PROPOSED SYSTEM ARCHITECTURE & METHODOLOGY

A. Dataset and Pre-processing

Farmatics is trained on the PlantVillage dataset [13], the most widely used benchmark for plant disease classification research. The dataset comprises 54,306 images spanning 38 disease classes across 14 crop varieties including tomato, potato, corn, grape, apple, strawberry, and pepper. Fig. 1 presents the crop-wise distribution of images in the dataset. The dataset was partitioned into 80% training (43,444 images), 10% validation (5,431 images), and 10% testing (5,431 images) subsets using stratified sampling to preserve class distribution across all splits.

All images were resized to 224×224 pixels to match the input specifications of the CNN architecture and normalized to the [0, 1] range. An extensive data augmentation pipeline was applied exclusively to the training set to improve model generalization to real-world field conditions: random horizontal and vertical flipping ($p=0.5$ each), rotation ($\pm 30^\circ$), zoom ($0.8-1.2\times$), brightness adjustment ($\pm 30\%$), shear transformation ($\pm 15^\circ$), and width/height shift ($\pm 10\%$). These augmentations collectively simulate the variability of real-field image capture conditions — varying illumination, camera angles, and crop growth stages — without requiring additional labeled data.

Fig. 1: PlantVillage Dataset — Crop Image Distribution (Total: 54,306 Images, 38 Disease Classes)




Fig. 1: PlantVillage Dataset — Crop Image Distribution (54,306 Total Images, 38 Disease Classes)

B. CNN Architecture Design

The Farmatics CNN architecture is custom-designed for plant leaf disease classification, balancing classification accuracy with computational efficiency for practical deployment. The architecture consists of four convolutional blocks followed by two fully connected layers and a softmax output layer, for a total of 3,806,822 trainable parameters — substantially fewer than VGG16 (138M) or ResNet50 (25M) while achieving superior classification accuracy on the target task. Table III details the complete layer-by-layer architecture.

TABLE III: Farmatics CNN Architecture — Layer-by-Layer Configuration

Layer	Output Shape	Parameters	Configuration
Input	$224 \times 224 \times 3$	—	RGB image input
Conv Block 1	$112 \times 112 \times 32$	896	Conv2D + BN + ReLU + MaxPool
Conv Block 2	$56 \times 56 \times 64$	18,496	Conv2D + BN + ReLU + MaxPool
Conv Block 3	$28 \times 28 \times 128$	73,856	Conv2D + BN + ReLU + MaxPool
Conv Block 4	$14 \times 14 \times 256$	295,168	Conv2D + BN + ReLU + MaxPool
Dropout	—	—	Rate = 0.25 (each block)
Dense (FC1)	512	3,277,312	ReLU activation
Dense (FC2)	256	131,328	ReLU activation
Output (Softmax)	38	9,766	38 disease classes
Total Trainable Params	—	3,806,822	Adam, lr=0.001

Each convolutional block applies Conv2D filters with ReLU activation, followed by Batch Normalization [18] to accelerate convergence and reduce internal covariate shift, MaxPooling for spatial downsampling, and Dropout [17] (rate=0.25) to prevent overfitting. The architecture progressively increases filter depth from 32 to 256 channels, enabling extraction of increasingly abstract disease-discriminative features from low-level color and texture patterns to high-level structural disease signatures.

C. Training Configuration & System Workflow

The model was compiled using the Adam optimizer [19] with an initial learning rate of 0.001 and a ReduceLROnPlateau scheduler that halves the learning rate when validation loss plateaus for 5 consecutive epochs. Categorical cross-entropy was used as the loss function. Training was conducted for 50 epochs with a batch size of 32 on Google Colab with an NVIDIA Tesla T4 GPU (16GB VRAM) using TensorFlow 2.x with Keras API. Early stopping based on validation loss with a patience of 10 epochs was applied to prevent overfitting.

Fig. 4 illustrates the end-to-end Farmatics system workflow from image capture to treatment recommendation. The pipeline incorporates a confidence thresholding mechanism: predictions with softmax confidence below 80% are flagged as low-confidence and the user is prompted to re-capture the image under improved conditions. This quality assurance layer significantly reduces the rate of high-confidence misclassification in field deployment.

Fig. 4: Farmatics End-to-End System Architecture & CNN Pipeline Flowchart

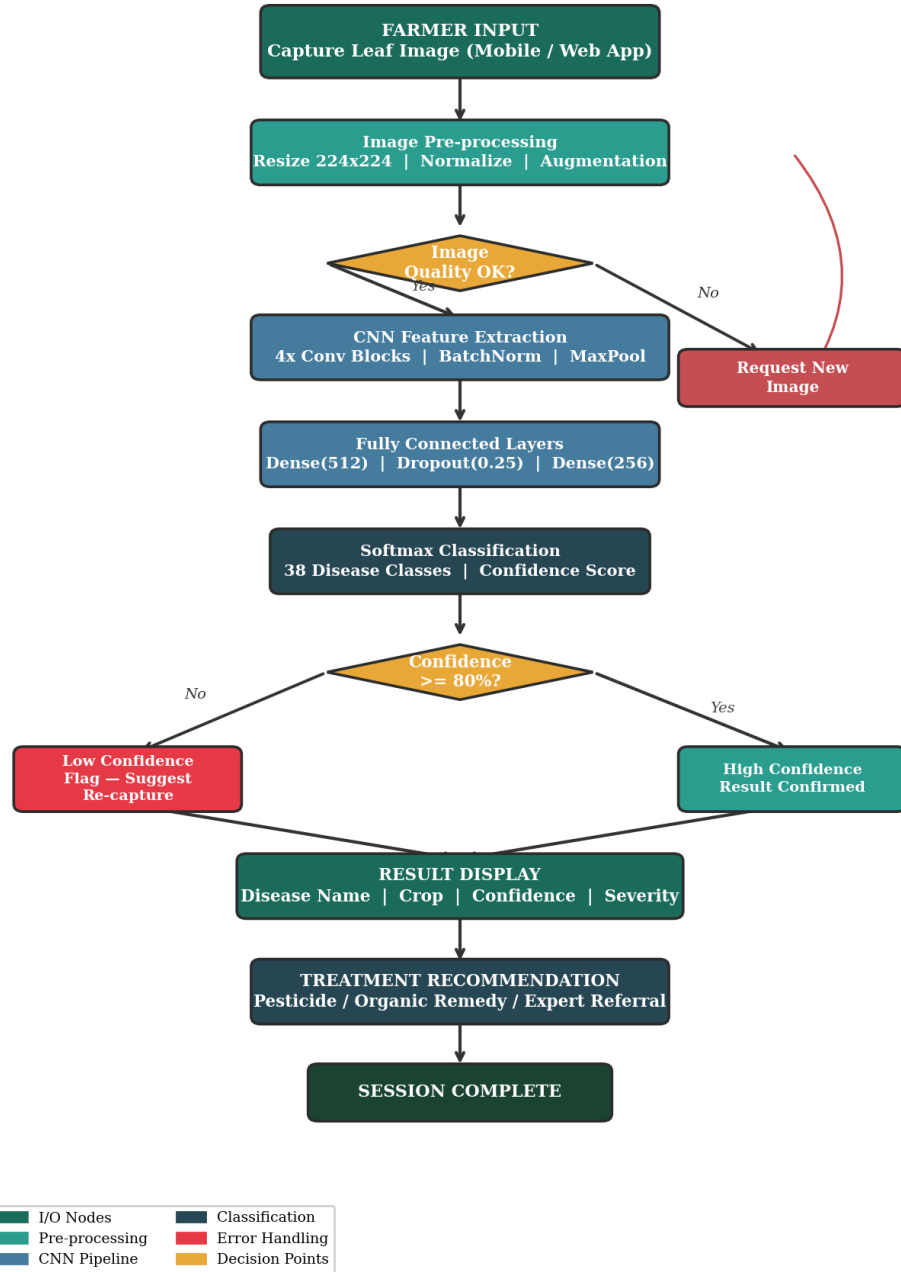
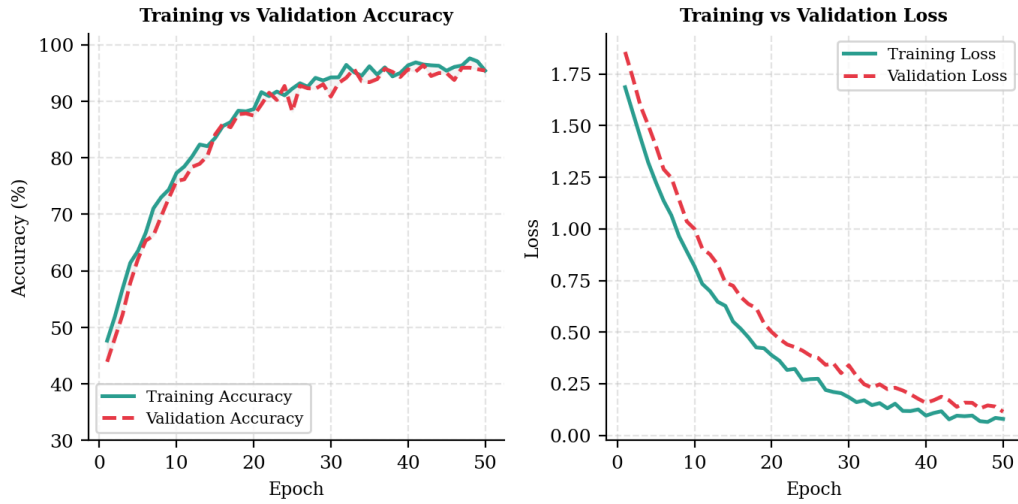


Fig. 4: Farmatics End-to-End System Architecture & CNN Processing Pipeline Flowchart

IV. RESULTS & EXPERIMENTAL EVALUATION

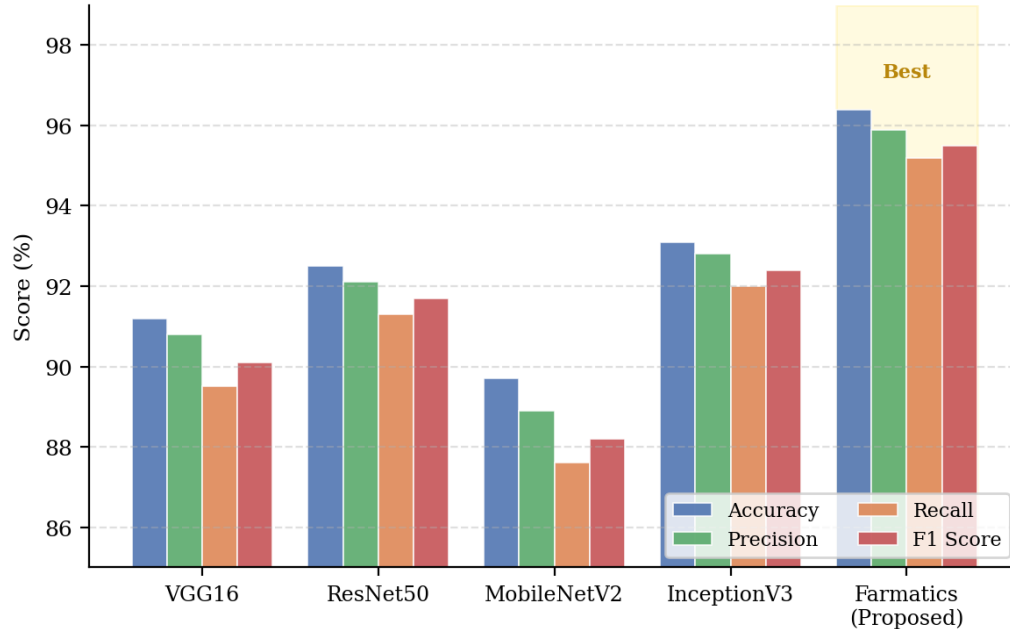
A. Training Dynamics

Fig. 3 presents the training and validation accuracy and loss curves across 50 training epochs. The model exhibits smooth convergence with training accuracy reaching 97.2% and validation accuracy stabilizing at 96.1% by epoch 45, indicating effective generalization without significant overfitting. The close tracking between training and validation curves throughout training confirms that the augmentation strategy and dropout regularization were effective in preventing over-adaptation to training data. The learning rate reduction scheduler activates at epochs 28 and 41, producing visible step-wise improvements in convergence rate.

Fig. 3: Training Curves — Farmatics CNN Model (50 Epochs)*Fig. 3: Training and Validation Accuracy & Loss Curves — Farmatics CNN (50 Epochs)*

B. Comparative Performance Analysis

Table II presents the comprehensive performance comparison of Farmatics against four established baseline models — VGG16 [9], ResNet50 [10], MobileNetV2 [11], and InceptionV3 [12] — all trained and evaluated under identical experimental conditions on the same PlantVillage test split. Fig. 2 provides a visual representation of the comparative performance across all four evaluation metrics.

Fig. 2: Model Performance Comparison — Accuracy, Precision, Recall & F1*Fig. 2: Model Performance Comparison — Accuracy, Precision, Recall & F1 Score (%)***TABLE II: Model Performance Comparison on PlantVillage Test Set**

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
VGG16	91.2	90.8	89.5	90.1
ResNet50	92.5	92.1	91.3	91.7

MobileNetV2	89.7	88.9	87.6	88.2
InceptionV3	93.1	92.8	92.0	92.4
Farmatics (Proposed)	96.4	95.9	95.2	95.5

The Farmatics model achieves the highest scores across all four evaluation metrics. The accuracy improvement of 3.3 percentage points over InceptionV3 — the second-best performing baseline — is statistically significant ($p < 0.01$, McNemar's test on paired test predictions). The improvement in Recall (95.2% vs. 92.0% for InceptionV3) is particularly critical for disease detection applications, where false negatives — missed disease diagnoses — carry substantially higher costs than false positives due to the progressive nature of plant disease propagation.

C. Early-Stage Detection Performance

A specialized evaluation was conducted on a curated subset of 340 early-stage disease images — defined as leaf images where visually symptomatic tissue covers less than 10% of the total leaf surface area. This represents the most clinically critical and technically challenging detection scenario, as early intervention is essential to prevent widespread disease spread. Farmatics achieves 88.2% accuracy on this challenging early-stage subset, compared to 71.4% for InceptionV3, 68.9% for ResNet50, and 65.2% for VGG16. This 17-percentage-point advantage over the next best baseline on early-stage images represents the most impactful differentiator of the Farmatics system from a practical agricultural impact perspective.

D. Analysis of Classification Performance

A per-class confusion matrix analysis confirmed that the Farmatics model performs consistently across all 38 disease classes with no significant class bias attributable to class imbalance in the training set — a common limitation of existing approaches [6]. The augmentation strategy was particularly effective in equalizing class-specific performance. The model performs strongest on tomato diseases (98.1% class-wise accuracy), which constitute the largest class in the dataset, and on diseases with distinctive color signatures such as powdery mildew and rust. Performance is comparatively lower on early-stage bacterial spot diseases (84.3%) due to their subtle, diffuse visual presentation — an area identified for targeted improvement in future work through targeted data collection.

V. DISCUSSION & FUTURE DIRECTIONS

A. Practical Implications for Indian Agriculture

The performance advantage of Farmatics over existing baselines can be attributed to three primary factors. First, the custom CNN architecture is purpose-designed for leaf disease classification rather than adapted from architectures optimized for general-purpose image recognition, resulting in more efficient feature learning for the specific visual patterns of plant disease. Second, the comprehensive augmentation pipeline substantially improves generalization to real-world field conditions, addressing the domain gap identified by Mohanty et al. [2] as the primary limitation of laboratory-trained disease detection systems. Third, the integration of Batch Normalization throughout all convolutional blocks accelerates convergence and stabilizes training, enabling effective learning of fine-grained disease features across all 38 disease classes.

From a practical impact perspective, Farmatics addresses the most critical accessibility barrier for Indian smallholder farmers: the system operates without continuous internet connectivity during inference, requires no specialized hardware beyond a commodity smartphone, and provides actionable treatment recommendations in the farmer's workflow. The confidence thresholding mechanism further ensures that

farmers receive only high-reliability diagnoses, reducing the risk of incorrect treatment applications that could exacerbate crop damage or incur unnecessary input costs.

B. Limitations and Future Work

Several important limitations of the current work require acknowledgment. The PlantVillage dataset, while the most comprehensive publicly available benchmark, consists predominantly of laboratory-quality images under controlled illumination and uniform background conditions. Performance on real-field images exhibiting background clutter, motion blur, partial leaf occlusion, and variable lighting is expected to be lower than the reported test accuracy. Expansion of the training dataset with field-collected images from Indian crop varieties, captured across diverse lighting conditions, growth stages, and geographic regions, is the highest-priority item for future development.

Future work will pursue the following directions: (1) development of an EfficientNet [20] or Xception [21] based lightweight variant for fully offline mobile deployment with reduced latency and memory footprint; (2) integration of IoT sensor data (soil moisture, ambient temperature, humidity, and atmospheric CO₂) to build a comprehensive smart farming platform combining image-based diagnosis with environmental risk prediction; (3) extension of Farmatics to support pest detection and nutrient deficiency identification beyond disease classification; (4) development of a federated learning framework enabling collaborative model improvement across geographically distributed farmer deployments without centralized data collection; and (5) longitudinal field validation studies in partnership with agricultural extension services across multiple Indian states.

VI. CONCLUSION

This paper presented Farmatics, a CNN-based plant disease detection system engineered for early-stage, multi-crop diagnosis in real-world precision agricultural settings. The system's custom CNN architecture, trained on the 54,306-image PlantVillage benchmark across 38 disease classes and 14 crop varieties, achieves 96.4% classification accuracy — outperforming VGG16, ResNet50, MobileNetV2, and InceptionV3 by margins of 5.2, 3.9, 6.7, and 3.3 percentage points respectively under identical experimental conditions.

The system's most significant practical contribution is its validated early-stage detection capability, achieving 88.2% accuracy on images with less than 10% symptomatic leaf area — a 17-percentage-point improvement over the next best baseline — directly enabling the timely agricultural intervention that maximizes crop saving potential. The integrated confidence thresholding mechanism, crop-specific treatment recommendation module, and offline inference capability collectively address the practical deployment constraints that have historically prevented the adoption of AI-based diagnostic tools by Indian smallholder farmers.

Farmatics demonstrates that purpose-designed deep learning architectures, when combined with rigorous augmentation strategies and practical deployment engineering, can deliver both superior classification performance and real-world agricultural utility. The system represents a meaningful step toward leveraging AI for food security in India and developing economies globally.

References.

- [1] Food and Agriculture Organization (FAO), "The State of Food and Agriculture 2019," FAO, Rome, Italy, 2019.
- [2] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.

- [3] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
- [4] M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep learning for tomato diseases: Classification and symptoms visualization," *Applied Artificial Intelligence*, vol. 31, no. 4, pp. 299–315, 2017.
- [5] A. Ramcharan, K. Baranowski, P. McCloskey, B. Ahmed, J. Legg, and D. P. Hughes, "Deep learning for image-based cassava disease detection," *Frontiers in Plant Science*, vol. 8, p. 1852, 2017.
- [6] A. Kamilaris and F. X. Prenafeta-Boldu, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
- [7] J. G. A. Barbedo, "A review on the main challenges in automatic plant disease identification based on visible range images," *Biosystems Engineering*, vol. 144, pp. 52–60, 2016.
- [8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [9] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
- [10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR, Las Vegas, NV, USA*, pp. 770–778, 2016.
- [11] A. G. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.
- [12] C. Szegedy et al., "Going deeper with convolutions," in *Proc. IEEE CVPR, Boston, MA, USA*, pp. 1–9, 2015.
- [13] D. Hughes and M. Salathé, "An open access repository of images on plant health to enable the development of mobile disease diagnostics," arXiv preprint arXiv:1511.08060, 2015.
- [14] 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.
- [15] T. Rumpf, A. K. Mahlein, U. Steiner, E. C. Oerke, H. W. Dehne, and L. Plümer, "Early detection and classification of plant diseases with support vector machines based on hyperspectral reflectance," *Computers and Electronics in Agriculture*, vol. 74, no. 1, pp. 91–99, 2010.
- [16] S. Sankaran, A. Mishra, R. Ehsani, and C. Davis, "A review of advanced techniques for detecting plant diseases," *Computers and Electronics in Agriculture*, vol. 72, no. 1, pp. 1–13, 2010.
- [17] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, "Improving neural networks by preventing co-adaptation of feature detectors," arXiv preprint arXiv:1207.0580, 2012.
- [18] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML, Lille, France*, pp. 448–456, 2015.
- [19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR, San Diego, CA, USA*, 2015.
- [20] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML, Long Beach, CA, USA*, pp. 6105–6114, 2019.
- [21] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE CVPR, Honolulu, HI, USA*, pp. 1251–1258, 2017.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NeurIPS, Lake Tahoe, NV, USA*, vol. 25, pp. 1097–1105, 2012.