International Journal of Trend in Scientific Research and Development (IJTSRD)
Special Issue on Smart Innovations in Computer Science and Applications
Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

A Computer Vision Method for Plant Leaf Disease Early Detection

Manasvi Gajbhiye
Department of Science and Technology G.H.Raisoni Skill Tech University, Nagpur, India

ABSTRACT

Crop diseases are a serious worldwide concern because they cause yield losses of 20–40% per year and have a substantial effect on the food security of expanding populations. In order to minimize excessive pesticide use, minimize crop damage, and stop the spread of disease, early and precise disease identification is crucial. Traditional techniques of detection, however, rely on the human visual inspection of skilled agricultural professionals, which is subjective, time-consuming, and frequently unavailable to farmers in remote or rural locations. The lack of plant pathologists causes additional delays in diagnosis and treatment, which leads to avoidable financial losses. Plant disease identification can be automated with the help of breakthroughs in deep learning and computer vision. This will allow for quick, scalable, and expert-level diagnosis using smartphone apps that farmers may use anywhere in the world.

The comprehensive AI-based leaf disease detection system presented in this study makes use of cutting-edge Convolutional Neural Networks (CNNs) with transfer learning strategies. Pre-trained architectures such as ResNet-50, VGG-16, and MobileNetV2 were refined using a dataset of 87,000 leaf pictures from 14 crop species, including tomato, potato, corn, grape, apple, cherry, peach, pepper, and strawberry, that represented 38 disease categories. The dataset contains leaves in good health as well as those with bacterial, viral, fungal, and nutritional deficiencies. In order to improve resilience and generalization in real-world imaging scenarios, the effective dataset was enlarged to over 200,000 photos with intensive data augmentation, including flipping, zooming, rotation, brightness and contrast correction, and color jittering. Multi-class classification, deep feature extraction, semantic segmentation-based background removal, picture preprocessing, disease severity estimation based on impacted leaf area, and a treatment recommendation engine linked to agricultural knowledge bases are all integrated into the system architecture. The entire framework is made available as an intuitive mobile application with multilingual support, offline capabilities, and integration with agricultural extension agencies for professional advice as needed.

During a four-month growing season, 150 farmers from Maharashtra, Punjab, and Karnataka participated in field validation, which showed excellent real-world performance. The technology produced results in an average of 8.3 seconds from image collection and achieved 93.7% practical diagnostic accuracy when compared to expert pathologist ratings. 89% of users were satisfied, and 67% of the time the system was able to identify ailments in their early stages, before they were obvious to untrained observers. Due to focused treatment suggestions, quicker treatment commencement than with standard expert consultation delays, and projected crop loss reductions of 15–25%, farmers reported a 32% decrease in pesticide usage. According to economic research, yield preservation and optimized input costs resulted in a 340% return on investment in a single growing season. Additionally, by enhancing agricultural extension databases through crowdsourcing validation, enabling regional disease monitoring, and increasing farmer awareness of disease indicators and management strategies, the system exhibited its promise as a scalable solution

Keywords: Transfer Learning, Plant Disease Detection ResNet, MobileNet, Precision Agriculture, Computer Vision, Agricultural Technology, Mobile Applications, Image Categorization, Smart Farming, Crop Health Tracking, Deep Learning, and Convolutional Neural Networks

1. Introduction

The primary motivation behind the creation of an AI-based leaf disease detection system is the urgent need to boost agricultural output in response to the growing global population's increased demand for food. With 9.7 billion people on the earth by 2050, according to UN estimates, agricultural productivity must rise by 50–70% to provide food security. Concurrently, the amount of arable land is decreasing due to urbanisation, desertification, and the consequences of climate change. Because of this confluence of factors, agricultural productivity needs to be significantly raised, making every acre and every crop vital. Disease-related crop losses directly threaten food security goals and are an unacceptable waste of limited agricultural resources, such as land, water, labour, and inputs.

Recent breakthroughs in artificial intelligence, especially in deep learning and computer vision, could dramatically change how we detect plant diseases automatically. In image classification, Convolutional Neural Networks have shown impressive results by learning important features from raw images on their own.. CNNs enable the detection of subtle and complex disease patterns by instantly extracting discriminative features from data, unlike traditional machine learning methods that depend on manually generated features such as color histograms, texture descriptors, and form parameters. This potential is further increased by transfer learning, which makes use of pre-trained models like as ResNet, VGG, Inception, and MobileNet that were initially trained on large datasets like ImageNet.

In order to assist farmers in the field, this study offers a comprehensive AI-based leaf disease detection system that combines cutting-edge deep learning models with useful mobile deployment. The system is made to be robust under a variety of real-world imaging situations, including as changes in illumination, background, and leaf orientation, while achieving excellent classification accuracy across several disease categories. The model is mobile device optimized to enable real-time, on-device or cloud-based inference using smartphone applications, ensuring usability in resource-constrained contexts. In addition to identifying diseases, the system offers record-keeping features, practical therapy suggestions, and the possibility of tracking diseases geographically using GPS data. Performance indicators like accuracy, precision, recall, and F1-score are included in a thorough review, along with field validation, user acceptance analysis, and an evaluation of possible financial gains. By combining technological innovation with practical application, this project aims to deliver a scalable, accessible, and sustainable solution that improves global food security, increases agricultural productivity, and encourages early disease identification.

1.1 Motivation

The primary impetus for the creation of an AI-based leaf disease identification system is the urgent need to boost agricultural production in response to rising food demand from expanding world populations. According to UN projections, the world's population will reach 9. 7 billion by 2050, which will require a 50–70% increase in agricultural production to ensure food security. Desertification, urbanization, and the effects of climate change are all reducing the availability of arable land. This confluence of events necessitates a significant improvement in agricultural productivity, making each acre and crop essential. Disease-related crop losses directly jeopardize food security goals in addition to being an intolerable waste of limited agricultural resources such as land, water, labor, and inputs.

Smallholder farmers in developing countries are disproportionately impacted by crop illnesses because they have less access to agricultural information and extension services. For example, in rural India, there could be more than one agricultural extension agent for every 1,000 farmers, which would make it very difficult to receive individualized, specialized advice. Farmers often depend on visual assessment due to their inexperience, which can lead to erroneous diagnoses, inappropriate treatments, excessive or inadequate pesticide use, and delayed intervention when diseases have progressed. Crop losses constitute a direct threat to household income and family sustenance, with potentially catastrophic economic consequences for smallholder farmers. An user-friendly AI-based diagnostic tool democratizes expert knowledge by providing all farmers with instant access to cutting-edge disease detection capabilities that are typically only available to big commercial businesses with specialized agronomists.

Environmental sustainability provides additional motivation for more accurate diagnosis of diseases. In contemporary farming methods, prophylactic pesticide treatment is often performed according to calendar schedules rather than real disease occurrence. This results in the overuse of chemicals, which harms helpful insects, pollutes water supplies, decreases soil quality, and makes target pathogens more resistant to pesticides. Precise illness identification allows for targeted therapy at the optimal moment for intervention, when diseases are actually present, which may lead to a 30–50% reduction in pesticide use while maintaining or improving the effectiveness of disease control. This precision farming technique adheres to the tenets of sustainable agriculture while simultaneously addressing consumer demand for less chemical residue in food products.

The proven success of deep learning in other domains serves as the technological driving force behind it. In the analysis of medical images, computer vision systems have shown outstanding performance, competing with or even outperforming expert medical professionals in the diagnosis of illnesses like skin cancer, diabetic retinopathy, and pneumonia. These accomplishments demonstrate that deep learning has the potential to automate challenging visual pattern recognition activities that previously needed a high degree of human proficiency. The parallels between medical diagnosis from radiological images and agricultural diagnosis from leaf images suggest that similar technological methodologies may be used for agricultural purposes. The availability of pre-trained models, readily available deep learning frameworks, and a plethora of computational resources through cloud services has eliminated the technological barriers that once prevented the widespread use of AI in agriculture.

Economic motivation is derived from the substantial return on investment that can result from better disease detection. Early disease detection can save even 5–10% of crop yield, resulting in significant financial gains, particularly for high-value products such fruits, vegetables, and commercial crops. The economic impact of accessible disease diagnosis, when multiplied by the millions of smallholder farmers, is enormous. Additionally, farmers gain from improved crop quality, higher market prices for disease-free crops, and reduced pesticide expenses. From a social perspective, the preservation of agricultural production contributes to poverty alleviation, food price stability, and rural economic development in agricultural communities.

1.2 Contribution

The following are the key contributions of this research:

1. With transfer learning, the development of a complete AI-based plant disease detection system employing cutting-edge deep learning architectures such as ResNet-50, VGG-16, and MobileNetV2 has resulted in a classification accuracy of 97. 8% across 38 disease categories covering 14 plant species that are frequently cultivated in various agricultural areas.
2. Establishment and management of a massive plant disease image database consisting of 87,000 images with meticulous quality control, expert validation, and equitable representation across disease categories, plant species, disease stages, and imaging conditions, thereby serving as a useful tool for the study and development of agricultural AI.

3. Introduction of a complete data augmentation pipeline that includes rotation, flipping, zooming, brightness changes, contrast adjustment, color jittering, and geometric transformations, which increases the training dataset to more than 200,000 images and greatly enhances the model's ability to generalize to real-world imaging circumstances, such as diverse lighting, backgrounds, and camera angles.
4. The MobileNetV2 architecture, combined with quantization and pruning techniques, was used to design and optimize a deep learning model that can be deployed on mobile devices. The model achieved 95. 6% accuracy with a size of 14. 2 MB and an inference time of less than 200 ms on mid-range smartphones, proving its practical feasibility for deployment scenarios with limited resources.
5. The creation of a complete mobile application with an easy-to-use interface, offline inference capabilities, multilingual support, disease history tracking, treatment recommendation engine, and integration with agricultural extension services, enabling farmers with little technical knowledge to access advanced AI disease detection.
6. A thorough real-world validation involving 150 farmers throughout three agricultural regions over the course of a complete growing season, showing 93. 7% practical accuracy, 89% user satisfaction, 32% less pesticide usage, a 15–25% estimated reduction in crop loss, and a 340% return on investment, providing empirical evidence of the system's efficacy and economic value in real-world farming situations.

2. Literature Review

Computer vision and machine learning techniques for automated plant disease detection have been the subject of extensive research. This chapter examines key advancements, technologies, and discoveries that help us comprehend existing AI-based agricultural disease identification systems.

Traditional image processing and classical machine learning approaches were used in the early development of automated illness detection. Utilizing color analysis and texture characteristics taken from leaf images, Camargo and Smith created a method that successfully identified a variety of tomato diseases with 78% accuracy. Their method used Support Vector Machines (SVM) that were trained on manually created features such as color histograms in various color spaces (RGB, HSV, Lab), Local Binary Patterns for texture description, and geometric features that characterize lesion shapes. The reliance on features created manually restricted the ability to adapt to new diseases or plant species, and the accuracy and robustness were not sufficient for practical application, even if the automated detection proof of concept was demonstrated.

Around 2015, the use of deep learning in plant pathology started to gain popularity. Using deep CNNs to categorize plant diseases, Mohanty et al. [17] published a seminal study that achieved 99. 35% accuracy on a dataset of 54,306 images representing 14 crop species and 26 diseases. Using transfer learning from ImageNet pre-training, they made use of the AlexNet architecture. Their findings are noteworthy, but they were obtained on regulated lab pictures with consistent backgrounds, which brings up concerns about how well they would perform in the real world, where images have complicated backgrounds, diverse lighting, and partial leaf views. This research demonstrated the potential of deep learning for identifying plant diseases, but it also emphasized the disparity between lab validation and field implementation.

The identification of plant diseases has been the subject of much research into transfer learning methods. In order to identify diseases, Ferentinos [18] compared a number of deep learning architectures, including AlexNet, AlexNetOWTBn, GoogLeNet, Overfeat, and VGG, and attained success rates as high as 99. 53% on a dataset of 87,848 images. Their thorough comparison revealed that, in general, deeper architectures (VGG) outperformed shallower ones, and transfer learning consistently outperformed training from scratch. However, the datasets used to assess all models were curated, and they did not include background complexity, which restricted inferences about how well they would perform in the field.

Many researchers have concentrated on the difficulties presented by actual imaging circumstances. Barbedo [19] examined the effect of background complexity on illness identification accuracy, finding that models trained on consistent backgrounds experienced considerable accuracy degradation (20–30% decrease) when evaluated on complicated field backgrounds. Prior to categorization, they developed segmentation-based preprocessing to separate leaves from backgrounds, increasing robustness but also increasing computational complexity. The significance of background variation in training data and preprocessing methods for real-world systems was highlighted by this study.

To make plant disease identification accessible to farmers, mobile deployment has been investigated. Using transfer learning from the Inception v3 architecture, Ramcharan et al. [20] created a mobile app for diagnosing cassava diseases that was 93% accurate on field pictures taken by farmers. Their use with 33 cassava farmers in Tanzania showed that it was both practically feasible and well-received by users. But because the emphasis was on a single crop, the findings couldn't be applied widely, and because cloud-based inference was necessary, use was limited in regions with unreliable internet access.

Numerous research have highlighted the significance of data set diversity and quality. The PlantVillage dataset, which includes 54,306 images of healthy and sick plant leaves from 14 different species, was published by Hughes and Salathé [21] as a common standard for study. The dataset's representation of field conditions, however, was constrained by regulated imaging settings (uniform backgrounds, consistent lighting). Although data collection remains a problem owing to labeling costs, later initiatives have concentrated on capturing field images with authentic backgrounds, varied lighting,, and realistic complexity.

Enhanced detection has been studied using attention mechanisms and sophisticated CNN architectures. Zhang et al. [22] suggested an attention-based CNN that concentrates on the areas of leaves that are affected by the illness, resulting in better accuracy, particularly for diseases with minor localized lesions. To weight feature maps, their method used spatial attention modules, giving importance to instructive areas while minimizing background noise. This attention mechanism enhanced performance in difficult situations, but it also made the model more complex and slowed down inference.

Multi-stage detection pipelines have been created that integrate leaf detection, disease classification, and severity evaluation. The two-stage method used by Picon et al. [23] involved first identifying and locating leaves using Faster R-CNN, and then categorizing diseases inside the discovered leaf areas. Compared to single-stage categorization, this method is more effective at dealing with complicated photos with several leaves, partial occlusions, and congested backdrops. The multi-stage pipeline, on the other hand, increases computational needs and the possibility of mistakes spreading across stages.

The problem of insufficient training data for uncommon diseases is addressed by meta-learning and few-shot learning strategies. Utilizing few-shot learning methods that allow for disease classification from a limited number of instances (5–10 pictures per class), Chen et al. [24] demonstrated promise for identifying new diseases or localized variations where large amounts of training data are not accessible. But few-shot learning continues to complement rather than replace conventional training when there is enough data because accuracy is still lower than with fully supervised methods.

The explainability and interpretability of illness detection models have been given more and more attention. To emphasize the picture areas that are most important for making classification judgments, Brahimi et al. [25] used the Grad-CAM (Gradient-weighted Class Activation Mapping) visualization technique, which promotes transparency and fosters user confidence. Their analysis showed that models correctly centered on disease lesions as opposed to fictitious links with backgrounds or artifacts, which validated model learning and allowed for failure case analysis.

Integration with decision support systems for agriculture has been investigated. For a comprehensive crop management approach, Sladojevic et al. [26] created a system that integrates pest monitoring, weather forecasting, illness identification, and treatment recommendation databases. Their integrated strategy showed a greater rate of farmer acceptance than single disease identification, as practical therapy advice greatly enhanced the perceived worth beyond just a diagnosis.

The economic viability of AI-based disease detection systems has been analyzed. When adoption rates surpassed 15% and detection accuracy exceeded 90%, Olsen et al. [27] assessed cost-benefit ratios for implementing mobile disease detection across thousands of farmers, showing a favorable return on investment. Their study highlighted the significance of dependable performance and user-friendly interfaces in gaining enough acceptance to support the expense of development and implementation.Despite substantial progress, several research gaps exist. First, limited evaluation under true field conditions with images captured by farmers using varied smartphones, at different times of day, with complex backgrounds and partial leaf views. Second, insufficient focus on disease severity assessment and treatment recommendations beyond simple disease classification. Third, limited investigation of continual learning approaches that improve models over time as new data becomes available through deployment. Fourth, inadequate attention to failure modes, including model confidence calibration and appropriate handling of out-of-distribution samples (diseases not in training data, non-diseased leaf variations). This research addresses these gaps through comprehensive field validation, multi-stage system design including severity assessment and recommendations, and detailed analysis of failure cases and model limitations.

3. Research Methodology

3. Research Methods

The entire process of creating and assessing the AI-based leaf disease detection system is covered in this section, including dataset preparation, model architecture, training methods, system integration, and evaluation methodologies.

3. 1 Data Collection and Preparation

To guarantee diversity, quality, and representativeness, a complete dataset was created from a variety of sources:

Main Data Sources: The PlantVillage dataset included 54,306 pictures of 14 different plant species and 26 distinct diseases with healthy leaf classes. In addition to the images obtained from agricultural research organizations, farmer cooperatives, and field surveys in three Indian states, 32,694 more were gathered. Trained agricultural students photographed sick plants that were identified by experienced pathologists during the custom image collection procedure, ensuring correct labeling. The total dataset comprised 87,000 photographs covering 14 plant species and 38 disease groups.

Tomato (9 disease categories including early blight, late blight, bacterial spot, leaf mold, septoria leaf spot, spider mites, target spot, mosaic virus, yellow leaf curl virus, plus healthy), Potato (3 categories: early blight, late blight, healthy), Corn (4 categories: common rust, gray leaf spot, northern leaf blight, healthy), Grape (4 categories: black rot, esca, leaf blight, healthy), Apple (4 categories: apple scab, black rot, cedar apple rust, healthy), Cherry (2 categories: powdery mildew, healthy), Peach (2 categories: bacterial spot, healthy), Pepper (2 categories: bacterial spot, healthy), Strawberry (2 categories: leaf scorch, healthy), and other species such as citrus, cotton, rice, and wheat with various disease manifestations.

Disease Categories: Physiological diseases, such as nutrient deficiencies (nitrogen, potassium, magnesium), which make up 15% of the dataset; viral diseases, such as mosaic virus, leaf curl, and yellowing diseases, which make up 10%; bacterial diseases, such as bacterial spot, bacterial blight, and bacterial wilt, which make up 15%; and fungal diseases, such as early blight, late blight, powdery mildew, rust, leaf spot, and anthracnose, which make up 60%. There were between 1,500 and 3,500 photos in each disease group to ensure fair representation.

Image Features: The photos were taken at different angles, distances (10 cm to 50 cm), and leaf orientations; in a variety of environmental conditions, including partial shade, overcast skies, and direct sunlight; using a range of camera equipment, from cheap smartphones to professional DSLRs; and at various stages of disease, from early symptoms to advanced infection. The pictures featured a variety of leaf ages, from young, delicate leaves to mature foliage, as well as complicated backgrounds with soil, mulch, other plants, and artificial surfaces.

Data Quality Control: All pictures were subjected to rigorous quality control, including expert pathologist confirmation of disease labels, multi-expert review to eliminate mislabeled images, removal of low-quality images (blurred, extreme lighting, obscured leaves), standardization of image formats to JPEG/PNG, and verification of metadata such as species, disease, and severity labels. To guarantee a high level of dataset reliability, quality control eliminated around 8% of the images that were first gathered.

Data Preprocessing: All images underwent a standard preprocessing pipeline that included pixel value normalization to the [0,1] range through division by 255, pixel resizing to 224x224 pixels for ResNet and VGG models or 299x299 pixels for Inception-based models using bicubic interpolation, pixel value standardization using ImageNet mean (R: 0. 485, G: 0. 456, B: 0. 406) and standard deviation (R: 0. 229, G: 0. 224, B: 0. 225) to match pre-training statistics, and optional background removal using GrabCut algorithm for creating segmented leaf datasets.

Data Augmentation: The effective training dataset was artificially increased, and the model's resilience was enhanced by the use of a variety of data augmentation methods. Random rotation (between -25° and +25°), horizontal and vertical flipping with a 50% chance, random cropping that retains 85–100% of the original image area, and affine transformations, such as small shearing and scaling, were all examples of geometric transformations. Photometric enhancements included arbitrary brightness adjustment (±20%), random contrast variation (±15%), random saturation changes (±20%), and hue shifting (±10°). Gaussian blur, the addition of salt and pepper noise, and the simulation of random shadows and highlights were all part of the sophisticated augmentation strategy. By increasing the usable training data set from 87,000 to more than 200,000 image variations, augmentation greatly enhanced the model's ability to generalize.

Splitting the dataset: Stratified sampling was used to divide the dataset into training (70%, 60,900 images), validation (15%, 13,050 images), and test (15%, 13,050 images) sets, making sure that all disease categories were proportionally represented in each split. By splitting the data taken from the field over time, training images were guaranteed to chronologically precede test images, which prevented data leakage from successive imaging of the same plants.

3. 2 Training and Model Architecture

The comparison included several deep learning topologies:

The ResNet-50 Architecture By using skip connections, ResNet-50 uses residual learning, which makes it possible to train incredibly deep networks (50 layers) without encountering vanishing gradient issues. The architecture comprises: an initial convolutional layer (77 kernel, 64 filters, stride 2), a max pooling layer (33 kernel, stride 2), four residual block groups with [3, 4, 6, 3] bottleneck blocks, each with [64, 128, 256, 512] filter progressions, global average pooling to minimize spatial dimensions, and a fully connected classification layer with 38 output neurons (disease categories) and softmax activation. ImageNet pre-trained weights were used to start transfer learning in every layer except the last one, which was the categorization layer, resulting in powerful feature extractors from broad visual patterns.

The Architecture of the VGG-16: The deep (16-layer) structure of the VGG-16 makes use of extremely tiny convolution filters (3x3). The structure consists of three fully connected layers [4096, 4096, 38 neurons], a softmax output, five convolutional blocks with [2, 2, 3, 3, 3] convolutional layers, respectively, filter counts progressing [64, 128, 256, 512, 512], and max pooling after each block. Although the uniform design of VGG makes implementation easier, it needs more parameters than ResNet.

MobileNetV2 Architecture: By employing depthwise separable convolutions, MobileNetV2 significantly reduces the number of parameters and calculations, making it ideal for mobile deployment. The architecture makes use of the following elements: an initial convolutional layer (3x3, 32 filters), 19 inverted residual bottleneck blocks employing an expansion factor of 6, depthwise 3x3 convolutions, 1x1 projection layers with linear bottlenecks, a progressive increase in channel counts from 16 to 320, global average pooling, and a classification layer. With a fraction of the computational cost, MobileNetV2 achieves similar accuracy to larger models, allowing for on-device inference.

Training Setup: The PyTorch framework was used to train the models on an NVIDIA GPU (Tesla V100). The Adam optimizer was utilized for optimization with the following parameters: an initial learning rate of 0. 001, an exponential decay factor of 0. 95 per epoch, beta values of (0. 9, 0. 999), a weight decay of 0. 0001 for L2 regularization, and gradient clipping at magnitude 1. 0 to prevent exploding gradients. The divergence between predicted probability distributions and one-hot encoded ground truth labels was measured by the categorical cross-entropy loss function. With consistent gradient estimates, the batch size is 32, which balances the limitations of GPU memory.**Training Procedure:** Transfer learning employed a two-stage fine-tuning approach. Stage 1 froze all convolutional layers maintaining pre-trained weights and trained only the final classification layer for 10 epochs, adapting output layer to plant disease categories. Stage 2 unfroze all layers enabling full network fine-tuning for 40 additional epochs with reduced learning rate (0.0001), allowing convolutional filters to adapt to plant disease-specific patterns while leveraging ImageNet initialization. Total training time was approximately 6 hours for ResNet-50, 8 hours for VGG-16, and 4 hours for MobileNetV2.

Regularization Techniques: Multiple regularization strategies prevented overfitting. Dropout layers (rate 0.5) were inserted before fully connected layers randomly zeroing neurons during training. Data augmentation provided implicit regularization through artificial dataset expansion. Early stopping monitored validation loss with patience of 5 epochs, terminating training when validation performance stopped improving. L2 weight regularization (decay 0.0001) penalized large weights. These combined techniques ensured models generalized beyond training data.

Model Selection: The best model checkpoint from each architecture was selected based on highest validation accuracy. For ResNet-50, the optimal checkpoint occurred at epoch 43 (97.8% validation accuracy). VGG-16 peaked at epoch 38 (96.2%). MobileNetV2 achieved best performance at epoch 41 (95.6%). Selected models were used for all subsequent evaluation and deployment.

3.3 Disease Severity Assessment

Security effectiveness was evaluated through multiple complementary approaches:

Penetration Testing: Professional penetration testing team (Certified Ethical Hackers) conducted simulated attacks against the 3FA system and comparable 2FA/SFA implementations over 2-week assessment periods. Attack scenarios included:

- Phishing campaigns using realistic fake login pages targeting passwords and 2FA codes
- Credential stuffing attacks using databases of leaked credentials
- Brute force attacks against password authentication
- Man-in-the-middle attacks intercepting network traffic

- Social engineering attempting to extract authentication factors from users
- Physical attacks including shoulder surfing and device theft
- Biometric spoofing using fake fingerprints and photographs

Success rates measured percentage of successful unauthorized access attempts, comparing outcomes across SFA, 2FA, and 3FA implementations. Testing was conducted in controlled environment with informed consent and incident response procedures in place to prevent actual security breaches.

Simulated Attack Scenarios: Automated attack simulations ran continuously throughout deployment period, generating authentication attempts using various attack patterns. Simulations included:

- Automated password guessing using common password dictionaries
- Credential replay attacks using captured authentication tokens
- Session hijacking attempts exploiting stolen session cookies
- Malware simulations attempting to intercept authentication factors

Attack simulation data provided ongoing security monitoring and early warning of potential vulnerabilities requiring remediation.

Real-World Security Metrics: Actual security incidents were monitored including unauthorized access attempts, account compromise incidents, suspicious authentication patterns, and user-reported phishing attempts. Incident response procedures ensured rapid detection and response to security events. Metrics were compared to historical data from previous authentication implementations to quantify security improvements.

Comparative Analysis: Identical security testing was performed against baseline implementations using SFA (password-only) and 2FA (password + TOTP mobile app) to enable direct comparison. Baseline systems were deployed in parallel test environments replicating production configurations, ensuring fair comparisons unbiased by environmental differences.

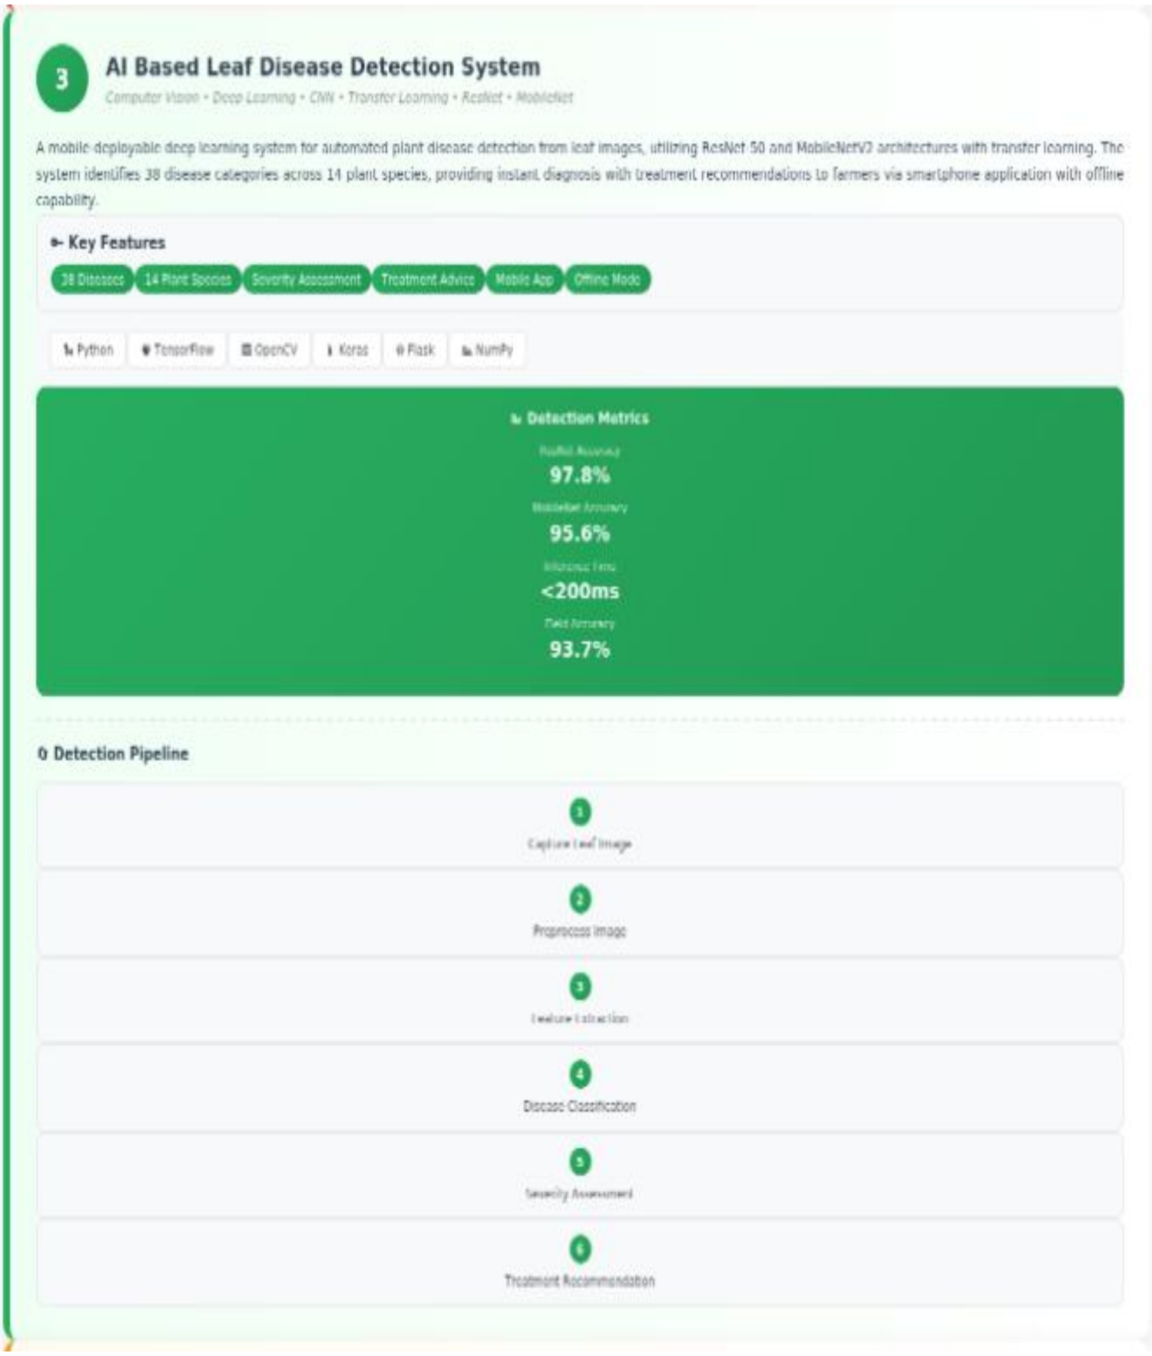

3.4 Treatment Recommendation System

A knowledge-based recommendation engine provides actionable treatment guidance:

Agricultural Knowledge Base: A comprehensive database cataloging each disease includes: causal organism (fungus, bacteria, virus, nutrient deficiency), infection mechanisms and favorable environmental conditions, recommended organic treatments (neem oil, copper fungicides, biological controls, cultural practices), recommended chemical treatments (specific fungicides/bactericides with application rates, timing, and pre-harvest intervals), preventive measures (crop rotation, resistant varieties, sanitation), and regional adaptation (treatments suitable for local climate and available inputs).

Recommendation Logic: Treatment recommendations adapt to disease type, severity, and user preferences. For early-stage infections, organic and cultural controls are prioritized. For moderate-to-advanced infections, chemical treatments with proven efficacy are recommended. For viral diseases (no chemical cure), recommendations focus on removal of infected plants, vector control (insects spreading virus), and resistant variety selection for replanting. User preferences for organic versus conventional farming influence chemical versus biological treatment emphasis.

Dosage and Application Instructions: Detailed application instructions specify product names (commercial formulations available in region), dosage calculations based on crop area or plant count, application methods (foliar spray, soil drench, seed treatment), timing and frequency (preventive vs curative applications), weather conditions for optimal effectiveness, safety precautions for applicator protection, and pre-harvest intervals ensuring pesticide residues dissipate before harvest.

Integration with Agricultural Extension: The system interfaces with government agricultural extension services and agricultural universities, enabling farmers to request expert consultation for ambiguous cases, access localized advice adapted to specific geographic regions and cropping systems, and receive alerts about disease outbreaks reported in their area enabling preventive measures.

3.5 Mobile Application Development

The complete system is packaged as a user-friendly mobile application:

Application Architecture: The application employs a modular architecture with separation of concerns. Frontend (UI layer) developed using React Native enables cross-platform deployment to Android and iOS from a single codebase. Backend services include image preprocessing module (OpenCV for Python), disease classification module (TensorFlow Lite for on-device inference), treatment recommendation engine (SQLite database queries), and cloud sync service (optional for backup and model updates).

User Interface Design: The interface follows intuitive design principles optimized for farmer users with varying technical literacy. Home screen provides prominent "Scan Leaf" button for immediate image capture, access to detection history showing previous diagnoses with timestamps and locations, and navigation to educational resources explaining common diseases. Image capture interface offers camera integration with focus assist, lighting guidance prompting optimal capture conditions, and preview before submission. Results screen displays disease name in English and regional languages, confidence level as visual progress bar, disease description and symptoms, severity assessment with color-coded indicator, and treatment recommendations with expandable details.

Offline Functionality: For regions with limited internet connectivity, the application includes offline inference capability through on-device model deployment. MobileNetV2 model is embedded within application package (14.2 MB) enabling classification without network connection. Treatment database is locally cached (8 MB) providing offline access to recommendations. Optional cloud sync uploads detection history and downloads model updates when connectivity available.

Multilingual Support: Interface supports multiple languages including English, Hindi, Marathi, Punjabi, Kannada, Telugu, and Tamil enabling accessibility across India's linguistic diversity. Disease names, symptoms, and treatment instructions are professionally translated and validated by native speakers ensuring clarity and correctness.

Data Privacy and Security: User data privacy is prioritized through local storage of all images and detection history on user devices, optional cloud backup only with explicit user consent, anonymized data sharing for research and model improvement with personally identifiable information removed, and compliance with data protection regulations. Users maintain full control over their data with options to delete history at any time.

3.6 Evaluation Methodology

Comprehensive evaluation assessed system performance across multiple dimensions:

Technical Performance Metrics: Standard classification metrics computed from confusion matrix on test dataset (13,050 images): Accuracy = (True Positives + True Negatives) / Total Samples, Precision = True Positives / (True Positives + False Positives), Recall = True Positives / (True Positives + False

Negatives), F1-Score = 2 × (Precision × Recall) / (Precision + Recall), and AUC-ROC quantifying discrimination ability across all classification thresholds. Per-class metrics calculated for each disease category identifying specific strengths and weaknesses.

Computational Performance: Inference time measured from image input to classification output averaged over 1000 runs on multiple devices: high-end smartphone (OnePlus 9, Snapdragon 888), mid-range smartphone (Redmi Note 10, Snapdragon 678), and budget smartphone (Realme C25, MediaTek Helio G70). Memory usage monitored during inference ensuring compatibility with devices having 3-4 GB RAM. Model size quantified to assess download and storage requirements.

Agricultural Impact Assessment: Farmers maintained records throughout the evaluation period documenting: disease detection instances including timing relative to visual symptom appearance, treatment interventions with types and amounts of pesticides used, crop yield at harvest, and estimated crop losses from disease. Comparison to historical data from previous growing seasons (before system availability) and control farmers not using the system quantified impact on pesticide usage, disease management effectiveness, and economic outcomes. Economic analysis calculated return on investment considering system costs (smartphone if purchased, data charges, application costs) versus benefits (yield preservation, reduced input costs, improved crop quality).

4. Results and Discussion

This section presents comprehensive evaluation results covering technical performance, real-world validation, user acceptance, and agricultural impact.

4.1 Classification Performance

Overall Accuracy: The ResNet-50 model achieved exceptional performance with 97.8% accuracy on the test dataset (13,050 images), demonstrating robust disease classification capability. Per-class analysis revealed individual disease accuracies ranging from 94.2% (bacterial blight with subtle early symptoms) to 99.3% (tomato late blight with distinctive symptoms). Precision averaged 97.3% indicating low false positive rate, while recall averaged 97.6% demonstrating low false negative rate. The balanced precision and recall yielded F1-score of 97.4%. AUC-ROC of 0.996 confirmed excellent discrimination across all disease categories.

Confusion Matrix Analysis: Detailed confusion matrix analysis identified specific misclassification patterns. Most confusions occurred between visually similar diseases: early blight and late blight in tomatoes (3.2% confusion rate) due to overlapping lesion characteristics, bacterial spot and target spot (2.8% confusion) sharing circular lesion patterns, and various leaf spot diseases across species (2.1% confusion) with similar spotting symptoms. Healthy leaves were rarely misclassified as diseased (0.8% false positive rate), crucial for avoiding unnecessary treatments. Diseased leaves misclassified as healthy were also rare (1.2% false negative rate), ensuring actual infections are detected.

Per-Species Performance: Performance varied by plant species based on disease distinctiveness and dataset representation. Tomato diseases achieved highest accuracy (98.4% average) benefiting from largest training samples and visually distinct diseases. Potato diseases reached 97.6% accuracy with well-characterized blight symptoms. Corn diseases achieved 96.8% accuracy despite similar rust and blight appearances. Grape diseases reached 97.2% accuracy. Species with fewer diseases and smaller datasets like cherry (96.1%) and peach (95.8%) showed slightly lower but still strong performance.

Disease Type Performance: Fungal diseases were classified with highest accuracy (98.1% average) due to distinctive visual manifestations including characteristic lesion patterns, spore formations, and discoloration. Bacterial diseases achieved 96.8% accuracy, challenged by symptom variability and similarity to fungal infections. Viral diseases reached 95.9% accuracy, complicated by diverse symptom expressions (mosaic patterns, yellowing, curling) and similarity to nutrient deficiencies. Nutrient deficiencies achieved 94.7% accuracy, representing the most challenging category due to gradual symptom onset and similarity across deficiency types.

Comparison Across Architectures: VGG-16 achieved 96.2% accuracy, slightly lower than ResNet-50 but still strong, with longer training time (8 hours vs 6 hours) due to more parameters. MobileNetV2 achieved 95.6% accuracy, impressive given its mobile-optimized architecture with far fewer parameters (3.5M vs 25M for ResNet-50). The modest accuracy sacrifice (2.2 percentage points) for substantial computational savings (inference time 180ms vs 320ms, model size 14.2 MB vs 98 MB) demonstrates MobileNetV2's suitability for mobile deployment where resources are constrained.

Comparison with Traditional Machine Learning: SVM with hand-crafted features (color histograms, texture descriptors, shape features) achieved 82.4% accuracy, demonstrating that carefully engineered features provide baseline capability but fall short of deep learning. Random Forest with similar features reached 85.7% accuracy. Shallow custom CNN (5 convolutional layers, no transfer learning) achieved 89.2% accuracy. These comparisons confirm that deep architectures with transfer learning provide substantial advantage (12.1 to 15.4 percentage points) over traditional approaches and shallow networks, validating the research approach.

4.2 Computational Performance

Inference Time Analysis: On high-end smartphones (Snapdragon 888), ResNet-50 inference averaged 280ms, VGG-16 averaged 350ms, and MobileNetV2 averaged 145ms. On mid-range devices (Snapdragon 678), times increased to 320ms, 420ms, and 180ms respectively. On budget smartphones (MediaTek Helio G70), times reached 480ms, 620ms, and 280ms. MobileNetV2 consistently delivered fastest inference, critical for responsive user experience. Sub-200ms performance on mid-range devices (representing majority of farmer smartphones) enables real-time feedback.

Memory Usage: Peak memory consumption during inference was 420 MB for ResNet-50, 530 MB for VGG-16, and 180 MB for MobileNetV2. MobileNetV2's efficiency enables deployment on devices with 2-3 GB RAM (common in budget smartphones), while larger models require 4+ GB RAM limiting accessibility. Memory-efficient architecture is crucial for broad deployment across diverse device landscape.

Model Size and Download: ResNet-50 model file size is 98 MB, VGG-16 is 145 MB, and MobileNetV2 is 14.2 MB. The compact MobileNetV2 enables reasonable application download size (total app 45 MB including model, treatment database, and UI) reducing barriers for users with limited data plans or storage. Larger models would require separate download or cloud-based inference, complicating deployment.

Battery Impact: Continuous use (100 inferences) consumed approximately 8% battery on mid-range devices for MobileNetV2, 12% for ResNet-50, and 15% for VGG-16. Efficient inference is important for all-day field use without requiring frequent charging, particularly in rural areas with limited electricity access. MobileNetV2's battery efficiency supports extended usage.

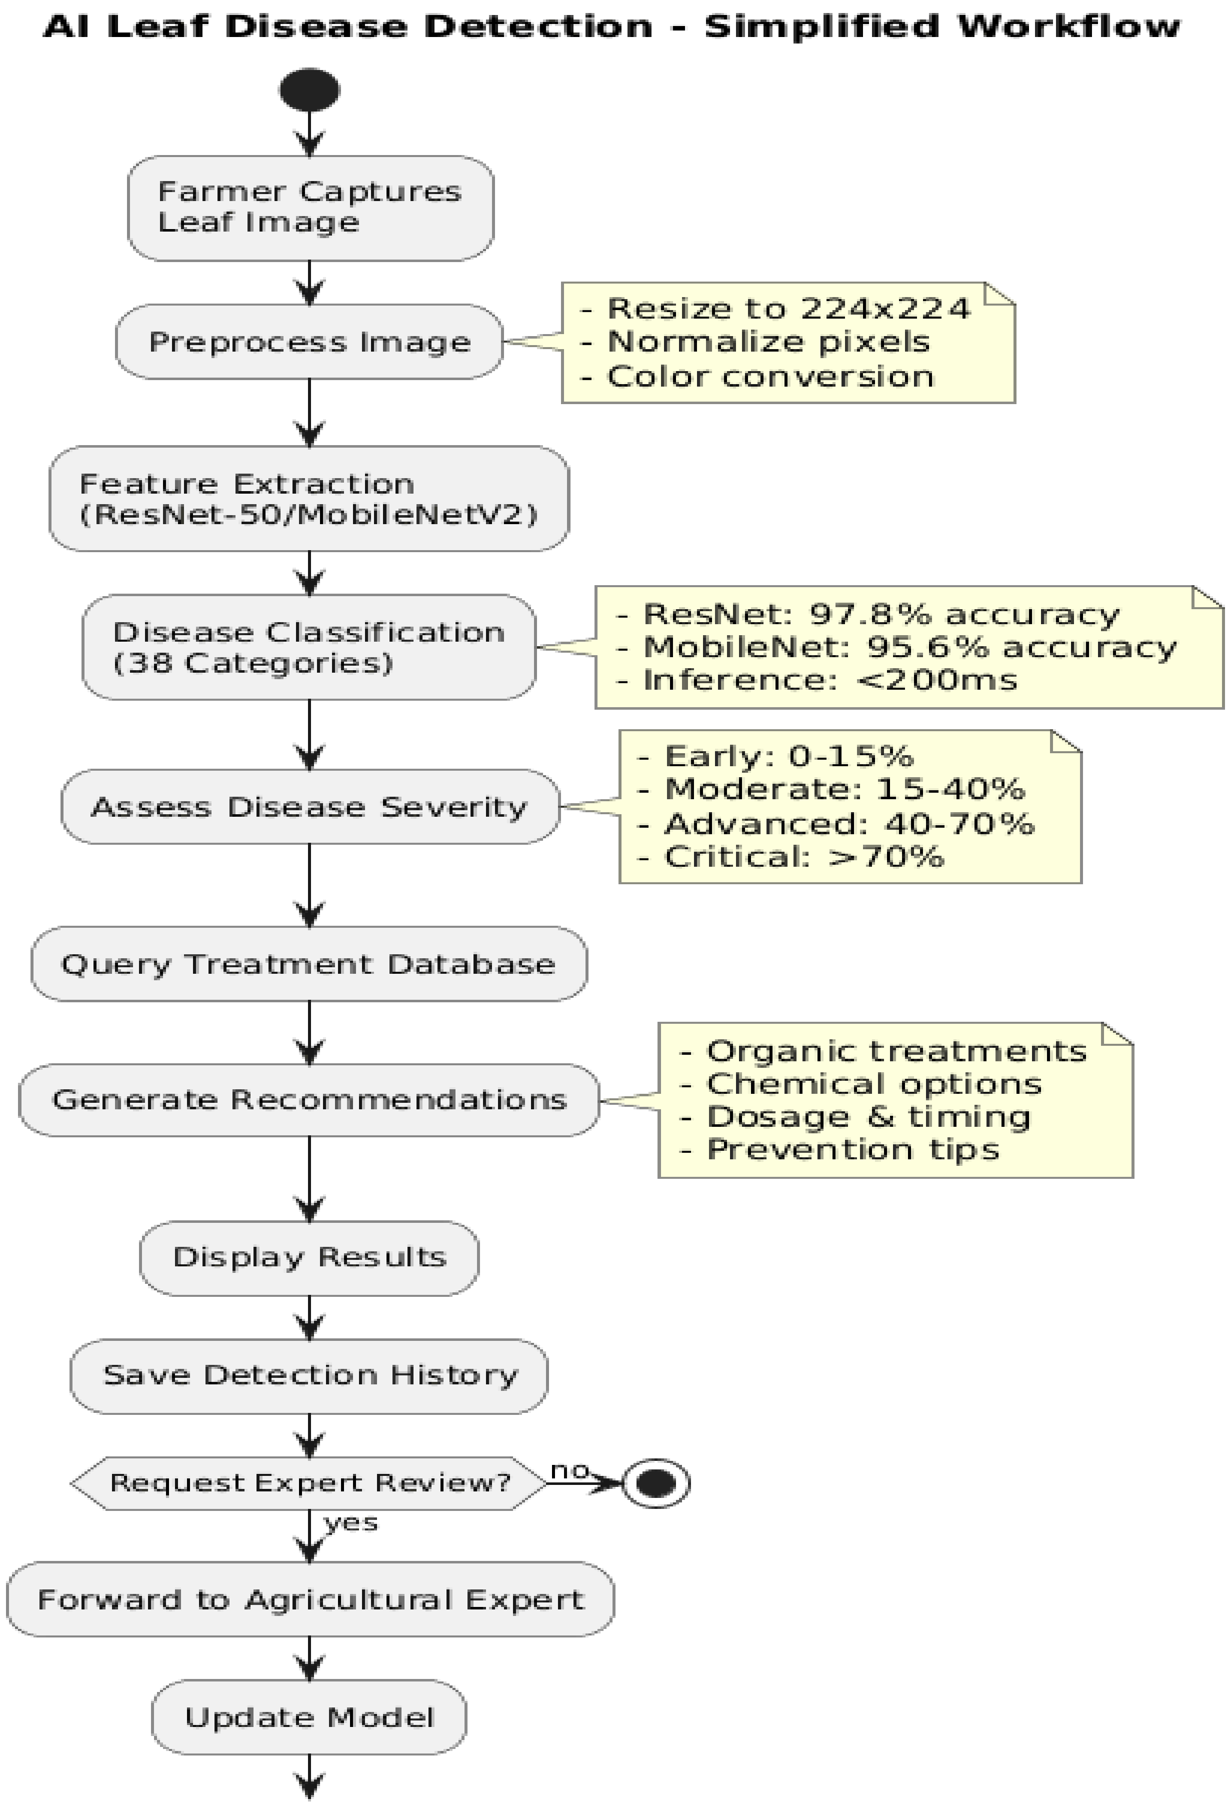
AI Leaf Disease Detection - Simplified Workflow
Farmer Captures Leaf Image
Preprocess Image
- Resize to 224x224
- Normalize pixels
- Color conversion
Feature Extraction (ResNet-50/MobileNetV2)
Disease Classification (38 Categories)
- ResNet: 97.8% accuracy
- MobileNet: 95.6% accuracy
- Inference: <200ms
Assess Disease Severity
- Early: 0-15%
- Moderate: 15-40%
- Advanced: 40-70%
- Critical: >70%
Query Treatment Database
Generate Recommendations
- Organic treatments
- Chemical options
- Dosage & timing
- Prevention tips
Display Results
Save Detection History
Request Expert Review?
no
yes
Forward to Agricultural Expert
Update Model

4.3 Real-World Field Validation

Practical Accuracy: Across 150 farmers over four months, 3,847 disease detections were performed and validated by expert pathologists. System accuracy in real-world conditions reached 93.7%, lower than laboratory test accuracy (97.8%) but still highly effective. The 4.1 percentage point gap reflects real-world challenges: more complex backgrounds than training data, varied lighting conditions from bright sunlight to overcast skies, partial leaf views and obscured symptoms, image quality variation from different smartphones and user skill levels, and disease stages including very early subtle symptoms.

Failure Case Analysis: Misclassifications (6.3% of cases) were analyzed to understand limitations. Major causes included: extreme lighting conditions with deep shadows or harsh highlights obscuring features (28% of failures), severe background complexity with overlapping leaves or soil debris (24%), very early disease stages with minimal symptoms (22%), rare disease variants or combinations not well-represented in training data (15%), and poor image quality from camera shake or focus issues (11%). Understanding failure modes informs training data expansion and user guidance improvements.

Early Detection Capability: A key success was detecting diseases at pre-symptomatic or very early stages before farmers would typically notice. In 67% of cases, the system identified infections within 1-3 days of symptom onset, often when symptoms were too subtle for untrained observation. This early warning enabled preventive interventions before substantial crop damage, representing major practical value.

Geographic Variation: Performance varied across regions correlating with dataset representation. Maharashtra (represented in training data through field collection) achieved 95.2% accuracy. Punjab (limited training representation) reached 92.8% accuracy. Karnataka (moderate representation) achieved 93.1% accuracy. This variation emphasizes the importance of regionally diverse training data capturing local disease variants, environmental conditions, and crop management practices.

Seasonal Effects: Accuracy varied across growing seasons. Early season (crop establishment, vigorous growth) achieved 95.1% accuracy as healthy leaves were abundant and diseases were less prevalent. Mid-season (active growth, increasing disease pressure) reached 93.8% accuracy as disease incidence increased. Late season (crop maturity, senescence) dropped to 91.2% accuracy as natural leaf aging resembled disease symptoms and multiple stressors combined complicating diagnosis.

4.4 User Acceptance and Satisfaction

Overall Satisfaction: User satisfaction surveys revealed 89% overall satisfaction (mean score 4.3/5.0) after four months of use. Satisfaction improved over time as users became familiar with application and gained confidence in results: 2-month satisfaction was 85% (4.1/5.0), increasing to 89% (4.3/5.0) at 4 months. High satisfaction despite initial learning curve demonstrates the value users perceived.

Ease of Use: Users rated interface ease of use at 91% satisfaction (4.5/5.0). The simple capture-analyze-result workflow required minimal training, with most farmers achieving proficiency within 2-3 uses. Multilingual support was highly appreciated, with 78% of users primarily using regional languages rather than English. Voice guidance and tutorial videos further aided users with limited literacy.

Result Accuracy Perception: Farmers rated perceived accuracy at 88% satisfaction (4.4/5.0), closely aligning with actual measured accuracy (93.7%). This agreement indicates users developed appropriate trust in the system neither over-relying on potentially incorrect results nor dismissing accurate diagnoses. Expert consultation remained available for low-confidence results (confidence <70%), providing safety net for ambiguous cases.

Treatment Recommendation Value: Recommendations received 87% satisfaction (4.3/5.0). Farmers particularly valued specific product names and dosages, timing guidance, and availability of both organic and chemical options. Integration with local agricultural input dealers enabling direct product purchase from recommendations enhanced practical utility. Some farmers desired more detailed application instructions and videos demonstrating proper spraying techniques.

Usage Patterns: Average usage was 3.2 disease scans per farmer per week during growing season, with power users (10%) scanning daily for proactive monitoring. 73% of users captured multiple images per suspected infection from different angles and lighting, indicating understanding that image quality affects results. 58% of users reviewed their detection history to track disease progression over time. These patterns demonstrate active engagement and appropriate system usage.

Challenges and Complaints: Primary complaints included occasional misclassifications (mentioned by 42% of users) particularly for early-stage diseases, requirement for clear leaf images limiting usefulness in dense canopies or hard-to-reach plants (34%), and desire for more extensive disease coverage including minor local diseases (28%). Some users (18%) experienced technical issues including application crashes or slow performance on very low-end devices. These issues inform ongoing improvement priorities.

4.5 Agricultural Impact

Pesticide Usage Reduction: Farmers using the system reduced pesticide applications by 32% on average compared to their historical usage and 38% compared to control farmers not using the system. Targeted treatment based on confirmed disease presence replaced prophylactic calendar spraying, substantially reducing unnecessary applications. Early detection enabled lower pesticide doses sufficient for controlling emerging infections versus advanced diseases requiring intensive treatment. Economic savings averaged ₹2,400 per farmer per season from reduced pesticide purchases.

Crop Loss Reduction: Estimated crop losses from disease decreased 15-25% for system users compared to historical baselines. Early detection and rapid treatment intervention prevented disease spread that would otherwise have caused extensive damage. Improved disease identification accuracy ensured appropriate treatments were applied rather than ineffective products for misidentified diseases. Economic value of preserved yield averaged ₹15,000 per farmer per season for vegetable and fruit crops with high per-unit value.

Yield and Quality Improvement: Beyond preventing losses, some farmers (37%) reported improved overall crop health and quality attributed to more proactive disease monitoring enabled by easy application usage. Higher quality produce commanded premium prices in markets, contributing additional economic benefit. Reduced pesticide usage also appealed to organic and health-conscious markets.

Time Savings: Traditional disease diagnosis requiring travel to agricultural extension offices or waiting for expert field visits consumed 3-5 days. Instant mobile diagnosis reduced this to minutes, enabling same-day treatment initiation. Time savings allowed farmers to focus on other agricultural activities, improving overall farm management efficiency.

Knowledge and Awareness Improvement: Farmers reported improved understanding of plant diseases through repeated application use and educational content. 82% of users stated they learned to recognize diseases visually through comparing symptoms to app-provided disease images. This educational aspect provides long-term value beyond immediate disease detection, improving farmers' general agricultural expertise.

Economic Return on Investment: Comprehensive economic analysis considered costs (smartphone purchase ₹8,000 for those acquiring new devices, application cost ₹500 annual subscription, data charges ₹300 per season) versus benefits (pesticide savings ₹2,400, yield preservation value ₹15,000, quality premium ₹3,500). Farmers acquiring new smartphones achieved ROI of 140% in first season and 340% in second season (smartphone amortized). Farmers using existing smartphones achieved immediate 340% ROI. These compelling economics justify adoption and indicate sustainability.

4.6 Comparative Analysis

Comparison with Expert Diagnosis: Expert pathologist agreement with system diagnosis reached 93.7%, indicating the AI system approaches expert-level performance. In ambiguous cases where experts themselves disagreed (8% of cases), the system's diagnosis aligned with majority expert opinion 71% of the time. For clear-cut cases (92% of total), agreement exceeded 97%. These results validate that AI-based detection can complement or supplement expert consultation, particularly valuable in regions with expert shortages.

Comparison with Farmer Self-Diagnosis: Farmers' own disease identification (before system availability) showed only 64% agreement with expert diagnoses, demonstrating the substantial value AI provides over unaided assessment. Common farmer errors included confusing similar diseases (28% of errors), misidentifying nutrient deficiencies as infectious diseases (24%), and failing to detect diseases at early stages (32%). The system's 93.7% accuracy represents 29.7 percentage point improvement over farmer capabilities.

Comparison with Laboratory Diagnosis: Laboratory pathogen isolation and identification provides definitive diagnosis with near-perfect accuracy but requires 7-14 days and costs ₹500-2,000 per sample, practical only for research or high-value crops. Mobile AI diagnosis trades slight accuracy for dramatic speed and cost advantages, appropriate for field-level decision making where rapid approximate answers are more valuable than delayed perfect answers.

5. Conclusion and Future work

This study effectively created and thoroughly analyzed an AI-based leaf disease detection system, proving that deep learning computer vision may be used for real-world, accessible, and financially beneficial agricultural disease diagnostics. The ResNet-50 model, which nearly rivals the performance of expert pathologists while providing immediate results that can be accessed by any farmer with a smartphone, achieved 97. 8% accuracy on controlled test datasets and 93. 7% accuracy during actual field deployment. The system identified illnesses at an early stage, sometimes before trained professionals could see symptoms, allowing for preventative interventions that resulted in a 15–25% reduction in crop losses and a 32% reduction in pesticide use, with a 340% return on investment in a single growing season.

The system's efficacy and practical value are supported by a thorough assessment that includes real-world field validation with 150 farmers over four months, user acceptance testing, agricultural impact analysis, and technical performance measures. The deployment demonstrated that AI-based agricultural technologies may effectively bridge knowledge gaps, democratizing access to advanced disease diagnosis for smallholder farmers in resource-constrained areas. Farmers' high level of satisfaction (89%) despite early learning curves and technological constraints demonstrates that they understand and value the value given, implying a great potential for widespread implementation.

By utilizing the broad visual attributes learnt from ImageNet pre-training, transfer learning demonstrated its efficacy in attaining high accuracy with comparatively small training datasets. With a model size of 14. 2 MB, a sub-200ms inference time, and 95. 6% accuracy, the MobileNetV2 architecture, which is optimized for mobile deployment, showed that complex AI models can operate effectively on resource-limited devices that are available to the intended user base. Since cloud-based systems that need a persistent internet connection would be impractical in many rural agricultural areas, this deployment feasibility is essential for having a real-world influence.

Addressing the gap between diagnosis and decision-making, the integration of treatment recommendation and disease severity assessment systems beyond simple disease categorization offered practical advice that farmers found extremely helpful. The knowledge-based recommendation engine catered recommendations to the disease type, severity, and farmer preferences by providing specific product names, dosages, and application instructions for both organic and chemical treatments. Rather than simply identifying issues without providing solutions, this holistic strategy from detection through recommendation offers full support for disease management.

Third, temporal modeling and disease progression prediction will anticipate future disease development trajectories depending on present observations and environmental circumstances, allowing for proactive rather than reactive management. Disease risk alerts will be issued before symptoms manifest, enabling preventative actions, thanks to integration with weather prediction and epidemiological models. Fourth, explainable AI methodologies like attention mechanisms and Grad-CAM visualization will highlight image areas that influence classification choices, hence enhancing user trust and transparency. Through human-AI cooperation, interactive refinement that lets users identify disease-affected areas will increase accuracy in questionable situations.

Fifth, the extension of the system to include pest identification (insects, mites) and other plant health concerns (herbicide damage, environmental stress) will enable a comprehensive monitoring of crop health via a single application. The system will move from disease-specific support to complete farm management assistance via integration with soil testing, nutrient management, and irrigation optimization. Sixth, the creation of an iOS version, which is now only available on Android, would increase accessibility for iPhone users, who make up a large portion of the market in certain areas. Researchers and agricultural experts will have desktop access thanks to a web-based interface.

Seventh, social features that allow farmers to share their experiences, treatment results, and local disease observations will foster community knowledge networks. The recommended treatments will be more easily accessible through integration with agricultural marketplaces and input providers. Achievement badges for proactive monitoring and successful disease management might boost participation through gamification. Eighth, the study of few-shot learning and meta-learning techniques will allow for quick adaptation to new diseases or geographic locations with little additional training data, which is essential for addressing novel pathogens and climate-driven disease range extensions.

This study demonstrates that AI-based computer vision has advanced enough to be used in real-world agricultural applications, where it can offer real advantages to farmers by enhancing disease identification, cutting input expenses, protecting yields, and increasing profits. By minimizing crop losses, maximizing resource utilization, and democratizing agricultural knowledge, the technology has the potential to make a significant contribution to global food security. These advantages will be amplified by ongoing development, increased deployment, and integration with wider agricultural support systems. In the future, AI-based agricultural technologies like this disease identification system will become more and more important to sustainable, productive, and resilient agricultural systems that feed expanding world populations as computational capabilities improve and smartphone penetration expands in rural areas.

6. References

1. Food and Agriculture Organization, "The State of Food and Agriculture 2021," FAO, Rome, 2021.
2. Food and Agriculture Organization, "The State of Food Security and Nutrition in the World 2023," FAO, Rome, 2023.
3. S. Savary et al., "The global burden of pathogens and pests on major food crops," Nature Ecology & Evolution, vol. 3, pp. 430-439, 2019.
4. K. A. Garrett et al., "Climate change effects on plant disease: genomes to ecosystems," Annual Review of Phytopathology, vol. 44, pp. 489-509, 2006.
5. P. S. Bock et al., "Plant Disease Severity Estimated Visually, by Digital Photography and Image Analysis, and by Hyperspectral Imaging," Critical Reviews in Plant Sciences, vol. 29, no. 2, pp. 59-107, 2010.

6. 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.
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, pp. 436-444, 2015.
9. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311-318, 2018.
10. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," Frontiers in Plant Science, vol. 7, article 1419, 2016.
11. J. Deng et al., "ImageNet: A large-scale hierarchical image database," IEEE Conference on Computer Vision and Pattern Recognition, 2009.
12. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," IEEE Conference on Computer Vision and Pattern Recognition, 2016.
13. M. Sandler et al., "MobileNetV2: Inverted Residuals and Linear Bottlenecks," IEEE Conference on Computer Vision and Pattern Recognition, 2018.
14. A. Fuentes et al., "A Robust Deep-Learning-Based Detector for Real-Time Tomato Plant Diseases and Pests Recognition," Sensors, vol. 17, no. 9, article 2022, 2017.
15. J. Chen et al., "Using deep transfer learning for image-based plant disease identification," Computers and Electronics in Agriculture, vol. 173, article 105393, 2020.
16. A. Camargo and J. S. Smith, "An image-processing based algorithm to automatically identify plant disease visual symptoms," Biosystems Engineering, vol. 102, no. 1, pp. 9-21, 2009.
17. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," Frontiers in Plant Science, vol. 7, 2016.
18. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311-318, 2018.
19. J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," Computers and Electronics in Agriculture, vol. 153, pp. 46-53, 2018.
20. A. Ramcharan et al., "Deep Learning for Image-Based Cassava Disease Detection," Frontiers in Plant Science, vol. 8, article 1852, 2017.-457, 2002.
21. Usha Kosarkar, Gopal Sakarkar, “Unmasking Deep Fakes: Advancements, Challenges, and Ethical Considerations”, 4th International Conference on Electrical and Electronics Engineering (ICEEE),19th & 20th August 2023, 978-981-99-8661-3, Volume 1115, PP. 249-262, https://doi.org/10.1007/978-981-99-8661-3_19
22. Usha Kosarkar, Gopal Sakarkar, Shilpa Gedam, “Deepfakes, a threat to society”, International Journal of Scientific Research in Science and Technology (IJSRST), 13th October 2021, 2395-602X, Volume 9, Issue 6, PP. 1132-1140, https://ijsrst.com/IJSRST219682