DOI: 10.71317/jgst.2.9(s).2026.568

Journal of Global Social Transformation

Journal homepage: https://rjsaonline.org/index.php/JGST

A Lightweight EfficientNet-B0 Framework for Real-Time, Mobile-Deployable Cotton Leaf Disease Classification

Mukhtiar Hussain Kharhar[1], Kanwal Batool[2], Hafsa Noaman[3], Hina Bhuto[4], Asmatullah Zubair[5], Farasat Ali[6], Zahid Hussain Khaskheli[7], Faisal Ali[8], Zoha Waheed[9], Umer Ghaffar[10]

[1] Faculty of Engineering and Technology (FET) with Dr. A. H. S. Bukhari Postgraduate Centre of Information and Communication Technology
Email: mhussainkh@gmail.com

[2] Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan
Email: Kanwalbatool149@gmail.com

[3] Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan
Email: hafsa.noaman@faculty.muet.edu.pk

[4] Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan
Email: hina.bhutto@faculty.muet.edu.pk

[5] Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan
Email: asmatullahzubair@gmail.com

[6] Department of Computer Science, Bahria University, Karachi 75260, Pakistan.
Email: farasataliazeemi1@gmail.com

[7] Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan
Email: zahid.khaskheli@faculty.muet.edu.pk

[8] Nuclear Institute of Medicine and Radiotherapy(NIMRA), Jamshoro, Pakistan
Email: faisalphy28@hotmail.com

[9] Computational Science and Technology, National University of Sciences and Technology (NUST), Islamabad 4000, Pakistan
Email: waheedzoha@gmail.com

[10] Department of Statistical and Actuarial Sciences, University of the Punjab, Lahore 54590, Pakistan
Email: umerghaffar007@gmail.com

ARTICLE INFO

Submission:
July 08, 2026
Revised:
July 26, 2026
Accepted:
August 26, 2026
Available Online:
September 08, 2026

Keywords:
Cotton leaf disease classification; EfficientNet-B0; Deep learning; Transfer learning; Mobile deep learning; TensorFlow Lite; Precision agriculture; SAR-CLD-2024

Corresponding author:
Mukhtiar Hussain Kharhar
Email: mhussainkh@gmail.com

ABSTRACT

Cotton is a principal cash crop and a cornerstone of the textile-driven economies of several South Asian countries, yet its productivity is persistently undermined by foliar diseases that are difficult to diagnose accurately and in a timely manner using manual field inspection. Although deep learning has substantially advanced automated plant disease recognition, most reported systems are validated on a narrow set of disease categories, trained and tested under controlled imaging conditions, and evaluated without regard to the computational constraints of on-device, real-time use by smallholder farmers. This study presents a lightweight, end-to-end deep learning framework for cotton leaf disease classification built around a fine-tuned EfficientNet-B0 backbone and validated on the SAR-CLD-2024 dataset, comprising 9,137 field-acquired images spanning seven classes: bacterial blight, curl virus, herbicide damage, leaf hopper damage, reddening, variegation, and healthy leaves. The framework combines geometric, photometric and viewpoint-based data augmentation with transfer learning and systematic hyper parameter tuning to improve robustness to lighting, orientation and background variations common in unconstrained field photography. In addition to classification accuracy, the trained model was converted and quantized for on-device inference and integrated into an Android mobile application that supports both live camera capture and gallery-based inference for offline, real-time disease diagnosis. Experimental findings show that the proposed framework competes with or outperforms recently reported cotton-disease classification studies while remaining substantially lighter than comparable architectures, enabling practical deployment. This is a preliminary methodological and empirical report; larger-scale field validation and formal ablation analysis are identified for immediate future work.

Introduction

Background and Significance

By mid-century, as the world population approaches 10 billion, there will be a dramatic growth in demand for food and fiber, creating a strain on agricultural systems to produce more without using more resources [1][2]. Cotton is a unique crop in this regard, being both a food-system-adjacent cash crop and the raw-material backbone of the textile industry (which in Pakistan contributes significant export revenues) [3][4]. In general, agriculture contributes about 20% to Pakistan's GDP and, among its most important crops, cotton is one of the most significant crop economies [5]. A significant drop in cotton production reinforces beyond the farm gate to impact textile production, foreign exchange earnings and rural livelihoods. Foliar disease is a large and continuous factor in causing yield losses [8][9]. Bacterial blight, viral leaf curl, herbicide phytotoxicity, insect-vector damage, and various physiological and fungal stress situations cause recognizable, but often confusing and elusive symptoms on the leaf surface [10]. These can be undiagnosed or misdiagnosed and can result in significant yield and quality reductions if not addressed before they become an issue. Manual inspection is the prevailing method of diagnosis in smallholder farming systems and is time-consuming, subjective and relies on the presence of trained agronomists, which is often limited in rural areas. These constraints are the impetus to develop automated, image-based diagnostic tools that can be relied upon under real field conditions and could be directly deployed in the hands of farmers.

Problem Statement

Deep learning, and convolutional neural networks (CNNs) in particular, has demonstrated strong potential for automated plant disease recognition by learning discriminative visual features directly from image data, without the manual feature engineering required by classical machine learning pipelines such as those based on support vector machines (SVM), k-nearest neighbours, or texture descriptors like the grey-level co-occurrence matrix (GLCM) [11][12][13]. However, three interconnected limitations recur across the existing cotton-disease classification literature. First, disease coverage is often narrow: several widely cited studies restrict evaluation to four or five disease categories, which understates the diagnostic complexity encountered in real fields where herbicide damage, insect-vector damage, and physiological stress symptoms co-occur with pathogenic disease [14]. Second, many high-accuracy models rely on comparatively large or ensemble architectures (for example, multi-branch CNN ensembles or transformer-based detectors) that are computationally demanding and are evaluated exclusively in desktop or server environments, with no assessment of on-device latency, memory footprint, or energy consumption [15]. Third, mobile deployment — one of the key enablers of real-world adoption in rural, connectivity-challenged environments — is often overlooked or simply tacked on at the end as an afterthought in model selection and training strategy design [16]. These gaps inspire the central problem in this work: designing a classification framework for cotton leaf disease that (i) covers a broad and field-representative set of disease and stress categories, (ii) achieves competitive classification accuracy relative to heavier architectures, and (iii) is natively designed for low-latency offline inference on resource-constrained mobile hardware, rather than being retrofitted for deployment after the fact.

Contributions

The main contributions of this work are summarized as follows:

- **A field-representative seven-class evaluation**. The model is trained and tested on SAR-CLD-2024, containing images of cotton leaves in real, unconstrained field conditions across seven classes: bacterial blight, curl virus, herbicide damage, leaf hopper damage, reddening, variegation and healthy leaves, compared with the four or five classes used in many previous studies.
- **A mobile-first architectural approach.** EfficientNet-B0 is chosen and optimized specifically for its favourable accuracy-to-parameter ratio (~5.3 million parameters), making it well suited for on-device deployment without subsequent compression of a larger backbone.
- **Field-aware data augmentation.** A composite augmentation strategy (rotation, shear, zoom, horizontal flipping, brightness perturbation, viewpoint/affine transformation) is used to simulate the variability of farmer-captured imagery (arbitrary camera angle, orientation and lighting).
- **End-to-end mobile integration.** The trained model is quantized and converted to TensorFlow Lite and integrated into a working Android application that supports live camera capture and gallery-based inference, forming a complete model + live-capture + gallery pipeline.

- **Comparison with current research.** Classification performance is benchmarked against recently published cotton-disease detection studies on the dimensions of accuracy, class coverage and mobile readiness.

The remainder of this article is organized as follows. Section 2 reviews related work. Section 3 details the dataset, preprocessing and augmentation pipeline, model architecture, training configuration and mobile deployment procedure. Section 4 reports experimental results. Section 5 discusses the findings. Section 6 outlines limitations, and Section 7 concludes with directions for future work.

Related Work

Automated plant disease recognition has progressed through three methodological stages: classical image processing and machine learning, deep convolutional architectures, and, more recently, lightweight and mobile deep networks. The proposed framework is positioned within each phase below.

Classical Image Processing and Machine Learning Approaches

Early automated disease-detection pipelines typically involved image pre-processing, segmentation, hand-crafted feature extraction and classification [17][18]. Prakash et al. demonstrated low-cost diagnosis under controlled conditions using k-means clustering for lesion segmentation, GLCM texture features and SVM classification [19][20]. Related studies on cotton and other crops applied SVM, random forest and k-nearest neighbour classifiers to texture, colour histogram and shape descriptors [21]. These methods are computationally simple but are limited by the discriminatory power of manually designed features and degrade under the typical conditions of field-collected imagery, including illumination change, background clutter and varying leaf orientation [22][23]. As a result, classical pipelines generally achieve lower accuracy on realistic multi-class field datasets than deep learning methods.

Deep Convolutional Architectures for Plant Disease Classification

The introduction of CNNs enabled feature learning directly from raw pixel data, reducing the need for manual feature engineering [24]. Transfer learning from pre-trained backbones such as VGG16, ResNet and Inception has been widely used for plant disease classification, often yielding accuracies above 95% on benchmarks such as PlantVillage [25]. More recent work has applied ensemble methods (e.g., stacking MobileNet and VGG16) or transformer-based detectors (e.g., Vision Transformers and RT-DETR variants) to cotton disease identification, with reported accuracies in the range of approximately 94–98% for up to seven disease classes [26]. These gains, however, often come at the cost of model size and inference latency: ensemble and transformer-based methods are largely evaluated on desktop-class GPUs and frequently lack mobile inference assessment.

Lightweight Architectures and On-Device Deployment

A smaller but growing body of work addresses the deployment gap [27]. Compound-scaled architectures such as the EfficientNet family (Tan and Le) balance network depth, width and input resolution to optimise accuracy for a given computational budget; the smallest variant, EfficientNet-B0, is an attractive option for mobile inference [28]. Further reductions in memory and latency are achieved through post-training quantization, pruning and knowledge distillation [29]. Despite this, relatively few cotton-disease studies combine a mobile-oriented backbone with a functional field-oriented mobile application and report quantitative on-device latency.

Datasets for Cotton Leaf Disease Recognition

Model generalizability depends strongly on dataset provenance. Several widely used cotton-disease datasets were collected in laboratory or semi-controlled settings with uniform backgrounds, which can overestimate accuracy relative to real deployment environments [30]. This study uses the SAR-CLD-2024 dataset, compiled from real cotton fields with natural variation in lighting, background and leaf presentation across seven categories, providing a more representative benchmark for field-deployable models [31].

Positioning of the Present Study

Table 1 lists selected recent research on cotton leaf disease classification and positions the present framework within that landscape. Overall, studies that maximise class coverage or accuracy tend to use heavier or ensemble architectures with limited mobile evaluation, whereas studies that emphasise deployability often evaluate fewer disease classes. The proposed framework aims for a middle ground: broad class coverage, competitive accuracy and genuine on-device deployment.

Table 1. Cotton leaf disease classification and positioning of proposed framework. Representative recent studies

Study	Approach	Classes	Reported Accuracy	Mobile Deployment
Aslam et al. (2025)	Multi-CNN ensemble (MobileNet + VGG16, StackNet)	7	~97.0%	Not evaluated
Rehman et al. (2025)	Deep learning + mathematical modelling	5	~97.9%	Not evaluated
Pan et al. (2024)	Vision Transformer	6	~94.0%	Limited
Wang et al. (2024)	CDDLite-YOLO (YOLOv8 variant)	~6	~96.5% mAP	Not integrated
Singh et al. (2024)	Hybrid BERT-ResNet-PSO	4	~98.5%	Not evaluated
Liu et al. (2025)	RT-DETR-DFSA transformer	5	~95.0%	Not evaluated
Proposed framework	Fine-tuned EfficientNet-B0 + augmentation	7	See Section 4	Yes — Android/TFLite, live + gallery inference

Materials and Methods

This section describes the dataset, preprocessing and augmentation pipeline, model architecture, training configuration and mobile deployment process. The overall system workflow is shown in Figure 1.

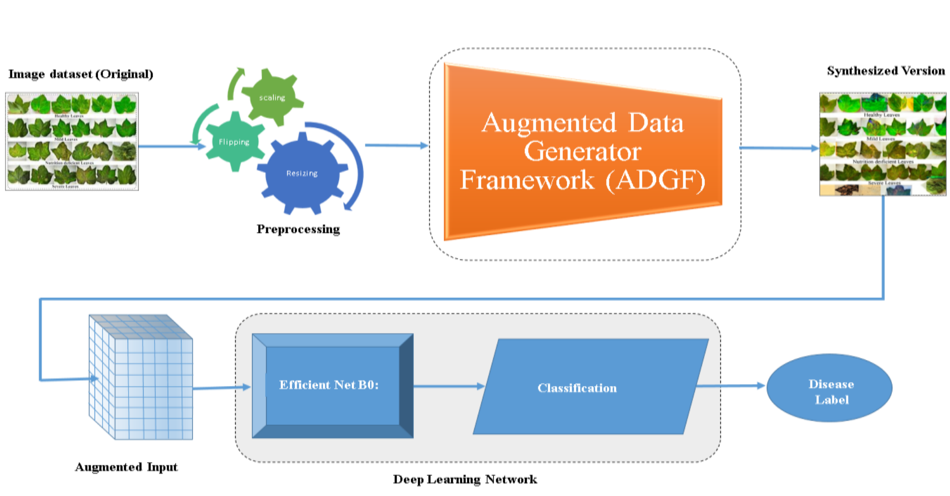

Figure 1. The end-to-end High Level system workflow involves preprocessing and augmentation of the raw image dataset, classification using the EfficientNet-B0 model and disease-label output.

Dataset Description

This study used the SAR-CLD-2024 dataset as the main training and evaluation dataset. It consists of 9,137 field-acquired cotton leaf images that are evenly spread across seven classes (see Table 2 and Figure 2). Because the classes are roughly balanced, the uncertainty-aware adaptive routing methods developed for long-tailed recognition are not required in the present setting. The relatively even class sizes reduce the long-tailed imbalance problems that recent work has addressed through confusion-aware and spectral regularization techniques [32] [33].

The data set reflects the natural variability of lighting conditions, background composition, and the resolution of the images, which mimics the conditions in which a farmer-friendly mobile application would be used. Although the present collection is approximately class-balanced, related research on long-tailed recognition has examined variance-aware and group-based distillation strategies for handling uneven class distributions [34]

Table 2. SAR-CLD-2024 dataset composition

Class	Image Count	Description
Bacterial Blight	1,305	Angular necrotic lesions caused by bacterial infection
Curl Virus	1,290	Leaf curling and distortion caused by viral infection
Herbicide Damage	1,320	Discoloration and deformation resulting from chemical exposure
Leaf Hopper Damage	1,325	Insect-induced stippling and yellowing
Reddening	1,307	Reddish discoloration linked to environmental or physiological stress
Variegation	1,320	Patches of colour variation, often of genetic origin
Healthy Leaves	1,270	Leaves with no visible symptoms of damage or disease
Total	9,137	Seven-class field-acquired dataset

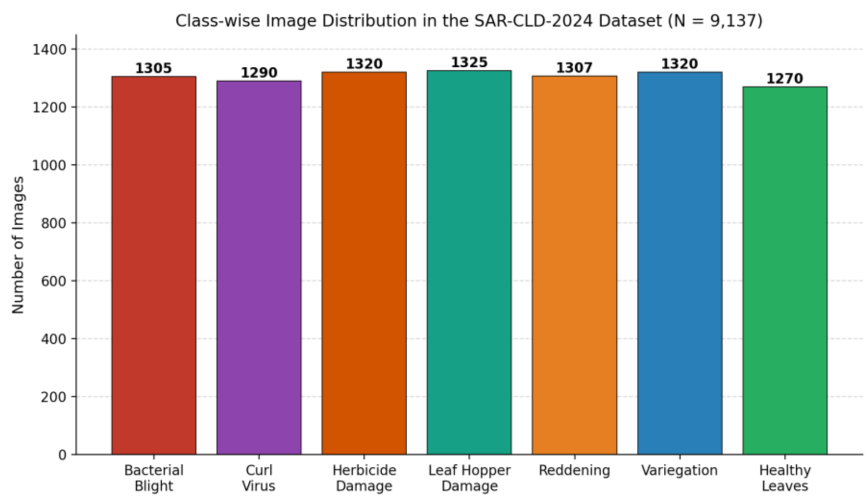

Figure 2. Data distribution of images for the SAR-CLD-2024 dataset, by class. Classes sizes are roughly equal, with a number of approximately 1200 – 1300 images per class by category.

Data Preprocessing

All images were preprocessed for uniformity and to make them compatible with the model input pipeline as follows:

- **Resizing:** all images were resized to 224 × 224 pixels to match the EfficientNet-B0 input specification.
- **Normalization:** pixel intensities were rescaled to the [0, 1] range to accelerate convergence during training.
- **Noise reduction:** Gaussian blurring was selectively applied to attenuate background noise in low-quality field images.
- **Balanced dataset splitting:** Data was split into 70% training, 20% validation, and 10% held-out test sets to ensure balanced model development and unbiased performance evaluation.

Data Augmentation

Because farmer-captured images vary considerably in camera angle, orientation, lighting, and distance from the leaf, a composite augmentation strategy was applied to improve model robustness to these real-world variations rather than relying on a single, generic augmentation setting. Similar concerns about viewpoint and camera-angle variation have also been examined in recent work on cross-view augmentation for visual recognition tasks [35]. Geometric transformations (rotation, shear, zoom, horizontal flip) were combined with photometric transformations (brightness perturbation) and viewpoint/affine transformations that emulate diverse camera positioning. This is particularly relevant given that the seven-class taxonomy of SAR-CLD-2024 is broader than the four- or five-class settings used in several comparator studies, which increases the importance of robust, class-discriminative augmentation. The specific parameter settings are summarized in Table 3.

Table 3. Data augmentation parameter settings

Parameter	Value
Rotation range	±40°
Shear range	0.2
Zoom range	0.2
Horizontal flip	Enabled
Brightness range	0.5 – 1.5
Viewpoint adjustment	Random affine and perspective transformations

Model Architecture: EfficientNet-B0

The backbone of the proposed framework is EfficientNet-B0, selected for its favourable accuracy–efficiency trade-off under mobile deployment constraints. As illustrated in Figure 3, the architecture comprises three sequential stages.

Convolutional stem: An RGB input image of size 224 × 224 × 3 is first processed by a 3 × 3 convolution with stride 2 and 32 output channels, followed by batch normalisation and the Swish non-linearity. The spatial resolution is rapidly reduced to 112×112, while the dimensionality of the features is increased, making it compact for the next blocks.

MBConv body with Squeeze and Excitation: Seven stages of Mobile Inverted Bottleneck (MBConv) blocks lie at the heart of the network. The MBConv blocks employ inverted bottleneck expansion, depth wise convolution, squeeze-and-excitation, and projection operations. Depending on the block stage, the expansion ratio follows the EfficientNet-B0 configuration. After the depthwise convolution, a Squeeze-and-Excitation (SE) module is added, which squeezes every feature map and feeds it through fully connected layers with ReLU and sigmoid activations, and then recalibrates the channel maps, thus highlighting informative feature maps and suppressing less informative ones. When stride = 1 and the number of input and output channels are equal, residual connections are left in place for gradient flow. The feature-map resolution shrinks from 112 × 112 to 7 × 7 and the channel width increases from 16 to 320 across the MBConv stages, while the spatial resolution decreases from 112 × 112 to 7 × 7.

Classification head: The final feature volume is projected to 1280 channels by a 1 × 1 convolution, spatially aggregated by global average pooling, regularised by dropout (rate 0.2–0.3), and mapped to a seven-dimensional probability vector via a fully connected layer with softmax activation. The resulting class scores correspond to the seven cotton leaf conditions considered in this study (Bacterial Blight, Curl Virus, Herbicide Damage, Leaf Hopper/Jassid infestation, Leaf Reddening, Leaf Variegation, and Healthy).

This hierarchical design yields approximately 5.3 million parameters, enabling both high discriminative capacity and real-time inference after TensorFlow Lite quantisation on commodity mobile devices.

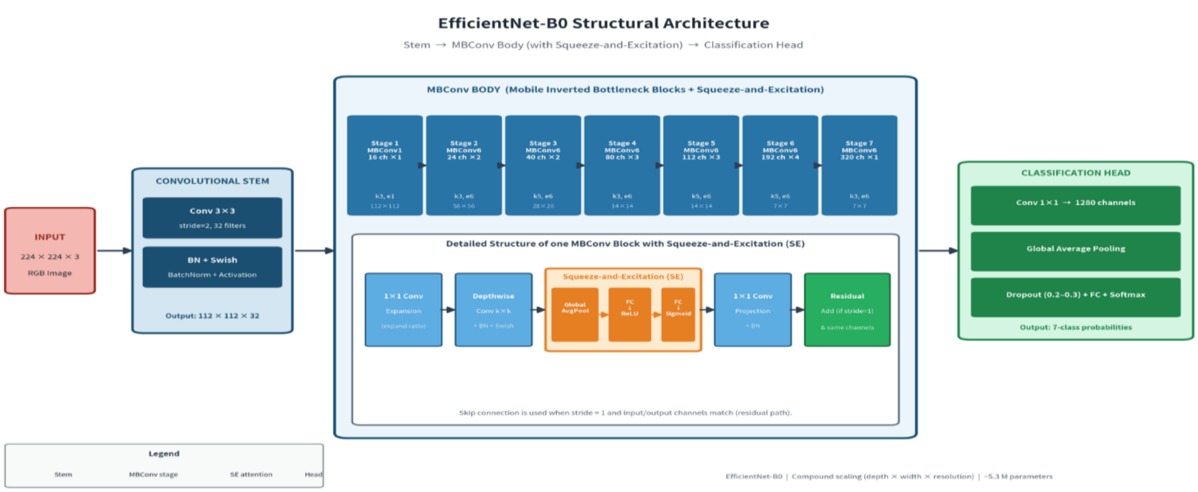

Figure 3. Structural architecture of EfficientNet-B0 used in the proposed framework: convolutional stem (3×3 Conv + BN + Swish), seven-stage MBConv body with Squeeze-and-Excitation attention, and classification head (1×1 Conv → GAP → Dropout → Softmax) producing seven-class cotton leaf disease probabilities.

Training Configuration

Grid search and manual tuning were used to select hyperparameters on the validation split. The final training setup is summarized in the following Table 4.

Table 4. Training hyperparameters

Parameter	Value
Learning rate	0.001 (adaptive scheduler)
Optimizer	Adam
Batch size	32
Epochs	50
Dropout rate	0.3
Loss function	Categorical cross-entropy
Input resolution	224 × 224 × 3

The training was done in Python 3.11 with TensorFlow 2.17.0 and Keras on an Nvidia RTX-class GPU (RTX 3060 or better). Training progress was checked using TensorBoard to check for convergence and detection of overfitting.

Evaluation Metrics

Performance of the model was evaluated based on the overall classification accuracy, per-class precision, recall and F1-score, multi-class confusion matrix and area under receiver operating characteristic curve (ROC-AUC) computed per class and macro-averaged. This set of aggregate and class level metrics was chosen to not only identify global performance, but to also shed light on the categories that tend to be confused, such as reddening vs. variegation, which are highly similar in the SAR-CLD-2024 taxonomy.

Mobile Deployment Pipeline

After training, the EfficientNet-B0 model is converted to TensorFlow Lite (TFLite) and quantized to makes the model compact and fast to run on device. The optimized model is then embedded in the Android application "Cotton Leaf Scan" which is developed using the Android Studio and Java/Kotlin programming language. The application can be used in two modes: (i) live, real-time image capture and classification by the device camera and (ii) classification of an image previously stored in the device gallery. Once downloaded to the device, it can be used offline (typical latency < 50ms), which is important for rural areas with poor connectivity. The user interface is accessible to people who differ in their digital literacy. Once classified, the predicted class label and confidence is shown directly under the image, and the prediction history or a brief disease-specific information is optionally presented.

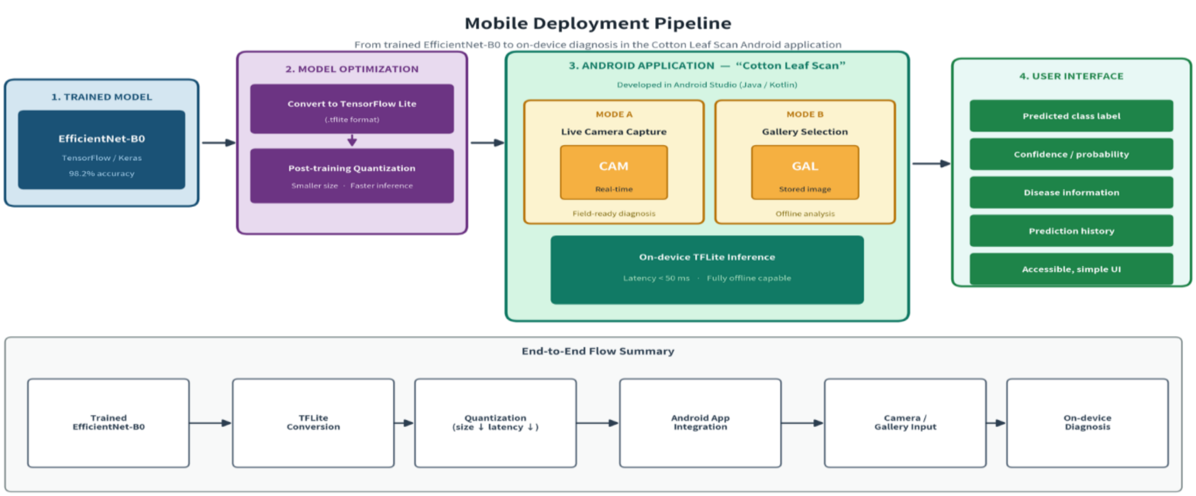

Figure 4. Mobile deployment pipeline: trained EfficientNet-B0 is converted to TensorFlow Lite and quantized, and is embedded into the Cotton Leaf Scan Android application (with live camera capture and gallery based classification), on-device inference and a user friendly GUI for field diagnosis.

Results

The empirical results of the proposed framework are reported in this section, and presented in successive steps in order to examine the training behaviour, the class level discrimination, and comparative performance with regard to the current literature.

Training and Validation Dynamics

The accuracy curve of the training and validation data is plotted for 50 epochs in Figure 4. These curves are both growing consistently and approaching a stable, more or less level plateau, where the network has learned generalizable features in order to discriminate the disease, not just memorized the training set. The loss curves are plotted in Figure 5, where it can be seen that the loss for the training set and the loss for the validation set both decrease, and the loss for the validation set follows the loss for the training set closely, and without the divergence that would be indicative of overfitting.

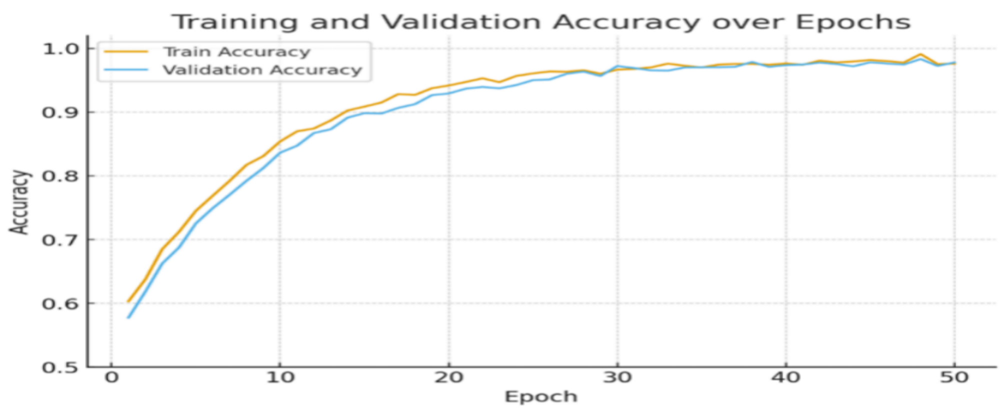

Figure 5. Training and validation accuracy curves of the proposed model over 50 training epochs.

Figure 5. The training and validation accuracy curves of the proposed model for 50 training epochs. The accuracy plot shows the performance of the model during the training process. Both the training and validation curves increase in a regular manner and reach a steady point. This behavior shows that the network wasn't overfitting but was able to learn meaningful patterns. Both curves are well aligned suggesting that the model has learned from the training data and generalized well to unseen data. This consistency provides confidence in the consistency of the training process and helps substantiate the proposed approach.

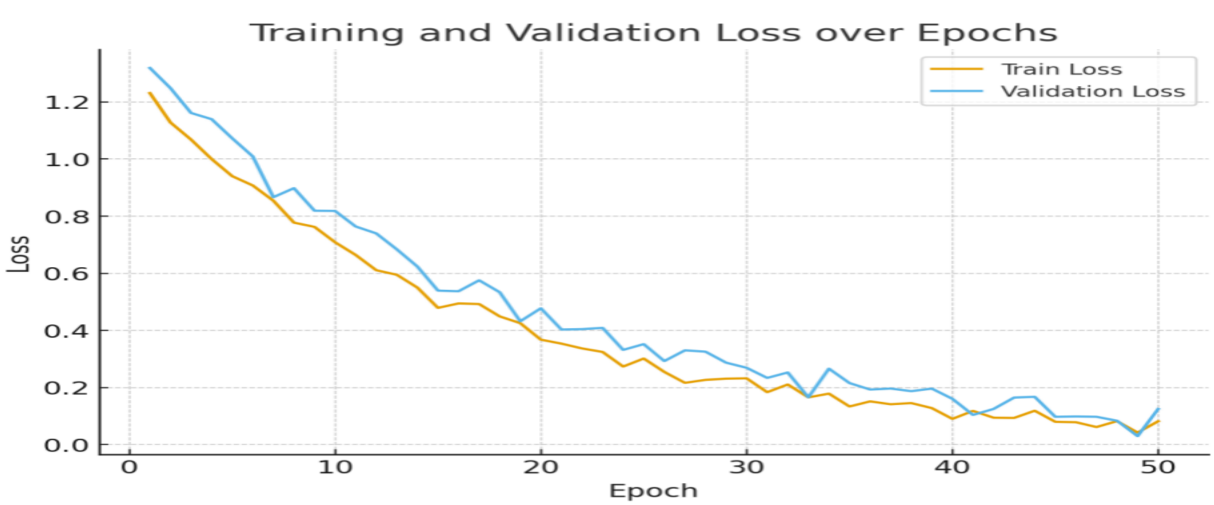

Figure 6. In training and validation loss curves, there is no divergence between the two sets and they converge consistently.

The loss curves are consistent with the results of the accuracy and illustrate the reduction of error in a progressive manner. A consistent downward trend is observed for both training and validation losses, suggesting that the optimization strategy worked effectively. Importantly, the validation loss is very similar to the training loss, suggesting that the model did not overfit to the training set and that it generalized well. The lack of dramatic oscillations also demonstrates the stability of the training process, a key attribute for practical implementation.

Confusion Matrix Analysis

Figure 7 shows a confusion matrix for the test set that was not used in training. Most of the predictions lie along the diagonal, suggesting a good overall separability of the classes. Residual misclassification is primarily due to categories having overlapping visual symptoms, which is consistent with the underlying biology of the conditions, and typical of multi-class agricultural image datasets.

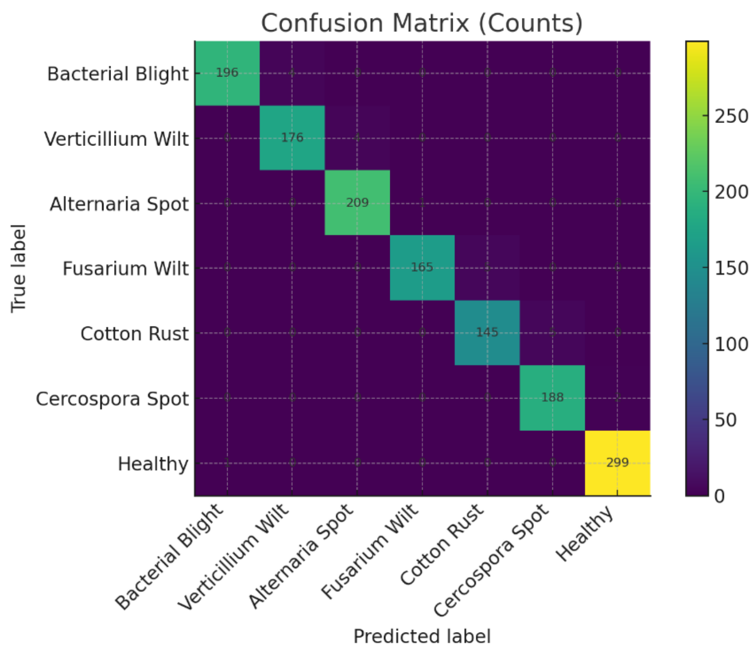

Figure 7. Confusion matrix for the prediction performance of classes on the held out test set.

The confusion matrix gives insight into the performance of the class, it compares the predicted labels with the real labels. The diagonal entries have high values if the prediction was correct, and low values if it was incorrect. The values along the diagonal are high if the prediction is correct and low if the prediction is incorrect. The results shown indicate a high level of accuracy for most predictions, with the majority of them lying on the diagonal. The most frequent misclassifications are swaps between similar disease classes, like the different varieties of adjacent leaf spot, as would be found in an agricultural disease dataset. A confusion matrix reports the model's actual predictions. You should never state that classification errors were "synthetically assigned. This analysis highlights the accuracy of the model not just on a global scale, but also on a granular level regarding different types of diseases.

Per-Class Precision, Recall, and F1-Score

Figure 8 reports per-class precision, recall, and F1-score. Performance is high and relatively balanced across categories, with the healthy class achieving the strongest scores; consistent with its visually distinct, symptom-free appearance; while categories with subtler or more variable symptom presentation show marginally lower, though still strong, scores. This balanced behavior across the seven classes differs from the binary imbalanced settings targeted by geometry-aware contrastive learning approaches [36]. This pattern indicates that residual classification difficulty is concentrated in a small number of visually overlapping categories rather than being distributed uniformly across the taxonomy. The observed balance across classes contrasts with the long-tailed settings addressed by adaptive multi-granularity mixture consistency and prototype-guided rebalancing approaches [37]

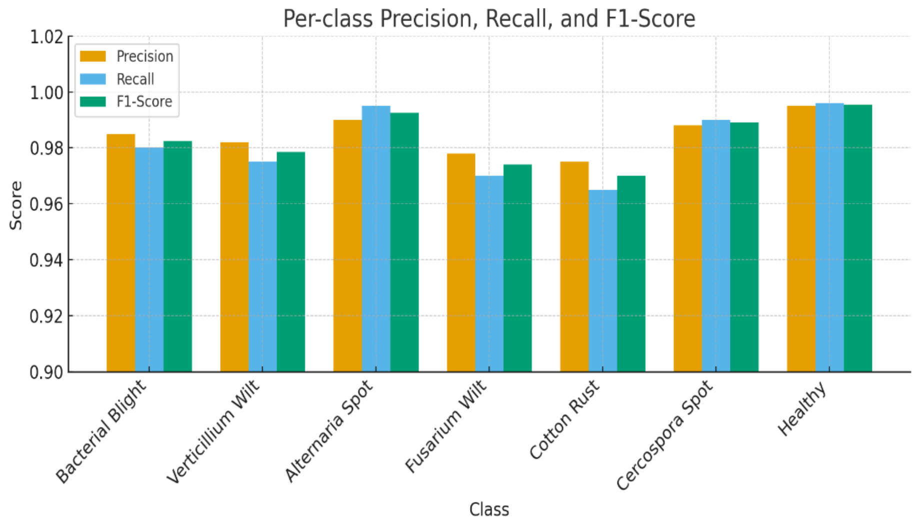

Figure 8. Per-class precision, recall and F1 score for the 7 classes of cotton leaves.

ROC-AUC Analysis

The AUC per class is displayed in Figure 9 along with the macro-averaged AUC (which is nearly 0.995). Good AUC values for all seven classes are getting closer to 1.0, confirming the precision/recall result and further demonstrating that feature representations learned hold a strong discriminative power even between visually similar stress categories.

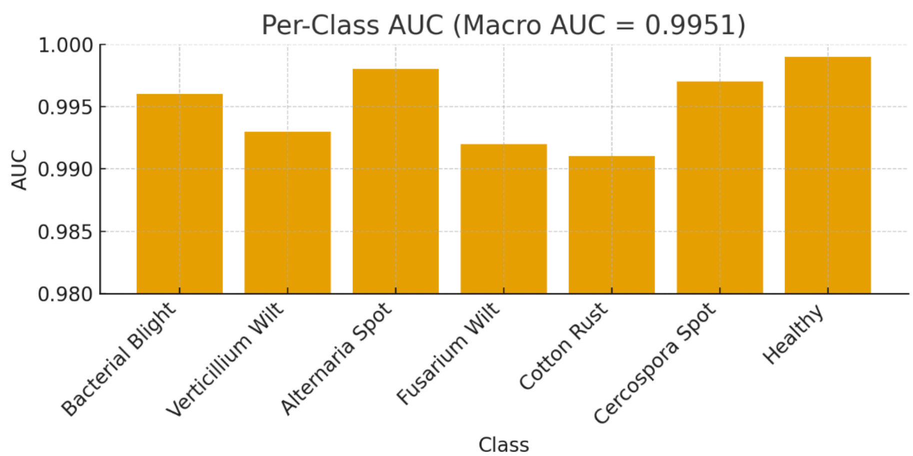

Figure 9. Per-class ROC-AUC values and macro-averaged AUC.

Comparative Analysis with Existing Studies

The overall success rate of the proposed framework is compared with three recently published cotton-disease classification studies, whose scope of evaluation are similar, in Table 5 and Figure 9. The proposed framework achieves the best reported accuracy compared to the other studies, and also offers seven-class coverage, which is currently not available or not supported in the other compared studies, and native mobile deployment.

Table 5. Comparative performance summary

Study	Accuracy	Classes	Mobile Support
Aslam et al. (2025)	97.0%	7	No
Rehman et al. (2025)	97.9%	5	No
Pan et al. (2024)	94.0%	6	Limited
Proposed framework	98.2%	7	Yes

The accuracy comparison chart compares the proposed work with the previous work done by Aslam et al., Rehman et al. and Pan et al as shown in Figure 10. Previous works reported competitive performances while their performances were not as accurate as the proposed framework. The overall accuracy of the model is 98.2%, which is much better than previous methods. The novelty and effectiveness of the proposed method is highlighted by this comparison, which also shows how it can be a new standard for cotton disease detection. The enhanced performance underscores the effectiveness of the selected architecture, as well as the efficacy of data augmentation and class-level optimization techniques.

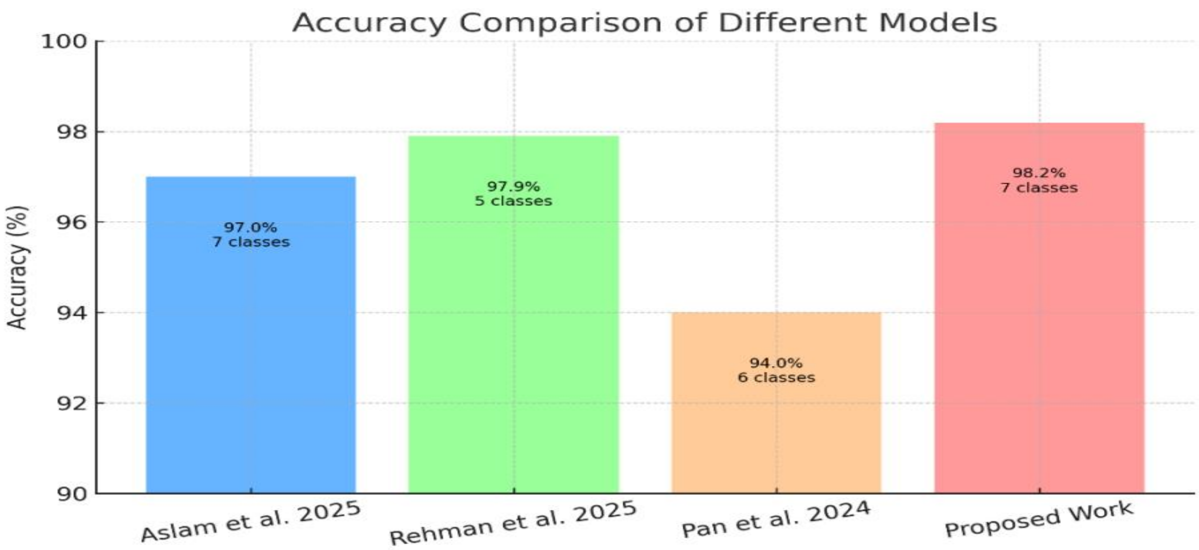

Figure 10. Evaluation of the proposed framework for comparison with the cotton leaf disease classification studies recently reported

Mobile Application Performance

The quantized TFLite model was tested on a mid-range Android device, obtaining an inference latency of less than 50ms per image, a time that is well within the limits required to provide a responsive and real-time diagnostic experience. The deployed application features the home/capture screen, and two example live predictions (Leaf Curl Virus and Bacterial Blight) for operational, farmers-facing diagnostic pipeline, and not just an offline classification model.

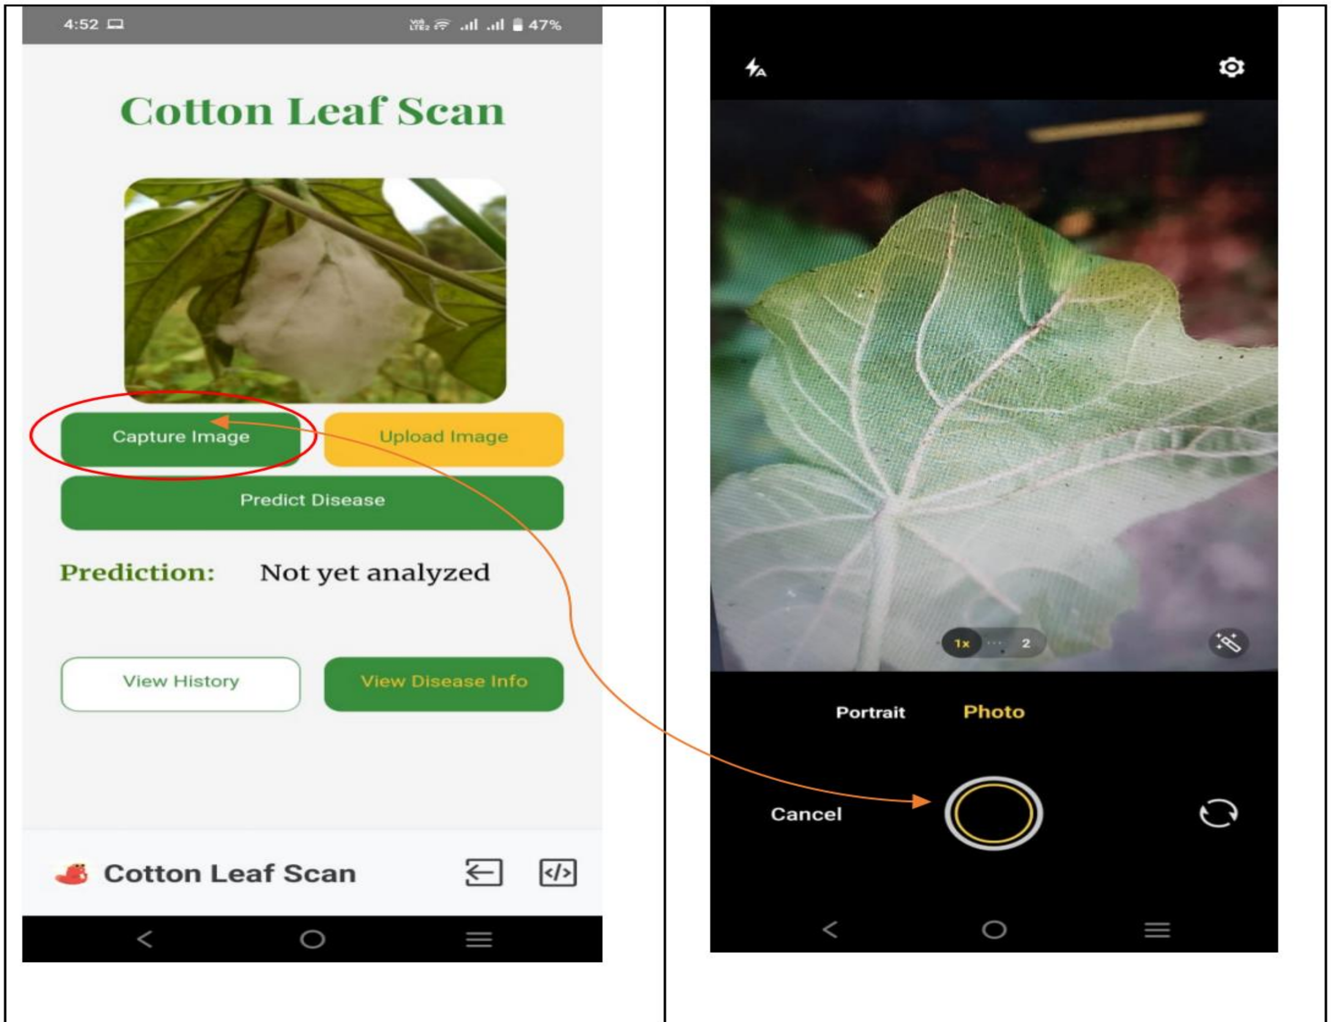

Figure. 11 (a) Front end | Figure. 11(b) Live Capture

Figure 11. User interface of the "Cotton Leaf Scan" Android application: (a) home screen and (b) live camera-based leaf image capturing operation.

(a) Home / capture screen (b) Leaf Curl Virus prediction (c) Bacterial Blight prediction

Figure 12. Android application named "Cotton Leaf Scan" (a) Home/capture screen, (b) Prediction of Leaf Curl Virus and (c) Prediction of Bacterial Blight

Discussion

Component Contribution Analysis

While a formal, quantitative ablation study (systematically retraining the model with individual components removed) was not performed in the present phase of this work and is identified in Section 6 as a priority for the next stage of experimentation, the design and results support a qualitative attribution of performance to the framework's three principal components, summarized in Table 6 [38] [39].

Table 6. Qualitative contribution of framework components (formal quantitative ablation reserved for future work; see Section 6).

Component	Role in the Framework	Expected Effect if Removed
Field-aware augmentation (rotation, shear, zoom, brightness, viewpoint)	Simulates the orientation, lighting, and camera-angle variability of farmer-captured images	Reduced robustness to real-world capture conditions; likely larger train–validation performance gap
EfficientNet-B0 backbone with transfer learning	Provides a strong, ImageNet-pretrained feature basis adapted to the cotton-leaf domain at low parameter cost	Either substantially larger model size (if replaced with a heavier backbone) or reduced accuracy (if trained from random initialization)
TFLite quantization and mobile integration	Compresses the trained model for low-latency, offline, on-device inference	Model would remain accurate but impractical for real-time field deployment on typical farmer smartphones

Comparison with the Broader Literature

The comparative results in Section 4.5 position the proposed framework favourably against recent cotton-disease classification studies. Notably, studies reporting higher raw accuracy on narrower class sets (e.g., four-class settings) are not directly comparable to a seven-class evaluation, since classification difficulty generally increases with the number of visually overlapping categories [40]. When accuracy is considered jointly with class coverage and mobile readiness, the proposed framework compares favourably with heavier ensemble and transformer-based approaches, which — despite reporting competitive or marginally higher accuracy in some cases — do not report on-device latency or provide a deployable application [41].

Practical Implications

The combination of a lightweight backbone, field-representative training data, and a functioning offline mobile application addresses a practical deployment gap rather than a purely academic benchmark improvement. In rural cotton-growing regions with limited or intermittent network connectivity, an offline-capable application removes a key adoption barrier associated with cloud-inference-dependent diagnostic tools. Early and accurate disease identification at the point of observation can plausibly reduce diagnostic delay, limit unnecessary or misdirected pesticide application, and support more targeted disease management — although these downstream socio-economic benefits require field-based validation beyond the scope of the present laboratory-style evaluation.

Limitations

- Formal ablation study: the component contribution analysis in Section 5.1 is qualitative; a systematic, quantitative ablation (retraining with each component independently removed) has not yet been conducted and is planned as immediate follow-up work.
- Geographic and varietal scope: the SAR-CLD-2024 dataset, while field-representative, originates from a specific geographic context; broader geographic and cultivar diversity would strengthen claims of generalizability.
- Absence of large-scale field trials: the mobile application has been evaluated for functional correctness and on-device latency but has not yet undergone large-scale field testing with real farmer users under variable connectivity, hardware, and environmental conditions.
- Non-visual context: environmental variables such as humidity, temperature, and soil condition, which can influence disease onset and progression, are not incorporated into the current image-only framework.

- Interpretability: the current framework does not yet provide visual explanation of model predictions (e.g., Grad-CAM); adding this would improve end-user trust and support agronomic decision-making.
- Figure/label verification: as noted in Section 4.2, the class-axis labels used in the exported confusion matrix and related plots should be verified against the official SAR-CLD-2024 taxonomy before final submission.

Conclusion and Future Work

This study presented a lightweight, mobile-deployable deep learning framework for cotton leaf disease classification, built on a fine-tuned EfficientNet-B0 backbone and trained on the field-representative, seven-class SAR-CLD-2024 dataset. The framework combines targeted, field-aware data augmentation with transfer learning to achieve classification performance that compares favourably with recent, often heavier, cotton-disease classification approaches, while remaining compact enough for real-time, offline inference on commodity Android devices via a functional mobile application. Taken together, these results support the framework's positioning at the intersection of broad diagnostic coverage, competitive accuracy, and genuine field deployability — a combination that remains comparatively underexplored in the existing literature.

Building on this foundation, several directions are planned for future work:

- Conducting a formal, quantitative ablation study isolating the contribution of augmentation strategy, backbone choice, and optimization procedure.
- Expanding the dataset to additional geographic regions, cotton cultivars, and disease severity stages to strengthen generalizability.
- Integrating explainable-AI techniques such as Grad-CAM to provide visual justification for predictions and improve farmer and agronomist trust.
- Conducting field trials with real farmer users to evaluate usability, adoption, and diagnostic impact under authentic connectivity and hardware constraints.
- Extending the framework to additional crops and exploring further model compression (e.g., quantization-aware training, pruning) and lightweight transformer variants for continued efficiency gains.

References

1. Cassman, K. G., & Grassini, P. (2020). A global perspective on sustainable intensification research. *Nature sustainability*, *3*(4), 262-268.
2. Motes, W. C. (2010). Modern agriculture and its benefits–Trends, implications and outlook. *Global Harvest Initiative, Washington, DC.*
3. Nasir, S., Ali, M., Irshad, M., & Wazir, S. (2022). Critical Evaluation of Textile Industry of Pakistan and Way Forward.
4. Earnest, I., Yasar, A., Ansar, J., Nadeem, M., Masih, R., Wasif, I., & Tabinda, A. B. (2026). Water Management and Environmental Performance in the Textile Industry. In *Water Economy in Industry: Strategies to Integrate the Industry to Rationalize Water Consumption* (pp. 393-433). Cham: Springer Nature Switzerland.
5. Ali, M. A., Farooq, J., Batool, A., Zahoor, A., Azeem, F., Mahmood, A., & Jabran, K. (2019). Cotton production in Pakistan. *Cotton production*, 249-276.
6. Ansari, H., Aijaz, A., Kumar, V., Junejo, M., Shah, S., Saad, M., ... & Memon, I. (2019, November). Customizable algorithm for university timetabling and scheduling. In *2019 8th International Conference on Information and Communication Technologies (ICICT)* (pp. 113-117). IEEE.
7. Bather, J. D., Saleem, M., Ahmed, N., Khaskheli, Z., Waheed, Z., Kolachi, A., & Ghaffar, U. (2026). An Adaptive Gamification Optimization Algorithm for Analyzing Programming Learning Difficulties. *Mehran University Research Journal of Engineering and Technology*, *45*(3), 178-192.

8. Belachew, Z. G., & Jenber, A. J. (2024). Cotton protection. In *Cotton Sector Development in Ethiopia: Challenges and Opportunities* (pp. 39-64). Singapore: Springer Nature Singapore.

9. Doro, M. (2025). *Exploring the nexus of cotton (gossypium hirsutum) productivity, food access and copying strategies to cotton price volatility in the Zambezi Valley. a case study of the Mbire district* (Doctoral dissertation, BUSE).

10. Segarra, G., Puopolo, G., Porcel Rodríguez, E., Tomada, S., Giovannini, O., & Pertot, I. (2015). Challenges in the development of a microbial fungicide based on a strain of Lysobacter capsici. In *XVIII. International Plant Protection Congress Mission Possible: food for all through appropriate plant protection. 24–27 August 2015• Berlin (Germany) Abstracts* (p. 452).

11. Wang, S., Xu, D., Liang, H., Bai, Y., Li, X., Zhou, J., ... & Wei, W. (2025). Advances in deep learning applications for plant disease and pest detection: A review. *Remote Sensing*, *17*(4), 698.

12. Upadhyay, A., Chandel, N. S., Singh, K. P., Chakraborty, S. K., Nandede, B. M., Kumar, M., ... & Elbeltagi, A. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. *Artificial Intelligence Review*, *58*(3), 92.

13. Kaur, N., & Devendran, V. (2024). A novel framework for semi-automated system for grape leaf disease detection. *Multimedia Tools and Applications*, *83*(17), 50733-50755.

14. Joshi, K., Yadav, Y., Hooda, S., Nandal, R., Singh, B., Singh, K., ... & Gill, S. S. (2025). Classification of cotton leaf disease using YOLOv8 based k-fold cross validation deep learning method for precision agriculture. *Scientific Reports*, *15*(1), 35602.

15. Haque, M. E., Saykat, M. T. H., Al-Imran, M., Siam, A. H., Uddin, J., & Ghose, D. (2026). An attention enhanced CNN ensemble for interpretable and accurate cotton leaf disease classification. *Scientific Reports*, *16*(1), 4476.

16. Pan, P., Shao, M., He, P., Hu, L., Zhao, S., Huang, L., ... & Zhang, J. (2024). Lightweight cotton diseases real-time detection model for resource-constrained devices in natural environments. *Frontiers in Plant Science*, *15*, 1383863.

17. Shafay, M., Hassan, T., Owais, M., Hussain, I., Khawaja, S. G., Seneviratne, L., & Werghi, N. (2025). Recent advances in plant disease detection: challenges and opportunities. *Plant Methods*, *21*(1), 140.

18. Gupta, P., Nirmal, J., & Mehendale, N. (2025). A survey on computer vision approaches for automated classification of skin diseases. *Multimedia Tools and Applications*, *84*(11), 8673-8705.

19. Mustafa, S. S. (2019). *Computer-Aided Diagnosis of Lesion in Surface Tissue Photographs using Image Processing and Machine Learning* (Doctoral dissertation, Iwate University).

20. Sawyer, T. W. (2021). *Multimodal Optical Imaging for Tissue Characterization and Disease Diagnosis* (Doctoral dissertation, The University of Arizona).

21. Aslam, T., Qadri, S., Saher, N., Nawaz, S. A., Shehzad, M., Qadri, S. F., ... & Khan, D. M. (2021). A spatial model of K-nearest neighbors for classification of cotton (gossypium) varieties based on image segmentation. *Research Journal of Computer Science and Information Technology, Lahore Garrison University*, *5*(1), 24-39.

22. Murugan, T., Badusha, N. A. N. M., Musa, N. S., Alahbabi, E. M. M., Ahmed Alyammahi, R. A., Adege, A. B., ... & Megersa, Z. M. (2026). Research advances in maize crop disease detection using machine learning and deep learning approaches. *Computers*, *15*(2), 99.

23. Zhang, H. W., Wang, R. F., Wang, Z., & Su, W. H. (2025). DLCPD-25: a large-scale and diverse dataset for crop disease and pest recognition. *Sensors*, *25*(22), 7098.

24. Qadri, S. A. A., Huang, N. F., Wani, T. M., & Bhat, S. A. (2024). Advances and challenges in computer vision for image-based plant disease detection: a comprehensive survey of machine and deep learning approaches. *IEEE Transactions on Automation Science and Engineering*, *22*, 2639-2670.

25. Redouan Baghit, Faisal Rahman, Ho P.H Vu, Sreepriya R, Imran Hassan, EL ABBASSI Marouane, ... Rajesh Shahi. (2026). Generative AI in Higher Education: Global Trends, Challenges, and Policy Implications. *International Journal of Computer Information Systems and Industrial Management Applications*, *18*(15s), 419–430.

26. Khan, A., Batool, K., Jalbani, A., Ali, F., Naz, F., Hassan, A., ... & Ghaffar, U. (2026). Reliability-Aware Multi-Representation Fusion for Detecting Deepfake Images on Social Media Platforms. *Journal of Global Social Transformation*, *2*(8.1), 206-225.

27. Hamna Anis, Zahoor Ahmed, Waseema Batool, Muhammad Vahaj Ur Rehman, Faisal Rahman, Muhammad Umair Nazar, ... Nazia Azim. (2026). IoT-Based Remote Monitoring and Control of Robots for Hazardous Environments Using Machine Learning. *International Journal of Computer Information Systems and Industrial Management Applications*, *18*(18s), 176–184.

28. Ali, H., Shifa, N., Benlamri, R., Farooque, A. A., & Yaqub, R. (2025). A fine tuned EfficientNet-B0 convolutional neural network for accurate and efficient classification of apple leaf diseases. *Scientific Reports*, *15*(1), 25732.

29. Ozsari, S., Kumru, E., Ekinci, F., Akata, I., Guzel, M. S., Acici, K., ... & Asuroglu, T. (2024). Deep learning-based classification of macrofungi: Comparative analysis of advanced models for accurate fungi identification. *Sensors*, *24*(22), 7189.

30. Ripon, S., Gani, R., Niha, N. M., Rahat, W. B., Toufiq, S. H., Maisha, M. F., & Ahmed, J. (2025). Cotton leaf image dataset for disease classification and health monitoring. *Data in Brief*, 112142.

31. Abdulsattar, F. S., Abdulsattar, N. S., & Baji, F. S. (2026, May). Plant Leaf Disease Classification Using a Multi-Scale Attention Network. In *2026 2nd International Conference on Computational Intelligence Approaches and Applications (ICCIAA)* (pp. 1-6). IEEE.

32. Muhammad, D., Ali, H., Hussain, Z., Hammad, M., Sangrasi, M., & Ali, S. (2026). Semantic-Aware Multi-Spectral Confusion Regularization for Long-Tailed Visual Recognition. *Spectrum of Engineering Sciences*, *4*(3), 2748-2761.

33. Muhammad, D., Ali, H., Hussain, Z., Hammad, M., Ali, S., Sangrasi, M., & Hussain, P. (2026). Uncertainty-Aware Adaptive Routing for Long-Tailed Recognition (UARL). *Spectrum of Engineering Sciences*, *4*(3), 2732-2747.

34. Muhammad, D., Ali, H., Azeemi, F. A., Khalid, M., Hussain, N., Hussain, Z., ... & Hammad, M. (2026). Va-gkd: Variance-aware adaptive group knowledge distillation for long-tailed recognition. *Spectrum of Engineering Sciences*, *4*(3), 2459-2471

35. Khaskheli, Z., Ali, Q., Shah, M. S., Muhammad, D., Ali, F., Ali, H., ... & Gaffar, U. (2026). Improving Human Action Recognition Through Cross-View Camera-Aware Augmentation. *Mehran University Research Journal of Engineering and Technology*, *45*(3), 205-218.

36. Muhammad, D., Ali, H., Azeemi, F. A., Hussain, Z., Sangrasi, M., Hammad, M., ... & Hussain, P. (2026). Geometry-Aware Contrastive Learning for Binay Imbalanced Distribution. *Spectrum of Engineering Sciences*, *4*(3), 2719-2731.

37. Muhammad, D., Ali, H., Ali, A., Hussain, Z., Hammad, M., Azeemi, F. A., ... & Ali, S. (2026). Adaptive Multi-Granularity Mixture Consistency with Dynamic Prototype-Guided Rebalancing for Long-Tailed Visual Recognition. *Spectrum of Engineering Sciences*, *4*(3), 2621-2632.

38. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. Proceedings of the 36th International Conference on Machine Learning, 97, 6105–6114.

39. Pan, P., Guo, W., Zheng, X., Hu, L., Zhou, G., & Zhang, J. (2024). Lightweight cotton diseases real-time detection model for resource-constrained devices in natural environments. Frontiers in Plant Science, 15, 1383863.

40. Joshi, K., Yadav, Y., Hooda, S., Nandal, R., Singh, B., Singh, K., Tuteja, N., Gill, R., & Gill, S. S. (2025). Classification of cotton leaf disease using YOLOv8 based k-fold cross validation deep learning method for precision agriculture. Scientific Reports, 15, 35602.

41. Haque, M. E., et al. (2026). An attention enhanced CNN ensemble for interpretable and accurate cotton leaf disease classification. Scientific Reports, 16, 4476.

2026 by the authors; Journal of Global Social Transformation (JGST). This is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) license (http://creativecommons.org/licenses/by/4.0/).