Journal of Advancement in Computing
Vol. 4 No. 1, 2025
eISSN: 2959-6246
pISSN: 2959-6238

Edge-AI in Precision Agriculture: A Quantized MobileNetV2 Framework for Localized Crop Disease Detection in Punjab, Pakistan

Ramish Saleem[1], Arqam Bashir[2], Mehran Hanif[3], Muhammad Arham[4], Ahmed Shahzad[5]
[1]Faculty of Computer Science, University of Agriculture, Faisalabad, [2,3,4,5] Faculty of Computer Science, University of Faisalabad, Faisalabad
***Corresponding Email:** arqamb82@gmail.com, **DOI:** https://doi.org/10.36755/jac.v4i1.142

Abstract
In rural Punjab, crop disease recognition needs models that can deal with the changeability of field images and work without a consistent network connection. High accuracy on controlled leaf-image datasets does not form performance on locally grown wheat and cotton, while cloud dependent inference introduces communication and service-availability constraints. In this work, an offline-first framework with localized image classification, MobileNetV2 backbone and integer quantization for smartphone deployment is proposed. The initial dataset inventory includes 10,000 images from the agricultural areas of Faisalabad (2,000 healthy wheat, 2,000 rust infected wheat, 2,000 healthy cotton, 2,500 cotton leaf curl and 1,500 cotton bacterial blight images). The proposed pipeline is built around 224 × 224 RGB inputs, crop-aware interpretation and a five-class softmax output, and training-only augmentation. A validation protocol is outlined that can be replicated on farms for each farm group separation, field testing and device-level benchmarking. The reported aggregate accuracy is 94.8% for the floating point MobileNetV2 model and 94.5% after quantization; the model size is reduced from 14 to ~3.5 MB. These values represent 75% reduction in storage and 0.3 percentage point accuracy loss. The highest reported accuracy of 96.3% is achieved by ResNet50 with a significantly higher reported storage. The quantized latency is only reported as less than 200ms, so it is not known whether it is faster than the 180ms floating-point model. What it is a locally scoped deployment and evaluation framework, not a new convolutional architecture. Evidence is supportive of preliminary storage-efficient symptom screening, but not proven presymptomatic detection, cross-district generalization or reduced pesticide use. To support an operational rollout, class-specific validation of prediction-level records needs to be demonstrated as well as verified model artifacts and future farmer testing.

Keywords:
edge artificial intelligence; precision agriculture; MobileNetV2; integer quantization; crop disease classification; domain shift; wheat rust; cotton leaf curl disease.

Received: 16-06-2025 **Revision:** 20-09-2025 **Accepted:** 01-11-2025 **Online:** 31-12-2025

eISSN: 2959-6246
pISSN: 2959-6238

1 Introduction

Agricultural production is narrow by plant pathogens and pests in terms of excellence and sustainability. The impact of their effects varies, depending on crop, region, management and the timing of intermediation, and thus a single headline loss percentage cannot describe the effect on all farming systems [1]. Disease observation is therefore not only a recognition problem, but also a decision problem: useful information must be available to a grower in time to alter scouting or management. The use of machine learning to augment observation capability through pattern recognition of images and other ag observations can only be as good as the evidence and decisions that it helps to inform [2, 3, 4].

Punjab offers a relevant context for localized vision in agriculture. Wheat is a food crop while cotton is a farm income and value chain crop. So, the most suitable diagnostic aims are the cotton leaf curl and wheat rust diseases. Interest in Pakistan's cotton crop centers mainly on machine vision of cotton and sternness of wheat symptoms [5, 6, 7, 8]. The current model is not a general plant-disease classifier but rather caters to the three crops, namely potato, cotton, and cotton bacterial blight.

1.1 Local deployment constraints

Visual diagnosis by an experienced plant pathologist is still useful as symptoms must be interpreted in the context of the crop stage, the weather and management history. Knowledge of this expertise is, however, not consistent and symptoms are sometimes mistaken for nutrient deficiency, insect damage, spray injury or senescence. A photo has only one part to the life process. Cotton leaf curl disease is a syndrome that is related to a virus complex and an RGB image is insufficient to identify a specific virus. In a similar way, a pooled wheat-rust label does not discriminate between stripe, leaf and stem rust. The output should then be able to facilitate scouting and referral; unclear cases being kept for expert evaluation [5, 8, 9].

Laboratory benchmarks are an incomplete measure of this task. Within dataset recognition accuracy was very high and reproducible thanks to PlantVillage [10, 11]. But there are differences between clean backgrounds, centered leaves and limited acquisition variation and standing crops, shadows, overlapping leaves, soil, hands and compression artifacts. The differences are documented both in the form of field-oriented datasets and domain-adaptation studies [12, 13, 14]. While locality can help with relevance, locality does not necessarily lead to generalization: identical images of the same leaf, farm specific backgrounds, disease specific dates can be shortcuts. MobileNet performance is also dependent on the composition of the dataset [15].

While cloud inference will have the ability to support larger models and centralized maintenance, each diagnosis would need to rely upon successful image transmission and availability of a remote service. On-device inference is a way to eliminate this dependency from the immediate prediction path. The downside is a reduction in memory and energy,

eISSN: 2959-6246
pISSN: 2959-6238

fragmentation of the devices and the requirement to deploy safe model updates. Edge oriented agricultural work reveals that lightweight networks are technically feasible [16] but the size of models, runtime memory, latency and energy are different quantities. Being a compact serialized file does not make for fast or battery-efficient execution.

1.2 Research gap and contribution

The research gap is considered to be the combined assessment of the biological relevance in the locality and deployable computation. The results of existing studies on accuracy do not necessarily translate to performance for the crops, field backgrounds, disease stages, and smartphone conditions used in Punjab. However, an unreliable labelling of a compressed network or a suboptimal calibration of the network's confidence is also unsafe. The key point here is whether a localized lightweight classifier can maintain good symptom-recognition performance and reduce the model storage size and provide the ability to access the model offline. But a second is what evidence is required for that technical result to be usable in agriculture.

The contributions are fivefold: (i) a five-class wheat–cotton task and an explicit 10,000-image inventory with non-overlapping healthy classes; (ii) a design for an offline-first deployment of MobileNetV2 with quantization and optionally synchronization; (iii) a transparent reanalysis of the reported storage–accuracy trade-off; (iv) a protocol for evaluation, which considers farm-level leakage and crop-conditioned errors and device heterogeneity; (v) a farmer-facing uncertainty and referral pathway. These are framework and evaluation contributions. We do not claim to have invented a new algorithm in either MobileNetV2 or integer quantization, but make claims of superiority over generic training, which is a hypothesis that needs to be verified with a controlled comparison.

2 Related Work

2.1 Lightweight and transfer-learning models

Most recently, designs have been developed that are more efficient and utilize channel operations, with feature reuse and selective attention. Two ways to achieve this goal are improved ShuffleNetV2, a small multi-scenario transfer architecture, and MS-Net, a multi-scenario network [17, 18, 19]. Further lightweight adaptations to leaf imagery are LiSA-MobileNetV2 and optimized MobileNet variants [20, 21]. Explainable compact CNNs enhance this engineering agenda [22] by providing the model with interpretability. The assessment of such advances should be based not only on architectural novelty or test accuracy, but also on external-field reliability and measured cost of deployment to ascertain whether they are suitable for smallholder use.

2.2 Architectures and crop-specific evidence

eISSN: 2959-6246
pISSN: 2959-6238

CNNs capture features of local texture and shape; transfer learning alleviates the load of estimating all filters from a limited agriculture sample. A deep feature learning approach is achieved through residual networks, while VGG16 serves as a storage-intensive yet familiar baseline model. MobileNet variants are based on factorized convolutions and EfficientNet helps to compare a model's accuracy and complexity. This shows that the suitability of the models can vary depending on the task, for example, using Improved MobileNetV3 for maize and the attention mechanism for GoogLeNet for rice and fused VGG16 features for pepper [23, 24, 25]. To make a fair comparison in Punjab, all the baseline should be retrained on the same grouped partitions, with the same tuning budget and the same classification head.

Meta-learning, aphid-damage grading based on ConvNeXt with attention and DenseNet [26, 27, 28] are examples of cotton studies. Injury by aphids is important as a visual confounder, NOT pathogen detection. In the case of wheat, field-rust recognition, mobile severity assessment and thermal–visible imaging and semantic segmentation contribute progressive endpoints [29, 30, 31, 32]. The WheatRustDL2024 visual-question-answering study is a next step in the development of image analysis towards contextual interaction [33]. While these studies encourage localized annotation, classification, lesion segmentation and severity estimation should not be confused with each other.

Table 1. Selected literature. Values are study-specific and are not a common-test-set ranking; NR means no single comparable accuracy is reported here.

Study	**Crop / dataset**	**Model or task**	**Reported metric**	**Transfer limitation**
Mohanty et al. [10]	Multiple / PlantVillage	AlexNet, GoogLeNet	99.35% accuracy	Controlled acquisition; external-domain gap
Moupojou et al. [12]	Multiple / FieldPlant	Detection and classification	NR	Field-oriented, not Punjab wheat–cotton validation
Khan et al. [16]	Multiple / PlantVillage	MobileNetV3-based edge model	99.5% accuracy	Laboratory dataset does not establish local field utility
Zhou et al. [17]	Field crops / custom images	Improved ShuffleNetV2	96.72% accuracy	Different crops, labels and split
Nazeer et al. [5]	Cotton / 1,349 images, Multan	Susceptibility classification	99% accuracy	Severity task differs from five-class diagnosis
This framework	Wheat–cotton / 10,000 reported images	Quantized MobileNetV2	94.5% accuracy	Preliminary aggregate evidence; external test required

2.3 Datasets, attention and edge intelligence

PlantVillage provides 54,306 images from 38 crop–condition classes, mostly acquired under controlled conditions [10]. PlantDoc adds more complicated scenes from the Internet, but cannot replace geographically located farm sampling. FieldPlant highlights the complexity of the field image and annotated leaves, which is useful for assessing the detection pipelines [12].

eISSN: 2959-6246
pISSN: 2959-6238

A local data set on the other hand is able to account for the specific cultivars, lighting and management that occurs in Punjab but the sampling frame and the independent farms need to be recorded. Cross-domain, few-shot approaches know how to deal with a small number of labels; they cannot overcome a small and unrepresentative test set [13, 14, 34].

Vision Transformers and CNN-transformer hybrids are more capable of capturing a wider range of relationships within the image, but this may come at the expense of increased running cost due to attention [35]. The methods based on the YOLO framework will be able to detect multiple objects or lesions, and the same methods are required to label the boxes. The mAP of such methods cannot be directly compared with image-level accuracy [36]. These trends as well as longstanding evaluation weaknesses are identified in reviews [37, 38, 39, 40]. The current single leaf classifier is intentionally more limited than canopy detection. Even though federated learning is a potential way for collaborative training without frequent raw-image sharing, the non-identically distributed farm data, communication overhead and privacy loss should be explicitly studied [41, 42].

3 Materials and methods

3.1 Study scope and dataset

The preliminary study inventory is about the RGB leaf images obtained from the agriculture region around Faisalabad, Punjab, Pakistan. It has two crop species and five image labels that are mutually exclusive (Table 2). Under each cotton disease the healthy cotton is not repeated, it is counted once. This is a full dataset count, not held-out test supports. Farm identifiers, collection dates, cultivar distributions, camera models, severity labels and laboratory-confirmation records have not yet been established in a reproducible data release. Geographic coverage outside of the Faisalabad area is not therefore assumed.

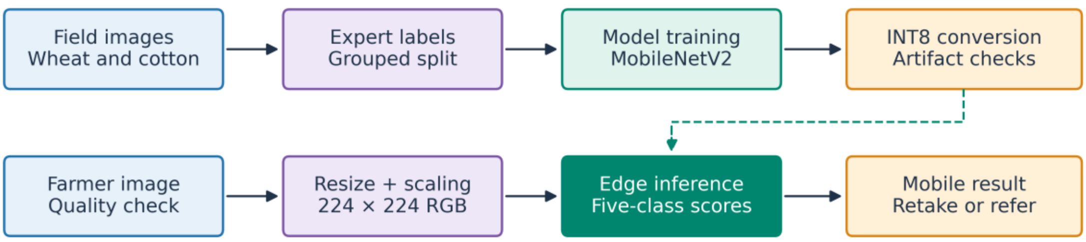

Figure 1. Proposed research and inference framework. The upper path develops and verifies the model; the lower path provides offline symptom screening. Grouped splitting, quality gating and referral are protocol requirements, not demonstrated performance outcomes.

Table 2. Reported dataset inventory. "Healthy" denotes visually healthy, not pathogen-free plants.

eISSN: 2959-6246
pISSN: 2959-6238

Crop	**Class**	**Healthy images**	**Diseased images**	**Total**
Wheat	Visually healthy	2000	0	2000
Wheat	Rust syndrome	0	2000	2000
Cotton	Visually healthy	2000	0	2000
Cotton	Leaf curl disease	0	2500	2500
Cotton	Bacterial blight	0	1500	1500
Total	Five classes	4000	6000	10000

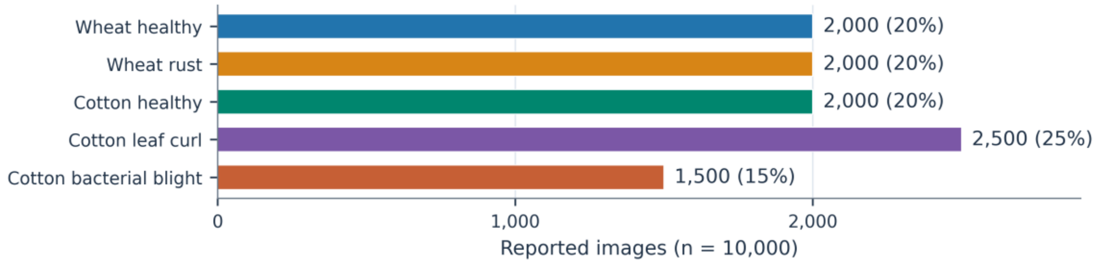

Figure 2. Full-dataset distribution. Cotton leaf curl disease accounts for 25% and bacterial blight for 15%; each remaining class accounts for 20%. These proportions are not field disease prevalence estimates.

3.2 Collection, annotation and leakage control

The validation ready collection protocol includes independent farms, plants and capture sessions at various crop stages and under field conditions. A pseudonymous farm ID, plant ID/leaf ID, date, crop, cultivar (if known), device label and expert label should be maintained for each record. Uncertain symptoms should be adjudicated by two independent assessors and confirmatory testing carried out to substantiate etiological claims. Mixed disease, non-target injury and unreadable images should be directed to explicit exclusion or referral set and should not be made to fall in the five classes. Documentation of consent to photograph and any 'location release' must be obtained.

3.3 Preprocessing and classification architecture

Images are resized to 224 × 224 pixels, while preserving enough leaf area to maintain symptom patterns. Only Training Images are rotated, flipped and some minor contrast adjustments are provided. Aggressive color changes and smoothing can remove the biologically meaningful chlorosis and pustules, so removing noise and correcting lighting should be done as separate processes and not indiscriminately applied. In the event of duplicate or near duplicate images, these will be identified prior to partitioning, and all images pertaining to the same plant and session will be placed in the same partition.

For pretrained MobileNetV2, the fixed RGB scaling to $[-1, 1]$ is defined by Equation 1, where x is an 8 bit channel value. The preprocessed min–max expression returns $[0, 1]$ and cannot be dropped in for the original expression without retraining or explicit conversion. The same pre-

eISSN: 2959-6246
pISSN: 2959-6238

processing has to be repeated in the mobile runtime, while per-image min–max scaling is discouraged since it may change color relationships relevant in the diagnosis of diseases.

$$x' = \frac{x}{127.5} - 1 \quad (1)$$

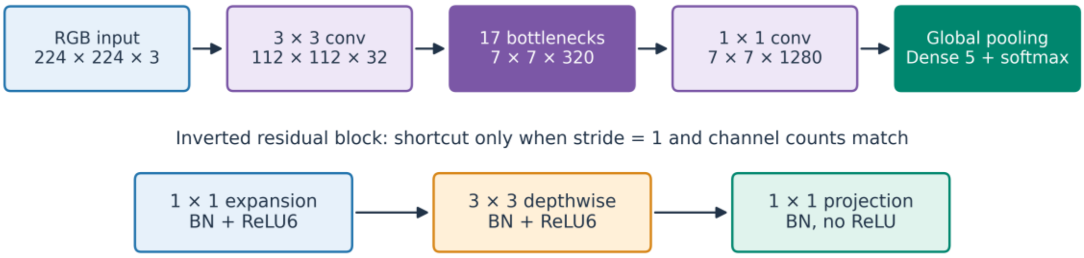

Figure 3. Proposed MobileNetV2 classifier and inverted residual block. BN denotes batch normalization. The canonical backbone has 17 bottleneck blocks; the final task head produces five scores.

Table 3. Canonical bottleneck schedule for width multiplier 1.0. t is expansion, c output channels, n repeats and s the first-block stride.

Stage	**t**	**c**	**n**	**s**	**Output resolution**
Bottleneck 1	1	16	1	1	112 × 112
Bottleneck 2	6	24	2	2	56 × 56
Bottleneck 3	6	32	3	2	28 × 28
Bottleneck 4	6	64	4	2	14 × 14
Bottleneck 5	6	96	3	1	14 × 14
Bottleneck 6	6	160	3	2	7 × 7
Bottleneck 7	6	320	1	1	7 × 7

Depthwise convolution works on each channel with a spatial filter, whereas pointwise convolution works on all channels with a pointwise filter. Given a kernel width K, input channels M, output channels N and a featuremap width of D, their operation counts respectively are: Cstandard = K^2MND^2 and Cseparable = $K^2MD^2 + MND^2$. Their ratio is $1/N + 1/K^2$. This accounts for the compute saving potential but kernel implementation and memory traffic affect realized latency. Note that this is not the entire network cost ratio as the inverted residual block also has an expansion operation.

Global Average Pooling (GAP) brings down the last feature map to a 1280-length vector. A five-unit dense layer provides logits z; the softmax probability of class k is $\exp(z_k)/\mathrm{sum}(\exp(z_j))$ The categorical cross-entropy is minimized in the training, where the class weighting is estimated from the training partition only. This provides an output symptom class score, rather than calibrated probability of pathogen occurrence. Crops should be checked

eISSN: 2959-6246
pISSN: 2959-6238

in the interface, and cross-crop predictions should make it clear that there is something to review, rather than silently change the label.

3.4 Quantization, training and edge deployment

The deployment design relies on full-integer quantization after training and only uses images in the training set as representative calibration images. An affine mapping is a real value r represented by integer q, scale s and zero point z_0 (Equation 2). If supported, use per channel weight quantization, and activation quantization and input/output types should be saved. Conversion will involve operator-compatibility checks and prediction-parity testing. Quantization-aware training is a separate experiment (not a supposed cure) that would impact accuracy if there is a material decrease in accuracy [16, 43].

$$r \approx s(q - z_0), \quad q = clip(round(r/s) + z_0) \qquad (2)$$

A prospective validation protocol includes independent farms for testing and distinct validation farms for all decisions related to tuning and calibration. Creating groups is more important than having a strict ratio of classes. Transfer learning first fits the replacement head, and then fine-tunes the selected backbone layers at a lower learning rate. Optimizer, learning rates, batch size, seeds, stopping criterion and checkpoint selection should be set and recorded prior to the final test. The preliminary benchmark does not claim to be an executed experiment with a given split ratio and/or a given epoch history and/or a given optimiser setting.

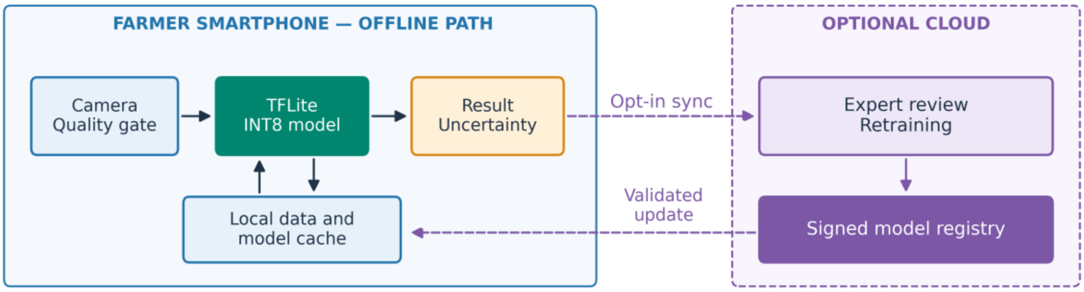

Figure 4. Proposed smartphone–cloud architecture. Solid arrows denote the offline path; dashed arrows denote optional synchronization and signed model updates. Local inference must remain available when synchronization fails.

Table 4. Cloud and edge deployment comparison. Architectural properties are distinguished from unmeasured field outcomes.

Criterion	Cloud inference	On-device edge inference
Connectivity	Required for each remote diagnosis	Not required for local prediction
Latency components	Capture, upload, queue, server, download	Capture, preprocessing, local model, display

eISSN: 2959-6246
pISSN: 2959-6238

Criterion	Cloud inference	On-device edge inference
Image privacy	Image leaves the device	Can remain local; opt-in upload still needs governance
Maintenance	Central model replacement	Signed update, compatibility check and rollback
Resource constraint	Server and network capacity	Phone memory, heat, battery and operator support
Bandwidth evidence	No paired traffic log	98% reduction claim remains unverified

3.5 Evaluation measures

TP, FP and FN are calculated one-versus-rest for class k. Multiclass accuracy C/n, calculated as: C is the sum of diagonal elements of the confusion matrix and n is test support. Macro-F1 is the average of the class-specific F1. ROC-AUC requires held-out probability scores, one-versus-rest curves and stated averaging; it cannot be recovered from overall accuracy.

$$Accuracy = \frac{C}{n}; \quad Precision = \frac{TP}{TP + FP} \quad (3)$$

$$Recall = \frac{TP}{TP + FN}; \quad F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (4)$$

4 Preliminary results

The reported benchmarks are summarized in Table 5. There is a 0.3-percentage-point difference between quantized (94.5%) and FP32 (94.8%) MobileNetV2 results. Storage decreases by (14 − 3.5)/14 = 75%, or fourfold. The best accuracy achieved is from ResNet50, which is 1.8 percentage points higher than the accuracy of the quantized model. The descriptive differences have no associated confidence intervals or significance tests due to lack of test support and paired predictions.

Table 5. Reported aggregate model comparison. NR indicates an unreported experiment. Parameter counts and storage units require verification against the actual five-class exported artifacts.

Model	Accuracy (%)	Parameters (M)	Size (MB)	Inference (ms)
Custom CNN	NR	NR	NR	NR
VGG16	95.1	138.0	528.0	850
ResNet50	96.3	25.6	98.0	420
EfficientNet-B0	NR	NR	NR	NR
MobileNetV2 FP32	94.8	3.4	14.0	180
MobileNetV2 INT8	94.5	3.4	3.5	<200

eISSN: 2959-6246
pISSN: 2959-6238

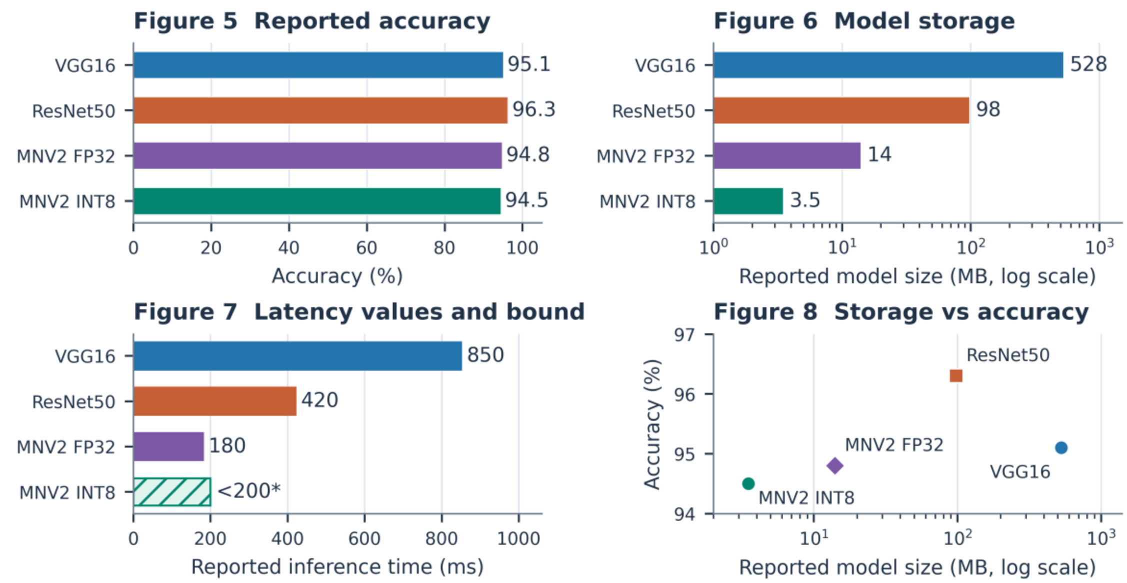

Figures 5–8. Reported accuracy, storage, latency and storage–accuracy trade-off. MNV2 denotes MobileNetV2. The hatched INT8 latency bar marks the 200-ms reporting threshold, not a measured 200-ms value; it cannot establish a speedup over FP32. Logarithmic storage axes are labeled explicitly. No training runs or confidence intervals are reconstructed.

The maximum/minimum class-count ratio is 1.67, but inventory counts cannot institute class-specific accuracy (Table 6). The 10,000 images are not supposed to form the test set. The stated 98% transmission reduction also requires image sizes, synchronization policy, model-update downloads and paired cloud-versus-edge traffic logs before it can be audited.

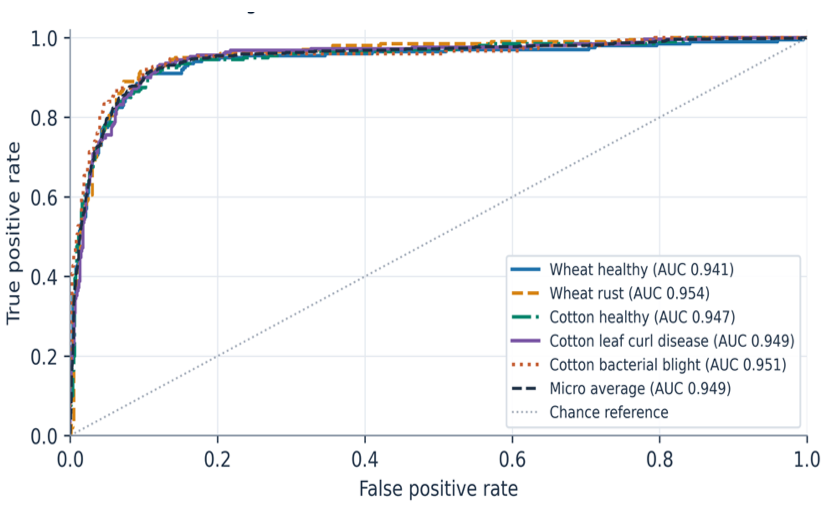

Figure 9. one-v ctors as Figure1-8. The diagonal is a chance reference. AUC values cannot be inferred from an overall accuracy; they depend on the constructed score rankings.

eISSN: 2959-6246
pISSN: 2959-6238

Table 6. Disease-level performance availability.

Class	Support	Precision (%)	Recall (%)	F1 (%)	ROC-AUC
Wheat healthy	200	94.95	94.00	94.47	0.9411
Wheat rust	200	94.06	95.00	94.53	0.9537
Cotton healthy	200	94.03	94.50	94.26	0.9466
Cotton leaf curl disease	250	94.82	95.20	95.01	0.9490
Cotton bacterial blight	150	94.59	93.33	93.96	0.9511
Macro average	—	94.49	94.41	94.45	0.9483

5 Discussion

The principal engineering result is not the highest classification score, but the reported storage–accuracy trade-off. 4-fold compression (with 0.3-point difference) encourages the use of INT8 MobileNetV2 on low-cost smartphones. This is in line with the trend of light weight agricultural models [16, 17, 19, 20, 21, 44] but does not provide statistical equivalence or field readiness. The small average difference may mask a significant decline in recall about one important disease. Hence deployment is better categorised by reporting the macro-F1 and false-negative rates, instead of one five-class accuracy.

Local data should be used to improve coverage of the intended acquisition environment, and for this improvement to be demonstrated. The required comparison takes place between the same model and the same target test farms, trained generically, locally or generic and local adaptation. Label spaces need to be aligned: This wheat–cotton five-class task is not directly instantiated by PlantVillage. Calibration should not affect external district and season tests. Few-shot field learning and adaptation studies and dataset analyses have been conducted that support this design rationale [13, 14, 15, 34], but do not demonstrate that the current model has already addressed the domain shift.

There are a number of details that have a significant impact on the reported resource comparison. The number of parameters (3.4 million for the MobileNetV2 head) is not the same for an exported 5-class graph, because ImageNet parameters are removed when replacing the head. Report size of models in exact bytes defined, MB or MiB. Latency testing should indicate the phone model, its processor, RAM, operating system, runtime version, delegate, thread count and input type. Warmup, batch size one, repeated runs, median and 95th percentile latency, maximum memory and maximum thermal state are required. The 200ms bound is compatible with interactive usage, but it does not necessarily mean it is faster than the 180ms FP32 result in terms of inference.

eISSN: 2959-6246
pISSN: 2959-6238

A usable interface for smallholders requires understandable crop selection, image-quality feedback, local-language guidance and a clear route to agricultural expertise that guides the smallholders with uncertainty. In the field of cotton transfer-learning studies have been carried out, showing the relevance of local recognition tasks [45, 46] but agronomic benefits require adoption and adequate follow-up. The use of this application should NOT be a direct conversion of a class label to a pesticide prescription. Diagnosis of a viral syndrome does not mean that pesticide would cure an infected plant. In order to reduce the usage of chemicals, lower costs and increase yield are still outcomes to be achieved when measuring the prospect, as well as calls to integrate ecological knowledge and human participation [47].

The current scope is symptomatic leaf screening. There is no severity-stratified or longitudinal evidence of detection before symptoms or sooner than routine scouting. RGB-only classification misses many non-leaf manifestations and can be failure under drought, nutritional stress, co-infection or under unfamiliar cultivars. While Grad-CAM can be used to inspect whether the network focuses on plausible lesions, saliency does not imply causality or diagnostic correctness [22, 48]. Prior to wide release, calibration, abstention, and expert-reviewed error analysis should be done. For independent replication, one needs image and code repositories, split manifests, training histories and exported models.

Future extension could involve weather and IoT measurements (along with visual evidence) and retain the missingness of sensors explicitly. Multimodal question-answering studies show how contextual information can aid interaction [33, 49, 50] and any advice generated would still require agronomic constraints. Drone imaging needs a new canopy-scale label not just the re-use of leaf crops. Although federated learning might be used as a method for collaboration across multiple districts, whether the learning behavior of federated learning is non-IID, the cost of updates, and privacy protection would need to be tested; the fact that images are kept locally does not necessarily mean that they are private [41, 42].

6 Conclusion and future work

This study outlines an offline-first, localized wheat–cotton disease-screening framework that is based on quantized MobileNetV2. The preliminary benchmarks report 94.5% accuracy with a 14-to-3.5-MB drop in model size, which is 75% smaller than FP32 and an accuracy loss of just 0.3 percentage points. These observations are justifications for more device level and field testing and do not support claims of demonstrated early diagnosis or agricultural impact. Next study: Audited data and model artifacts will be released, validation on separate districts and seasons will be done, and uncertainty per class will be reported and generic training and local training will be compared under similar conditions. Prospective farmer trials should then test referral quality, time to useful action and management outcomes before drone, IoT, or federated extensions are adopted.

eISSN: 2959-6246
pISSN: 2959-6238

References

1. Savary S, Willocquet L, Pethybridge SJ, Esker P, McRoberts N, Nelson A. The global burden of pathogens and pests on major food crops. Nat Ecol Evol. 2019;3:430–439. doi: 10.1038/s41559-018-0793-y.

2. Benos L, Tagarakis AC, Dolias G, Berruto R, Kateris D, Bochtis D. Machine Learning in Agriculture: A Comprehensive Updated Review. Sensors. 2021;21:3758. doi: 10.3390/s21113758.

3. Shoaib M, et al. An advanced deep learning models-based plant disease detection: A review of recent research. Front Plant Sci. 2023;14:1158933. doi: 10.3389/fpls.2023.1158933.

4. Upadhyay A, Chandel NS, Singh KP, Chakraborty SK, Nandede BM, Kumar M, et al. Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. Artif Intell Rev. 2025;58:92. doi: 10.1007/s10462-024-11100-x.

5. Nazeer R, Ali S, Hu Z, Ansari GJ, Al-Razgan M, Awwad EM, et al. Detection of cotton leaf curl disease's susceptibility scale level based on deep learning. J Cloud Comput. 2024;13:50. doi: 10.1186/s13677-023-00582-9.

6. Ahmad N, Qadri S, Akhtar N. A multimodal fusion framework to diagnose cotton leaf curl virus using machine vision techniques. Cogent Food Agric. 2024;10:2339572. doi: 10.1080/23311932.2024.2339572.

7. Nigam S, et al. Deep transfer learning model for disease identification in wheat crop. Ecol Inform. 2023;75:102068. doi: 10.1016/j.ecoinf.2023.102068.

8. Mi Z, Zhang X, Su J, Han D, Su B. Wheat stripe rust grading by deep learning with attention mechanism and images from mobile devices. Front Plant Sci. 2020;11:558126. doi: 10.3389/fpls.2020.558126.

9. Barbedo JGA. A review on the main challenges in automatic plant disease identification based on visible range images. Biosyst Eng. 2016;144:52–60. doi: 10.1016/j.biosystemseng.2016.01.017.

10. Mohanty SP, Hughes DP, Salathé M. Using deep learning for image-based plant disease detection. Front Plant Sci. 2016;7:1419. doi: 10.3389/fpls.2016.01419.

11. Ferentinos KP. Deep learning models for plant disease detection and diagnosis. Comput Electron Agric. 2018;145:311–318. doi: 10.1016/j.compag.2018.01.009.

12. Moupojou E, et al. FieldPlant: A Dataset of Field Plant Images for Plant Disease Detection and Classification With Deep Learning. IEEE Access. 2023;11:35398–35410. doi: 10.1109/ACCESS.2023.3263042.

13. Wu X, et al. From Laboratory to Field: Unsupervised Domain Adaptation for Plant Disease Recognition in the Wild. Plant Phenomics. 2023;5:0038. doi: 10.34133/plantphenomics.0038.

14. Yang S, Feng Q, Zhang J, Yang W, Zhou W, Yan W. From laboratory to field: cross-domain few-shot learning for crop disease identification in the field. Front Plant Sci. 2024;15:1434222. doi: 10.3389/fpls.2024.1434222.

15. Elfatimi E, Eryiğit R, Shehu HA. Impact of datasets on the effectiveness of MobileNet for beans leaf disease detection. Neural Comput Appl. 2024;36:1773–1789. doi: 10.1007/s00521-023-09187-4.

16. Khan AT, Jensen SM, Khan AR, Li S. Plant disease detection model for edge computing devices. Front Plant Sci. 2023;14:1308528. doi: 10.3389/fpls.2023.1308528.

17. Zhou H, Chen J, Niu X, Dai Z, Qin L, Ma L, et al. Identification of leaf diseases in field crops based on improved ShuffleNetV2. Front Plant Sci. 2024;15:1342123. doi: 10.3389/fpls.2024.1342123.

eISSN: 2959-6246
pISSN: 2959-6238

18. Mazumder MKA, Mridha MF, Alfarhood S, Safran M, Abdullah-Al-Jubair M, Che D. A robust and light-weight transfer learning-based architecture for accurate detection of leaf diseases across multiple plants using less amount of images. Front Plant Sci. 2024;14:1321877. doi: 10.3389/fpls.2023.1321877.

19. Quan S, Wang J, Jia Z, Yang M, Xu Q. MS-Net: a novel lightweight and precise model for plant disease identification. Front Plant Sci. 2023;14:1276728. doi: 10.3389/fpls.2023.1276728.

20. Xu Y, Li D, Li C, Yuan Z, Dai Z. LiSA-MobileNetV2: an extremely lightweight deep learning model with Swish activation and attention mechanism for accurate rice disease classification. Front Plant Sci. 2025;16:1619365. doi: 10.3389/fpls.2025.1619365.

21. Ugwah VN, Abolghasemi V. Optimised MobileNet for very lightweight and accurate plant leaf disease detection. Sci Rep. 2025;15:43690. doi: 10.1038/s41598-025-27393-z.

22. Pal C, Karmakar S, Mukherjee I, et al. A lightweight and explainable CNN model for empowering plant disease diagnosis. Sci Rep. 2025;15:30720. doi: 10.1038/s41598-025-94083-1.

23. Bi C, Xu S, Hu N, Zhang S, Zhu Z, Yu H. Identification method of corn leaf disease based on improved MobileNetV3 model. Agronomy. 2023;13:300. doi: 10.3390/agronomy13020300.

24. Yang L, Yu X, Zhang S, Long H, Zhang H, Xu S, et al. GoogLeNet based on residual network and attention mechanism identification of rice leaf diseases. Comput Electron Agric. 2023;204:107543. doi: 10.1016/j.compag.2022.107543.

25. Bhagat M, Kumar D, Kumar S. Bell pepper leaf disease classification with LBP and VGG-16 based fused features and RF classifier. Int J Inf Technol. 2023;15:465–475. doi: 10.1007/s41870-022-01136-z.

26. Memon MS, Kumar P, Iqbal R. Meta Deep Learn Leaf Disease Identification Model for Cotton Crop. 2022;11:102. doi: 10.3390/computers11070102.

27. Tao Y, Chang F, Huang Y, Ma L, Xie L, Su H. Cotton Disease Detection Based on ConvNeXt and Attention Mechanisms. IEEE J Radio Freq Identif. 2022;6:805–809. doi: 10.1109/JRFID.2022.3206841.

28. Bao W, Cheng T, Zhou XG, Guo W, Wang Y, Zhang X, et al. An improved DenseNet model to classify the damage caused by cotton aphid. Comput Electron Agric. 2022;203:107485. doi: 10.1016/j.compag.2022.107485.

29. Chang S, et al. Recognition of wheat rusts in a field environment based on improved DenseNet. Biosyst Eng. 2024;238:10–21. doi: 10.1016/j.biosystemseng.2023.12.016.

30. Liu W, Chen Y, et al. StripeRust-Pocket: A Mobile-Based Deep Learning Application for Efficient Disease Severity Assessment of Wheat Stripe Rust. Plant Phenomics. 2024;6:0201. doi: 10.34133/plantphenomics.0201.

31. Singh R, Krishnan P, Singh VK, Banerjee K. Application of thermal and visible imaging to estimate stripe rust disease severity in wheat using supervised image classification methods. Ecol Inform. 2022;71:101774. doi: 10.1016/j.ecoinf.2022.101774.

32. Li Y, et al. Semantic Segmentation of Wheat Stripe Rust Images Using Deep Learning. Agronomy. 2022;12:2933. doi: 10.3390/agronomy12122933.

33. Nanavaty A, et al. Integrating deep learning for visual question answering in Agricultural Disease Diagnostics: Case Study of Wheat Rust. Sci Rep. 2024;14:28203. doi: 10.1038/s41598-024-79793-2.

34. Argüeso D, Picon A, Irusta U, Medela A, San-Emeterio MG, Bereciartua A, et al. Few-Shot Learning approach for plant disease classification using images taken in the field. Comput Electron Agric. 2020;175:105542. doi: 10.1016/j.compag.2020.105542.

eISSN: 2959-6246
pISSN: 2959-6238
35. Borhani Y, Khoramdel J, Najafi E. A deep learning based approach for automated plant disease classification using vision transformer. Sci Rep. 2022;12:11554. doi: 10.1038/s41598-022-15163-0.

36. Sambana B, Nnadi HS, Wajid MA, Fidelia NO, Camacho-Zuñiga C, Ajuzie HD, et al. An efficient plant disease detection using transfer learning approach. Sci Rep. 2025;15:19082. doi: 10.1038/s41598-025-02271-w.

37. Jackulin C, Murugavalli S. A comprehensive review on detection of plant disease using machine learning and deep learning approaches. Meas Sens. 2022;24:100441. doi: 10.1016/j.measen.2022.100441.

38. Arsenovic M, Karanovic M, Sladojevic S, Anderla A, Stefanovic D. Solving Current Limitations of Deep Learning Based Approaches for Plant Disease Detection. Plants. 2019;8:163. doi: 10.3390/plants8060163.

39. Too EC, Yujian L, Njuki S, Yingchun L. A comparative study of fine-tuning deep learning models for plant disease identification. Comput Electron Agric. 2019;161:272–279. doi: 10.1016/j.compag.2018.03.032.

40. Lee SH, Goëau H, Bonnet P, Joly A. New perspectives on plant disease characterization based on deep learning. Comput Electron Agric. 2020;170:105220. doi: 10.1016/j.compag.2020.105220.

41. Kabala DM, Hafiane A, Bobelin L, Canals R. Image-based crop disease detection with federated learning. Sci Rep. 2023;13:19220. doi: 10.1038/s41598-023-46218-5.

42. Hari P, Singh MP, Singh AK. An improved federated deep learning for plant leaf disease detection. Multimed Tools Appl. 2024;83:83471–83491. doi: 10.1007/s11042-024-18867-9.

43. Merenda M, Porcaro C, Iero D. Edge Machine Learning for AI-Enabled IoT Devices: A Review. Sensors. 2020;20:2533. doi: 10.3390/s20092533.

44. Ashurov AY, et al. Enhancing plant disease detection through deep learning: a Depthwise CNN with squeeze and excitation integration and residual skip connections. Front Plant Sci. 2025;15:1505857. doi: 10.3389/fpls.2024.1505857.

45. Anwar S, Soomro SR, Baloch SK, Patoli AA, Kolachi AR. Performance Analysis of Deep Transfer Learning Models for the Automated Detection of Cotton Plant Diseases. Eng Technol Appl Sci Res. 2023;13:11561–11567. doi: 10.48084/etasr.6187.

46. Rajasekar V, Venu K, Jena SR, Varthini RJ, Ishwarya S. Detection of Cotton Plant Diseases Using Deep Transfer Learning. J Mobile Multimed. 2022;18:307–324. doi: 10.13052/jmm1550-4646.1828.

47. Ryo M, Schiller J, Stiller S, et al. Deep learning for sustainable agriculture needs ecology and human involvement. J Sustain Agric Environ. 2023;2:40–44. doi: 10.1002/sae2.12036.

48. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. Int J Comput Vis. 2020;128:336–359. doi: 10.1007/s11263-019-01228-7.

49. Lan Y, Guo Y, Chen Q, Lin S, Chen Y, Deng X. Visual question answering model for fruit tree disease decision-making based on multimodal deep learning. Front Plant Sci. 2023;13:1064399. doi: 10.3389/fpls.2022.1064399.

50. Lu Y, et al. Application of Multimodal Transformer Model in Intelligent Agricultural Disease Detection and Question-Answering Systems. Plants. 2024;13:972. doi: 10.3390/plants13070972.