

Field-Deployable Cassava Disease Detection on Low-End Android Devices

K. Boateng, F. Sow, A. Mensah, C. Okoye · GCT Exchange, Accra Hub · CVPR Workshop 2026

Abstract

Cassava is a staple crop for over 500 million people across sub-Saharan Africa. Three diseases — Cassava Mosaic Disease (CMD), Cassava Brown Streak Disease (CBSD), and Cassava Bacterial Blight (CBB) — together account for an estimated 30% of yield loss. We present a 14MB MobileNetV3 classifier trained on CassavaField-90k that achieves 91.7% top-1 accuracy on the held-out evaluation set and runs at 38ms per inference on a \$40 Android device. The model and dataset are released openly. We report on a 14-farmer pilot in the Ashanti region.

1. Background

Existing cassava classifiers are accurate but heavy: most published baselines exceed 100MB and require GPU acceleration that is unavailable in the field. Farmers in the regions where the disease burden is highest typically carry low-end phones with intermittent connectivity. Any model that requires an upload step is, in practice, not used.

2. Dataset

CassavaField-90k contains 89,412 labeled leaf images collected across 12 farms in Ghana, Côte d'Ivoire, and Nigeria between 2023 and 2025. Each image is double-annotated by trained agronomists; inter-annotator agreement is 0.87 Cohen's kappa. Images include realistic field conditions: glare, partial occlusion, varying soil backgrounds, and three growth stages.

3. Model

We use MobileNetV3-Small (alpha=0.75) with a 5-class head (healthy, CMD, CBSD, CBB, mixed). Training: 200 epochs, AdamW, cosine schedule, mixup + random erasing. INT8 quantization via TFLite post-training quantization with a 1k-image calibration set.

4. Results

Model	Acc ↑	Size	Latency
EfficientNet-B3 baseline	93.2%	47MB	412ms
MobileNetV3-S FP32	92.1%	21MB	94ms
MobileNetV3-S INT8 (ours)	91.7%	14MB	38ms

5. Field Pilot

Fourteen farmers across two cooperatives in the Ashanti region used the app for one growing season (June–November 2025). Of 1,847 inferences, 71% led to a documented intervention (rogueing, replanting, or extension officer consult). Self-reported confidence in disease identification rose from 4.1/10 to 7.6/10. Failure modes were dominated by mixed-disease presentations and poor lighting at dawn.

6. Limitations

The model is trained on West and Central African cultivars and we do not claim transfer to East African varieties without re-evaluation. The model is a screening tool, not a diagnostic; all UI surfaces are designed to suggest, not assert.

7. Release

Model: huggingface.co/gctexchange/cassava-mobilenet-v3 (Apache-2.0).

Dataset: huggingface.co/datasets/gctexchange/cassavafield-90k (CC-BY 4.0).

Android app source: github.com/gctexchange/cassava-field (Apache-2.0).

References

- [1] Howard et al. 2019. Searching for MobileNetV3.
- [2] Mwebaze et al. 2019. iCassava 2019 Fine-Grained Visual Categorization Challenge.
- [3] FAO 2022. Cassava production statistics, sub-Saharan Africa.