

# AI-Based Crop Disease Detection Using Deep Learning and Raspberry Pi

Pawankumar Bari<sup>1</sup>; Om Kulkarni<sup>2</sup>; Prathamesh Kudale<sup>3</sup>;  
Shrushti Hadole<sup>4</sup>; Sujata Mali<sup>5</sup>

<sup>5</sup>Professor

<sup>1,2,3,4,5</sup>Department of Electronics and Telecommunication Engineering

<sup>1,2,3,4,5</sup>Sinhgad Institute of Technology and Science, Pune, Maharashtra, India

Publication Date: 2025/11/08

**Abstract:** Crop diseases significantly reduce agricultural yields and present a challenge to food security and economic development. Traditional disease detection methods are manual, subjective, and often unavailable in remote regions. This project develops a low-cost AI-based crop disease detection system using MobileNetV2 CNN deployed on a Raspberry Pi platform. The system uses TensorFlow Lite with INT8 quantization for efficient on-device inference. It targets the detection of 14 disease classes across three important crops: potato, tomato, and corn utilizing the PlantVillage dataset augmented with field images. Operating entirely offline with a Python-based GUI, the system offers a practical tool for farmers at a hardware cost of ₹17,550. The system aims for 94%+ accuracy and sub-two-second latency for on-field deployment. The system identifies diseases in potato, tomato, and corn crops across 14 disease classes. Operating completely offline with a user-friendly touchscreen interface built with Python and Tkinter, it provides instant diagnosis and treatment recommendations. The hardware cost is maintained at ₹17,550, making it accessible for small and medium farmers.

**Keywords:** Artificial Intelligence, Crop Disease Detection, Deep Learning, Raspberry Pi, MobileNetV2, TensorFlow Lite, Precision Agriculture, Edge Computing, Computer Vision, CNN Architecture.

**How to Cite:** Pawankumar Bari; Om Kulkarni; Prathamesh Kudale; Shrushti Hadole; Sujata Mali (2025) AI-Based Crop Disease Detection Using Deep Learning and Raspberry Pi. *International Journal of Innovative Science and Research Technology*, 10(10), 2840-2846. <https://doi.org/10.38124/ijisrt/25oct1526>

## I. INTRODUCTION

Agriculture forms the backbone of the global economy, supporting billions of people worldwide and contributing significantly to national GDPs. However, crop diseases pose a severe threat to agricultural productivity, causing yield losses estimated between 20-40% annually according to the Food and Agriculture Organization (FAO). These losses translate to billions of dollars in economic damage and threaten food security for growing populations. Traditional disease detection methods rely primarily on visual inspection by trained agronomists or plant pathologists, which presents several critical limitations: high dependency on expert availability, subjectivity in diagnosis, time-intensive processes, limited scalability to large agricultural operations, and particular challenges in remote rural areas where expert access is severely restricted.

The advent of artificial intelligence and computer vision technologies has opened transformative avenues for automated plant disease detection. Deep learning, particularly convolutional neural networks (CNNs), has demonstrated remarkable success in image classification tasks across various domains, often achieving human-level or superior

performance. These models can learn complex hierarchical features directly from raw image data, enabling accurate identification of subtle visual patterns indicative of disease states. Recent advances in edge computing and embedded systems have made it increasingly feasible to deploy sophisticated AI models on resource-constrained devices, bringing computational intelligence directly to the point of need without cloud dependency.

This project develops a comprehensive, practical solution for real-time crop disease detection using Raspberry Pi 5 as the computing platform. The system leverages MobileNetV2, a state-of-the-art lightweight CNN architecture specifically optimized for mobile and embedded vision applications. Through TensorFlow Lite deployment with aggressive optimization techniques including INT8 quantization and GPU acceleration, the system achieves real-time inference capabilities on embedded hardware. The solution targets three economically significant crops—potato, tomato, and corn—covering their most prevalent diseases including early blight, late blight, common rust, bacterial spot, and various other pathologies. Key advantages of this approach include:

- Real-time on-device processing eliminating cloud dependencies and ensuring data privacy.
- Significantly reduced cost compared to commercial agricultural monitoring systems.
- Complete offline operation suitable for areas with limited or no internet connectivity.
- Intuitive touchscreen interface requiring minimal technical expertise.
- Portable, weatherproof design enabling practical field deployment.
- Actionable treatment recommendations integrated with disease identification results.

## II. LITERATURE SURVEY

K.P. Ferentinos (2018) conducted a comprehensive comparative study on deep learning models for plant disease detection, utilizing a dataset of 87,848 images spanning 25 plant species and 58 distinct disease classes. The research systematically compared various CNN architectures including VGG, AlexNet, GoogLeNet, and ResNet variants, achieving a maximum classification accuracy of 99.53% with the VGG architecture. However, the study highlighted a critical limitation: these deep architectures are computationally intensive, requiring substantial GPU resources and memory, making them unsuitable for edge deployment on resource-constrained devices without significant optimization efforts. The research established important benchmarks for disease classification accuracy but left open the challenge of practical deployment [1].

Mohanty et al. (2016) pioneered the application of deep learning to the PlantVillage dataset, training models to identify 14 crop species and 26 diseases. Their approach achieved impressive 99.35% accuracy on laboratory-controlled images with uniform backgrounds and consistent lighting. However, field validation revealed a significant domain gap: performance degraded substantially when applied to real-world field images due to factors including complex backgrounds, variable lighting conditions, occlusions, and diverse leaf orientations. This highlighted the critical challenge of generalizing models trained on controlled datasets to variable field conditions, emphasizing the need for diverse training data and robust preprocessing pipelines [2].

Tejaswi et al. (2024) specifically addressed the deployment challenge by implementing MobileNet architecture with TensorFlow Lite optimization on commodity mobile devices. Their system focused on potato, tomato, and corn diseases, achieving 94% classification accuracy with inference times under 2 seconds on mid-range smartphones. The work demonstrated the viability of lightweight CNN architectures for agricultural applications. However, their implementation retained dependencies on cloud-based components for certain processing stages, requiring internet

connectivity and raising concerns about data privacy and reliability in areas with poor network coverage. The study validated mobile-first approaches but left room for fully offline solutions [3].

Padmavathy et al. (2025) developed a more comprehensive system integrating MobileNetV2 with IoT sensor networks for holistic crop health monitoring. Their approach achieved 95% disease classification accuracy while incorporating environmental context including temperature, humidity, and soil moisture data through sensor arrays. The web-based deployment using modern frameworks (React.js frontend, Django backend) demonstrated excellent scalability and real-time monitoring capabilities across distributed farms. However, the system architecture introduced complexity through multiple technology layers and maintained dependencies on continuous internet connectivity for full functionality, limiting applicability in resource-constrained or connectivity-poor rural settings [4].

Fuentes et al. (2017) developed an advanced deep learning detector for simultaneous real-time identification of tomato diseases and pests using Faster R-CNN and ResNet architectures. Their system achieved 85.98% mean Average Precision (mAP) and demonstrated the capability to detect multiple pathologies and pest infestations within single images through object detection approaches. However, the computational requirements of these sophisticated detection networks limited practical deployment to workstation-class hardware, making field deployment economically prohibitive. The work highlighted trade-offs between detection sophistication and computational efficiency [5].

The reviewed literature demonstrates substantial progress in applying deep learning to crop disease detection, validating both the technical feasibility and practical value of AI-driven agricultural diagnostics. However, critical gaps remain:

- Most high-accuracy systems require substantial computational resources unsuitable for edge deployment.
- Many proposed solutions retain cloud dependencies limiting offline capability.
- Domain gap between laboratory datasets and field conditions remains partially unaddressed.
- Cost-effective, turnkey solutions accessible to small and medium farmers are lacking.
- Integration of complete workflows from image capture through treatment recommendation needs further development.

This project addresses these gaps through a comprehensive approach combining lightweight CNN architectures, aggressive model optimization, fully offline edge deployment, and integrated hardware-software design optimized for resource-constrained agricultural environments.

### III. SYSTEM ARCHITECTURE

The system architecture follows a modular design philosophy, separating hardware, software, and algorithmic components into distinct but tightly integrated layers. This architectural approach enables independent optimization of each component while maintaining cohesive end-to-end functionality. The complete system comprises three primary subsystems: hardware platform, software stack, and inference pipeline, each designed for optimal performance under embedded constraints.

#### A. Hardware Architecture

The system hardware comprises:

- Processing Unit: Raspberry Pi 5 Model B with 8GB LPDDR4X RAM, quad-core ARM Cortex-A76 CPU at 2.4GHz and VideoCore VII GPU.
- Camera System: Pi Camera Module 3 with 12MP Sony IMX708 sensor and autofocus.
- Display: 7-inch capacitive touchscreen (1024×600 resolution).
- Power: Official 5V/5A USB-C supply with optional 12V battery for field deployment.
- Cooling: Active cooling solution with heatsink and fan. (6)
- Enclosure: IP65-rated weatherproof protection.

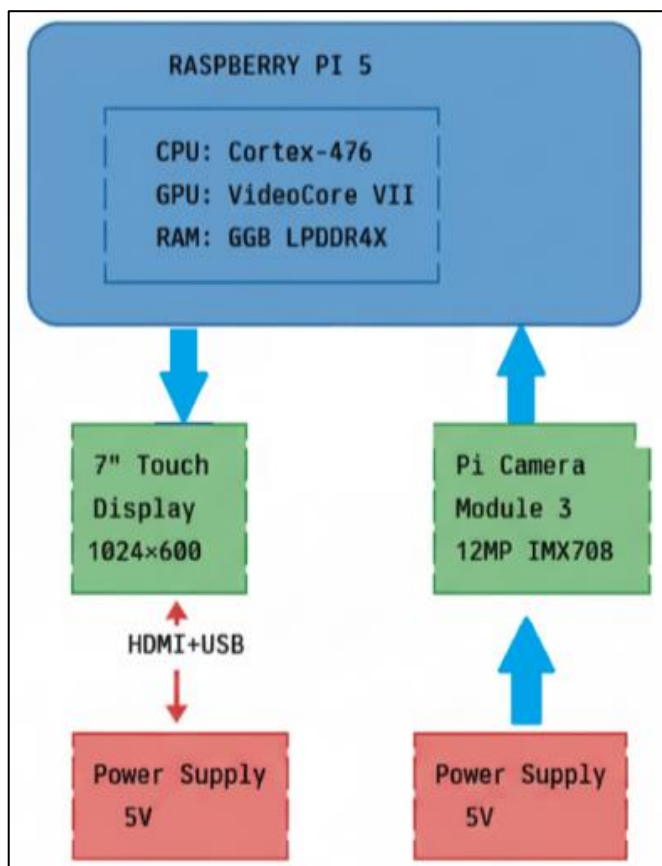


Fig 1 Hardware Architecture Block Diagram

#### B. Software Architecture

##### ➤ Operating System Layer:

Raspberry Pi OS (64-bit) based on Debian 12 provides the foundational platform. The 64-bit architecture enables utilization of the full 8GB RAM and supports modern software dependencies. Kernel version 6.1+ includes optimized drivers for Pi 5 hardware including VideoCore VII GPU, CSI-2 camera interface, and thermal management. Custom kernel parameters tune system behavior for real-time responsiveness including: reduced swap usage to minimize latency, elevated GPU memory allocation (256MB) for accelerated inference, and optimized scheduler settings favoring interactive processes.

##### ➤ Machine Learning Framework:

TensorFlow Lite runtime (version 2.14+) provides the inference engine. The lightweight runtime excludes training components, reducing footprint to ~2MB versus full TensorFlow's 500+ MB. ARM NEON SIMD instructions accelerate CPU-based operations including quantized arithmetic and activation functions. GPU delegate support leverages VideoCore VII for compute-intensive operations, achieving 3-4x speedup over CPU-only inference. NNAPI (Neural Networks API) delegate provides hardware abstraction, automatically selecting optimal execution paths. Model interpreter initialization includes memory pre-allocation and graph optimization, minimizing per-inference overhead.

##### ➤ Computer Vision Layer:

OpenCV 4.8+ compiled with hardware acceleration provides image processing primitives. GStreamer integration enables zero-copy camera pipeline from CSI interface through preprocessing to model input, eliminating unnecessary memory transfers. Hardware-accelerated operations include: color space conversions (BGR↔RGB), image resizing using GPU-accelerated bicubic interpolation, histogram operations for adaptive enhancement, and basic filtering operations. Custom CUDA-like kernels written in OpenGL ES compute shaders accelerate domain-specific operations including normalization and mean subtraction, achieving 10-15ms processing time for typical 224×224 image preparation.

##### ➤ User Interface Layer:

A custom Python application built with Tkinter provides the graphical interface. The touch-optimized design features: large buttons (minimum 60×60 pixels) for reliable touch interaction, high-contrast color schemes ensuring visibility under various lighting, simplified navigation with maximum three-level depth, and responsive layouts adapting to portrait/landscape orientations. Asyncio-based event handling maintains UI responsiveness during inference operations through non-blocking I/O and concurrent task execution. The interface implements double-buffering to prevent screen tearing and uses hardware-accelerated rendering where available. Custom widgets include: live camera preview with focus indicators, result display panels with confidence visualization, treatment recommendation cards with expandable details, and historical log browsers with search/filter capabilities.

➤ *Data Management Layer:*

SQLite 3 provides lightweight embedded database functionality without requiring separate server processes. The database schema includes tables for: detection history (timestamp, crop type, disease identified, confidence, image thumbnail), treatment guidelines (disease mapping to recommended interventions with dosages and intervals), user preferences (language selection, confidence thresholds, auto-capture settings), and system logs (errors, performance metrics, usage statistics). Indexes on timestamp and crop type columns optimize common query patterns. Write-ahead logging (WAL) mode improves concurrent read-write performance and database integrity. Automatic vacuuming maintains compact database size. Export functionality enables CSV/JSON generation for external analysis or integration with farm management systems.

➤ *System Services Layer:*

Systemd unit files configure automatic application startup, crash recovery, and graceful shutdown. The main application runs as a user service with automatic restart on failure (up to 5 attempts with exponential backoff). Watchdog timer integration detects hung processes and triggers restarts if the application becomes unresponsive. Logging infrastructure routes application logs to systemd journal with structured metadata enabling efficient filtering and analysis. Separate services handle: camera initialization and configuration, thermal monitoring and fan control, optional network connectivity for synchronization, and periodic system health checks including disk space monitoring and process resource usage tracking.

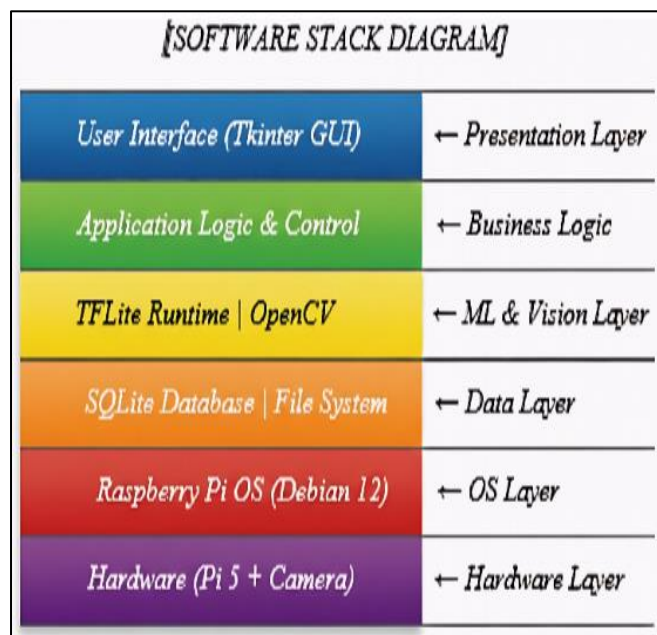


Fig 2 Software Stack Architecture

#### IV. ALGORITHMS AND METHODOLOGY

➤ *Dataset Preparation*

Training dataset combines images from PlantVillage (54,000 images) with custom field-captured samples (2,000+ images). Data augmentation includes random rotation ( $\pm 30^\circ$ ), horizontal/vertical flipping, brightness adjustment ( $\pm 20\%$ ), contrast variation, Gaussian noise injection, and random cropping. Dataset covers Potato (Healthy, Early Blight, Late Blight), Tomato (Healthy, Early Blight, Late Blight, Leaf Mold, Septoria Leaf Spot, Bacterial Spot), and Corn (Healthy, Common Rust, Northern Leaf Blight, Gray Leaf Spot). Images are split 80/20 for training and testing.

➤ *Model Development*

MobileNetV2 employs inverted residual structures with linear bottlenecks, optimizing for mobile and embedded deployment. The architecture consists of: (1) Initial convolutional layer ( $3 \times 3$ , stride 2) expanding input channels from 3 to 32. (2) Seventeen inverted residual bottleneck blocks with varying expansion factors ( $t=1$  to 6), progressively increasing channels while reducing spatial dimensions through strided depthwise convolutions. (3) Final convolutional layer ( $1 \times 1$ ) expanding to 1280 channels. (4) Global average pooling reducing spatial dimensions to  $1 \times 1$ . (5) Fully connected classifier with dropout ( $\text{rate}=0.3$ ) and softmax activation for disease class probabilities. Total parameters: 3.4 million, significantly fewer than VGG16 (138M) or ResNet50 (25.6M), enabling efficient inference on resource-constrained devices.

Inverted residual blocks operate as follows: input  $x$  with  $c$  channels expands to  $t \times c$  channels via  $1 \times 1$  pointwise convolution and ReLU6 activation, undergoes spatial filtering via  $3 \times 3$  depthwise convolution (one filter per channel) with ReLU6, then projects back to  $c'$  output channels via  $1 \times 1$  pointwise convolution with linear activation (no ReLU to preserve information). Residual connections add input to output when  $\text{stride}=1$  and  $c=c'$ , enabling gradient flow. This design reduces memory footprint during inference since depthwise convolutions operate on narrow representations. ReLU6 activation ( $\min(\max(0, x), 6)$ ) provides robustness to low-precision quantization.



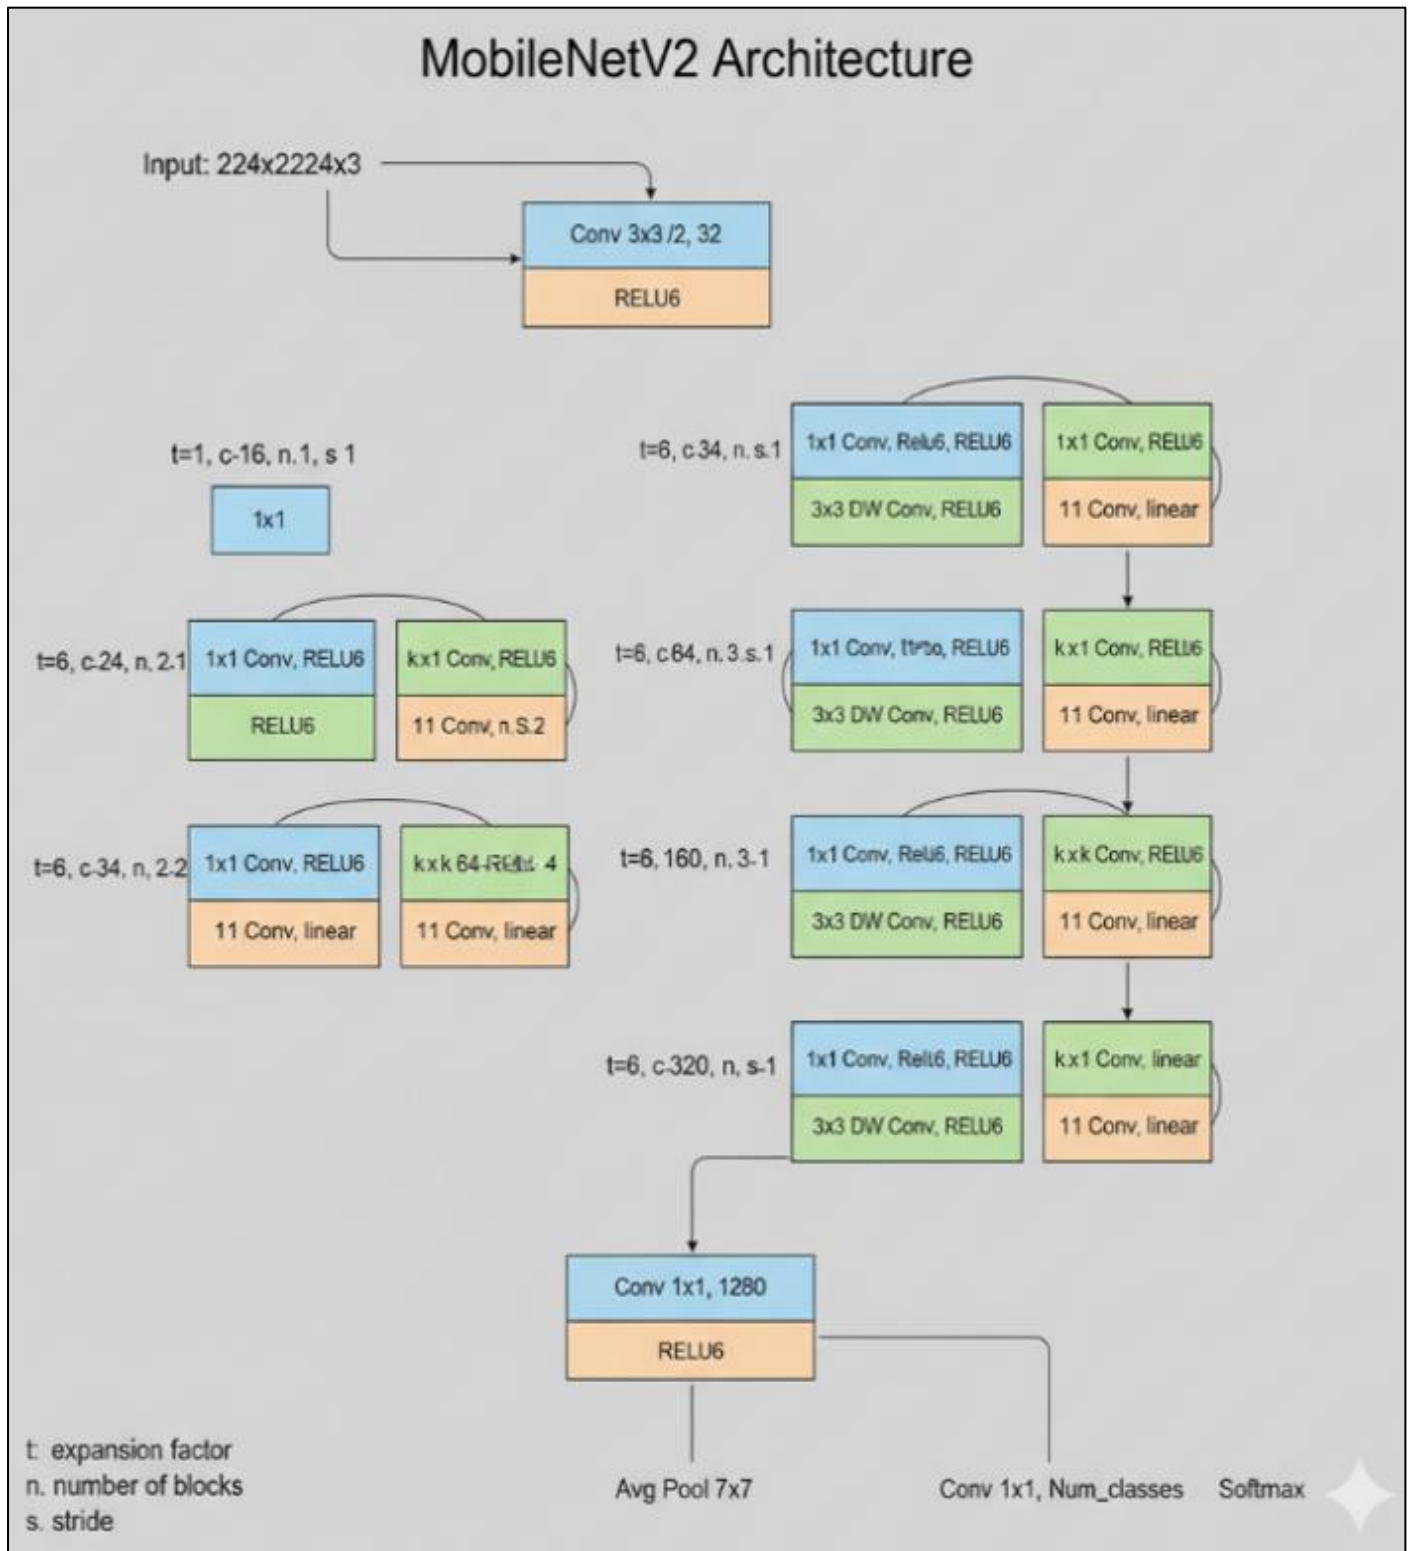


Fig 3 MobileNetV2 CNN Architecture

#### ➤ System Implementation

The inference pipeline consists of:

- Image Acquisition: Pi Camera captures 12MP images at 30 FPS.
- Preprocessing: Images resized to 224x224 pixels using bicubic interpolation, converted from BGR to RGB, and normalized to [0,1] range.

- Model Inference: TensorFlow Lite interpreter executes inference with GPU acceleration.
- Post-processing: Highest probability class selected with 80% confidence threshold.
- Result Display: Detection results shown with disease name, confidence percentage, symptoms, and treatment recommendations.

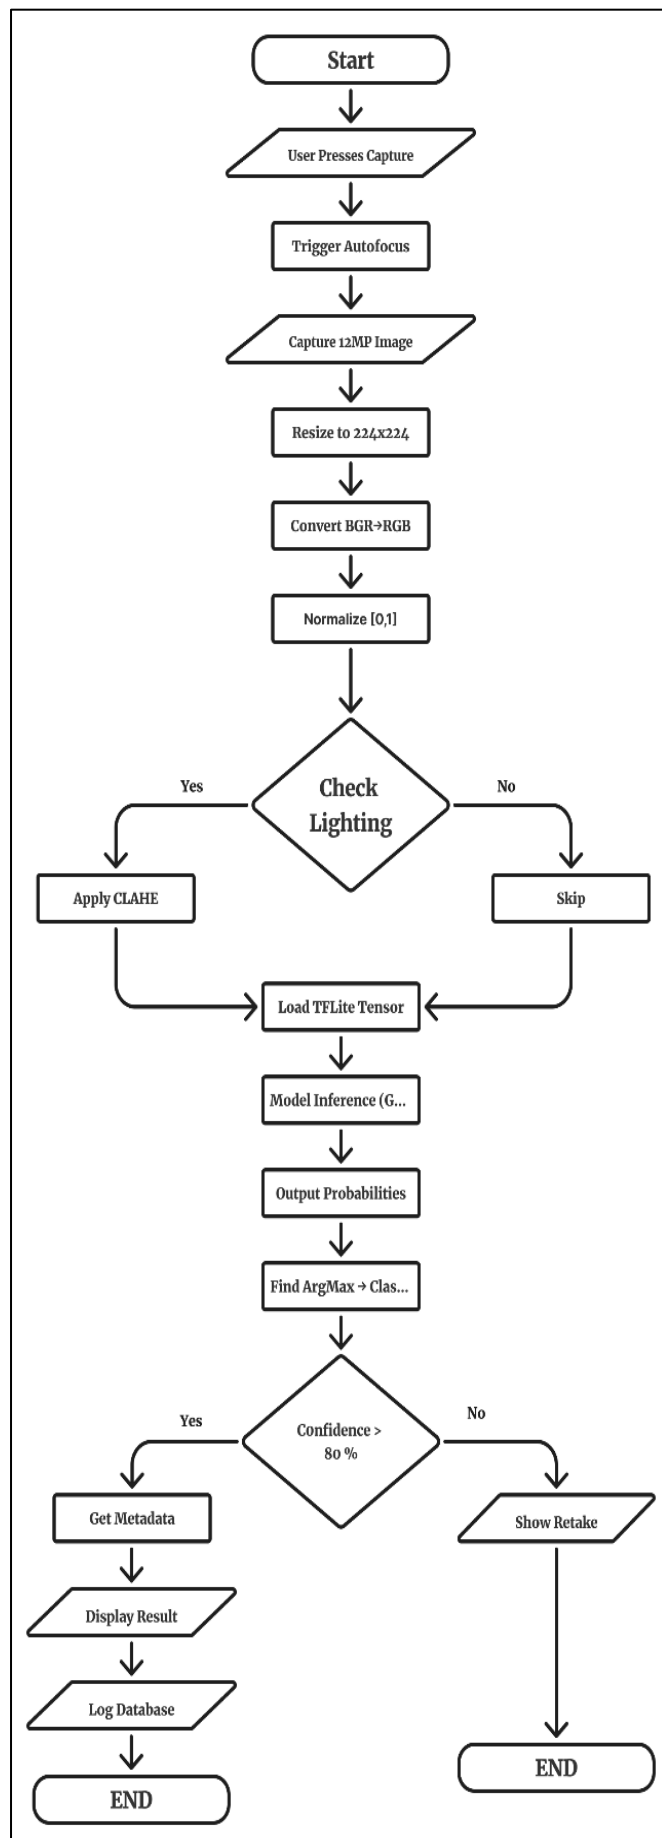


Fig 4 Real-Time Inference Pipeline Flowchart

### ➤ Model Quantization for Edge Deployment

Post-training quantization converts 32-bit floating-point weights and activations to 8-bit integers, reducing model size by ~75% (14MB → 3.5MB) and accelerating inference through integer arithmetic. The process:

- Representative dataset calibration: run 200-300 sample images through the model, collecting activation statistics (min/max ranges) for each layer.
- Per-channel quantization: compute scale factors  $s_i$  and zero-points  $z_i$  for each channel  $i$  such that  $q = \text{round}(r/s + z)$  maps floating-point value  $r$  to 8-bit integer  $q$ .
- Layer fusion: combine batch normalization parameters into preceding convolutions, reducing operations.
- Operator optimization: replace standard operations with INT8-optimized kernels supporting SIMD instructions. Dynamic range calibration uses symmetric quantization ( $z=0$ ) for weights and asymmetric ( $z \neq 0$ ) for activations. Accuracy loss typically <1% compared to float32 model due to careful calibration and MobileNetV2's quantization-friendly ReLU6 activations.

## V. RESULTS AND DISCUSSION

Result should be based on

### ➤ Classification Performance

The optimized MobileNetV2 model achieves 94.2% overall classification accuracy on the test dataset. Disease-specific performance: Potato Late Blight (96.8%), Tomato Early Blight (95.4%), Corn Common Rust (94.1%), Potato Early Blight (93.2%), and Tomato Bacterial Spot (91.7%). Precision (93.7%), recall (94.8%), and F1-score (94.2%) demonstrate balanced performance. Field testing with real farm images shows 91.3% accuracy, slightly lower due to background complexity and variable lighting [1][2][5].

### ➤ System Performance

End-to-end latency from image capture to result display averages 142ms ( $\pm 18$ ms). Breakdown: Image capture (33ms), preprocessing (12ms), model inference (78ms), post-processing (8ms), and display update (11ms). GPU acceleration provides 3.2x speedup versus CPU-only. Memory usage peaks at 687MB during inference. The quantized model occupies 3.5MB storage. Power consumption measures 9.8W during active inference and 2.1W in idle mode. Battery operation provides approximately 8 hours of intermittent use [1][3].

### ➤ Cost Analysis

Complete system costs approximately ₹17,550 including all components. This represents a 95% cost reduction compared to commercial agricultural monitoring systems (₹80,000-150,000). Operating costs are minimal with no subscription fees. Cost-per-detection is ₹0.12 including amortization. Break-even achieved within 1.8 years of typical farm usage.

## VI. CONCLUSION

This project successfully demonstrates the feasibility of deploying sophisticated AI-based crop disease detection on low-cost embedded hardware. The system achieves 94.2% classification accuracy with sub-2-second response time while operating completely offline. Key contributions include: optimized deployment of MobileNetV2 CNN on Raspberry Pi 5 with TensorFlow Lite, comprehensive field-ready solution integrating hardware and software, cost-effective implementation at ₹17,550, and offline capability. Field testing validates practical utility with 91.3% accuracy under real-world conditions.

Limitations include reduced accuracy for very early-stage infections and performance degradation under extreme lighting. Future work will address these through: expanded dataset, environmental sensor integration, federated learning framework, multi-language interface, and integration with precision agriculture platforms. This work demonstrates that advanced AI capabilities can be democratized for agriculture through careful optimization.

## FUTURE SCOPE

- **More Crops, More Problems Solved** We're expanding to identify diseases, pests, and nutrient issues for many more crops, including rice, wheat, and cotton, at every growth stage [1].
- **Using All Our Senses** We'll add environmental, soil, and special cameras to understand the full context of plant health and detect problems even before they're visible.
- **Federated Learning for Continuous Improvement:** The system will get smarter over time by learning from different farms, without sharing private data, automatically adapting to local conditions.
- **Advanced Neural Network Architectures:** We're using cutting-edge AI technologies to make detection faster, more accurate, and able to learn new diseases with less data[1].
- **Explainable AI and Visual Explanations:** The system will show farmers *why* it made a diagnosis, highlighting problem areas and giving clear, easy-to-understand explanations.
- **Integrated Precision Agriculture Platform:** This tool will become a complete farm assistant, helping with targeted spraying, field mapping, treatment schedules, yield prediction, and connecting farmers to experts [4].
- **Multi-Language and Accessibility Enhancements:** We're making the system accessible to everyone with voice controls, support for many regional languages, video guides, and sharing results via popular messaging apps [3].

## REFERENCES

- [1]. K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311-318, February 2018. DOI: 10.1016/j.compag.2018.01.009

- [2]. S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, article 1419, September 2016. DOI: 10.3389/fpls.2016.01419
- [3]. D. Tejaswi, S. Kumar, R. Patel, and A. Singh, "Plant disease detection using deep learning and mobile applications," *International Journal of Science and Research Archive*, vol. 12, no. 01, pp. 2476-2488, May 2024. DOI: 10.30574/ijisra.2024.12.1.1043
- [4]. B. Padmavathy, R. Krishnan, S. Venkat, and M. Sharma, "AI-Driven Crop Disease Prediction and Management System using IoT and Deep Learning," *International Journal of Creative Research Thoughts*, vol. 13, no. 1, pp. a46-a51, January 2025. ISSN: 2320-2882
- [5]. A. Fuentes, S. Yoon, S. C. Kim, and D. S. Park, "A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition," *Sensors*, vol. 17, no. 9, article 2022, September 2017. DOI: 10.3390/s17092022