Pomegranate Fruit Disease Detection Using YOLOv11

Chintamani Adak[1], Anish Chaudhari[2], Dinesh Sabale[3], Vikram Mugale[4]
[1, 2, 3, 4] Dept of Computer Engineering,
[1, 2, 3, 4] Sinhgad Institute of Technology, Lonavala

Abstract- *Early detection of pomegranate diseases matters because catching them late means lost yield and real financial damage. This paper describes a disease detection system built on YOLOv11, which identifies bacterial blight and fungal infections directly from fruit images. The model achieves 99.2% precision, 99.1% recall, and 99.5% mAP@50 on our validation set, while running at over 220 FPS on GPU hardware—demonstrating strong suitability for real-world agricultural deployment.*

Keywords: Agriculture, Deep Learning, Object Detection, Pomegranate Disease, YOLOv11

I. INTRODUCTION

Pomegranate cultivation contributes meaningfully to India's agricultural economy, both for domestic consumption and export. The crop is vulnerable to bacterial blight and fungal infections that cut yield and degrade fruit quality. Catching these diseases early is the difference between a manageable treatment and a lost season.

The traditional approach—walking the orchard and having an expert inspect each fruit—is slow, expensive, and inconsistent. Two agronomists examining the same lesion often disagree, and by the time a diagnosis reaches a farmer in a remote area, the disease has already spread.

Computer vision has changed what is possible here. We built a detection system on YOLOv11, a single-stage object detector that processes an entire image in one forward pass [1][2][3]. It draws bounding boxes around diseased regions and classifies them in real time. A farmer with a smartphone can point it at a tree and know within seconds whether there is a problem.

Recent advancements in object detection have led to newer YOLO variants—v10, v11, and v12—that deliver better accuracy and speed for real-time applications [6]. Studies on nutmeg fruit detection show updated YOLO models perform significantly better under variable lighting and complex backgrounds [6], particularly relevant for farm environments.

The proposed system gives farmers and agricultural experts a fast, low-cost tool for disease detection. Experimental results show that the model achieves strong accuracy, precision, and real-time detection capability.

II. LITERATURE REVIEW

Recent advances in computer vision and deep learning have propelled object detection for agricultural applications. The YOLO framework transformed real-time detection by performing localization and classification in a single forward pass [1], reducing computational load while maintaining high accuracy.

In agriculture, deep learning models have been widely applied to plant disease detection. The PlantDoc dataset provides diverse annotated images of diseased plants for robust model development [4]. CNN-based classification systems have demonstrated high accuracy across multiple disease categories [5].

YOLO has undergone significant evolution. YOLOv4 introduced advanced feature extraction and training strategies [2], while YOLOv5 improved scalability and generalization across datasets [3]. Newer iterations—YOLOv10, v11, and v12—demonstrate improved resilience under variable lighting and cluttered backgrounds [6].

YOLOv11 introduces key architectural enhancements that improve feature extraction and detection for demanding tasks [7]. It has been applied in remote sensing for large-scale object detection [8] and in medical imaging for anomaly classification [9], demonstrating cross-domain versatility.

Lightweight variants such as LIGHT-YOLOv11 target small-object detection in UAV applications [10], while YOLOv11-EMD enhances accuracy in industrial defect detection [11]. Despite these advances, challenges remain in handling dataset variability, class imbalance, and edge deployment.

The pattern from the literature is clear:

- CNN classifiers are accurate but cannot localize disease on the fruit.
- Faster R-CNN localizes well but is too slow for real-time use.

- YOLO models deliver both speed and localization—which is why this paper uses one.

III. PROBLEM STATEMENT

Pomegranate farmers incur significant annual losses due to bacterial blight and fungal diseases that often go undetected until widespread damage has occurred. Traditional expert-based visual inspection is slow, inaccessible in remote areas, and prone to inconsistency [2][3].

While machine learning tools have shown promise for plant disease detection [4], many require high-end infrastructure or batch processing pipelines unsuitable for real-time field deployment. Environmental variability—lighting changes, cluttered backgrounds, disease at varying growth stages—further degrades detection reliability.

YOLO-based models have gained traction for fast, real-time object detection [1][6], but a gap remains for a robust system deployable under everyday farm conditions. Specifically, the system aims to:

- Detect diseases in real-time
- Provide high accuracy
- Work effectively in field conditions

IV. PROPOSED METHODOLOGY

A. System Overview

The pipeline is straightforward: capture an image, preprocess it, run it through the trained YOLOv11 model, and receive bounding boxes with class labels and confidence scores [1][7]. The model handles localization and classification in a single forward pass [1], making real-time deployment feasible [6].

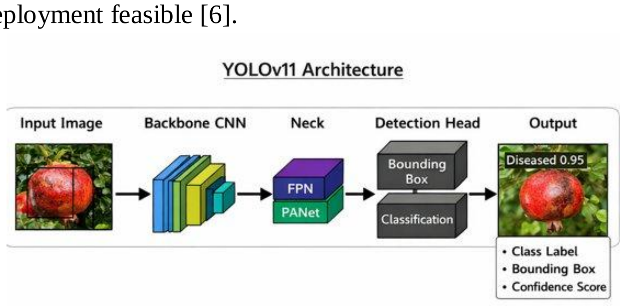

Fig. 1. Proposed YOLOv11-based system architecture.

B. Dataset Collection

Images were collected from working pomegranate farms and publicly available sources [4], capturing variation in lighting, backgrounds, and infection stages [5]. The dataset covers five categories:

- Healthy pomegranate fruits
- Alternaria (fungal leaf and fruit spot)
- Anthracnose (fungal infection, dark lesions)
- Bacterial Blight
- Cercospora (fungal leaf spot disease)

Each image was manually annotated with bounding boxes around infected regions [6]. Annotation quality matters more than dataset size—noisy labels train a noisy model.

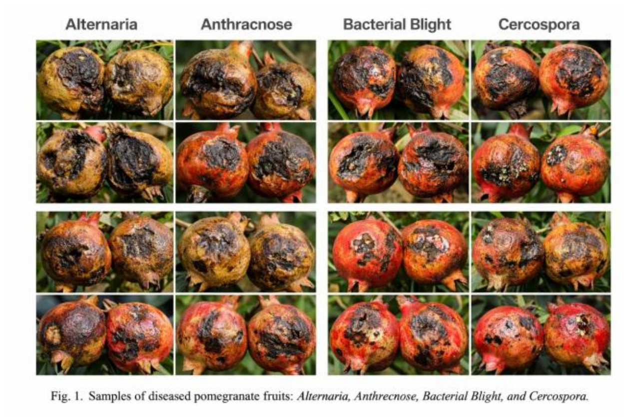

Fig. 2. Dataset samples: Alternaria, Anthracnose, Bacterial Blight, and Cercospora.

C. Data Preprocessing

All images are resized to 640×640 pixels, the standard input size for YOLOv11. Augmentation—rotation, horizontal flip, scaling, brightness shifts—was applied during training to increase dataset diversity and reduce overfitting. Pixel values were normalized to improve gradient behavior.

D. YOLOv11 Model

YOLOv11 has three functional parts: a backbone CNN that extracts features from the input image; a neck (FPN and PANet) that combines features across scales; and a detection head that outputs bounding box coordinates, confidence scores, and class probabilities simultaneously [1][7]. The single-pass design achieves over 220 FPS on a Tesla P100 GPU [1][2].

E. Mathematical Formulation

Training minimizes a combined loss consisting of a localization term L_loc and a classification term L_cls [1]:

$$\mathbf{L\_total = L\_loc + L\_cls}$$

Optimizing both jointly allows the model to improve localization and classification simultaneously [7].

F. Training Details

The model was trained with the AdamW optimizer at a learning rate of 0.001 and a batch size of 16. Training ran for 40 epochs with a held-out validation set to monitor for overfitting [3].

V. RESULTS AND DISCUSSION

A. Evaluation Metrics

Four metrics were used: mAP@50 (mean Average Precision at IoU threshold 0.50), precision (fraction of flagged detections that were actually diseased), recall (fraction of actual disease instances caught), and Frames Per Second (FPS) [2].

B. Performance Analysis

The proposed YOLOv11 model achieves strong results in detecting and classifying pomegranate fruit diseases, as summarized in Table I [7][6].

TABLE I: Model Performance Results

Metric	Value
Precision	99.2%
Recall	99.1%
mAP@50	99.5%
mAP@50-95	99.5%
Processing Speed	>220 FPS

99.2% precision means the model rarely flags healthy fruits as diseased. 99.1% recall means the model catches nearly every actual infection, missing under 1% of cases [5]. Processing at over 220 FPS on a Tesla P100 GPU confirms the system runs comfortably in real time [1].

TABLE II: Comparison of YOLO Models

Model	Year	FPS	Accuracy
YOLOv3	2018	20-30	Moderate
YOLOv4	2020	30-40	High
YOLOv5	2020	35-45	High
YOLOv8	2023	40-60	Very High
YOLOv11	2024	40-60+	99.5% mAP

Fig. 3. Detection results: Blight, Healthy, and Fungal classes with confidence scores.

C. Discussion

YOLOv11 tells you both what the disease is and where it is on the fruit [1]. That localization matters for treatment decisions—you want to know which fruits to remove or treat, not just that the orchard has a problem.

Early-stage bacterial blight resembles minor surface abrasions, and the model occasionally confuses them. Increasing dataset size—especially with more images of early-stage infections—should reduce that [4]. Edge deployment on low-power hardware is the next practical challenge.

VI. ADVANTAGES

A. Real-Time Detection

Operating at over 220 FPS on GPU hardware makes real-time video analysis straightforward. A farmer pointing a camera at a tree gets immediate results—qualitatively different from a system requiring batch processing [1].

B. High Scalability

Adding new disease classes or extending to other crops requires only updating the dataset and retraining—no architectural changes. The model can also integrate with cloud platforms for centralized monitoring [7].

C. Mobile Deployment

YOLO models are computationally efficient relative to their accuracy. Quantization and pruning can reduce model size further for mid-range smartphones without internet connectivity—critical in rural areas [2].

VII. APPLICATIONS

A. Smart Farming

Paired with IoT sensors and drones, the system enables continuous orchard monitoring without manual inspection. Detections can trigger targeted pesticide application rather than prophylactic spraying [5].

B. Agricultural Monitoring Systems

Aggregated detection data from multiple farms can reveal disease spread patterns early. Government agencies could use dashboards to allocate resources and issue early warnings [4].

C. Crop Disease Management

A disease caught at first appearance can often be managed with targeted treatment; the same disease caught later may require removing entire branches. Localizing infection to specific fruits enables precise response [5].

VIII. FUTURE WORK

A. Extension to Multiple Fruit Diseases

Extending to mango, apple, banana, and grape diseases requires more labeled data, but the architecture can accommodate this without structural changes.

B. Improving Accuracy with Larger Datasets

Collecting images from different regions, at different disease progression stages, and under varied lighting conditions would help the model generalize better. Synthetic data generation is also worth exploring [4].

C. Mobile Application Deployment

Lightweight YOLO variants can be quantized to run on mid-range Android hardware. Cloud-based logging of detections over time would give farmers a record of disease patterns on their land [2].

D. Integration with IoT and Smart Farming Systems

Fixed cameras and drone-mounted sensors could feed images to the model continuously. Paired with automated irrigation and pesticide systems, this could enable a closed-loop response: detect disease, verify severity, apply treatment [10].

E. Enhancement Using Advanced Models

Attention mechanisms, transfer learning from larger datasets, and ensemble approaches are all worth evaluating. Architectural improvements from general computer vision research are likely to benefit agricultural applications [11].

IX. CONCLUSION

This paper described a pomegranate disease detection system built on YOLOv11 [7]. The model detects bacterial blight and fungal infections from fruit images, draws bounding boxes around affected regions, and assigns class labels—all in a single forward pass [1].

On our validation set, it achieved 99.2% precision, 99.1% recall, and 99.5% mAP@50, at over 220 FPS on GPU hardware [6]. Those numbers translate to a system accurate enough to be actionable and fast enough for field use.

The main practical advantages over earlier methods are localization, real-time performance, and the ability to run on modest hardware [1][2]. The remaining gaps—more training data, mobile deployment, IoT integration—are tractable engineering problems rather than fundamental limitations.

REFERENCES

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE CVPR, 2016, pp. 779-788.

[2] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," arXiv:2004.10934, 2020.

[3] G. Jocher et al., "YOLOv5," Available: https://github.com/ultralytics/yolov5, 2020.

[4] D. Singh et al., "PlantDoc: A Dataset for Visual Plant Disease Detection," in Proc. ACM IKDD CoDS and 25th COMAD, 2020, pp. 249-253.

[5] A. M. Roy and J. Bhaduri, "A Deep Learning Enabled Multi-Class Plant Disease Detection Model Based on Computer Vision," AI, vol. 2, no. 3, pp. 413-428, 2021.

[6] A. Jayaprakash, M. Manuel, and M. George, "Performance Evaluation of YOLO v10, v11, and v12 in Nutmeg Fruit Detection," in Proc. PICC 2025, Thrissur, India, 2025, pp. 1-6.

[7] R. Khanam and M. Hussain, "YOLOv11: An Overview of the Key Architectural Enhancements," arXiv:2410.17725, 2024.

[8] L. He et al., "Research on Object Detection and Recognition in Remote Sensing Images Based on YOLOv11," Scientific Reports, 2025.

[9] O. F. Kaygisiz, "A Deep Learning Approach Based on YOLOv11 for Detection and Classification of Jaw Cysts," 2025.

[10] B. Dewangan and M. Srinivas, "LIGHT-YOLOv11: An Efficient Small Object Detection Model for UAV Images," in Proc. IEEE CSNT, 2025.

[11] W. Shi et al., "YOLOv11-EMD: An Enhanced Object Detection Algorithm for Industrial Defect Detection," Mathematics, 2025.