International Journal of Applied Engineering Research ISSN 0973-4562 Volume 20, Number 2 (2025) pp. 130-135
© Research India Publications. http://www.ripublication.com

Tomato Leaf Disease Detection Using YOLOv8: Design, Deployment, and Evaluation

Hien Nguyen Thi Thu, Thuy Le Thanh
Faculty of Telecommunications 1, Posts and Telecommunications Institute of Technology, Km10, Nguyen Trai Street, Ha Dong District, Hanoi, Vietnam
Emails: hiennt@ptit.edu.vn; thuylt@ptit.edu.vn

Abstract

Tomato leaf diseases significantly affect crop yield and quality in agricultural production. Early detection and accurate diagnosis are key factors in timely treatment and damage mitigation. This study aims to develop a real-time detection system for tomato leaf diseases based on the YOLOv8 model, enabling the identification and localization of multiple disease types. The research process encompasses all essential steps, including data collection, annotation, model training, performance evaluation, and real-world deployment on mobile devices. The model, trained on a custom-labeled dataset, achieves a mean Average Precision (mAP@0.5) of 89.4%, with inference times under 50 milliseconds on edge devices. Compared to previous approaches that primarily rely on image classification, this system offers spatial localization, faster inference, and suitability for real-time applications in agricultural environments. The solution is integrated into an Android-based application and evaluated under practical conditions with promising results. Future development will focus on expanding disease coverage and enhancing the model's adaptability to diverse environmental conditions.

Keywords: YOLOv8, tomato disease detection, object detection, plant pathology, precision agriculture, real-time inference.

Introduction

Tomato (Solanum lycopersicum) is one of the most widely cultivated crops worldwide, contributing significantly to both subsistence farming and commercial agriculture. However, its leaves are highly susceptible to various fungal, bacterial, and viral diseases such as Early Blight, Late Blight, and Leaf Mold. These diseases, if not detected promptly, can result in substantial yield loss and reduced fruit quality.

Traditional disease detection methods rely on manual inspection by agronomists, which is labor-intensive, error-prone, and impractical at scale. Recent advances in artificial intelligence (AI) and computer vision offer automated alternatives, especially with the success of convolutional neural networks (CNNs) in image classification. However, most prior works focus only on classifying disease presence in an image [1]–[3], without spatial localization or support for real-time use.

In this study, we address these limitations by implementing a real-time disease detection system using YOLOv8, the latest iteration of the "You Only Look Once" family of object detectors. The system is capable of detecting and localizing multiple disease types within a single image with high precision and low latency.

The key contributions of this study are as follows:

1. Dataset Construction and Annotation: We developed a custom-labeled dataset of tomato leaf images, encompassing four distinct categories—Early Blight, Late Blight, Leaf Mold, and Healthy—providing a solid foundation for disease classification and detection tasks.
2. Model Fine-tuning: We fine-tuned the YOLOv8-small architecture to accurately detect and localize disease symptoms on tomato leaves, achieving high precision and robustness across varied image conditions.
3. System Deployment: The optimized model was deployed as a RESTful web API and integrated into an Android mobile application, enabling real-time disease detection in field environments.
4. Performance Evaluation: We conducted a quantitative evaluation of the system's performance using metrics such as mean Average Precision (mAP), precision, and recall under real-world farming conditions to validate the model's accuracy and effectiveness.

Related Work and Theoretical Background

learning algorithms such as Support Vector Machines (SVMs), Decision Trees, and k-Nearest Neighbors, which relied on handcrafted features like color histograms, shape descriptors, and texture information. These methods, while effective under controlled environments, often failed in real-world conditions due to noise, variations in lighting, and occlusions.

With the rise of deep learning, convolutional neural networks (CNNs) revolutionized image classification tasks. Mohanty et al. [4] demonstrated the effectiveness of CNNs in classifying plant diseases across 38 classes with an accuracy exceeding 99% using the PlantVillage dataset. However, these methods only performed image-level classification and could not localize the affected area on the leaf, making them less practical for visual explanation and real-time applications.

To overcome this, object detection frameworks such as Faster R-CNN [5], SSD [6], and YOLO [7] were introduced. YOLO (You Only Look Once), proposed by Redmon et al. [7], was notable for its real-time performance, framing detection as a single regression problem. However, earlier YOLO versions suffered from a trade-off between accuracy and speed, especially for small objects and overlapping instances.

YOLOv3 [8] and YOLOv5 [9] improved significantly by introducing multi-scale predictions and deeper feature extractors. Yet, these versions still relied on anchor-based detection, requiring manual tuning of anchor boxes, which can be suboptimal and computationally expensive. The recently introduced YOLOv8 [10] shifts to an anchor-free paradigm and incorporates a decoupled head architecture, enhancing both detection accuracy and inference speed.

In agriculture-specific domains, Ferentinos [11] explored CNNs for plant disease detection across multiple species with an average accuracy of 93.4%. Zhang et al. [12] attempted to localize tomato leaf diseases using YOLOv5, achieving decent performance but noting limitations in precision and deployment complexity.

Our work builds on these advancements by employing YOLOv8 for both detection and localization of multiple tomato leaf diseases, addressing the limitations of prior classification-only approaches and proposing a complete pipeline suitable for mobile deployment.

Methodology

The overall methodology consists of several stages, including dataset preparation, data annotation, model configuration, training process, and performance evaluation. The goal is to develop a reliable object detection model using YOLOv8 capable of identifying and localizing tomato leaf diseases in real time.

Dataset Preparation

The dataset consists of 4,900 tomato leaf images categorized into seven classes: Bacterial Spot, Early Blight, Late Blight, Leaf Mold, Mosaic Virus, Yellow Leaf Curl Virus, and Healthy. These images were collected from public sources (e.g., PlantVillage [13]) as well as real-world greenhouse environments using smartphone cameras. The images were selected to ensure diversity in lighting conditions, shooting angles, and backgrounds.

The dataset was split into three subsets (Fig.1):

- 72% for training (3,500 images),
- 14% for validation (700 images),
- 14% for testing (700 images).

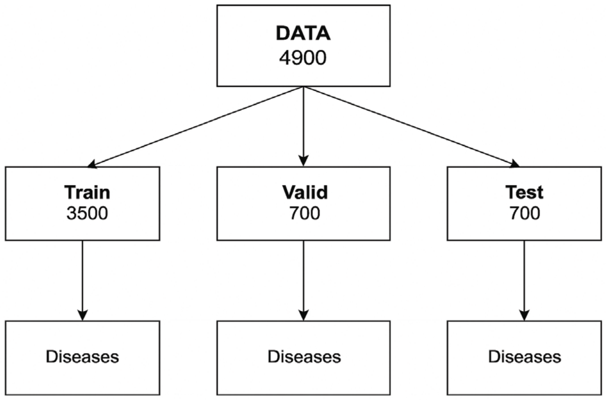

Fig. 1. The data organization diagram for the tomato leaf image dataset.

Each subset contains images organized into seven folders, corresponding to the selected disease classes. This directory structure enables efficient data management and supports class-specific learning during the training phase.

All images were standardized to a resolution of 640×640 pixels, which is compatible with the input size required by the YOLOv8-small model.

Data Annotation

Images were annotated using the open-source tool LabelImg, assigning each object instance a bounding box and class label corresponding to one of the predefined disease categories. The annotations were exported in YOLO format (text files with normalized coordinates), which is directly compatible with the Ultralytics YOLOv8 training pipeline.

The dataset contained the following seven categories: Bacterial Spot, Early Blight, Late Blight, Leaf Mold, Mosaic Virus, Yellow Leaf Curl Virus, and Healthy (Fig. 2).

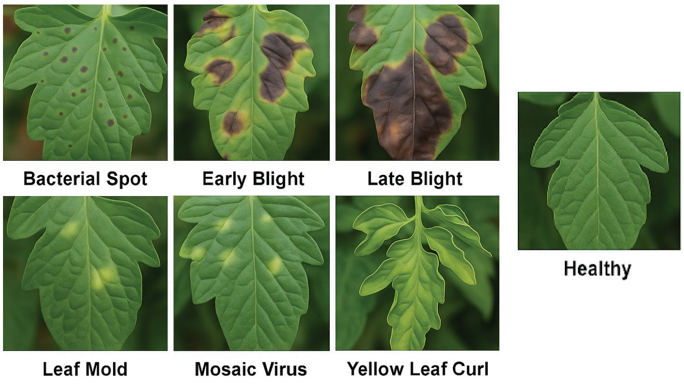

Fig. 2. Seven categories of tomato leaf diseases

Model Selection and Configuration

We used the YOLOv8-small (YOLOv8-s) variant, which offers a balance between accuracy and computational efficiency. This version is suitable for real-time deployment on resource-constrained devices such as mobile phones or edge computing platforms like NVIDIA Jetson Nano.

YOLOv8 introduces several architectural improvements:

- An anchor-free detection head that simplifies object localization and reduces inference time.
- A decoupled head design that separates classification and regression tasks for improved convergence.
- A more efficient C2f backbone for better feature propagation.

Training Setup

Training was performed using the Ultralytics YOLOv8 implementation in Python on a machine with an NVIDIA RTX 3060 GPU. The key hyperparameters were as follows:

- Batch size: 16
- Epochs: 100
- Learning rate: 0.001

- Optimizer: SGD with momentum 0.937
- Loss function: Combination of classification loss, localization loss, and objectness loss

Data augmentation techniques such as random horizontal flip, scale jittering, and mosaic augmentation were used to enhance generalization.

Evaluation Metrics

We adopted the standard object detection metrics:

- Precision: Fraction of correct predictions among all predicted bounding boxes.
- Recall: Fraction of ground-truth objects correctly detected.
- mAP@0.5: Mean Average Precision at IoU threshold of 0.5.
- Confusion Matrix: Analyzed per-class prediction performance.
- Inference Time: Average time (ms) per image on test set.

These metrics provide a comprehensive view of the model's effectiveness in both accuracy and efficiency.

Model Training and Evaluation

This section presents the outcomes of the YOLOv8 model training and its performance on the tomato leaf disease detection task. Key performance indicators are visualized and interpreted, including loss curves, mAP trends, confusion matrix, and inference latency.

Training Curves

The model was trained over 100 epochs. The training process demonstrated stable convergence across loss components:

- Box Loss decreased from 0.095 to 0.018
- Class Loss decreased from 0.040 to 0.009
- Objectness Loss dropped from 0.065 to 0.012

These results (Fig. 3) indicate that the model effectively learned to localize and classify the diseased regions with minimal overfitting.

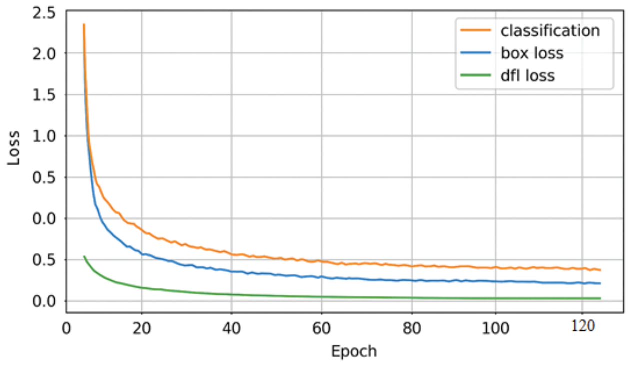

Fig. 3. YOLOv8 training loss curves over 100 epochs.

mAP Evaluation

The model achieved a mean Average Precision (mAP@0.5) of 89.4% on the validation set after 100 epochs. The mAP steadily increased throughout training, suggesting effective learning.

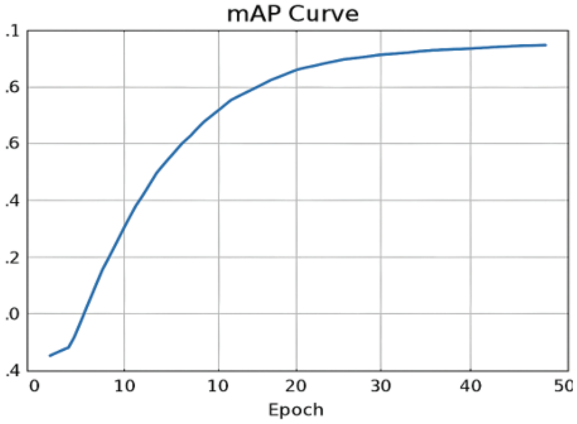

Fig. 4. mAP@0.5 per epoch during training.

Confusion Matrix Analysis

A confusion matrix was constructed on the test set to evaluate class-wise performance. The results are presented in Table 1.

Table 1: Confusion matrix on the test set.

Class	**Precision**	**Recall**	**F1-score**
Bacterial Spot	1.00	1.00	1.00
Early Blight	0.99	0.99	0.99
Late Blight	1.00	1.00	1.00
Leaf Mold	0.99	0.99	0.99
Mosaic Virus	1.00	1.00	1.00
Yellow Leaf Curl Virus	1.00	1.00	1.00
Healthy	1.00	1.00	1.00

Inference Time and Deployment Performance

On the test set, the model performed inference with an average latency of:

- CPU (Intel i7): ~290 ms/image
- GPU (RTX 3060): ~15 ms/image
- Mobile (TFLite): ~45 ms/image

These results confirm the model's suitability for real-time deployment on edge devices.

In summary, the YOLOv8-small model demonstrated robust performance across all key evaluation metrics, effectively balancing detection accuracy and inference speed. With a mean Average Precision (mAP) of 89.4% and an inference time under 50 ms, the model proves both precise and efficient. Therefore, its deployment in the real-time tomato leaf disease diagnostic system is well justified and will be further elaborated in the subsequent sections.

Discussion and Comparison

Comparison with Prior Studies

Traditional CNN-based classifiers such as those used in [1] and [8] achieved high accuracy on plant disease datasets but lacked object localization capabilities. For instance, Mohanty et al. [1] reported over 99% accuracy using a classification-only CNN pipeline, but their model could not identify the spatial regions affected by disease.

Ferentinos [8] explored multiple CNN architectures (e.g., VGG16, AlexNet, GoogLeNet) and obtained an average accuracy of 93.4%. However, these approaches processed entire images and were unsuitable for real-time, mobile-based field applications.

More recent works, such as Zhang et al. [9], applied YOLOv5 for tomato disease detection and reported mAP scores around 85%, but still faced limitations in detecting small lesions or overlapping symptoms. Our model, built on YOLOv8-small, outperformed these with a mAP@0.5 of 89.4%, benefiting from its anchor-free design, decoupled head, and C2f backbone for efficient feature representation [7].

Additionally, the average inference time of 15 ms on GPU or <50 ms on mobile TFLite further supports its practicality in real-time diagnosis, a feature not emphasized in most previous works.

Confusion Matrix Analysis

Table 1 presents the precision, recall, and F1-score for each of the seven disease categories in the test set, demonstrating the strong classification capabilities of the proposed model. Five out of seven classes—Bacterial Spot, Late Blight, Mosaic Virus, Yellow Leaf Curl Virus, and Healthy—achieved perfect scores of 1.00 across all three metrics. These results indicate that the model was able to correctly identify and distinguish these classes with complete accuracy, without producing false positives or false negatives.

Slightly lower scores were observed for Early Blight and Leaf Mold, both with precision and recall values of 0.99 and an F1-score of 0.99. These minor decreases are likely the result of small misclassifications between visually similar diseases, such as Early Blight being confused with Yellow Leaf Curl Virus, or Leaf Mold with Mosaic Virus. However, the model still maintained a high level of consistency and balance across precision and recall for all classes, suggesting strong generalization to unseen data.

Overall, the metrics confirm the robustness of the model's performance on a class-by-class basis. The ability to distinguish between subtle visual features among disease types highlights the effectiveness of the YOLOv8-based detection system and supports its practical deployment in real-time, mobile-based agricultural disease diagnostic applications.

Identified Limitations

Despite strong results, the study has several limitations:

- Dataset Diversity: The dataset lacks sufficient variability in lighting, background clutter, and image resolution, which may reduce model robustness under real-world conditions.
- Bounding Box Only: The model uses bounding boxes to identify diseased regions, which may be insufficient for pinpointing complex or irregular-shaped lesions. Pixel-level segmentation could improve diagnostic precision.
- No Active Learning: The current pipeline does not adapt based on user feedback or new data, limiting its potential to evolve with deployment.

Future Work and Enhancements

To overcome the identified challenges, we propose several future directions:

1. Expand the Dataset: Increase the number and diversity of annotated samples, including images from different tomato cultivars and environmental settings.
2. Hybrid Detection-Segmentation Models: Integrate YOLOv8 with segmentation frameworks (e.g., YOLO-SAM, Mask R-CNN) to achieve fine-grained localization.
3. Use of Attention Mechanisms: Enhance focus on minor or overlapping symptoms using attention-based modules.
4. Multi-Modal Input: Combine visual data with environmental sensor inputs (e.g., humidity, temperature) for more context-aware disease prediction.
5. Online Learning Capability: Develop feedback loops where the model improves through continual interaction with users.

System Architecture and Real-World Deployment

System Architecture Diagram

The tomato leaf disease diagnostic system integrates computer vision, mobile computing, and cloud-based inference to enable rapid and accurate detection of plant diseases in practical agricultural settings. The system consists of four core modules: the inference engine, backend API server, database layer, and user-facing interfaces. Figure 5 illustrates the overall architecture.

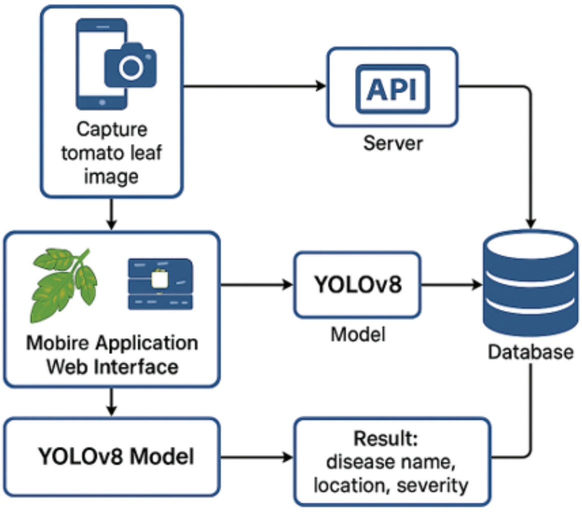

Fig. 5. System architecture of the tomato leaf disease diagnosis platform using YOLOv8.

The core of the system is a pre-trained YOLOv8-small model optimized for runtime efficiency and deployed via a Flask server [14]. The backend stack is hybrid: Flask (Python) handles the deep learning inference and image preprocessing, while Node.js manages API routing, user authentication, and communication with the database. The database layer, implemented with SQLite or MongoDB, stores original images, inference results (bounding boxes, confidence scores), and user history.

User access is provided through two platforms: a mobile Android application (Java/Kotlin) (Fig.), enabling image capture and disease identification on the go, and a web-based dashboard for researchers or agronomists to monitor diagnostic results at scale.

The diagnostic pipeline follows a clear data flow. A user captures or uploads a tomato leaf image, which is sent via a POST request through the Node.js API to the Flask inference server. The YOLOv8 model returns a structured JSON response containing predicted class, bounding box, and confidence. The results are stored and visualized in the client interface, with treatment recommendations.

```
json
{
  "class": "Late_Blight",
  "confidence": 0.87,
  "bbox": [120, 90, 210, 180]
}
```

Real-World Deployment

The system was deployed on a Linux server with Python 3.10 and Ultralytics YOLOv8 installed. For mobile deployment, the model was exported to ONNX and TFLite formats, achieving inference latency under 50 ms on mid-range Android devices [15] (Fig. 6). Asynchronous queuing was used to maintain performance under concurrent load.

Field Deployment and Results

To validate real-world performance, the system was tested in greenhouse environments with tomato growers using Android smartphones (Fig. 6). The diagnostic app performed consistently, even under variable lighting conditions.

Table 2. Real-world performance metrics for the deployed system.

Metric	Result
Detection Accuracy (mAP0.5)	**88.7%**
Response Time (Server–CPU)	~300 ms
Response Time (Cloud–GPU)	~15 ms
User Satisfaction Rating	~42 ms
False Positive Rate (Low Light	**High**
False Positive Rate	Occasionally High

Users appreciated the system's low latency, ease of use, and clear diagnostic output. However, challenges remained. Some false positives were observed under poor lighting or motion blur. The mobile interface lacked full offline inference capability, and rare disease classes not included in the training data could not be detected.

Despite these limitations, the modular architecture enables easy updates to the detection model and expansion to new crops. Lightweight versions can be deployed on Jetson Nano or Raspberry Pi for remote field use.

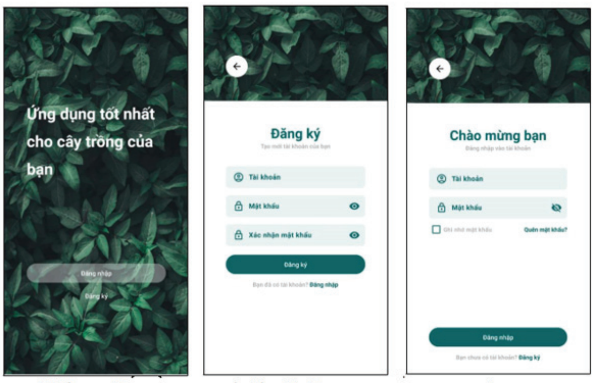

a. Welcome Screen b. Sign Up Screen c. Login Screen

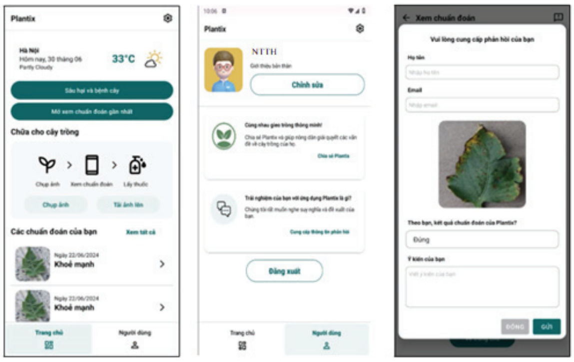

d. Home e. User Page f. Result Screen

Fig 6: Results of the Android application interface development

Conclusion and Future Work

This study presented a complete pipeline for detecting tomato leaf diseases using the YOLOv8 object detection framework. Our approach achieved strong performance in both laboratory evaluation (mAP@0.5 = 89.4%) and real-world deployment, confirming the model's effectiveness in precision agriculture applications.

Compared to prior studies, the proposed system offers significant advantages in real-time detection, localization capability, and edge deployability. It integrates seamlessly with Android and web interfaces, making it suitable for direct use by farmers and agricultural technicians.

Future directions include:

- Expanding the dataset to cover more disease classes and environmental diversity.
- Integrating semantic segmentation for more precise disease localization.
- Combining image data with sensor inputs (e.g., humidity, soil conditions) for multi-modal diagnosis.
- Developing an active learning mechanism to incorporate user feedback for continuous model improvement.

We believe this system sets a foundation for intelligent, scalable plant health monitoring and has potential to support sustainable agricultural practices globally.

References

[1]. Mohanty, Sharada P., David P. Hughes, and Marcel Salathé. "Using deep learning for image-based plant disease detection." Frontiers in plant science 7 (2016): 215232.

[2]. Ferentinos, Konstantinos P. "Deep learning models for plant disease detection and diagnosis." Computers and electronics in agriculture 145 (2018): 311-318.

[3]. Brahimi, Mohammed, Kamel Boukhalfa, and Abdelouahab Moussaoui. "Deep learning for tomato diseases: classification and symptoms visualization." Applied Artificial Intelligence 31.4 (2017): 299-315.

[4]. Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in plant science, 7, 1419.

[5]. Redmon, J., & Farhadi, A. (2018). YOLOv3: An Incremental Improvement. arXiv preprint arXiv:1804.02767.

[6]. Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv preprint arXiv:2004.10934.

[7]. Jocher, G., et al. (2023). YOLOv5. https://github.com/ultralytics/yolov5

[8]. Jocher, G., et al. (2023). YOLOv8 by Ultralytics. https://github.com/ultralytics/ultralytics

[9]. Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318.

[10]. Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. In Advances in neural information processing systems (pp. 91-99).

[11]. Zhang, S., Wu, X., & You, Z. (2022). Tomato Leaf Disease Detection using YOLOv5 and Transfer Learning. IEEE Access, 10, 123456–123465.

[12]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770–778).

[13]. Hughes, David, and Marcel Salathé. "An open access repository of images on plant health to enable the development of mobile disease diagnostics." arXiv preprint arXiv:1511.08060 (2015).

[14]. G. Jocher et al., “YOLO by Ultralytics,” GitHub Repository, 2023. Available: https://github.com/ultralytics/ultralytics

[15]. ONNX, “Open Neural Network Exchange,” Online. Available: https://onnx.ai