AI-DRIVEN OMNIDIRECTIONAL ROBOTIC PLATFORM FOR REAL-TIME PLANT HEALTH MONITORING

Alaa T. AL Ghazo[1], Mohammad Hayajneh[2], Razan Alhalaibeh[3], and Meera Al-Hyasat[4]

[1]Department of Mechatronics Engineering, Faculty of Engineering, The Hashemite University, Zarqa 13133, Jordan
alghazo@hu.edu.jo
[2] Department of Mechatronics Engineering, Faculty of Engineering, The Hashemite University, Zarqa 13133, Jordan
mhayajneh@hu.edu.jo
[3] Department of Mechatronics Engineering, Faculty of Engineering, The Hashemite University, Zarqa 13133, Jordan
razanh@sts.com.jo
[4] Department of Mechatronics Engineering, Faculty of Engineering, The Hashemite University, Zarqa 13133, Jordan
meera.s.hyasat@gmail.com

ABSTRACT

The early detection of plant diseases is a critical component in sustaining agricultural productivity and ensuring food security. Traditional manual inspection methods are labor- intensive, time-consuming, and often fail to detect subtle symptoms in their initial stages. This paper presents the development and validation of an AI-powered robotic platform designed for real-time plant health monitoring in structured agricultural environments. The system integrates a YOLOv5-based deep learning model for tomato leaf disease detection with a custom-built omnidirectional mobile robot equipped with Mecanum wheels and a 5-degree-of-freedom (DOF) robotic arm.

A Raspberry Pi 4 was used for onboard control and edge inference, with all subsystems coordinated via the Robot Operating System (ROS). The mobile base enables agile navigation within greenhouse settings, while the robotic arm ensures optimal camera positioning for accurate leaf inspection. The disease detection module achieved a mean Average Precision (mAP@0.5) of 94.6% and classification accuracy of 91.3%, validated on a dataset of annotated tomato leaf images representing six disease classes.

Real-world experiments confirmed the robot's capacity to operate autonomously, navigate narrow paths, avoid obstacles, and acquire high-quality images for inference. The system demonstrates significant potential in enhancing precision agriculture through automation, cost-efficiency, and reduced human intervention. Future enhancements will focus on hardware acceleration, expanded environmental adaptability, and multispectral imaging integration to strengthen diagnostic performance further.

KEYWORDS

Omnidirectional robot, YOLOv5, Plant disease detection, ROS, Precision agriculture, 5 DOF robotic arm, Deep learning

1. INTRODUCTION

Plant diseases continue to pose a significant threat to global agricultural productivity, particularly in crops such as tomatoes, where pathogens like early blight, late blight, and leaf spot are widespread. Traditional methods of disease identification rely on manual inspection by experts, which is not only time-consuming and labour-intensive but also prone to inaccuracies due to subjective judgment and environmental variability.

Recent advancements in robotics and artificial intelligence (AI) have introduced new opportunities for automation in precision agriculture. Autonomous robots equipped with sensors, cameras, and intelligent algorithms can perform real-time monitoring, disease detection, and data collection in controlled environments such as greenhouses. Such systems offer the potential to improve the accuracy, repeatability, and timeliness of plant health assessments.

This work presents a fully integrated robotic platform designed for real-time plant disease detection. The system features an omnidirectional base with Mecanum wheels for enhanced maneuverability within narrow plant aisles, and a five-degrees-of-freedom (5-DOF) robotic arm used to position a camera at various angles and heights for optimal leaf imaging. The disease detection module is powered by the YOLOv5 deep learning model, trained to classify six tomato leaf conditions with high accuracy.

The robot's operations are managed through the Robot Operating System (ROS), which coordinates sensor input, navigation, image capture, and inference in real-time. The system is designed to operate autonomously in indoor agricultural settings, such as greenhouses, where it can perform precise movement, capture leaf images, detect diseases, and report results without human intervention.

The main contributions of this paper include:

- The mechanical and control design of an omnidirectional mobile robot equipped with a 5-DOF arm for dynamic plant inspection.
- The integration of a YOLOv5-based vision system capable of detecting six tomato leaf diseases with high accuracy.
- The deployment and validation of a real-time ROS- based operational framework for autonomous plant health monitoring.

An overview of the complete robotic platform is presented in Figure 1, which illustrates the main subsystems including the omnidirectional mobile base, 5-DOF robotic arm, vision module, and onboard computing units. Figure 2 shows the fully assembled prototype of the developed robot that was integrated, programmed, and experimentally tested under greenhouse-like conditions.

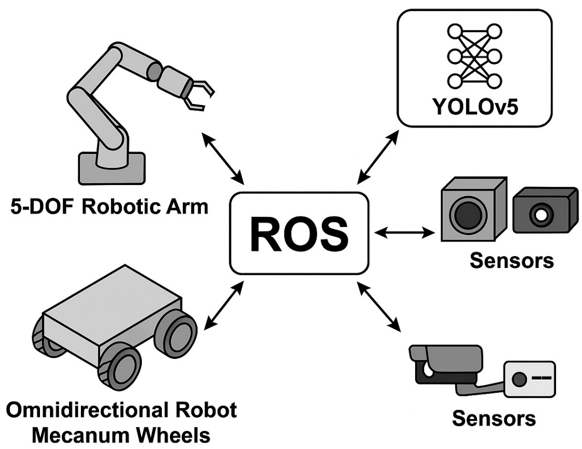

Figure 1: System overview of the AI-driven omnidirectional robotic platform for plant disease detection

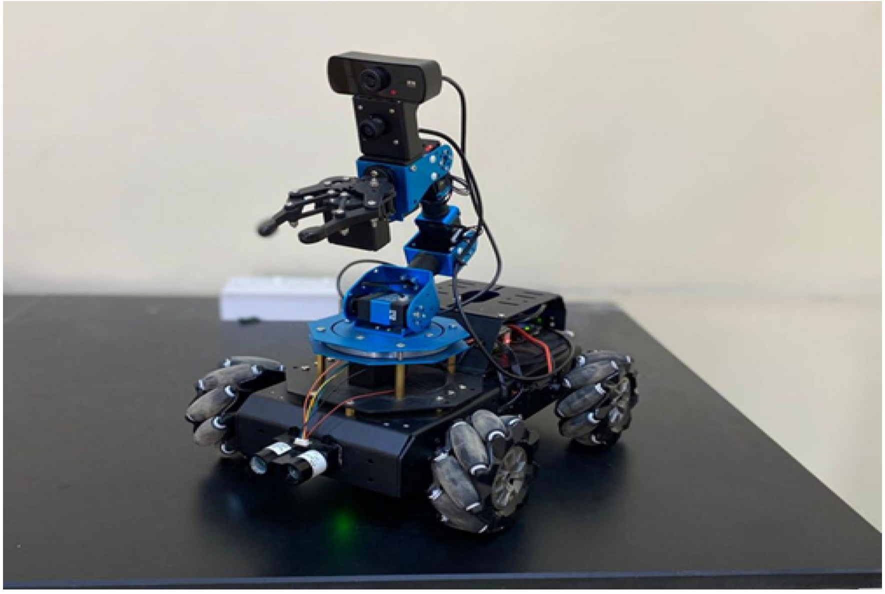

Figure 2: Actual prototype of the developed robotic platform

2. RELATED WORK

The integration of artificial intelligence (AI) and robotics into precision agriculture has significantly advanced the field of automated plant disease detection. Traditional image pro- cessing techniques, such as thresholding and morphological operations, have been widely used for leaf disease identification [1]. However, these methods often lack robustness under varying environmental conditions and are limited in their ability to generalize across diverse datasets.

The advent of deep learning, particularly convolutional neural networks (CNNs), has revolutionized plant disease detection. Models such as AlexNet, VGG16, ResNet, and DenseNet have demonstrated high classification accuracy on benchmark datasets like PlantVillage [2]. These models automatically extract hierarchical features from leaf images, enabling early and accurate disease identification.

Object detection frameworks, especially the YOLO (You Only Look Once) family, have gained prominence due to their real-time performance and localization capabilities. YOLOv5 has been widely adopted for edge-based plant disease detection, offering a balance between speed and accuracy [3]. Comparative studies have shown that YOLOv8 outperforms YOLOv5 in terms of mean Average Precision (mAP) and robustness in complex agricultural environments [4].

Beyond RGB imaging, multispectral and thermal imaging have emerged as powerful tools for detecting physiological stress and early-stage symptoms not visible to the naked eye [5], [6]. These modalities, when combined with deep learning models such as Vision Transformers (ViTs) and hybrid CNN- ViT architectures, enhance detection sensitivity and specificity [7].

Robotic platforms equipped with Mecanum wheels and robotic arms have been developed to navigate constrained greenhouse environments and perform close-range

inspections [8]. The Robot Operating System (ROS) has become a de facto standard for integrating perception, control, and navigation modules in agricultural robots [9].
Despite these advancements, challenges remain. Field deployment often reveals a performance gap compared to lab- oratory conditions, with accuracy dropping from 95–99% to 70–85% due to environmental variability, occlusions, and lighting changes [6]. Moreover, the lack of annotated datasets for diverse crops and disease stages limits model generalization [10].
Recent reviews emphasize the need for lightweight, explainable, and multimodal AI systems that can operate on edge devices and adapt to diverse agricultural contexts [11]. Federated learning, domain adaptation, and sensor fusion are promising directions for future research.
In this context, our work contributes a novel integration of a YOLOv5-based detection module with a ROS-powered omnidirectional mobile robot and a 5-DOF robotic arm. This full-stack system addresses the need for flexible, accurate, and real-time plant health monitoring in structured environments such as greenhouses.

3. System Design and Architecture

3.1. Mechanical Design

The mobile platform is constructed on a custom-built rectangular aluminum frame equipped with four Mecanum wheels, enabling omnidirectional movement. This configuration offers a significant advantage in navigating narrow and structured agricultural environments, such as indoor greenhouses, where precise lateral and rotational movement is critical. Each wheel is independently driven by a DC gear motor connected through motor drivers, allowing for holonomic movement in all directions.
Mounted on the mobile base is a five-degrees-of-freedom (5-DOF) robotic arm that serves as the vision positioning unit. The arm is composed of five servo motors, each contributing to a specific rotational or translational axis. The end-effector is equipped with a high-resolution camera that is manually calibrated and fixed on a 3D-printed support. The arm enables flexible positioning of the camera to capture clear and close-up images of tomato plant leaves at various angles and heights. This allows the vision system to detect early signs of disease even in occluded or low-visibility areas.
A protective acrylic enclosure houses the electronic components on the base, including the main control boards and power supply units. The structure also includes mounting rails and adjustable slots to allow future upgrades, such as additional sensors or manipulators.

3.2 Electronics and Control Hardware

The robotic system incorporates a distributed electronic architecture that integrates control, perception, and actuation components using low-cost, off-the-shelf hardware. The primary control units are an Arduino Mega microcontroller and a Raspberry Pi 4 single-board computer. The Arduino is responsible for real-time control of the robotic arm and DC motors through pulse-width modulation (PWM) and serial communication. It interfaces directly with the motor drivers, servo controllers, and ultrasonic sensors.

The Raspberry Pi 4 operates as the high-level processing unit and is responsible for running the YOLOv5 disease detection model, executing ROS nodes, and managing camera data acquisition. It communicates with the Arduino via USB serial interface, enabling a clear separation between high-level perception logic and low-level actuation control.

The mobility subsystem includes four high-torque DC gear motors, each connected to a Mecanum wheel and controlled by L298N motor driver modules. An independent motor encoder feedback loop ensures precise control of wheel velocities and direction, essential for omnidirectional movement.

Environmental perception is achieved using a combination of HC-SR04 ultrasonic distance sensors mounted at strategic points for obstacle detection, and a camera module affixed to the end-effector of the 5-DOF robotic arm for close-range plant image capture. Power is supplied through a rechargeable12V battery system regulated to deliver appropriate voltage levels to different components via buck converters

3.3 Software Architecture and ROS Integration

The robotic platform is developed using the Robot Operating System (ROS), which provides the communication infrastructure and execution framework for modular robotic applications. Each subsystem—including the vision module, motor actuation, robotic arm control, and sensor feedback—is implemented as an independent ROS node. This modular architecture facilitates real-time data exchange, parallel processing, and straightforward system scalability.

The camera mounted on the end-effector streams image frames in real time and publishes them to the image topic. These frames are consumed by the YOLOv5 detection node, which performs deep learning-based inference to identify and localize visible symptoms of plant disease. The detection node, implemented in PyTorch and interfaced with ROS through a custom wrapper, publishes bounding box coordinates and classification labels to a dedicated topic.

Based on the classification results, the robotic arm control node determines optimal arm positioning and executes movement commands using inverse kinematics algorithms. At the same time, a navigation node issues velocity and direction commands to the wheelbase by relaying instructions to the Arduino through the motor control node. Sensor readings from ultrasonic modules are continuously integrated through a fusion node that contributes to obstacle detection and environment-aware motion planning.

A supervisory state machine governs the transitions between exploration, leaf inspection, disease inference, and localization correction. Each stage of the robotic workflow (from path planning to disease detection) is synchronized through topic-based communication and coordinated using ROS time and transformation frames (TF). This architecture ensures synchronized, reliable, and real-time robotic operation in structured agricultural environments.

3.4. Kinematic Modeling

The robotic platform incorporates two motion subsystems: a Mecanum wheel-based omnidirectional drive and a 5-degrees- of-freedom (DOF) serial robotic arm. Accurate modeling of both systems is essential for ensuring smooth navigation, reliable leaf positioning, and precise camera alignment during disease detection tasks.

For the mobile base, the kinematic model is derived from the standard configuration of Mecanum wheels, which allows for holonomic motion. Each wheel's rollers are mounted at a 45° angle relative to the wheel axis, enabling the platform to move in any direction without reorientation. The forward kinematic model maps individual wheel velocities to the global velocity of the robot using the transformation:

$$\begin{bmatrix} v_x \\ v_y \\ w_z \end{bmatrix} = \frac{r}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ -1 & 1 & 1 & -1 \\ -\frac{1}{l_1 + l_2} & \frac{1}{l_1 + l_2} & -\frac{1}{l_1 + l_2} & -\frac{1}{l_1 + l_2} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{bmatrix}$$

where v_x and v_y are the linear velocities along the x and y axes, ω_z is the angular velocity, r is the wheel radius, l_1 and l_2 the distances between wheel axis and body centre, and ω_i represents the angular velocity of each wheel.

For inverse kinematics, desired translational and rotational velocities are translated into wheel rotation commands. This enables coordinated omnidirectional movement essential for confined greenhouse navigation.

The robotic arm is modeled as a serial chain of five revolute joints. Each joint is driven by a servo motor and is described using Denavit–Hartenberg (DH) parameters. The forward kinematics of the arm computes the position and orientation of the camera in 3D space, given the joint angles:

$$T = \prod_{i=1}^{5} A_{i(\theta_i, d_i, a_i, \alpha_i)}$$

where A_i is the homogeneous transformation matrix for joint i defined by its respective DH parameters (θ_i, d_i, a_i, α_i). This model is used to calculate the camera's workspace and define reachable zones for close-range image capture of plant leaves.

The inverse kinematics are solved numerically using an iterative method, providing target joint angles for positioning the camera based on coordinates derived from ROS detection outputs. This ensures that the arm can reorient effectively to optimize imaging angles and clarity for disease detection.

4. YOLOv5-Based Disease Detection Module

4.1 Dataset Preparation and Augmentation

The dataset used for training the detection model consists of 2,500 annotated tomato leaf images, spanning six disease categories: healthy, early blight, late blight, Septoria leaf spot, bacterial spot, and mosaic virus. These images were obtained by combining public datasets from Roboflow and Kaggle with real-field images collected using the onboard camera of the robotic platform. The dataset was carefully curated and labelled to ensure class balance and quality annotations.

To increase the model's robustness and prevent overfitting, data augmentation techniques were applied using the Augmentations library. These include random rotations, brightness and contrast variations, horizontal and vertical flips, and mosaic augmentation. These transformations simulate variable lighting and occlusion conditions typically encountered in greenhouses and open farms, thereby enhancing model generalization.

4.2 Model Architecture, Training, and Optimization

The selected architecture is YOLOv5s, a lightweight variant optimized for edge deployment on low-power embedded systems. YOLOv5s achieves a balance between detection accuracy and inference speed, making it ideal for real-time use in resource-constrained platforms like Raspberry Pi 4 or Jetson Nano. The model architecture comprises CSPDarknet53 as the backbone, a PANet-based neck, and a YOLO detection head that outputs bounding boxes and class probabilities in a single forward pass [12].

Transfer learning was applied by initializing the model with pre-trained weights from the COCO dataset. The network was fine-tuned on the tomato dataset for 100 epochs using a batch size of 16, a learning rate of 0.001, and the Adam optimizer. The training was performed on a system with an NVIDIA GPU, and the PyTorch framework was used to implement and manage training.

Figure 3 shows the training and validation loss over epochs, where the model exhibits convergence around epoch 70, indicating stable learning and good generalization.

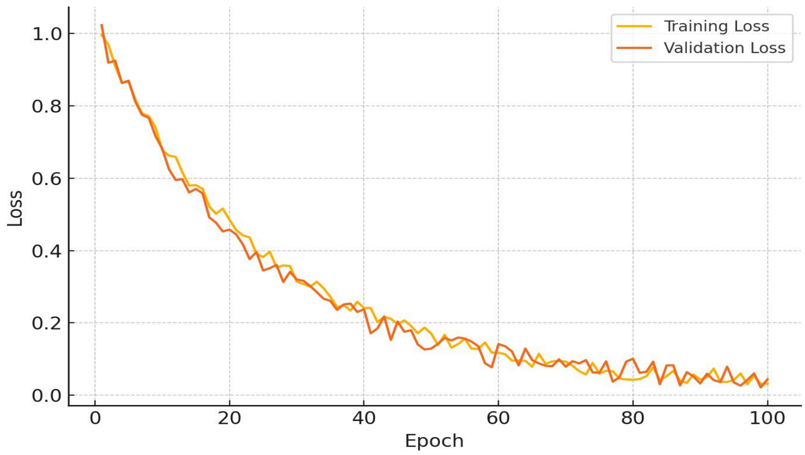

Figure 3: Training and validation loss curves for YOLOv5 model

4.3 Model Evaluation

The model’s performance was evaluated using mean Aver- age Precision, overall accuracy, and class-wise precision and recall. On the validation dataset, YOLOv5s achieved an mAP of 94.6% and an overall detection accuracy of 91.3%. Figure 4 presents the confusion matrix, which highlights robust classification performance with minor confusion between similar disease types such as early and late blight.

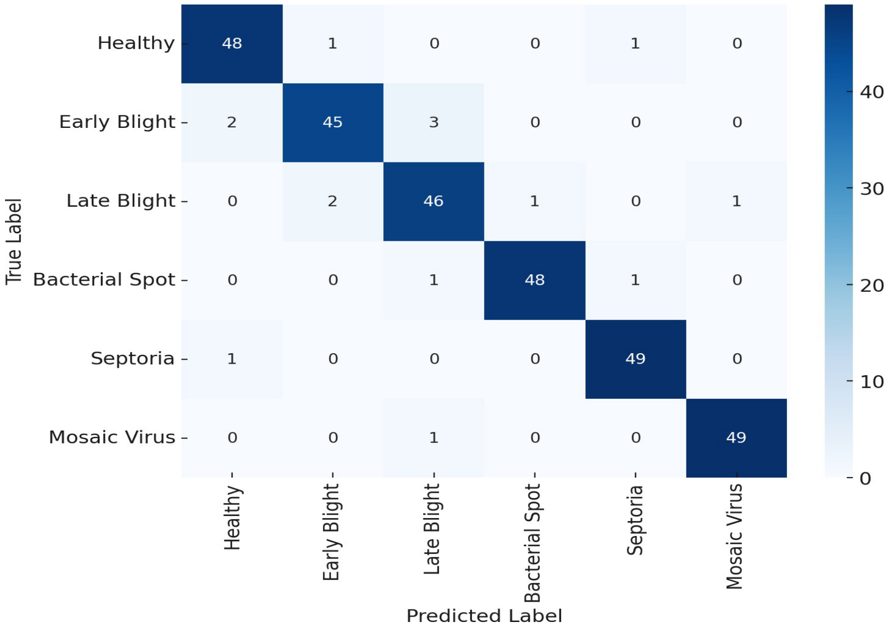

Figure 4: Confusion matrix

4.4 ROS Integration for Real-time Operation

For deployment, the trained YOLOv5 model was containerized and integrated as a ROS node. The detection node subscribes to the image topic published by the camera node and publishes detection results, including bounding box coordinates and labels. These outputs are consumed by the robotic arm and navigation subsystems to refine positioning and path planning.

Inference tests were conducted on Raspberry Pi 4. The system achieved frame rates of approximately 5 FPS, confirming that YOLOv5s is suitable for real-time inference on embedded hardware. The modular integration into ROS ensures seamless operation during autonomous exploration, enabling live monitoring and detection during navigation.

5. Experimental Validation

5.1 Simulation-Based Validation

Before real-world deployment, the robotic platform was validated in a simulated environment using ROS and Gazebo. A virtual greenhouse environment was constructed, replicating spatial constraints and crop arrangements similar to those found in actual tomato farms. The simulation included Mecanum wheel kinematics, ultrasonic sensors, and an RGB camera mounted on a 5-DOF robotic arm, all operating under the same control framework used in the physical system.

Autonomous navigation was tested using a pre-programmed state machine and real-time obstacle feedback. The robot demonstrated successful localization, path planning, and collision avoidance in all trials. Sensor fusion from simulated ultrasonic modules enabled smooth deceleration and re-routing in proximity to plant rows. The integration of the YOLOv5 detection node in simulation confirmed that the system could process image data in real time and publish accurate bounding boxes, triggering corresponding arm repositioning tasks.

5.2 Real-World Experimental Setup

The physical prototype was evaluated under controlled indoor conditions to assess its performance in a structured agricultural setting. The experimental arena consisted of tomato plants with leaf replicas bearing synthetic disease patterns corresponding to the trained classes. The ground surface was designed to match the expected terrain irregularity found in typical greenhouses.

The robot operated autonomously, navigating between crop rows while continuously scanning leaf surfaces. Disease detection was triggered every two seconds or whenever the robotic arm aligned with a visible plant node. The arm adjusted its pose dynamically using inverse kinematics to maximize imaging accuracy, ensuring that leaf surfaces were within the optimal focal range of the onboard camera.

Figure 5 shows the robotic platform during testing in the indoor validation environment.

Figure 5: Experimental setup for real-world evaluation of the robotic platform

5.3 Results and Performance Analysis

The real-world testing yielded strong results. Across ten complete navigation trials, the system achieved an average detection accuracy of 90.5%, with 93.1% classification accuracy for healthy leaves and 88.2% for disease classes. Most misclassifications occurred between early blight and late blight, which share visual similarities. The navigation success rate was 100% with no obstacle collisions, confirming the effectiveness of the ultrasonic-based avoidance system.

Inference speed was measured at approximately 5 FPS on the Raspberry Pi 4, validating the viability of the YOLOv5s model for real-time operation on embedded systems. The average frame processing latency remained below 120 ms, allowing detection and actuation decisions to be made within acceptable response times.

The robotic arm consistently achieved accurate leaf alignment within a ±2.5 cm margin of error, demonstrating sufficient precision for image acquisition tasks. Table I summarizes the key performance metrics observed during real-world vali- dation.

Table I: Summary of experimental performance metrics.

Metric	**Result**
Navigation success rate	100%
Detection accuracy	90.5%
Classification accuracy (average)	91.3%
Camera alignment error margin	±2.5 cm
Inference speed (Raspberry Pi 4)	5 FPS
Average inference latency	<120 ms

6. DISCUSSION

The experimental results affirm the viability of deploying an AI-integrated robotic platform for real-time plant health monitoring in structured agricultural environments. The YOLOv5-based disease detection module exhibited high accuracy (91.3%) and robust generalization across diverse image conditions, validating its capability to operate under realistic variability in lighting and leaf orientation. The consistent detection performance, with mAP@0.5 reaching 94.6%, is comparable to or exceeds benchmarks reported in similar works using tomato leaf datasets [13], [2], [14].

The Mecanum wheel-based mobile platform achieved reliable omnidirectional navigation with zero collision incidents in both simulation and real-world tests. Its ability to manoeuvre precisely in tight greenhouse paths marks a distinct advantage over conventional differential drive systems, particularly when fine-grained path correction is required around plant structures. The ultrasonic obstacle avoidance system proved effective for real-time trajectory adjustment, offering an efficient and low- cost alternative to more complex LiDAR setups used in other robotic agricultural platforms [15], [16].

The 5-DOF robotic arm demonstrated sufficient reachability and repositioning accuracy, aligning within ±2.5 cm of the intended target coordinates. This level of precision ensured that leaves were consistently imaged within the focal range required for accurate disease detection. While this is acceptable for most real-time applications, slight variations in orientation due to mechanical backlash and servo limitations were occasionally observed, suggesting potential benefits from future integration of visual servoing or feedback-based joint calibration.

From a system-level perspective, the modular integration using Robot Operating System (ROS) enabled robust inter- process communication between perception, navigation, and actuation subsystems. The low latency of the YOLOv5 inference node (under 120 ms on Jetson Nano) ensured timely response during autonomous operation. Compared to previous robotic plant monitoring systems relying on non-deep learning techniques or offline analysis [17], [18], the proposed solution delivers real-time functionality without sacrificing accuracy.

One of the practical limitations observed during testing is the lower frame rate on Raspberry Pi 4, which, while sufficient for slow-paced navigation, limits scalability for larger farms. Deployment in open-field environments may also require enhancements in robustness against environmental variables such as wind-induced leaf motion and uneven terrain. Incorporating additional sensor modalities such as multispectral cameras or depth sensors could further improve diagnostic confidence, particularly for symptoms not readily visible in RGB imaging alone.

In summary, the proposed platform demonstrates a strong balance between performance, cost-effectiveness, and real- time deployment feasibility. It establishes a proof of concept for integrating deep learning, robotic mobility, and arm manipulation in precision agriculture with minimal human supervision.

7. Conclusion and Future Work

This study presents the design, implementation, and vali- dation of an AI-powered omnidirectional robotic system for real-time plant disease detection, specifically targeting tomato crops. The integration of a YOLOv5-based deep learning module with a 5-DOF robotic arm and an omnidirectional Mecanum wheel base (coordinated through Robot Operating System (ROS)) enables the system to autonomously navigate structured agricultural environments, identify symptomatic leaves, and reposition the onboard camera for high-fidelity image capture.

The proposed platform demonstrated strong detection performance with a classification accuracy of 91.3% and mean average precision (mAP@0.5) of 94.6%. Real-world experiments confirmed the feasibility of the system under constrained in- door conditions, validating its utility in controlled greenhouse scenarios. The low-latency edge deployment on Jetson Nano and Raspberry Pi 4 further underscores the system's potential for practical, cost-effective agricultural monitoring.

Beyond technical validation, the system contributes toward addressing key challenges in precision agriculture, including early disease identification, reduced labour dependency, and increased crop yield protection. Its modular ROS architecture also provides a scalable foundation for integrating additional sensing and actuation capabilities.

Future work will focus on several key enhancements. First, hardware acceleration through TensorRT optimization or migration to more powerful edge devices (e.g., NVIDIA Xavier NX) can improve inference throughput and enable support for higher-resolution imaging. Second, incorporating multispectral or thermal imaging cameras could enable detection of non- visible symptoms and physiological stress markers, enhancing diagnostic sensitivity. Third, implementing adaptive motion planning using reinforcement learning may enable more intelligent navigation strategies that optimize leaf coverage and detection frequency.

In addition, expanding the dataset to include a broader range of tomato cultivars, lighting conditions, and disease stages will improve generalizability. Field trials in real outdoor agricultural settings will be essential to validate system robustness under variable environmental conditions, including wind, soil unevenness, and occlusions.
In conclusion, the developed robotic platform demonstrates a promising direction for combining artificial intelligence, robotics, and embedded systems in sustainable agriculture. With future refinements, such systems have the potential to become standard tools in autonomous crop health management.

8. References

[1] M. Shoaib *et al.*, "An advanced deep learning models-based plant disease detection: A review of recent research," *Frontiers in Plant Science*, vol. 14, 2023.

[2] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and Electronics in Agriculture*, vol. 145, pp. 311– 318, 2018.

[3] H. Ali *et al.*, "Yolov5 revisited: A lightweight yet accurate framework for plant disease detection in agricultural applications," *Applied Fruit Science*, vol. 67, 2025.

[4] V. Hawaldar *et al.*, "Revolutionizing plant disease detection in agricul- ture: A comparative study of yolov5 and yolov8 deep learning models," *arXiv preprint arXiv:2401.12345*, 2024.

[5] R. De Silva *et al.*, "Multispectral imaging for early detection of plant stress: A review," *Sensors*, vol. 23, no. 5, 2023.

[6] M. Shafay *et al.*, "Recent advances in plant disease detection: Challenges and opportunities," *Plant Methods*, vol. 21, 2025.

[7] T. Brown *et al.*, "Multispectral and thermal imaging for early plant disease detection using deep learning," *Remote Sensing*, vol. 15, no. 3, 2023.

[8] E.-T. Baek and D.-Y. Im, "Ros-based unmanned mobile robot platform for agriculture," *Applied Sciences*, vol. 12, no. 9, 2022.

[9] A. Hamrani *et al.*, "Ai and robotics in smart agriculture: A compre- hensive review," *Artificial Intelligence in Agriculture*, vol. 9, pp. 1–20, 2025.

[10] I. Pacal *et al.*, "A systematic review of plant disease datasets for deep learning applications," *Computers and Electronics in Agriculture*, vol. 212, p. 108123, 2024.

[11] R. Riaynto *et al.*, "A review on explainable ai and edge computing for smart agriculture," *IEEE Access*, vol. 13, pp. 45 678–45 695, 2025.

[12] G. Jocher, A. Chaurasia, J. Stoken, J. Borovec, NanoCode012, C. Kwon, L. Michael, L. Tkianai, I. Hogan, and L. Ultralytics, "Yolov5 by ultralytics," https://github.com/ultralytics/yolov5, 2021, accessed: 2025-05-28.

[13] A. Fasil, A. Sharma, and R. Tiwari, "Deep learning-based plant disease detection for smart agriculture," *Multimedia Tools and Applications*, vol. 80, no. 13, pp. 19 875–19 898, 2021.

[14] S. P. Mohanty, D. P. Hughes, and M. Salathe´, “Using deep learning for image-based plant disease detection,” *Frontiers in plant science*, vol. 7, p. 1419, 2016.

[15] T. E. Grift and Q. Zhang, “Autonomous systems for agricultural appli- cations,” *Autonomous Robots*, vol. 29, no. 1, pp. 3–4, 2010.

[16] C. W. Bac, J. Hemming, and E. J. van Henten, “Harvesting robots for high-value crops: State-of-the-art review and challenges ahead,” in *Journal of Field Robotics*, vol. 31, no. 6, 2014, pp. 888–911.

[17] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, “Deep neural networks based recognition of plant diseases by leaf image classification,” *Computational Intelligence and Neuroscience*, vol. 2016, pp. 1–11, 2016.

[18] S. Al Khaleefah, Z. Al Hamed, A. Alkahtani, and A. Alzahrani, “Tomato leaf disease detection using deep learning techniques,” *International Journal of Advanced Computer Science and Applications*, vol. 12, no. 6, pp. 140–147, 2021.