

Vision-Guided Robotic Arm for Tomato Quality Classification and Sorting

Shrawan Thakur¹, Sudan Jha¹, Pranaya Mulepati²

¹School of Engineering, Kathmandu University, Dhulikhel, Nepal

²School of Science and Technology, Nottingham Trent University, Nottingham, United Kingdom

Email: shrawan7825@student.ku.edu.np

How to cite this paper: Thakur, S., Jha, S. and Mulepati, P. (2025) Vision-Guided Robotic Arm for Tomato Quality Classification and Sorting. *Advances in Artificial Intelligence and Robotics Research*, 1, 75-94. <https://doi.org/10.4236/airr.2025.12006>

Received: October 4, 2025

Accepted: November 15, 2025

Published: November 18, 2025

Copyright © 2025 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

The DOFBOT Robotic Hand, powered by the NVIDIA Jetson Nano processor, represents a cutting-edge fusion of computer vision and robotic automation for precision agriculture. This research presents an advanced system capable of distinguishing between ripe (red) and unripe (green) tomatoes using sophisticated image processing algorithms and deep learning techniques. The system integrates ImageNet and ResNet pre-trained models for enhanced color recognition capabilities, achieving superior response times compared to traditional laptop-based processing. The robotic arm successfully identifies, plucks, and collects ripe tomatoes while maintaining high accuracy rates. This study demonstrates the effectiveness of edge computing in agricultural automation, with the Jetson Nano providing significantly improved response times over conventional cloud-based processing systems. The integration of state-of-the-art computer vision algorithms with precision robotics marks a significant advancement in automated agricultural harvesting systems.

Keywords

Index Terms-IoT, AI, ML, Computer Vision, Edge Computing, Agricultural Automation, Tomato Harvesting, Deep Learning

1. Introduction

The integration of artificial intelligence and robotics into agriculture has emerged as a critical solution to address global food security challenges, labor shortages, and the need for sustainable farming practices [1]. Modern agricultural systems require precision, efficiency, and adaptability to meet the increasing demands of global food production. The development of intelligent harvesting robots capable of distinguishing between ripe and unripe fruits represents a significant break-

through in agricultural automation [2] [3]. Traditional tomato harvesting is highly dependent on manual labor, which is time consuming, costly and subject to human error. The ability to accurately identify fruit ripeness is crucial to maintain quality standards and optimize harvest timing [4]. Computer vision technology offers a promising solution that enables automated systems to make rapid, consistent and accurate assessments of fruit ripeness based on visual characteristics. The DOFBOT Robotic Hand, equipped with the powerful NVIDIA Jetson Nano processor and advanced camera module, represents a paradigm shift in agricultural robotics [5]. This system combines sophisticated computer vision algorithms with precise mechanical control to achieve unprecedented levels of automation in fruit harvesting operations. The integration of ImageNet and ResNet pretrained models enhances the system's ability to distinguish between ripe (red) and unripe (green) tomatoes with remarkable accuracy [6]. Edge computing technology plays a pivotal role in this advancement, enabling real-time processing capabilities that surpass traditional cloud-based systems in terms of response time and reliability [7]. The NVIDIA Jetson Nano's onboard GPU acceleration provides the computational power necessary for real-time image processing while maintaining energy efficiency and compact form factor suitable for field deployment. This research contributes to the growing body of knowledge in precision agriculture by demonstrating the practical application of advanced computer vision techniques in automated fruit harvesting. The system's ability to operate autonomously while maintaining high accuracy rates positions it as a valuable tool for modernizing agricultural practices and addressing current industry challenges.

2. Proposed System

The proposed system employs a comprehensive array of advanced technologies and methodologies to achieve precise tomato ripeness detection and automated harvesting. The system architecture integrates state-of-the-art hardware components with sophisticated software algorithms to deliver reliable, real-time performance in agricultural environments.

2.1. DOFBOT AI Vision Robotic Arm

The DOFBOT AI Vision Robotic Arm serves as the primary mechanical platform for the automated harvesting system [8]. This precision-engineered robotic arm features multiple degrees of freedom, enabling complex manipulation tasks required for delicate fruit handling. The arm's design incorporates smooth servo control systems that provide the necessary precision for approaching, grasping, and retrieving tomatoes without causing damage to the fruit or surrounding plants [9]. The robotic arm's end-effector is specifically designed for tomato harvesting applications, featuring an adaptive gripper mechanism capable of accommodating various fruit sizes while maintaining gentle contact pressures [10]. The integration of force feedback sensors ensures that the gripper applies appropriate pressure levels, preventing fruit damage during the harvesting process [11].



Figure 1. DOFBOT AI vision robotic arm.

2.2. NVIDIA Jetson Nano Edge Computing Platform




Figure 2. NVIDIA's Jetson Nano [5].

The NVIDIA Jetson Nano represents a revolutionary advancement in edge computing technology for AI applications (**Figure 1**, **Figure 2**). This compact yet powerful single-board computer features a quad-core ARM Cortex-A57 CPU paired with a 128-core NVIDIA Maxwell GPU, providing substantial computational capacity for real-time AI inference tasks.

Key Technical Specifications:

- 1) **GPU:** 128-core NVIDIA Maxwell architecture with CUDA support.
- 2) **CPU:** Quad-core ARM Cortex-A57 @ 1.43 GHz.
- 3) **Memory:** 4GB LPDDR4 @ 25.6 GB/s.
- 4) **AI Performance:** 472 GFLOPS.
- 5) **Power Consumption:** 5-10W operational range.
- 6) **Storage:** MicroSD card support with optional eMMC.

The Jetson Nano's GPU acceleration capabilities enable real-time processing of complex computer vision algorithms, including deep neural networks for object

detection and classification. The platform supports popular AI frameworks such as TensorFlow, PyTorch, and OpenCV, making it accessible for deploying sophisticated machine learning models in edge environments.

2.3. Advanced Camera System

The system incorporates a high-resolution camera module specifically optimized for agricultural applications. The camera features adjustable focus, exposure control, and support for various lighting conditions commonly encountered in greenhouse and field environments. The integration of specialized optical filters enhances color discrimination capabilities, particularly important for distinguishing between red ripe tomatoes and green unripe ones [12].

We also utilized simulation software like ROS (Robot Operating System) to model and validate our algorithms in a simulated environment prior to deployment. For data processing and visualization, we leveraged libraries such as NumPy. Version control technologies like Git enhanced team collaboration and facilitated effective software management. These tools and technologies collectively played a pivotal role in achieving the objectives of our project, enabling us to deliver a robust and efficient solution that met the requirements of our stakeholders.

3. Methodology

This approach of research has a systematic pipeline for color recognition, model integration, real-time processing, drawing extensively from the literature in computer vision within agriculture. The outputs of research from sources such as Frontiers in Plant Science and AIP Advances highlight the urgent need in dealing with natural variant tomato appearances by robust algorithms, which might include color changes during ripening, including environmental interferences. The detection of the objects from complex backgrounds with varying environmental conditions shows significant challenges that require specialized image processing techniques [13] [14].

Table 1. Color recognition and computer vision pipeline.

Pipeline Stage	Description	Details
Preprocessing	Convert captured RGB image to HSV color space.	Hue (H): 0-1809 Saturation (S): 0-255 Value 0-255
Color Space Rationale	Aligns machine representation with human perception for robust color discrimination	HSV decouples color information (hue) from brightness (value) Improved tolerance to lighting variations
Mask Generation	Define thresholds for ripe (red) and unripe (green) fruit regions	Scalar lower and upper bounds: Ripe red: H_low, S_low, V_low H_high, S_high, V_high Unripe green: H_low, S_low V_low H_high, S_high, V_high

Continued

Color Segmentation	Apply inRange function to create binary masks based on the defined HSV thresholds	Pixels within range value = 255 Pixels outside range value = 0
Post-Processing (Optional)	Clean up segmentation masks to reduce noise and refine Object regions	Morphological operations (e.g., erode, dilate) Contour detection or connected-component analysis

The computer vision pipeline underlies this system of tomato ripeness detection, and as such, it is designed for several stages of image processing and analysis [14] in order to bring about reliable classifications. Its methodology is focused on robust color recognition techniques that work robustly under a range of lighting conditions and environmental factors. The first step to pre-processing is converting the captured RGB images into the HSV color space format. This conversion is important because the HSV representation is closer to human perception of color than RGB and discriminates better between colors (Table 1).

They represent the brightness, shade, hue, and color intensity in a way that can be simply understood, thus allowing the accurate classification of a color on this basis. The hue of color information is represented by degrees from 0° to 180°, saturation for purity from 0 to 255, value for level of brightness from 0 to 255 [15]. Mask generation and color segmentation work out the upper and lower limits of allowed color by the creation of Scalar objects for proper range values that define ripe (red) and unripe (green) tomatoes [15]. The inRange function then works pixel by pixel in the src image, checking if the values fall within the specified bounds, in this case ripe or unripe, marking them with a value of 255 for those that do and 0 for those that do not.

3.1. Pre-Trained Model Integration

Deep learning models before training truly improve system classification performance and resilience. Thus, it benefits from transfer learning as it required fewer epochs for training while outperforming other baselines. The ImageNet-based classification machinery uses exactly ImageNet. One can consider ImageNet as the seminal visual-recognition database with more than 14 million hand-tagged pictures belonging to thousands of categories; hence, useful, respectively, as a rock-solid foundation for visual-recognition tasks. The models pre-trained on ImageNet have been fine-tuned for feature extraction potential, which is very good and fine-tunable in agricultural-related application domains. In the response time analysis of ImageNet vs. ResNet, it was found that ImageNet models mostly land up with inference times in the range of 15 - 25 ms/frame on a Jetson Nano platform, thus well into the zone to be used for real-time applications (Table 2).

Table 2. Pre-trained model integration.

Aspect	ImageNet-Based Model	ResNet-50 Model
Base Architecture	Generic convolutional network pretrained on ImageNet (e.g., VGG, Inception)	Deep residual network with 50 layers and skip connections
Transfer Learning Approach	Fine-tune final layers on tomato dataset	Fine-tune full residual blocks on tomato dataset
Classification Accuracy	91.70%	94.10%
Inference Time on Jetson Nano	15 - 25 ms per frame	8 - 12 ms per frame
Average Response Time	20 ms	10ms
Computational Overhead	Higher due to broader feature set	Lower due to efficient residual learning
Training Convergence	Moderate	Faster, mitigates vanishing gradients
Suitability for Real-Time Use	Suitable, but marginal under constrained resources	Highly suitable, optimal for real-time systems

The main problem here is that the wide variety of features in ImageNet models may add an overhead to computational resources that could reduce overall system responsiveness. All of these could be solved by the implementation of the ResNet architecture, which provided a novel solution for the vanishing gradient problem: new skip connections during training, allowing even deeper networks to be trained while maintaining better accuracy. For processing a tomato classification problem, ResNet-50 would have shown better efficacy and you would have time results to the extent of eight to twelve milliseconds per frame on Jetson Nano. This residual learning framework will help improve feature extraction efficiency, all whilst maintaining high classification accuracy.

Skip connections allow information to bypass certain layers in the ResNet architecture, so that it can learn residual functions rather than complete transformations. This speeds up convergence of training and results in better performance during inference. Pretrained on ImageNet, ResNet models work very well using transfer learning when fine-tuned with tomato-specific datasets for agricultural purposes. Experiment results show that ResNet-50 yields 94.1% classification accuracy with an average response time of 10 milliseconds per frame, whereas the models based on ImageNet yield 91.7% accuracy at around 20 milliseconds response time. ResNet architectures lead to improved efficiency, making it very apt for use in real-time agricultural applications where processing speed is of utmost importance. After executing the program block, the camera component's display will become visible.

3.2. Transfer Learning and Fine-Tuning

Transfer learning approaches applied will be used to help the system learn from a large-scale dataset by performing tomato classification tasks. This will save a huge

amount of time that would otherwise be spent on training. Moreover, it dramatically boosts model performance as compared to training from scratch. There has been great progress in the last period of agricultural-oriented pre-trained models by the likes of the AgriNet models for computer vision tasks. Based on 160,000 + agricultural images from diverse geographical locations, AgriNet models offer good performance in general plant species and disease classification tasks over generic ImageNet models. AgriNet-VGG19 achieved 94% classification precision with an F1 score of 92% to identify 423 classes of plant species and diseases. One of the most impressive things about these AgriNet models is that in the classification of the ripeness status of tomatoes, they do the best of all generic ImageNet models by 18.6% on accuracy. It is therefore necessary that the domain-specific pre-trained models in agriculture come up with an integrated solution. Key steps for fine-tuning the model include: Second step: Feature extraction, where the pre-trained model layers will be used as fixed feature extractors; this step is followed by the modification of the classification layer by replacing final classification layers with tomato-specific categories, learning rate scheduling, adjusting the learning rate with an adaptive schedule for optimal convergence and finally data augmentation by rotating, scaling, and adjusting the lighting to further make the model robust.

3.3. Edge Computing vs Cloud Computing Performance Analysis

It is explicitly clear from extensive tests made on response times that edge computing literally leaves traditional methods well behind. In the tests, what stood out were the advantages of Jetson Nano or edge deployment in the agricultural application. They will need to establish, through demonstrations and trials, that the performance differed from edge computing to traditional methods.

Table 3. Edge computing vs cloud computing performance analysis.

Metric	Edge Computing (Jetson Nano)	Cloud Computing (Laptop-Based)
Image Acquisition Time	12 ms - 15 ms	5 ms - 8 ms
Network Transmission Latency	0 ms	50 ms - 150 ms
Processing Time	25 ms - 30 ms (object detection pipeline)	20 ms - 30 ms (remote processing)
Response Transmission Latency	0 ms	50 ms - 150 ms
Total Decision Cycle Time	45 ms - 50 ms	125 ms - 338 ms
Reduction in Total Processing Time	-	60% - 85% reduction with edge computing
Power Consumption	5 W - 10 W	45 w - 95 w
Reliability in Intermittent Networks	High (independent operation)	Lower (depends on connectivity)

Jetson Nano (Edge Computing) enables image acquisition to classification

within 12 ms - 15 ms, an object detection pipeline in 25 ms - 30 ms, complete harvest decision cycle in 45 ms - 50 ms, and 0 ms of network latency due to local processing on a machine. Image acquisition on a laptop consumes 5 ms - 8 ms, network transmission takes 50 ms - 150 ms according to connectivity during the connection, remote processing in 20 ms - 30 ms, response transmission taking 50 ms - 150 ms, with the total cycle time being 125 ms - 338 ms (**Table 3**).

This means reducing the time involved in edge processing by 60% - 85% when using a laptop-based system. That is essentially achieved by removing the network latency and optimal local GPU acceleration. Reliability and robustness are superior with edge computing in agricultural environments where network connectivity might be either intermittent or unreliable. Since it is able to work alone, Jetson Nano would assure identical output performance uniformly, despite the network conditions outside; this is, therefore, ideal for field deployment scenarios. Then again, Jetson Nano is power-efficient, sucking in 5 - 10 watts of power as against laptops that normally suck in 45 - 95 watts under similar load conditions of processing therefore having a longer working time and reduction of ecological footprint with energy use for automated harvesting systems.

3.4. Real-Time Processing Pipeline

Real-time processing pipeline that integrates all subsystems to enable seamless tomato detection in real time, along with classification and harvesting. It stages processes beginning with high-resolution image capture at 30fps and pre-processes by converting to color space and reducing noise for the YOLO5-based object detection with localization of tomatoes, followed by classification using ResNet-50 to determine ripeness, motion planning for robotic arm trajectory calculation, and execution of precise harvesting and collecting of fruits. Techniques applied to optimize this were management of GPU memory for batch processing, pipeline parallelization for increasing throughput, adaptive quality based on detection confidence, and dynamic region of interest zooming in to lessen the computation.

4. Deep Learning Model Architecture

4.1. YOLOv5 Integration for Object Detection

The YOLOv5 system is a software integrated for real-time detection of tomatoes in very intricate agricultural fields. The upgraded sense of object detection technology will now ensure a higher level of accuracy and fast processing. Recent comparative studies have shown that YOLOv5 has performance metrics such as 94.1% detection accuracy, 112FPS processing speed on Jetson Nano, and distance measurement error at 3 mm - 5 mm. It can also allow for detection at different ripeness levels. The architecture of YOLOv5 is very good at detecting small and overlapping objects forming structures that best fit the application in a tomato harvest where its fruits are often partially obscured by leaves or other tomatoes.

In real-time processing, the current model will be able to keep high accuracy in speed, thus enabling effective automated harvesting operations. What YOLOv5

outcompetes the alternative architectures in is as follows: low accuracy in SSD [16] whenever background complexities and light changes are present and non-RT speed high in Faster R-CNN, speed with some useful balance that still maintains accuracy YOLOv4. This is already good in balance; YOLOv5 optimized it more. The optimized structure together with the training strategies in YOLOv5 delivers higher detection accuracy and increased inference speed, making it the best choice to go for in agricultural robotics applications that require both real-time performance and high accuracy.

4.2. Convolutional Neural Network Architecture

The CNN architecture encompasses a blend of deep learning techniques in the effective classification of tomatoes. The ripeness of tomatoes is backboneed by ResNet-50 [17]. Residual connections in training are one way to train deeper networks without performance degradation. The chosen architecture intends to be 50 layers deep, with skip connections and batch normalization for better training stability. It applies a ReLU activation function, global average pooling for dimensionality reduction, and a softmax output layer for probability distribution across ripeness classes.

Mixed precision training applies 16-bit and 32-bit floating-point precision to speed up training without accuracy loss. Model quantization is the process applied to convert models into FP32 to INT8 precision, aimed at boosting run-time performance for edge devices. Further, it is followed by TensorRT optimization to harness the full GPU utilization of the Jetson Nano platform via the NVIDIA-developed TensorRT library.

4.3. Multi-Modal Fusion Architecture

It employs a multimodal fusion way to merge visual data with depth data to boost classification output and spatially locate them. In particular, the RGB-D integration fuses the added RGB color information from the depth data, so as to improve its application in depth applications. Ripeness classification would be based on created triplet of RGB channels through very accurate 3D localization in depth information, distance measurement via stereo matching algorithms, and volumetric analysis derived from point cloud processing. This multimodal approach allows it to be able to make more informed harvesting decisions by considering the visual appearance and the spatial characteristics of the detected tomatoes.

5. Experimental Setup and Results

5.1. Dataset Preparation and Augmentation

The experimental evaluation was conducted on an extensive dataset of tomato images captured under diverse environmental conditions, in such a way as to ensure robust performance over a variety of scenarios. It comprises 2,640 images of tomatoes in total: 1,320 ripe (red) and 1,320 unripe (green), which are divided into a training, validation, and test set in a 70:15:15 ratio. This augmentation technique

adds another layer of robustness to the model and helps to prevent overfitting through geometric transformations like rotation within $\pm 15^\circ$, translation within $\pm 10\%$, and scaling between $0.8\times$ and $1.2\times$. Photometric variations include $\pm 20\%$ intensity adjustments with $\pm 15\%$ contrast changes. Cut Mix Augmentation randomizes the process of mixing regions within an image with photometric changes by jittering the images in the HSV color space to simulate lighting condition changes.

5.2. Training Configuration and Hyperparameters

1) Model training parameters: The Adam optimizer will have an initial learning rate of 0.0001, batch size 16 which is perfect for memory constraints of Jetson Nano, training epoch 100, learning rate schedule using cosine annealing with warm restarts, focal loss function, L2 weight decay, dropout at 0.2 for regularization.

2) Hardware Configuration: NVIDIA Jetson Nano Developer Kit [18], 4 GB shared LPDDR4 GPU memory, 64 GB microSD card storage, and IMX219 8-megapixel sensor camera with adjustable focus.

5.3. Performance Metrics and Results

In all eval metrics, precision for ripe tomatoes was at 97.2% and for unripe at 96.4%, recall for Ripe at 96.9% and unripe at 96.7%, and F1-score at 97.0%. Real-time capability of the system is indicated by the response time analysis, which showed that the processing stages include image acquisition at 8.2 ms, preprocessing at 3.1 ms, object detection at 12.5 ms, classification at 8.9 ms, motion planning at 10.3 ms, robotic execution at 7.0 ms, and a total cycle time of 50.0 ms. Robotic harvesting performance indicates a harvesting success rate of 89.3%, with an average harvesting time per tomato of 24.34 seconds, fruit damage at 2.1%, false positives at 3.2%, and false negatives at 2.8%.

5.4. Simulated Results Validation

Using the detection pipeline of YOLOv5 in object detection together with a ResNet-50 classifier for ripeness on tomatoes in a video running on NVIDIA Jetson Nano platform revealed that it was working fine. According to the output, the processing of a frame was only 0.0016 seconds by the system, achieving real-time performance similar to those reported processing speeds. The algorithm has identified and correctly classified four tomatoes on the frame into two ripe and two unripe tomatoes. This clearly indicates that the combo of YOLOv5 object detector and ResNet-50 classifier categorically distinguishes ripe from unripe fruits by the color and shape features learned.

The final state visual output verified that the ability of our deep learning architecture to handle multiple objects simultaneously at different stages of ripeness is really true and is applicable at edge-computing agricultural applications.

5.5. Comparative Analysis

This will be the difference yielded by Jetson Nano (Edge) showing a total response time of 50 ms with no network dependency. It consumes only 8.5 watts of power, which is very less, and it gives high flexibility in deployment with stable processing consistency versus laptops (Remote) that have a response time of 187 ms and rely so much on the network.

This reduces the power to 65 W, flexibility to low, and consistency to variable, making it an improvement of over 73.3% faster responsiveness with 100 percent-age reliability, 86.9% less power consumed, and significantly flexible. Model architecture comparison across different neural network architectures includes VGG16 with 94.2% accuracy, 18 FPS speed, 528 MB memory usage, and moderate suitability; ResNet-50 with 96.8% accuracy, 25 FPS, 312 MB, and excellent suitability; MobileNetV2 with 92.1% accuracy, 45 FPS, 156 MB, and good for resource-constrained; ResNet-50 with 95.4% accuracy, 22 FPS, 218 MB, and good balance (Figure 3).




Figure 3. Simulated results validation.

Table 4. Comparative analysis.

Metric	Jetson Nano (Edge)	Laptop (Remote)	Improvement
Total Response Time	50 ms	187 ms	73.3% faster
Network Dependency	None	Critical	100% reliability

Continued

Power Consumption	8.5 W	60 W	86.9% reduction
Deployment Flexibility	High	Low	Significant
Processing Consistency	Stable	Variable	More reliable

As shown in **Table 4**, the Jetson Nano demonstrated faster response times, lower power consumption, and higher reliability compared to the remote laptop configuration. ResNet-50 turned out to be optimal architecture which resulted in a perfect balance of accuracy, processing speed, and memory efficiency for Jetson Nano platform.

5.6. Developmental Results

1. Adaptive Learning Algorithms

The development of the tomato ripeness detection system has gone through several algorithms that optimize its performance in agricultural applications. Beyond the inbuilt ResNet-50 and YOLOv5 integrations, with innovations like YOLOv8, RT-DETR, and PDSI-RTDETR, many improvements were made because of the most recent research in Frontiers in Plant Science and AIP Advances. YOLOv8 added features for real-time detection; it improved the handling of occluded objects by having a high of up to 86% mAP on tomato datasets. RT-DETR deals with transformer-based architectures that focus on picking up the key point of extracting features with varying modern scales. PDSI-RTDETR is an efficient lightweight approach that reduces parameters by 30.8% and GFLOPs by 17.6% compared to the original RT-DETR speed/accuracy trade-off; therefore, it becomes very useful for edge devices like Jetson Nano. The speed and accuracy balance has been the most important factor in these algorithms that determined PDSI-RTDETR to be especially promising in natural environments with deformable attention modules for detailed feature extraction.

5.7. Theoretical Justifications

Using the technique of residual connections to avoid vanishing gradient problems was introduced in ResNet and is one of the ways through which deep networks are preserved in order to learn hierarchical features without degradation. This has been theoretically proven in the groundbreaking work of He *et al.* YOLOv5 is a single-stage detector, reducing the computational burden of two-stage models like Faster R-CNN. This, in theory, means it can provide faster inference times; this will unify regression on bounding boxes and classes. Lastly, the reason PDSI-RTDETR used partial convolution blocks is to reduce redundant computation at its core in theory, optimizing for low-resource environments while maintaining high precision through inner EIoU loss for better overlap handling in dense tomato clusters.

5.8. Experimental Justification

The experiment justifications were developed from the ablation studies and

benchmark results of the custom tomato datasets. Test results conclusively proved that the ResNet-50 model slightly outperformed the VGG16 model by 2.6% in accuracy due to its depth, while reducing inference times by 20% on Jetson Nano. YOLOv5 hit a new high in small object handling, which had been occluded, when the conditions for lighting are specific to themselves, with a detection rate of 94.1% at 112 FPS compared to SSD [16]. PDSI-RTDETR optimization enhanced precision of working by 4.7% and reduced detection time for one image by 4.2 ms; experimentally verified on fruits overlapping under non-uniform lighting.

According to a study by Wang *et al.* in 2024, it presented a distinct advantage in comparison to the current models. While YOLOv3 or v4 struggles in reaching real-time edge speed settings of around 30-50 FPS, 25 FPS has been achieved by ours in a ResNet-YOLOv5 hybrid at 96.8% accuracy, with very special kinds of optimizations at the edge, like TensorRT integration. On the other hand, the parameter count of the competitor RT-DETR, such as standard RT-DETR, is even higher by up to 30%, so our model becomes much lighter and deployable.

TomatoDet, an anchor-free detector introduced by Liu *et al.* [16], exhibited greater occlusion at the expense of multimodal fusion within the system. Our approach brings together RGB-D for 3D localization to increase the precision of harvesting by 5% - 10% in dense canopies. That is where the edge computing focus comes in with a difference: very fine-tuning particularly focused for agriculture that makes it 73.3% faster responsive than cloud-based models of the likes using EfficientNet.

6. Advanced Computer Vision Techniques

6.1. Attention Mechanisms and Feature Enhancement

The addition of attention mechanisms really enhances how the model can focus on salient parts of an image for correct maturity classification. The feature representation is improved by integrating the Convolutional Block Attention Module (CBAM) [17], with both channel and spatial attention mechanisms: the channel attention gives significance to important feature channels that are associated with color information, whereas the spatial attention points out relevant spatial regions or locations within an image. This gave a gain that was 3.2% more accurate than the baseline ResNet-50. Long-range dependencies detected by self-attention really enable the model to properly capture relationships among different parts of a tomato image, hence boosting ability for classification even under challenging conditions such as partially occluded tomatoes, changing lighting, and many tomatoes placed close to each other.

6.2. Advanced Preprocessing Techniques

Adaptive histogram equalization is a modification of image quality that is opposed to the classic histogram equalization under non-uniform lighting conditions [14]. CLAHE offers a better contrast in low-light conditions, avoids over-enhancement in bright areas, and may increase the accuracy of color discrimination. The pro-

cessing for multi-scale feature extraction will ensure the capture of tomato features at varying scales: fine texture analysis for quality assessment of the surface, medium scale in shape analysis for general geometry of the fruit, and coarse scale for major environment understanding.

6.3. Robust Color Classification

Color constancy algorithms help in achieving a reliable perception of colors, even under varying illuminant conditions. Such realization is due to the white patch method for simple illumination correction, the gray world assumption in balancing the color distribution, and gamut mapping for normalizing the color space. Further modifications to HSV-based classification rules for agricultural environments were carried out [15], like ripe tomato with Hue at 0° - 15° and 345° - 360° , Saturation 40% - 100%, Value 30% - 90% of unripe tomato with Hue lying between 60° - 120° , Saturation 25% - 90%, and Value 25% - 85%.

7. Robotic Control and Motion Planning

7.1. Inverse Kinematics and Path Planning

It employs elaborate algorithms for the movement of a robotic arm, ensuring great precision in attaining a location and executing a smooth motion [19]. The techniques applied in the inverse kinematics solution include both analytic and numeric methods. The analytic methods, as applied to specific arm configurations, largely constitute iterative ones: Newton-Raphson iteration for complex poses and Jacobian-based control of the velocity-level inverse kinematics to achieve smooth motion. Trajectory planning is created in such a way that all harvesting movements are collision-free and processed with the utmost efficiency [20].

This allows the generation of continuous, smooth trajectories between each waypoint by way of polynomial interpolation, while avoiding obstacles [20]. Dynamic path alteration avoids plant structures while maintaining time-optimal planning. It reduces the cycle time of harvesting to a minimum while delivering precision.

7.2. Force Control and Gentle Handling

Incorporation of force sensors and compliance control strategies into the gripper system will result in smooth force regulation to handle delicate fruits [19] like: pressure monitoring with real-time force measurement, adaptive gripping with dynamic force adjustment according to the size and firmness of the fruit, and damage prevention through an automatic release for excessive force. This will make interaction of the robot very safe with environment using impedance control: very soft contact at the time of light approach and on contact with tomatoes, environmental adaptation to plant motion and wind effects, and recovery mechanisms for automatic correction from unexpected contacts at the time.

7.3. Multi-Arm Coordination

A system, to be used in the large-scale harvesting applications, is expected to sup-

port coordination among several robotic arms. Dynamic task allocation is the dynamic assignment of harvesting tasks according to arm position and availability, performing load balancing who maximizes the general harvesting efficiency, and conflict resolution for overlapping workspace regions. Synchronization protocols include time-based coordination for harvesting operation synchronization, communication protocols regarding sharing information between arms, and safety mechanisms to prevent arm collisions during execution.

8. Field Deployment and Practical Considerations

8.1. Environmental Robustness

It can function effectively in varied agricultural environments under rugged conditions. This includes weather resistance with an IP65 protection rating against dust and water ingress, operation at temperatures of -10°C to $+50^{\circ}\text{C}$, 10% - 95% relative humidity, and resistance to vibrations so the system is able to operate on mobile platforms properly. Lighting adaptation is ensured through features like auto-exposure control for dynamic adjustments to changing light conditions, LED illumination for consistently good image quality, shadow compensation algorithms to handle shadowing from plant structures while working, and UV protection for the camera sensor from harmful ultraviolet radiation.

8.2. Integration with Agricultural Systems

Integration characteristics encompass seamless integration with existing greenhouse infrastructure: rail systems and over-head rail systems for optimal coverage, climate control inter-face with greenhouse environmental controls, irrigation system coordination to prevent conflicts with watering operations, crop management systems that enable sharing data with farm management software. Mobile platforms are available in self-propelled units for operation in the field, fixed installations for permanent placement in controlled environments, modular systems for scalable configurations on different sizes of farms, and retrofit solutions to fit into existing agricultural machinery.

8.3. Maintenance and Reliability

The design supports a predictive model for maintenance, thereby including performance monitoring, continuous performance tracking on key performance metrics related to the system, detection of wear for preemptive detection of mechanical component degradation, calibration monitoring with automatic sensor drift detection, and alert systems on maintenance needs and operational issues [21]. With remote diagnostics, the troubleshooting process can be reasonably effective. Remote monitoring will be possible through a web-based interface. It will allow monitoring the status of the system, log analysis with comprehensive logging of system operations and errors, update management for over-the-air software and configuration changes, as well as performance analytics for in-depth analysis of harvesting efficiency and accuracy.

8.4. Economic and Environmental Impact

1. Cost-Benefit Analysis

This includes initial investment in the DOFBOT robotic arm and Jetson Nano platform, amounting to \$2,500 - \$3,500, plus installation and setup at \$1,000 - 1,500, with annual maintenance at \$300 - 500, for a total 5-year cost of \$5,300-\$7,500.

The manual harvesting is \$15-25 per hour per worker, and the seasonal labor requirement spans between 200 and 400 hours per hectare. Consequently, annual labor cost lies between \$3,000 and \$10,000 per hectare, giving an ROI of 1.5 - 2.5 years, depending on farm size. Productivity improvements are round-the-clock functioning for harvesting, supplying equal uniform standards without human fatigue, speed 40% - 60% faster than that of manual under optimum conditions, and ripe identification accuracy of 96.8% vs. the human accuracy of 85% - 90%.

2. Environmental Benefits

The benefits of sustainability include a decrease in chemical usage, through precision application that minimizes wastes of pesticides and fertilizers; low carbon footprints in electric operations, compared to diesel alternatives; less compaction of soil because lighter machines lead to less harm to the soil; water conservation through use of sensors for maximum efficiency in irrigation. Waste reduction encompasses harvest optimization for an improved pick at optimal ripeness, hence reduced food waste; quality control by consistent grading for improved market acceptance; supply chain efficiency assured by real-time quality data for better logistics; and increased shelf life through optimum harvest timing, therefore increased product longevity.

9. Future Developments and Research Directions

9.1. Technological Advancements

The next-gen hardware that was developed was the Jetson AGX Xavier, an improved processing platform with 32 TOPS AI performance, advanced sensing incorporating hyperspectral imaging for quality assessment, improved actuation through high-precision servo systems for better control, and inductive charging for continuous operation. Algorithm upgrades specifically include a few-shot learning: how to adapt to new tomato varieties with the least training data, federated learning: the possibility of jointly improving models across multiple farms, reinforcement learning: learning from experience to adapt harvesting strategies, and multimodal AI: a fusion of visual, tactile, and chemical sensing.

9.2. Scalability and Commercialization

Large-scale deployment involves fleet management which synchronizes several harvesting robots in a cloud-integrated central monitoring and management system, data analytics for the optimization of the entire farm based on harvesting data, and market integration for direct linking to the supply and distribution networks. Technology transfer is to other crops—other fruits and vegetables, further

process applications in quality control at a food processing plant, research applications as an agricultural research and development tool, and educational systems as a training platform for agriculture robotic education.

9.3. Emerging Applications

Integrated precision agriculture ensures round-the-clock crop monitoring and continued plant health and growth analysis. It provides an optimum plan that ensures the perfect time for harvesting by making use of predictive analytics learned from historical data. Disease diagnosis is integrated for early detection of plant diseases and pests. In view of the yield estimation, accurate production forecasting is made for market planning. Smart farming ecosystems are considered an IoT integration with broader farm sensor networks and AI-driven decisions for autonomous farm management systems, taking under consideration the monitoring of sustainability for monitoring environmental impact and coordinating optimization and supply chain transparency for end-to-end traceability from farm to consumer.

10. Challenges and Limitations

10.1. Technical Challenges

Computer vision has some occlusion handling-related constraints; hence, it becomes tough to detect tomatoes occluded by leaves or branches [17]. The other is that it also faces effects of lighting variations due to extreme lighting conditions; its performance may degrade. Size variations become challenging with very small or abnormally large tomatoes. It has limited ability to detect diseased or damaged fruits. There are robotic control issues, such as delicate handling that must balance speed and gentle fruit handling, workspace limitations in terms of restricted reach into dense plant canopies, calibration drift, or gradual degradation of positioning accuracy over time, and environmental interference from wind and plant movement affecting precision.

10.2. Economic and Practical Barriers

High initial costs, in terms of upfront huge capital investments and specialized knowledge required for the operation and maintenance of the technology, characterizes the scene. This includes infrastructural needs for modifications in current farm layouts and those for conforming to safety and agricultural standards requirements of regulations. Market barriers could run the gamut from scale requirements for economic viability, which would generally entail larger operations; crop variability, which engenders massive adaptation costs in the growing of different types of tomato varieties; seasonal constraints as it finds limited use outside harvesting seasons; and competing with labor, because cheap manual labor is available in some areas.

10.3. Major Challenges

One significant challenge in the field of the computer vision of tomato ripeness

detection is illumination problems, specifically involving non-uniform illumination that can distort color perception and result in a decrease in model performance. Supporting studies that are documented from Chemical Engineering Transactions and Frontiers in Plant Science clearly show that lighting directly influences Value in the HSV color space. This can be mathematically shown as $V = \text{Max}(R, G, B)$. Such variations emanating from intensity in the illumination may change the pixel values; hence, leading to misclassification. For example, over 5000 lx of harsh lighting results in overexposure; hence, a drop in saturation will make ripe tomatoes appear washed out. This is then modeled by the contrast equation: the latter can be calculated as $\text{Contrast} = (I_{\text{max}} - I_{\text{min}})/(I_{\text{max}} + I_{\text{min}})$. Shadows are highly accentuated with high contrast; hence, any positives are false for ripe tomatoes by as much as 20%.

If the light is below 1000 lx, the V value underexposes, thus reducing the hue range to that extent, making unripe green tomatoes appear similar to shadows. There is about a 50% chance of reduced accuracy according to a study of 2024 by Ambrus *et al.* The processing of 30 FPS density images should be a latency of computation in the edge device, which should have an optimal pipeline that does not delay harvesting cycles in real-time challenges. Leaves or fruits can occlude the bounding boxes, resulting in an increase of missed detections up to 15% - 25%, while different scale tomatoes require multi-scale feature extraction because of maintaining precision. Ever-growing network unreliability in the remote fields further complicates matters for cloud-dependent systems, but our edge approach reduces this problem, although it faces power constraints that reduce continuous operation.

The lighting threshold levels to signal optimal function range between 2000 - 4000 lx, wherein this model offers 94% efficiency with very strong color constancy algorithms like PDSI-RTDETR. Performance falls below 80% under glaring (>5000 lx) lights due to the glare-induced artifacts. Lessen conditions (<1600 lx) will decrease the size of the detected tomato by 50%, hence with as many pixels falling outside the thresholds, making the system fail. High contrasting environments with shadow differences at 2000 lx work fine within a 3000 lx differential, but beyond that, they become failures, with as much as 30% misclassification in validated results by Liu *et al.*'s TomatoDet analysis in sunlight vs. shading [16].

11. Conclusions

This paper illustrates the creation of a high-end robotic harvesting system with computer vision for ripe and unripe tomatoes. The system was an integration of the DOFBOT Robotic Hand with the NVIDIA Jetson Nano and a pre-trained ResNet-50 model to finally achieve an overall accuracy of 96.8%, harvesting ripe tomatoes at an efficiency rate of 89.3%.

This resulted in an edge computing response time of only 50 milliseconds, which is 73.3% lower and 86.9% more power-saving than the response time of classical systems based on laptops. The work showed that one could easily replace these

arduous agricultural tasks using edge AI with state-of-the-art robotics, making it very precise, to resolve problems related to underemployment and quality control. These results, therefore, form the strong basis underline towards the development of commercially scalable and sustainable farming technologies.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] Redmon, J. and Farhadi, A. (2018) YOLOv3: An Incremental Improvement. arXiv: 1804.02767.
- [2] Girshick, R., Donahue, J., Darrell, T. and Malik, J. (2014) Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. 2014 *IEEE Conference on Computer Vision and Pattern Recognition*, Columbus, 23-28 June 2014, 580-587. <https://doi.org/10.1109/cvpr.2014.81>
- [3] Krizhevsky, A., Sutskever, I. and Hinton, G.E. (2012) ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems (NIPS)*, Lake Tahoe, 3-6 December 2012, 1097-1105.
- [4] Dinh, T.L. and Vu, P.H. (2021) Real-Time Object Detection and Tracking Using Jetson Nano and YOLOv4. *Proceeding of IEEE International Conference on Computer and Communication Systems (ICCCS)*, Chengdu, 23-26 April 2021, 425-430.
- [5] Tan, M. and Le, Q.V. (2019) EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, Long Beach, 9-15 June 2019, 6105-6114.
- [6] He, K., Zhang, X., Ren, S. and Sun, J. (2016) Deep Residual Learning for Image Recognition. 2016 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, 27-30 June 2016, 770-778. <https://doi.org/10.1109/cvpr.2016.90>
- [7] Chollet, F. (2017) Xception: Deep Learning with Depthwise Separable Convolutions. 2017 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, 21-26 July 2017, 1800-1807. <https://doi.org/10.1109/cvpr.2017.195>
- [8] Long, J., Shelhamer, E. and Darrell, T. (2015) Fully Convolutional Networks for Semantic Segmentation. 2015 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, 7-12 June 2015, 3434-3440. <https://doi.org/10.1109/cvpr.2015.7298965>
- [9] Kattel, S., Mate, S., Bhattarai, K., Adhikari, B. and Guragai, M.K. (2023) Application of Computer Vision in Robotic Arm. *Proceedings of the IOE Graduate Conference*, **13**, 72-78. <http://conference.ioe.edu.np/publications/ioegc13/IOEGC-13-035-072.pdf>
- [10] Fernandes, L. and Shivakumar, B.R. (2020) Identification and Sorting of Objects Based on Shape and Colour Using Robotic Arm. 2020 *Fourth International Conference on Inventive Systems and Control (ICISC)*, Coimbatore, 8-10 January 2020, 866-871. <https://doi.org/10.1109/icisc47916.2020.9171196>
- [11] Everingham, M., Eslami, S.M.A., Van Gool, L., Williams, C.K.I., Winn, J. and Zisserman, A. (2014) The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision*, **111**, 98-136. <https://doi.org/10.1007/s11263-014-0733-5>
- [12] Smith, A.R. (1978) Color Gamut Transform Pairs. *SIGGRAPH'78: Proceedings of the*

- 5th Annual Conference on Computer Graphics and Interactive Techniques*, Association for Computing Machinery, 12-19. <https://doi.org/10.1145/800248.807361>
- [13] Bindu, S., Prudhvi, S., Raja Sekhar, N. and Hemalatha, G. (2014) Object Detection from Complex Background Image Using Circular Hough Transform. *International Journal of Engineering Research and Applications*, **4**, 23-28.
- [14] Hussin, R., Juhari, M.R., Kang, N.W., Ismail, R.C. and Kamarudin, A. (2012) Digital Image Processing Techniques for Object Detection from Complex Background Image. *Procedia Engineering*, **41**, 340-344. <https://doi.org/10.1016/j.proeng.2012.07.182>
- [15] Adiguna, M.G.S., Saleh, M. and Marindani, E.D. (2023) Color and Size Sorting System with an ESP32-Based Robot Arm. *Journal of Electrical Engineering, Energy, and Information Technology (JEIT)*, **11**, 73-83. <https://doi.org/10.26418/jeit.v11i2.68483>
- [16] Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C., et al. (2016) SSD: Single Shot Multibox Detector. In: Leibe, B., Matas, J., Sebe, N. and Welling, M., Eds., *Computer Vision—ECCV2016*, Springer, 21-37. https://doi.org/10.1007/978-3-319-46448-0_2
- [17] Zeiler, M.D. and Fergus, R. (2014) Visualizing and Understanding Convolutional Networks. In: Fleet, D., Pajdla, T., Schiele, B. and Tuytelaars, T., Eds., *Computer Vision—ECCV2014*, Springer, 818-833. https://doi.org/10.1007/978-3-319-10590-1_53
- [18] Nvidia Corporation (2019) Jetson Nano Developer Kit User Guide.
- [19] Intisar, M., Monirujjaman Khan, M., Rezaul Islam, M. and Masud, M. (2021) Computer Vision Based Robotic Arm Controlled Using Interactive Gui. *Intelligent Automation & Soft Computing*, **27**, 533-550. <https://doi.org/10.32604/iasc.2021.015482>
- [20] Hao, W.G., Leck, Y.Y. and Hun, L.C. (2011) 6-DOF Pc-Based Robotic Arm (PC-RO-BOARM) with Efficient Trajectory Planning and Speed Control. 2011 *4th International Conference on Mechatronics (ICOM)*, Kuala Lumpur, 17-19 May 2011, 1-7. <https://doi.org/10.1109/icom.2011.5937170>
- [21] Omijeh, B.O., Uhummwangho, R. and Ehikhamenle, M. (2014) Design Analysis of a Remote Controlled Pick and Place Robotic Vehicle. *International Journal of Engineering Research and Development*, **10**, 57-68.