

# Automated Fruit Ripeness Classification in Robotic Harvesting Using a GAN-Augmented Deep Learning Framework

Hakim Ayesha<sup>1\*</sup>, Saleem Kanwal<sup>2</sup>

<sup>1</sup> School of Electrical Engineering and Computer Science (SEecs), National University of Sciences and Technology (NUST), 44000, Islamabad, Pakistan, [ayesha.hakim@seecs.edu.pk](mailto:ayesha.hakim@seecs.edu.pk)

<sup>2</sup> Institute of Computing, MNS University of Agriculture, 60000, Multan, Pakistan, [2022-uam-2887@mnsuam.edu.pk](mailto:2022-uam-2887@mnsuam.edu.pk)

Manual fruit harvesting is labor-intensive and inconsistent, limiting efficiency in large-scale agricultural operations. This study proposes a Generative Adversarial Networks (GAN)-augmented dataset integrated with a deep learning framework for automated fruit ripeness classification. Four commercially important fruits, mango, strawberry, tomato, and sweet pepper, were classified into ripe and unripe categories. As collecting large numbers of labeled fruit images is difficult due to seasonal, environmental, and labor constraints, GANs were employed to expand the dataset by creating realistic fruit images under diverse conditions. DenseNet201, ResNet50, and a CNN were trained and evaluated, with DenseNet201 achieving the best performance (99.41% training, 98.01% validation, and 95.7% testing accuracy). GAN-based augmentation improved generalization under occlusion and illumination variability. The results confirm that DenseNet201 trained on a GAN-augmented dataset provides a reliable framework for automated ripeness classification, supporting the development of robotic harvesting systems.

**Keywords.** Transfer Learning, Synthetic Image Generation, Agricultural Automation, Precision Agriculture, Computer Vision

## Introduction

Fruit harvesting is one of the most labor-intensive operations in agriculture, with costs and inefficiencies driven by reliance on manual methods. Leaf occlusion, uneven illumination, and variable ripeness further complicate yield estimation and harvesting (Liu *et al.*, 2025). Deep learning and computer vision have advanced fruit detection and classification, with CNNs such as VGG16, ResNet, and EfficientNet achieving accuracies above 95% (Tomas *et al.*, 2024). Object detection frameworks like YOLO and Mask R-CNN have been widely applied for localization, while transformer models are emerging for ripeness detection (Xiao *et al.*, 2023).

A persistent limitation is the availability of large, balanced datasets. Field-based data collection is costly and seasonal, often leading to overfitting. Generative Adversarial Networks can address this by producing realistic synthetic fruit images across ripeness stages and conditions, improving training diversity (Lu *et al.*, 2022). Building on this, we propose a GAN-augmented DenseNet201 framework for multi-fruit ripeness classification, benchmarked against ResNet50 and CNN. To our knowledge, this is the first study applying GAN-augmented DenseNet201 to classify multiple fruit species under field-like conditions.

## Materials and Methods

A hybrid dataset was created using field-captured images from orchards in Multan, Pakistan, and publicly available datasets. It included 3,988 real images of mango, strawberry, tomato, and sweet pepper. GANs generated 53,351 synthetic images, producing a final dataset of 57,339 samples. Images were resized to 224×224 pixels, normalized to [0,1], and split into 70% training, 20% validation, and 10% testing sets.

DenseNet201, ResNet50, and a CNN were implemented for classification. DenseNet201 was chosen for its densely connected layers, ResNet50 for skip connections, and CNN as a lightweight baseline. Transfer learning with ImageNet weights was applied to DenseNet201 and ResNet50, with additional dense layers for binary classification. All models were trained with Adam optimizer ( $\text{lr}=0.001$ ), categorical cross-entropy loss, ReLU activations, Softmax output, dropout=0.2, and early stopping over 50 epochs. Performance was evaluated using accuracy, precision, recall, and F1-score.

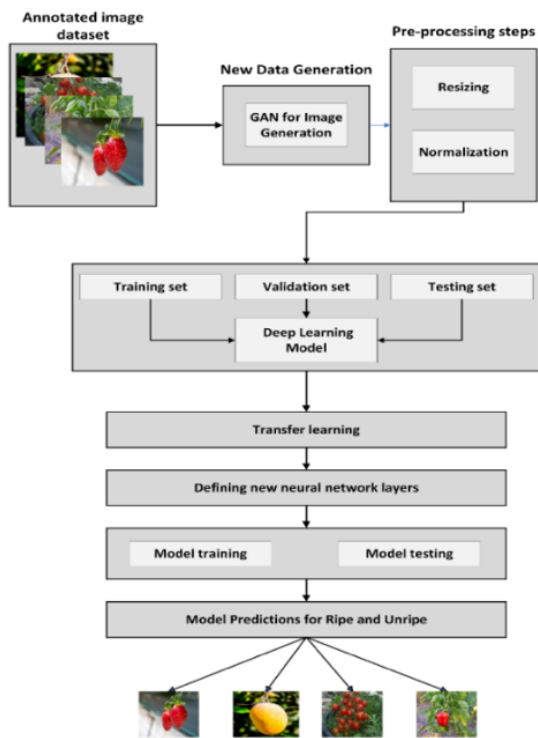


Figure 1. Flow chart of the proposed framework for classification of fruits into ripe and unripe stages

## Results and Discussion

The performance of DenseNet201, ResNet50, and the custom CNN was evaluated on the hybrid dataset containing both original and GAN-generated images. DenseNet201 outperformed ResNet50 and CNN across all metrics. It achieved 99.41% training accuracy, 98.01% validation accuracy, and 95.7% testing accuracy, with precision, recall, and F1-scores above 95%. ResNet50 reached 90% test accuracy, while CNN lagged at 81.6%. GAN-based augmentation proved critical, improving generalization under occlusion and lighting variability. Training DenseNet201 without early stopping led to overfitting, with poor test accuracy despite high training performance. Compared with prior approaches such as VGG16 and transformer models (Tomas *et al.*, 2024; Xiao *et al.*, 2023), the proposed GAN-augmented DenseNet201 provided competitive accuracy with lower annotation requirements. The results, as shown in Table 1, confirm its effectiveness as a robust framework for multi-fruit ripeness classification in robotic harvesting.

GAN-based augmentation significantly improved the model's ability to generalize under conditions of occlusion and varying illumination. DenseNet201 combined with early stopping effectively prevented overfitting, whereas training without early stopping led to poor generalization despite high training accuracy. Compared with prior works such as YOLO and transformer-based models, the proposed approach achieved higher classification accuracy while requiring fewer annotation efforts (Xiao *et al.*, 2023).

Table 1. Comparative evaluation of four deep learning models DenseNet201(with and without early stopping), ResNet50, and CNN

Metric (%)	DenseNet 201 (early stopping)	ResNet50	CNN	DenseNet201 (no early stopping)
Training acc	<b>99.41</b>	97.57	99.34	99.81
Validation acc	<b>98.01</b>	82.33	79.31	99.67
Testing acc	<b>95.70</b>	90.00	81.60	33.00
Precision	<b>97</b>	89	84	31
Recall	<b>96</b>	85	79	43
F1-score	<b>95</b>	87	81	32
Training loss (0-1)	0.02	0.08	0.04	0.03
Validation loss (0-1)	0.02	0.32	0.79	0.00

## Conclusions

GAN-augmented DenseNet201 provides a robust and accurate framework for automated fruit ripeness classification. Tested on mango, strawberry, tomato, and sweet pepper, it achieved superior results compared to ResNet50 and CNN. GAN-based augmentation proved essential for improving generalization, and DenseNet201 with early stopping offered the most reliable performance. Future work will extend the framework to more fruit species and integrate real-time detection for robotic harvesting.

## References:

- Liu Q, Zhao Y, Chen D, Huang Z. 2025. Advances in object detection and localization techniques for fruit harvesting robots. *Agronomy* 15(1): 145.
- Lu Y, Chen D, Olaniyi E, Huang Y. 2022. GANs for image augmentation in agriculture. *Comput Electron Agric* 200: 107208.
- Tomas MC, Beltran AJ, Aranez YE, Britanico E. 2024. Tomato ripeness classification using VGG16 CNN. *Telecommun Comput Electron Control* 22: 165–171.
- Xiao B, Nguyen M, Yan WQ. 2023. Fruit ripeness identification using transformers. *Applied Intelligence* 53: 22488–22499.