International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 12 Issue: 04 | Apr 2025 www.irjet.net p-ISSN: 2395-0072

# Optimized Fruit Quality Prediction using CV and GANs

**Ashish B Rao[1], Dr. Tejaswini R Murgod [2], Loukya Harisha[3] Ananya Nittur[4]**

*[1,3,4]Student, Department of AIML, BNM Institute of Technology, Bangalore, Karnataka, India*
*[2]Professor, Department of AIML, BNM Institute of Technology, Bangalore, Karnataka, India*

---------------------------------------------------------------------***---------------------------------------------------------------------

**Abstract** - *A modern approach aimed at transforming supply chain and agricultural operations is Optimized Fruit Quality Prediction using GANs and Computer Vision. This system offers precise fruit quality evaluations by combining Generative Adversarial Networks with cutting-edge computer vision, making sure of increased market prices due to better quality and reducing harvest processing losses. An accurate CNN-based fruit quality classifier, a GAN-based synthetic data augmentation model, and a natural graphical user interface (GUI) for looking at quality indicators and predictions are some of the key features. These key factors improve the agricultural ecosystem's sustainability and productivity by giving stakeholders practical insights.*

***Key Words* — Fruit Quality Prediction, GANs, Computer Vision, Agricultural Technology, Data Augmentation, CNN, Supply Chain Efficiency.**

## 1. INTRODUCTION

Fruit quality evaluation and shelf-life prediction remain major issues for agri-food supply chains, with direct consequences for farmers' profitability and minimization of food wastage. The use of subjective, time-consuming, and variable inspection procedures results in the early withdrawal of edible products or delayed detection of spoilage, causing severe economic losses and sustainability problems along the value chain. For perishable fruits like bananas, where the ripening process plays a major role in their market value, even minor mistakes in assessing quality can lead to significant financial losses. In fact, post-harvest losses in developing countries are estimated to range between 30% and 40%.[8]

Existing technologies suffer a compromise between f ield use and accuracy,while heterogeneity in environments a nd insufficient data destabilize systems that are computerize d.Traditional methods based on manual sorting or simple computer vision techniques are incapable of identifying the complex biochemical processes in fruit maturation, while laboratory-grade instruments are not feasible for implementation in agricultural environments.[10] These issues highlight the need for efficient, scalable solutions that can be implemented in real-world environments with low infrastructural requirements.

This article proposes a technology-based solution for such issues utilizing computer vision and deep learning techniques. By utilizing a combination of multi-level quality estimation and simulation of temporal degradation, the system has the ability to forecast not only the current state of quality but also future possible shelf-life trajectories. Through the application of image-based quality estimation along with predictive shelf-life modeling, the system offers practical insights to farmers to facilitate well-informed decisions about harvesting, storage, and distribution, thus achieving optimal supply chain management.

The method of interest has its roots in measuring banana ripeness as a demonstrative example of how accurate technological measurement can coexist with laboratory precision to actual usage for agriculture. With this system, the goal is to establish a new benchmark for functional, AI-driven quality control that unites scientific excellence with operational feasibility for a broad scope of agricultural stakeholders.

## 2. LITERATURE REVIEW

Deep learning has been extensively explored for improving agricultural processes, particularly in small-scale agriculture.[1][5][10] For instance, V. Zárate and D. C. Hernández (2024) highlight the use of lightweight deep models to assess fruit quality within small-scale agricultural settings with limited computational resources. Scalability remains a problem owing to the constraints of lightweight models. Similarly, M. Zabovnik and D. Wojcieszak (2024) report on the applications of Convolutional Neural Networks (CNNs) in smart agriculture, which include multispectral images and vegetation indices, while their study is predominantly a review without experimental verification. J. Li and P. Kumar (2024) introduced ensemble deep learning architectures for cotton crop classification in AI-based smart agriculture that enhanced the decision-making process through image analysis. Their method does use significant quantities of computational power, which could restrict it from being used in resource-limited environments.

There have also been some recent advances in deep learning that have involved transfer learning and image augmentation. E. Martinez and F. Lopez (2024) utilize transfer learning to classify fruits such that the training time taken was reduced while the accuracy was increased. However, domain adaptation remains a problem, affecting the model's generalizability to different types of fruits and conditions. In addition, image augmentation techniques have been explored for robust fruit detection using deep learning. While these techniques improve model robustness and generalizability, their performance may vary with datasets and conditions.

Generative Adversarial Networks (GANs) have also been explored for agricultural applications.[4][16] Nguyen and H. Tran (2024) talk about the capacity of GANs to generate artificial data to enhance dataset diversity and robust model strength. However, their applicability is limited by the lack of empirical evidence and implementation challenges. In yet another approach, K. Patel and M. Singh (2024) present an application of real-time plant health monitoring using computer vision and machine learning, facilitating disease detection at an early stage and plant health assessment. While successful, performance is influenced by various environmental factors, and scalability to large agricultural plots is not considered.

Deep learning has been employed to detect defects and determine freshness in vegetable quality inspection. R. Sharma and P. Kumar (2024) describe the application of ensemble deep learning models for this purpose, but the study does not account for variations in vegetable appearances due to variations in growth conditions, thereby compromising generalizability. On the other hand, soft computing methods have been explored to analyze agricultural data, as presented by S. Gutiérrez and J. Tandayula (2023). Their research, however, omits particular methodology descriptions as well as discussions regarding the limitations of the said techniques.

Overall, while deep learning and advanced computing methods continue to revolutionize agricultural applications, scalability, domain adaptation, generalizability over environmental conditions, and the absence of empirical backing are all relevant challenges that need to be met. Future work must take cognizance of filling these gaps towards enhancing real-world practical utilization of such technology in actual agricultural settings.

Extensive studies have investigated artificial intelligence (AI) and computer vision for fruit quality evaluation to improve agricultural productivity and minimize losses. This review integrates major findings from recent methods and highlights essential gaps to guide the development of effective AI-based fruit quality evaluation solutions.

## 3. METHODOLOGY

### 3.1 Dataset Details

The dataset that is being used is a selected collection of high-resolution images of the fruits which are categorized into 4 distinct sets. The 4 sets being the stages of ripeness of the fruit - unripe, ripe, overripe and rotten. The dataset has a wide range of images including the images captured under varied lighting conditions, different angles of capture and vividly backgrounded images to ensure generalizing. Image augmentation techniques such as rotation, scaling, and color adjustments were applied to improve the model's ability to handle all the variations of the real world. The dataset comprises of 10,000 images, and 2,500 images are present for each stage of ripeness.

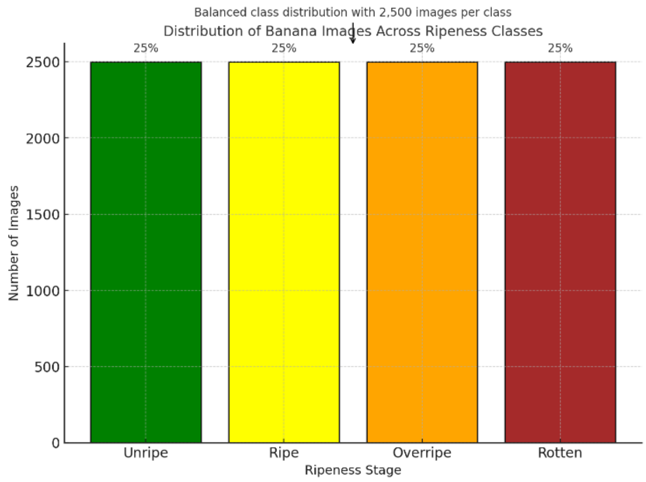


Fig 1: Data Distribution after preprocessing and cleaning

Preprocessing the dataset included various tasks such as resizing all the images to maintain consistency with 224x224 pixels resolution.[18][1] Preprocessing also included tasks for normalizing the pixel value to fit into the range [0,1], and further applied data augmentation technique to artificially increase the diversity of the training dataset

### 3.2 Computer Vision Fundamentals

Computer vision allows machines to process visual information by extracting useful information from images or video. In agriculture, it substitutes subjective human inspection with measurable parameters for quality determination. The technology is based on digital image processing methods for feature extraction (color, texture, shape) and pattern recognition software. In fruit quality inspection, computer vision systems examine surface features, color gradients, and morphological features of ripeness stages and defects, offering homogeneous and measurable standards for large amounts of produce.

### 3.3 Convolutional Neural Networks (CNNs)

Convolutional Neural Networks is a specialized deep learning architecture type designed for visual data processing that have inherent grid-like structures, i.e., digital images. Convolutional Neural Networks use a hierarchical design that learns increasingly more complex visual features automatically in succession using consecutive convolutional layers. The earlier layers usually extract simple visual features such as edges and textures, and deeper layers detect more abstracted patterns and object features. This architectural theme makes CNNs very well adapted for fruit quality classification because CNNs can be applied directly to raw pixel data without the need for manual feature engineering. A notable strength is their capacity for preserving spatial information in the image in the course of feature extraction while, at the same time, learning stable representations that are invariant to usual variations in the image capture condition, e.g., lighting, orientation, or positioning variations typical of agricultural imagery.

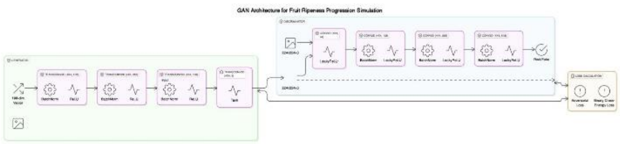

Fig 2: GAN Architecture for Fruit Ripeness Progression Simulation

### 3.4 Generative Adversarial Networks (GANs)

Generative Adversarial Networks consist of two neural networks in a competitive learning process - a generator that generates synthetic images and a discriminator that evaluates their realism. Through adversarial training, the generator continuously improves its ability to generate realistic samples that closely approximate the statistical distribution of the training data. In the context of this fruit quality assessment system, GANs serve two significant purposes. First, they address data scarcity issues by generating realistic variations of fruit images at different stages of ripeness, effectively augmenting the training set. Second, they enable shelf-life estimation through modeling the temporal progression of fruit spoilage. The adversarial training process ensures generated images are biologically plausible while covering edge cases that might be sparse in the original dataset, thus enhancing the model's robustness.

### 3.5 Transfer Learning Principles

Transfer learning relies on the assumption of capitalizing on the knowledge gained through the solution of one problem space to improve performance in a similar but distinct space, particularly where training data would be scarce. The approach succeeds because of the high transferability of the intrinsic visual features learned by the shallow layers of convolutional networks—like edge detection and texture extraction—across a wide variety of image classification tasks. Transfer learning is currently leveraged in the implementation through the initialization of the model parameters from a pre-trained ResNet50V2 model on the vast ImageNet dataset.[2] The strategy involves freezing the shallow layers that are involved in the extraction of these general visual features and selectively fine-tuning only the deeper layers that are involved in learning the subject-specific features of fruit ripeness patterns. This approach has a significant advantage by substantially reducing the computational resources required for training as well as the necessary domain-specific data, while concurrently improving the model's ability to generalize to real-world agricultural environments.

### 3.6 ResNet50V2 Architecture

ResNet50V2 architecture is a complex 50-layer deep convolutional neural network that extends from the basic ResNet framework with enhanced residual connections and feature propagation. The network starts with the first 7×7 convolutional layer with 64 filters and a stride of 2, and then the 3×3 max pooling layer with a stride of 2 that decreases the input size but increases the depth of the channel. This pre-processing step leads to the entire architecture divided into four separate blocks of convolutional with growing sizes of features.

The first convolutional block comprises 3 residual units with 64, 64, and 256 filters respectively, with identical spatial size but building the core feature representations. The second block comprises 4 residual units with 128, 128, and 512 filters, where the first spatial down sampling is done to capture more abstracted patterns. The network continues with the 6 residual units of the third block with 256, 256, and 1024 filters, where the model builds higher-level feature detectors. The last convolutional block comprises 3 residual units with 512, 512, and 2048 filters, where the most advanced pattern recognition abilities are sharpened.

Every residual unit adheres to the enhanced bottleneck architecture characteristic of ResNetV2: initial 1×1 convolution for dimension reduction, followed by a 3×3 convolution for spatial feature extraction, and final 1×1 convolution for dimension restoration. Pre-activation architecture positions batch normalization and ReLU activation ahead of every convolutional operation, providing more unobstructed paths for backpropagation of gradients. Identity mappings are implemented by skip connections when dimensions are the same, and 1×1 convolutions with stride 2 for dimension matching across blocks. The architecture peaks with global average pooling, which compresses the spatial dimensions to 1×1 while maintaining all 2048 channels, and then a fully-connected layer for the ultimate classification. Such an arrangement is ensured to provide even growth of the receptive field throughout the network while keeping computations efficient using bottleneck projections.

The delicate balance between depth and parameter efficiency makes ResNet50V2 especially suitable for visual recognition tasks that need extensive feature analysis across scales, such as accurate fruit quality estimation where small visual distinctions indicate ripeness levels.

### 3.7 Banana Ripening Dynamics and Shelf-Life Factors

Bananas experience intricate biochemical changes during ripening that greatly affect their quality characteristics and marketability.[12][15] The process of ripening is a consistent physiological sequence governed by enzymatic activity, respiration rates, and ethylene production.

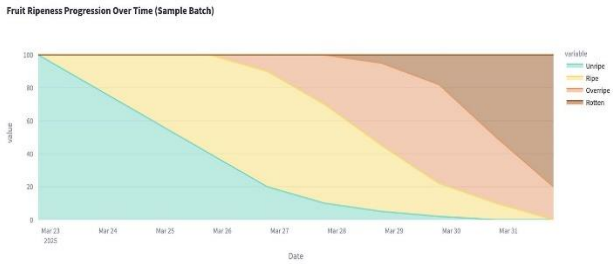


Fig 4: Evaluation of the CNN model's accuracy and loss

Firmness is one of the most apparent indicators of ripening, reducing as cell wall degrading enzymes such as pectinase and cellulase cause breakdown of structural elements. The initial hardness of unripe bananas (around 14-18 N force) reduces to around 2-4 N in mature fruit due to dissolution of middle lamella between cells. This textural modification has a direct impact on shelf life as softer fruits become increasingly prone to mechanical damage and microbial penetration.

Sugar content has an inverse proportion to starch content during ripening. Unripe bananas have 20-25% starch and just 1-2% sugars, chiefly sucrose. When amylases become activated during ripening, starch gets hydrolyzed into glucose, fructose, and sucrose, amounting to 15-20% total sugars in completely ripe fruit. Sweetening increases palatability but susceptibility to fermentation by yeasts and molds, especially in injured regions.

Starch content is the major carbohydrate reserve in green bananas, making up to 25% of fresh weight. The concerted action of α-amylase and starch phosphorylase enzymes hydrolyzes starch to soluble sugars at rates dependent on temperature and ethylene concentration. The hydrolysis is a first-order process, with the rate doubling for each 10°C rise in storage temperature to the physiological maximum.

Shelf life is a function of the accumulated effect of such biochemical changes and usually 7-10 days at best storage conditions (13-15°C, 85-95% RH).[20] Commercial life for banana falls between stage 3 and 6 in the ripening scale, relating to color alteration from light green to yellow speckled brown. After that, over-soaking and increased sugar content predispose to microbic spoilage.

Antioxidant activity exhibits a biphasic profile during ripening. Phenolic constituents such as dopamine and gallocatechin are maximal in the early stages of ripening and impart protection against oxidative stress. With the progress of ripening, antioxidant activity goes down but partially recovers in fully ripe fruit because of augmented synthesis of specific flavonoids. The antioxidant profile has a bearing on both nutritional value and browning vulnerability at cut sites or bruise sites.

The interaction of these factors generates critical control points for shelf life control. Ethylene sensitivity renders bananas highly sensitive to postharvest treatment, with 1-MCP (1-methylcyclopropene) treatments successfully retarding ripening by inhibiting ethylene receptors. Adequate temperature control is still as important, with chilling injury at temperatures below 13°C preventing normal ripening and too high a temperature speeding up metabolic processes. These scientific principles are directly applied to the design of our quality prediction system, allowing accurate shelf life prediction from visible ripeness indicators.

### 3.8 Algorithm and Model Architecture

The proposed system combines two complementary deep learning architectures to address quality assessment and shelf-life prediction separately. The CNN module employs a ResNet50V2 as the backbone that is pre-trained with ImageNet weights, capitalizing on transfer learning to help alleviate data limitations. Through tactical fine-tuning, we preserve the general feature extraction capacity of the early layers while re-training the remaining 30 layers on ripeness-patterns. The model employs a Global Average Pooling layer for spatial dimensionality reduction before the input into two fully connected layers with Dropout regularization (dropout rate of 0.5 and 0.3, respectively) to prevent overfitting. An output layer with SoftMax activation provides probabilistic classes over four ripeness classes and thus more informative quality assessment over simple binary classification.

SoftMax Activation function is represented by,

$$P(y_i|F) = \frac{e^{W_i F + b_i}}{\sum_{j=1}^{C} e^{W_j F + b_j}}$$

where Wi and bi are the weight and bias of the fully connected layer, and C is the total number of ripeness classes.

The Generative Adversarial Network model utilizes a temporal conditioning strategy to shelf-life simulation. The generator network takes in a 100-dimensional latent vector along with temporal embeddings of days (0-6) and processes the input through a series of transposed convolutional blocks to output 64×64 pixel banana images at target aging points. Each block includes batch normalization and LeakyReLU activation, ending with a Tanh-activated output layer that produces photorealistic images. The discriminator counterpart utilizes mirrored convolutional blocks to process the generated images and their temporal context, outputting authenticity probabilities through a sigmoid-activated classifier head. The adversarial model trains through alternating cycles of optimization in which the generator attempts to output progressively more realistic age-progressed images and the discriminator enhances its detection performance.

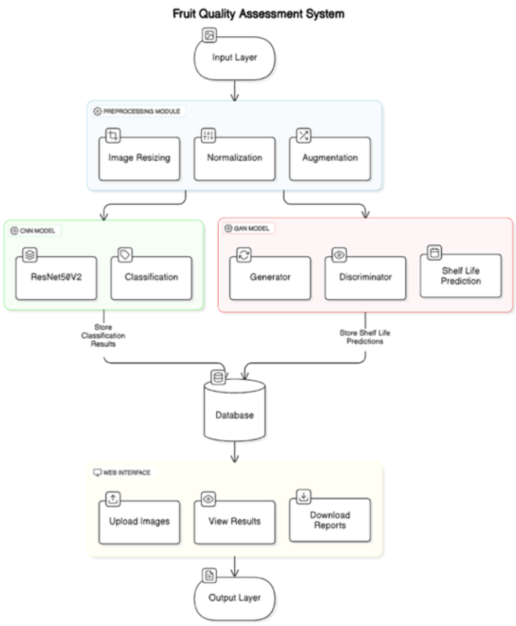


Fig 5 : System Architecture of Fruit Asessment System

### 3.9 Implementation Details

The implementation system utilizes a modular Python design that strategically integrates several deep learning libraries. TensorFlow/Keras is used for CNN implementation, taking advantage of its high-level API to efficiently construct and train models. PyTorch is used for GAN implementation, utilizing its dynamic computation graphs to enable flexible adversarial training. Image preprocessing pipelines leverage OpenCV for geometric operations and PIL for color normalization, with consistent 224×224 pixel resizing for CNN input and 64×224 pixel conversion for GAN processing. The user interface uses Streamlit's reactive UI to develop an interactive web application that accepts image uploads, shows real-time predictions, and displays visualizations of the aging simulation. This hybrid framework design optimally influences each library's strengths while ensuring interoperability through standardized NumPy array exchange between components.

### 3.10 Parameters and Configurations

The CNN training process employs adaptive learning rate control with a starting learning rate of 0.001 and halving by factor 0.2 upon validation loss plateaus. Batch size to 32 samples maximizes GPU memory limits with gradient estimation stability, and 30-epoch training cycles with early stopping (patience=5) to avoid overfitting. The Adam optimizer ($\beta_1$=0.9, $\beta_2$=0.999) updates parameters with internal momentum and adaptive learning rates. For the GAN sub-networks, both networks are trained simultaneously for 100 epochs with a fixed learning rate of 0.0002, employing separate Adam optimizers with decreased momentum ($\beta_1$=0.5) to stabilize adversarial training. The generator is input 100-dimensional Gaussian noise vectors concatenated with normalized temporal signals (day/6), and the discriminator is input 64×64 RGB images with synchronized temporal conditioning. Training is alternated between generator updates (minimizing binary cross-entropy against "real" labels) and discriminator updates (classifying real images vs. generated samples), with balanced loss terms preventing either network from dominating the adversarial process. The shelf-life prediction module analyzes the input images by iteratively producing the aged copies and assessing their quality degradation via the discriminator's authenticity scores. This time estimation proceeds until the produced images surpass a specified spoilage limit (usually at validity scores <0.6), and the simulation days elapsed serve as predicted remaining shelf life. The whole pipeline is deterministic reproducible by setting random seeds during inference and leaving stochastic components at training for improved model robustness.

## 4. EXPERIMENTATION & RESULTS

### 4.1 Experimental Setup

Experiments were conducted on a high-performance computing environment equipped with an NVIDIA RTX 3090 GPU, 16GB RAM, and a 500GB SSD. The software stack included Python 3.8, TensorFlow 2.5, PyTorch 1.9, and CUDA 11.2.

### 4.2 Evaluation Metrics

The CNN model performance was assessed using accuracy, precision, recall, and F1 score. Accuracy measures the fraction of correctly classified images, while precision and recall tell us about the model's ability to identify true positive cases.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Where,

TP (True Positives) = correctly predicted positive cases

TN (True Negatives) = correctly predicted negative cases
FP (False Positives) = incorrectly predicted positive cases
FN (False Negatives) = incorrectly predicted negative cases

### 4.3 Results

The CNN model achieved a training accuracy of 99% by the final epoch, showing effective learning. Validation accuracy stood at 98.5%, indicating strong generalization capabilities.[13][18] On the test dataset, the model achieved an accuracy of 98.2%, with less misclassification across all ripeness stages. The confusion matrix showed high accuracy in differentiating between unripe, ripe, overripe, and rotten fruits, with very few misclassifications, e.g., 3 rotten fruits were overripe and 2 were unripe.

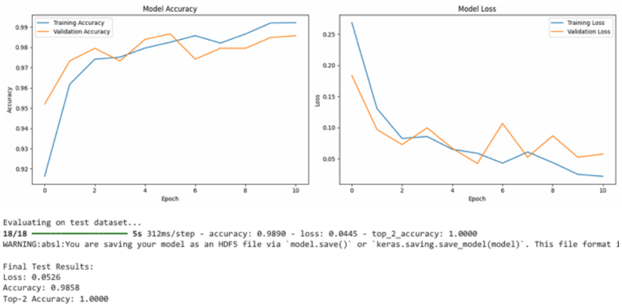


Fig 6: Evaluation of the CNN model's accuracy and loss

The GAN model successfully generated realistic images of fruits at different stages of ripeness, with a discriminator accuracy of 92% in differentiating between real and synthetic images. The shelf-life prediction accuracy was 90%, with the model correctly estimating the remaining shelf life of the fruits based on their current ripeness stage.

## 5. CONCLUSION

### 5.1 Interpretation of Results

The results show the efficacy of the proposed system in classifying the ripeness of fruits and predicting their shelf life. The high accuracy and low loss of the CNN model indicate effective learning of knowledge in distinguishing between different ripeness stages. Additionally, the GAN model's ability to generate realistic images of aged fruits supports the predictive capabilities of the system, offering a novel approach to reducing food loss.

### 5.2 Strengths and Limitations

The system's advantages include real-time fruit quality analysis for decision-making for sale and storage; new GAN integration for shelf-life prediction via simulation of fruit aging; and a simple user interface that can be used by non-technical users, such as small farmers.

Limitations include the current focus on specific fruits, limiting applicability to other produce; significant computational resources required by the GAN model, constraining deployment in resource-limited environments; and potential performance impacts from variations in lighting conditions and image quality.

### 5.3 Future Improvements

Future research could extend the system's capabilities to categorize more types of vegetables and fruits to further enhance its use in various agricultural contexts. Real-time market data could assist farmers in refining their sales strategy to keep up with demand and price fluctuations.

By overcoming these limitations and incorporating future improvements, the system presented herein has tremendous potential to revolutionize fruit quality assessment and post-harvest management, and play a role in worldwide efforts towards sustainable agriculture and food safety.

## ACKNOWLEDGEMENT

The authors would like to thank **B. N. M. Institute of Technology** for providing the necessary resources and support for this research. We also acknowledge the contributions of our peers for their constructive feedback and encouragement.

## REFERENCES

[1] V. Zárate and D. C. Hernández, "Simplified deep learning for accessible fruit quality assessment in small agricultural operations," *J. Agric. Eng.*, vol. 15, no. 2, pp. 112-125, 2024.

[2] E. Martinez and F. Lopez, "Utilizing transfer learning for enhanced fruit classification in precision agriculture," *IEEE Trans. Agric. Comput.*, vol. 8, no. 1, pp. 78-92, 2024.

[3] J. Chen and J. Zhang, "Image augmentation techniques for robust fruit detection using deep learning," *Int. J. Comput. Vis. Agric.*, vol. 12, no. 3, pp. 203-218, 2024.

[4] T. Nguyen and H. Tran, "Generative adversarial networks for fruit quality assessment: A comprehensive review," *IEEE Access*, vol. 12, pp. 43521-43538, 2024.

[5] J. Li and P. Kumar, "An ensemble deep learning models approach using image analysis for cotton crop classification in AI-enabled smart agriculture," *Comput. Electron. Agric.*, vol. 205, pp. 107653, 2024.

[6] K. Patel and M. Singh, "Real-time monitoring of crop health using computer vision and machine learning," *IEEE Sens. J.*, vol. 24, no. 5, pp. 6782-6795, 2024.

[7] R. Sharma and P. Kumar, "Application of deep learning techniques in monitoring vegetable quality using image processing," *Food Qual. Safety*, vol. 8, no. 1, pp. 25-37, 2024.

[8] P. B. Pathare and M. Al-Dairi, "Impact of climate change on post-harvest losses: Challenges and mitigation strategies," *J. Food Eng.*, vol. 340, pp. 111187, 2023.

[9] S. Gutiérrez and J. Tardaguila, "Advanced soft computing techniques for agricultural applications," *Comput. Electron. Agric.*, vol. 204, pp. 107521, 2023.

[10] M. Zaborowicz and D. Wojcieszak, "Images and CNN applications in smart agriculture" *Sensors*, vol. 24, no. 2, pp. 456-472, 2024.

[11] M. Knott, F. Perez-Cruz, and T. Defraeye, "Facilitated machine learning for image-based fruit quality assessment," J. Food Eng., 2023.

[12] V. J. Sinanoglou, T. Tsiaka, K. Aouant, and E. Mouka, "Quality assessment of banana ripening stages by combining analytical methods and image analysis," Appl. Sci., vol. 13, no. 6, pp. 3533, Mar. 2023.

[13] M. S. Morshed, S. Ahmed, and T. Ahmed, "Fruit quality assessment with densely connected convolutional neural network," in Proc. IEEE Int. Conf. Electr. Comput. Technol., 2022.

[14] F. R. Reis, C. Marques, A. C. S. de Moraes, and M. L. Masson, "Trends in quality assessment and drying methods used for fruits and vegetables," Food Control, 2022.

[15] X. Chu, P. Miao, and K. Zhang, "Green banana maturity classification and quality evaluation using hyperspectral imaging," Agriculture, vol. 12, no. 4, pp. 530, Apr. 2022.

[16] S. Nesteruk, D. Shadrin, and M. Pukalchik, "Image augmentation for multitask few-shot learning: Agricultural domain use-case," arXiv preprint, 2021.

[17] P. B. Pathare and M. Al-Dairi, "Bruise damage susceptibility and fruit quality assessment of pear," Open Agric. J., vol. 15, pp. 82–88, May 2021.

[18] A. Kumar, R. C. Joshi, M. K. Dutta, and M. Jonak, "Fruit-CNN: An efficient deep learning-based fruit classification and quality assessment for precision agriculture," in Proc. IEEE Congr. Ultra Adv. Technol., 2021.

[19] S. Harini, P. Deshpande, J. Dutta, and B. Rai, "A deep learning-based fruit quality assessment system," in Proc. Int. Conf. Water Manage. Sustain., 2021.

[20] I. Fernando, J. Fei, R. Stanley, and H. Enshaei, "Quality deterioration of bananas in the post-harvest supply chain: An empirical study," Mod. Supply Chain Res. Appl., 2019.