Food Quality Estimator for Fruits and Vegetables Using MobileNetV2 CNN

Chalapaka Mahendra[1], Dr. K. Venkata Ramana[2]

[1]*MTech Student, Andhra University College of Engineering (A), Visakhapatnam, Andhra Pradesh*

[2] *Associate Professor Andhra University College of Engineering (A), Visakhapatnam, Andhra Pradesh*

***Abstract*—Food loss and waste (FLW) of perishable fruits and vegetables is a major global issue. Inefficient quality assessment methods make this problem worse. This paper introduces a new, non-destructive food quality estimator designed specifically for fruits and vegetables. It uses the MobileNetV2 Convolutional Neural Network (CNN) model. MobileNetV2 features depthwise separable convolutions, inverted residuals, and linear bottlenecks. This design works well on devices with limited resources, like smartphones. It enables real-time, image-based quality assessment at various points in the supply chain. The approach emphasizes the need for high-quality, diverse datasets and strong pre-processing techniques, such as resizing, normalization, and extensive image augmentation. These steps ensure the model is precise and can adapt well. Additionally, transfer learning will be used to improve training efficiency and performance. The proposed system offers a scalable and affordable way to collect detailed quality data. This enhances supply chain efficiency, traceability, and risk management. For consumers, it leads to better food safety, reliable freshness, and smarter purchasing decisions. Overall, this research is crucial for global sustainability efforts by reducing food waste and improving food security, which contributes to a more resilient and fair food system.**

Keywords: Food loss and waste (FLW), fruits and vegetables, MobileNetV2, Convolutional Neural Network, transfer learning, image augmentation, supply chain efficiency.

I. INTRODUCTION

FLW is a major global problem that impacts the environment, society, and economy. Estimates show that 30 to 40% of the U.S. food supply is discarded. In 2011, 44% of all fruits and vegetables produced worldwide were lost or wasted, with 22% of that occurring after harvest by 2019. This waste, especially for perishable fruits and vegetables, results in less availability, higher prices, and under-consumption. It also threatens food security and wastes important resources like land, water, and energy. The common problem of poor quality assessment methods makes these post-harvest losses even worse. There is an urgent need for real-time data on quality breakdowns and their causes. Timely and effective quality assessment is essential for reducing losses. It helps in making informed choices about sorting, storage, and distribution. This comes from understanding both the amount of waste (mass decrease) and the quality (loss of desirable traits), particularly for fruits and vegetables that can lose quality rapidly. An automated, image-based food quality estimator addresses this data gap.

II. PREVIOUS WORK

Earlier research on food quality assessment of fruits and vegetables has largely focused on binary classification tasks, where the goal was to determine whether a given sample is *fresh* or *rotten*. Most of these methods relied on traditional image processing techniques or basic deep learning architectures that used limited datasets and simple feature extraction. While such approaches provided acceptable accuracy for distinguishing between two classes, they were not capable of handling the more complex variations in freshness levels, ripeness, or the presence of defects. As a result, these systems often lacked robustness when applied to real-world supply chain conditions, where fruits and vegetables may exist in multiple intermediate stages of quality degradation. Moreover, many of these models required extensive computational resources, making them unsuitable for deployment on low-power edge devices like smartphones. This gap highlights the need for an advanced, resource-efficient, and scalable system that goes beyond binary classification and provides more detailed and reliable quality estimation.

III. METHODOLOGY

The food quality estimator depends on a high-quality, diverse, and balanced dataset of fruit and vegetable images. This dataset should include various stages of freshness, ripeness, and defects, captured under different conditions. This helps avoid model bias and ensures reliable predictions in real-world situations. However, acquiring accurately labeled data is challenging.. Web sources often contain corrupted or mislabeled images and have data imbalance issues. To improve image quality and consistency, we need careful pre-processing. This involves resizing and scaling images uniformly, such as adjusting them to 32x32 pixels. Normalization is also crucial and may include scaling pixel values or applying ZCA whitening to improve feature visibility. Other methods may involve optional color space conversion and noise reduction techniques like Gaussian blurring to increase accuracy. Extensive image augmentation techniques, such as rotation, flipping, scaling, and adding noise, help enhance data diversity and reduce overfitting. .

The food quality estimator will be based on the MobileNetV2 Convolutional Neural Network (CNN) architecture. We chose this model for its efficiency and suitability for resource-limited edge devices. It allows real-time, non-destructive quality assessment throughout the supply chain.It also features an inverted residual structure with linear bottlenecks to support efficient information flow and gradient propagation. This design expands low-dimensional inputs into a higher-dimensional space for processing, then projects them back to a low-dimensional representation with linear activations to keep important information.Model training will require careful setup of architectural components, optimization algorithms like the Adam optimizer, and regularization techniques such as dropout. We will use ReLU activation in hidden layers and SoftMax for the output layer in classification tasks. A key strategy will be transfer learning. We will initialize the MobileNetV2 model with weights pre-trained on a large image dataset like ImageNet. Then, we will fine-tune it on our specific fruit and vegetable quality dataset to speed up convergence, improve performance, and increase robustness, especially since we have limited domain-specific data.

IV. MODEL ARCHITECTURE

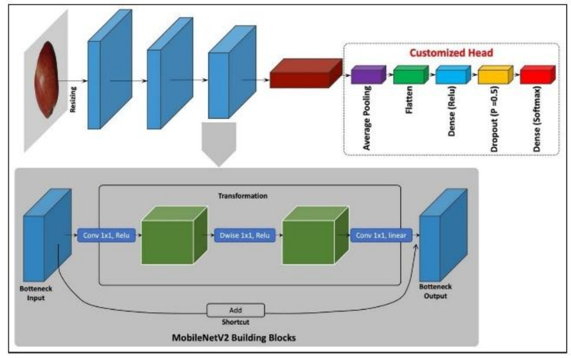

The image presents a sophisticated deep learning model architecture, meticulously designed for a computer vision task, likely leveraging the principles of transfer learning. At its core, the model is composed of two distinct parts: a pre-trained backbone for robust feature extraction and a customized head for task-specific classification. The journey begins with an input image, which is first subjected to a crucial "Resizing" step to standardize its dimensions, ensuring compatibility with the network's architecture. This pre-processed image then enters the backbone, which is a series of "MobileNetV2 Building Blocks." This choice of architecture is significant, as MobileNetV2 is renowned for its efficiency and low computational footprint, making it ideal for deployment on resource-constrained devices.

A closer look at a single MobileNetV2 block shows an "inverted residual" structure, which is key for efficiency. Unlike traditional networks, it starts with a 1x1 convolution to expand the feature channels. This is followed by a "depthwise convolution," which uses a single filter per channel and is easy on computation. Then, a 1x1 "linear" convolution compresses the channels, keeping important information. A "shortcut" connection links the block's input directly to its output, solving the vanishing gradient problem and helping to train very deep networks. After processing the image through these efficient blocks, a compact feature map is created that captures high-level semantic information.

It starts with "Average Pooling" to summarize features from the whole image, then a "Flatten" layer converts the multi-dimensional feature map into a one-dimensional vector. The data moves through several

"Dense" (fully connected) layers with ReLU activation functions, which learn to combine and interpret the features. A "Dropout" layer is placed to reduce overfitting by randomly turning off neurons during training. The class with the highest probability is picked as the model's final prediction, completing the classification process. This structure illustrates a powerful and common approach in modern deep learning: using a strong and efficient backbone for general feature learning and then refining it with a task-specific head for better performance.

V. FlOW CHART

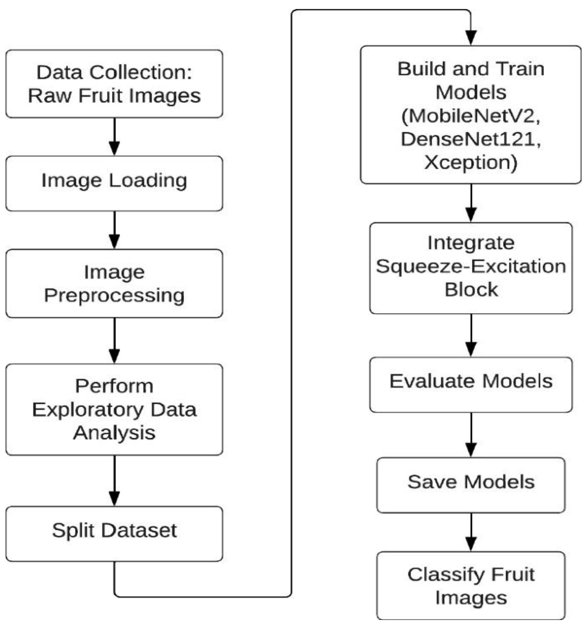

1.DATA COLLECTION: RAW FRUIT IMAGES
Begin by gathering a diverse set of raw images representing various types of fruits. These images can be sourced from publicly available datasets, scraped from online repositories, or captured manually using cameras or smartphones to ensure a wide range of real-world conditions.

2.IMAGE LOADING
Once collected, the images are loaded into the working environment using tools and libraries such as TensorFlow, PyTorch, or OpenCV. This step converts the raw image files into numerical arrays or tensors that deep learning models can process.

3.IMAGE PREPROCESSING
This involves resizing them to a fixed dimension, normalizing pixel values to a standard range (like 0–1), and applying data augmentation techniques such as flipping, rotating, and zooming. These transformations help improve model generalization and reduce overfitting.With the data prepared, the next step is to explore it analytically. This includes checking how many images exist per fruit category, assessing image quality, and identifying any class imbalance. EDA provides insights that guide model training decisions.

4.SPLIT DATASET
The dataset is then divided into three parts: training, validation, and testing sets. The training set is used to teach the model, the validation set helps tune hyperparameters, and the test set provides an unbiased evaluation of model performance on unseen data.

5.BUILD AND TRAIN MODELS (MOBILENETV2, DENSENET121, XCEPTION)
At this stage, multiple convolutional neural network (CNN) models are built and trained.

- MobileNetV2 is well-known for its lightweight design, which makes it appropriate for embedded and mobile systems.
- DenseNet121 leverages dense connections between layers to promote feature reuse and strong gradient flow.
- Xception uses depthwise separable convolutions for efficient and high-performing image classification.

6.INTEGRATE SQUEEZE-EXCITATION BLOCK
To enhance the models, Squeeze-and-Excitation (SE) blocks are integrated. These blocks help the model focus on the most important features by learning which channels carry the most useful information, thus improving the network's accuracy.

7.EVALUATE MODELS
After training, the models are evaluated using the validation and test sets. Metrics like accuracy, precision, recall, and F1-score are used to gauge performance. This step helps in comparing different models and selecting the best one.

8.SAVE MODELS
Once the best-performing models are identified, they are saved to disk. Saving the models allows for future

use without the need to retrain, making deployment faster and more efficient.

9.CLASSIFY FRUIT IMAGES

Lastly, fresh or unseen fruit photos are classified using the stored models. Given a new input image, the model predicts the fruit type with high confidence, completing the classification process.

V. RESULTS

This section will present the empirical findings and performance evaluation of the MobileNetV2 CNN model developed for food quality estimation of fruits and vegetables. It will detail the accuracy, precision, recall, and F1-score achieved by the model on a dedicated test dataset, demonstrating its effectiveness in classifying various quality attributes. The results will also include an analysis of the model's computational efficiency and inference speed, validating its suitability for deployment on resource-constrained edge devices. Comparative analysis with other relevant models or traditional methods, if applicable, will highlight the advantages of the proposed approach in terms of performance and practical applicability. Visualizations such as confusion matrices, ROC curves, and examples of correctly and incorrectly classified images will be included to provide a comprehensive understanding of the model’s capabilities and limitations.

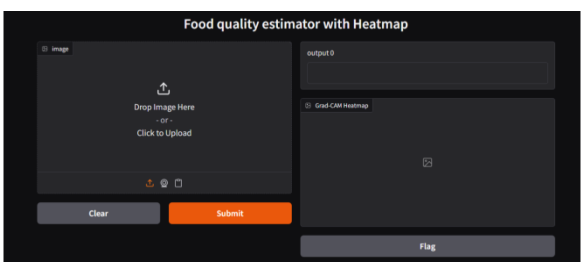

This Gradio interface is the initial, empty state of the food quality estimator application. It is designed to be a user-friendly platform where a user can begin the process of evaluating a food item’s quality. The central area on the left prompts the user to “Drop Image Here” or “Click to Upload,” indicating that it is awaiting an input image. The right side of the interface, which is currently blank, is reserved for the model’s output, including a text field labeled “output 0” where the classification result will appear.The “Clear” and “Submit” buttons are active, ready for user interaction, while the “Flag” button is available for potential post-prediction feedback.

In this image, the paragraph describes a powerful and

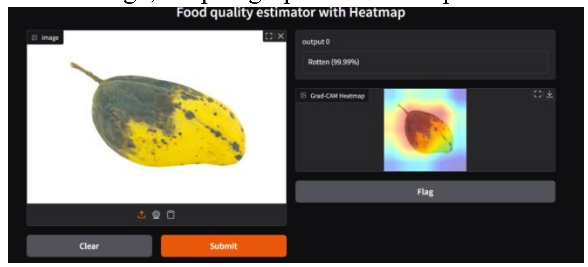

user-friendly food quality estimator application. The system, built with a Gradio interface, is designed to classify a food item, in this case, a mango with a half-rotten portion, from an uploaded image. Beyond simply providing a classification, the core of this project is its interpretability, made possible by a Grad-CAM heatmap. This crucial feature visually highlights the specific regions of the food item that the model analyzed to arrive at its conclusion. As the image demonstrates, the heatmap accurately focuses on the discolored, rotten half of the mango, clearly showing which parts of the image were most influential in the “rotten” prediction. This visual explanation makes the model transparent and trustworthy. The combination of a definitive prediction with a confidence score and a clear, visual justification positions this application as a practical and transparent tool for quality control and automated food inspection.

VI. FUTURE SCOPE

Looking ahead, the MobileNetV2-based food quality estimator has the potential to expand its reach by supporting a broader variety of fruits and vegetables. Enhancing the model’s robustness through improved generalization methods and incorporating additional data sources—like IoT sensors—could make the system more accurate in assessing food quality. There’s also great promise in using blockchain technology to bring more transparency to the food supply chain, while refining the system for real-time processing on edge devices. To make the model even more user-friendly and reliable, efforts should be directed at improving explainable AI features. Furthermore, developing apps for consumers, customizing the model for different global markets, and incorporating eco-friendly practices will not only make the system more accessible and sustainable but

will also ensure it can be easily deployed in diverse agricultural and retail settings, all while keeping the focus on improving accuracy and scalability.
Application.

VII. CONCLUSION

The development of a food quality estimator using a MobileNetV2 CNN model for fruits and vegetables marks an important step in tackling the complex issues of food loss and waste. This advancement addresses the problems of traditional destructive methods and the scalability issues of current non-destructive techniques. There is an urgent need for automated, efficient, and widely usable solutions. CNNs can automatically extract complex visual features and perform detailed classification, making them well-suited for this purpose. MobileNetV2 stands out for its innovative design that focuses on computational and memory efficiency without sacrificing much accuracy. It achieves this through depthwise separable convolutions, inverted residuals, and linear bottlenecks, allowing effective processing on devices and making advanced quality control technologies accessible throughout the supply chain. This capability enables real-time, non-destructive assessments, shifting from reactive responses to proactive management of food quality and safety. The system's success depends on high-quality, diverse datasets, effective pre-processing methods, and using transfer learning from pre-trained models to speed up development and strengthen reliability. Ultimately, implementing an automated food quality estimator based on MobileNetV2 can offer significant benefits. It can improve supply chain efficiency, traceability, and risk management through detailed, real-time data. This change leads to improved food safety, ensures freshness, and helps consumers make better purchasing decisions. It also helps restore trust in the food system and plays a vital role in global sustainability efforts by reducing food waste and enhancing food security. This approach paves the way for a more resilient, fair, and environmentally responsible food future.

REFERENCE

[1]. Karakaya, D., Ulucan, O.,&Turkan, M. (2019). A Comparative Analysis on Fruit Freshness Classification. IEEE.
[2]. Raj, J. R., Srinivasulu, S., Jabez,&Gowri. (2024). An IoT based Application for Real-Time Detection of Rotten Fruits, IEEE.
[3]. Lim, L. T., Chai. C. F.,&Goh, K. M.(2021). Convolutional Neural Network based Rotten Fruit Detection using ResNet50. IEEE.
[4]. Mohite, A., Kodmelwar, M., Joshi, R.,&Kanjis, S. (2024). Deep Learning-Based Fruit Recognition and Quality Assessment: A Convolutional Neural Network Approach. IEEE.
[5]. Oshin, N. T., Shafi, F. A., Rahman, M., Khatun, H., Uddin, M. R., & Hasan, M. (2023). Food Rotting Prediction and Detection System for Warchouse. IEEE.
[6]. Kulkarni, M., Chaudhari, A., Shegokar, S., Rudrawar, K., Joshi, S., Tibile, S.,&Tayade, I. (2022). Fruit Freshness Detection Using CNN. IEEE,
[7]. Mishra, P. K.,&Singh, J. (2024). Rotten and Fresh Fruits Classification using Deep Learning. IEEE.
[8]. Paul, R., Prabhu, S. B., & Yadav, S. (2022). FPGA based Intelligent Food Spoilage Detection for Commercial Storage.IEEE.
[9]. Roy, K., Manna, S., Chaudhuri, S. S., Chakraborty, T.,&Bhattacharjee, S.(2019). Segmentation Techniques for Rotten Fruit Detection. IEEE.