

Classification vs. Regression for Real-Time Fruit Shelf-Life Prediction: A Transfer Learning Approach with MobileNetV2 for Edge AI

Sai Sri Chandana Jada
M.Tech Student
QIS College of Engineering and
Technology
Ongole, India

M. Senthil, PhD
Professor, Department of CSE
QIS College of Engineering and
Technology
Ongole, India

Kolla Vivek
Assistant Professor
QIS College of Engineering and
Technology
Ongole, India

ABSTRACT

The agricultural industry faces significant challenges in managing perishable goods, with a substantial portion of produce being wasted due to spoilage. This paper presents a comparative study of two deep learning approaches for predicting the shelf life of fruits using Convolutional Neural Networks (CNNs) and edge computing. This study developed and evaluated both a classification model and a regression model, both based on the MobileNetV2 architecture for a fair comparison. The classification model achieved a test accuracy of 74.75%, while the regression model provided more granular predictions with a mean absolute error of 1.44 days. Both models were converted to the TensorFlow Lite (TFLite) format and evaluated on the test set, achieving identical performance to their Keras counterparts while significantly reducing prediction latency. This research explores the advantages and disadvantages of both classification and regression approaches, demonstrating the potential of deep learning and edge computing to create scalable and efficient solutions for real-time shelf-life prediction, which can help to reduce food waste and optimize supply chain management.

General Terms

Deep Learning, Computer Vision, Image Processing, Pattern Recognition, Machine Learning, Edge Computing, Internet of Things (IoT), Predictive Modeling, Classification, Data Science, Algorithms, Food Technology, Agricultural Technology.

Keywords

Shelf-Life Prediction, Fruit Spoilage, Convolutional Neural Networks (CNN), Edge Computing, TensorFlow Lite (TFLite), Image Classification, Computer Vision, Deep Learning, Food Waste, Supply Chain Management, Banana Ripeness.

1. INTRODUCTION

The prediction of shelf life for fruits and vegetables is a crucial aspect of modern agriculture, food distribution, and supply chain management. This paper explores two distinct deep learning methodologies for shelf-life prediction: a classification approach to categorize produce into ripeness stages, and a regression approach to predict the remaining shelf-life in days. Fruits, being perishable goods, are highly susceptible to spoilage. Traditional methods of estimating shelf life typically rely on manual inspection, which can be labor-intensive and subjective. Consequently, there is a growing need for innovative technologies to predict the shelf life of fruits more accurately and efficiently [1]. Convolutional Neural Networks (CNNs) have proven highly effective in processing and analyzing visual data, making them ideal for this task [2,

3]. By training CNNs on large datasets of fruit images, it is possible to create predictive models that can automatically estimate the remaining shelf life of fruits based on their appearance [4, 5].

This paper compares a classification model and a regression model, both based on the MobileNetV2 architecture, to ensure a fair comparison. The classification model assigns a ripeness category (e.g., "fresh", "ripe"), while the regression model predicts a continuous value for the remaining shelf life. The classification model offers a simpler, more interpretable output, while the regression model provides more granular and actionable data for fine-tuning supply chain logistics.

Furthermore, this research investigates the deployment of these models on edge devices using TensorFlow Lite (TFLite), enabling real-time, on-site analysis without reliance on cloud infrastructure [6]. By combining CNNs with Edge Computing, the aim is to provide a scalable and efficient solution that addresses the challenges associated with fruit quality monitoring, storage, and distribution [7]. This comparative study will shed light on the trade-offs between classification and regression models for this application, offering valuable insights for the development of future food technology systems.

2. LITERATURE SURVEY

Fruits and vegetables shelf-life estimation is a critical area of research in food science and technology, aiming to enhance food quality, reduce wastage, and improve supply chain efficiency. Traditional methods for determining fruit freshness rely on physical, chemical, and biological assessments, which are often time-consuming and subject to human error. The introduction of deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized fruit shelf-life estimation by enabling automated, accurate, and scalable predictions [8]. Despite these advancements, significant research gaps persist, including limitations in dataset availability, variability in environmental conditions, computational constraints, and ethical concerns. This literature survey explores existing research on fruit shelf-life estimation using CNNs and highlights key gaps that need to be addressed for more effective implementation.

A study titled "A Review on Automated Detection and Assessment of Fruit Damage Using Machine Learning" by Yonsai Safari, Joyce Nakatumba-Nabende, Rose Nakasi, and Rose Nakibuule, comprehensively examined 32 research papers from the past 13 years, focusing on automated fruit damage detection using machine learning and deep learning techniques [9]. The review highlights a significant shift towards deep learning models, particularly Convolutional

Neural Networks (CNNs), for their superior performance in processing large datasets and achieving high accuracy. For instance, studies using deep CNNs achieved accuracies up to 98.6% in apple and banana disease detection and 98% in citrus disease severity detection, while Faster R-CNN reached 88.1% in mango fruit detection. The most commonly extracted features for damage detection were color, texture, shape, and size. While deep learning models demonstrated high accuracy, particularly with models like ResNet-50 and VGG-16, traditional machine learning remains relevant due to its lower computational demands, especially for smaller datasets [9]. However, a major drawback is the limited availability of large, open-source datasets, hindering reproducibility and further research. The review also identified challenges such as occlusion in orchards, the detection of defects like chilling injuries that do not alter color or texture, and the manual selection of CNN hyperparameters. Furthermore, most of the research focused on post-harvest damage detection, with limited studies addressing in-field damage assessment. The review concludes that while deep learning shows significant promise for automated fruit damage detection, future research should focus on developing lightweight CNN architectures for mobile and embedded devices, addressing occlusion challenges, exploring alternative imaging modalities like thermal imaging, and developing domain adaptation methods for broader applicability [9].

Another study, "Comparative Analysis for Predicting Shelf life of Fruits Using Advanced Deep Learning Approaches" by Sanath S Shenoy, Radhika Mishra, and Ruchi Chaturvedi, explored the application of deep learning, specifically Faster R-CNN and YOLOv5, to predict the shelf life of bananas, aiming to reduce food waste by accurately assessing fruit maturity [10]. The research involved creating a dataset of 254 images of Cavendish bananas, categorized into five shelf-life stages, and training both models using 164 images, with 45 images each for validation and testing. The models were trained with a batch size of 16, and their performance was evaluated using mean Average Precision (mAP) at 50% Intersection over Union (mAP50) and mAP between 50% and 95% IoU (mAP50-95) across 250, 300, 500, and 1000 epochs or iterations. The results indicated that both models achieved mAP50 and mAP50-95 scores exceeding 80%, demonstrating their potential for shelf-life prediction. However, Faster R-CNN exhibited better performance at higher iterations but demanded more computational resources, while YOLOv5 offered model variants with lower hardware requirements [10]. The study highlights the feasibility of object detection models for fruit shelf-life prediction, but the dataset size of 254 images is relatively small, which could limit the generalizability of the findings. Additionally, the paper does not specify the exact mAP values obtained for each model and epoch, hindering a precise quantitative comparison.

The work done by Saeed S. Alahmari, and Tawfiq Salem titled "Food State Recognition Using Deep Learning", introduced a novel cascaded multi-head deep learning approach to simultaneously recognize food ingredient type and state, addressing a gap in existing research that primarily focused on ingredient type recognition [11]. The proposed method utilizes a cascaded architecture, where features learned for food ingredient type are fused with image deep representations to enhance food state recognition. The authors trained and evaluated their approach on a benchmark dataset, comparing it to a non-cascaded deep learning method. The cascaded approach achieved an 87% accuracy in food ingredient state recognition, demonstrating a significant improvement over the 81% accuracy obtained by the non-cascaded method. The

discussion highlighted that the non-cascaded approach struggled with distinguishing between similar food states, such as diced and sliced, or creamy-paste and grated, due to the lack of food type information. By incorporating the food type feature vector, the cascaded approach significantly reduced these misclassifications [11]. The paper acknowledges limitations, including dataset imbalances and instances where food state labels were missing, necessitating data preprocessing. Future work aims to address these issues by balancing the dataset and incorporating image segmentation to further improve recognition accuracy. The conclusion emphasizes the efficacy of the proposed cascaded multi-head deep neural network in jointly learning food ingredient type and state, showcasing its potential for applications in automated food handling, elderly care, and other intelligent systems.

Another study titled "Remaining Shelf-Life Estimation of Fresh Fruits and Vegetables During Transportation" by Arwa Abougharib, Mahmoud Awad and Malick Ndiaye, introduced a novel Shelf-Life Estimation Model (SLEM) designed to predict the Remaining Shelf-Life (RSL) of Fresh Fruits and Vegetables (FFVs) during transportation, leveraging real-time environmental data from IoT sensors [12]. The model estimates the decay rate based on respiration, integrating it over time, and offers a non-destructive, non-invasive approach without requiring pre-deployment accelerated shelf-life experiments. To facilitate real-time applications, a surrogate model (QnD) was developed to reduce computational intensity. Experimental validation using strawberries, apricots, and spinach in a domestic refrigerator demonstrated that the original model achieved a maximum prediction error of 1.3 days, while the surrogate model's error reached 2.95 days. Statistical analysis, including ANOVA tests, confirmed that both models produced predictions not significantly different from observed shelf lives, even at a 0.01 significance level, validating the model's accuracy despite simplifying assumptions. However, the surrogate model exhibited higher errors, particularly with sealed spinach, likely due to accumulated approximation errors and challenges in accurately estimating parameters like headspace air volume and porosity. The model consistently overestimated the RSL for apricots, potentially due to unaccounted exogenous ethylene influence. The study concluded that the original model is more accurate for sealed products and prolonged experiments, while both models perform well for unsealed products like strawberries and apricots. The research suggests that respiration rate, coupled with time, temperature, and CO₂ concentration, are primary decay predictors. The model, though validated in a refrigerator setting, is deemed applicable to transportation, cold storage, and inventory management due to its temperature-dependent nature [12].

2.1 Research Gaps

2.1.1 Limited availability of High-Quality Datasets

One of the primary challenges in CNN-based fruit shelf-life estimation is the lack of large, high-quality, labeled datasets. Most existing datasets focus on fruit classification rather than shelf-life prediction. The limited availability of diverse datasets hinders the generalizability of CNN models, leading to reduced accuracy when applied to different fruit varieties, climates, and storage conditions. To address this gap, researchers must develop standardized, publicly accessible datasets that

incorporate a wide range of fruit types and environmental factors.

2.1.2 Challenges in Model Generalization

CNN models trained on specific datasets often struggle to generalize to new environments. Factors such as lighting conditions, camera angles, and fruit positioning can significantly impact model accuracy. Existing studies have primarily focused on controlled environments, which do not reflect real-world supply chain conditions. Future research should emphasize transfer learning techniques and domain adaptation strategies to improve model robustness across diverse scenarios.

2.1.3 Ethical and Regulatory Challenges

Despite the benefits of CNN-based fruit shelf-life estimation, ethical concerns regarding data privacy, algorithmic bias, and food safety compliance remain unresolved. There is a lack of standardized regulatory frameworks governing the use of AI in food quality assessment. Addressing these ethical and legal gaps requires interdisciplinary collaboration between AI researchers, food scientists, and policymakers.

2.1.4 Computational and Hardware Constraints

CNN models require significant computational resources, making them less accessible for small-scale farmers and retailers. Existing research has largely focused on cloud-based implementations, which may not be feasible for regions with limited internet connectivity. Developing lightweight, edge-computing-based CNN models could enhance accessibility and scalability [13]. Although CNNs have demonstrated significant potential in fruit shelf-life estimation, various research gaps must be addressed to improve their effectiveness. The lack of high-quality datasets, challenges in model generalization, ethical considerations, and computational limitations present significant hurdles to widespread adoption. Future research should focus on developing robust, adaptable, and ethically sound AI-driven fruit shelf-life estimation models to optimize food supply chains and reduce wastage. By bridging these research gaps, CNN-based solutions can contribute to a more efficient and sustainable food industry.

2.2 Research Objectives

Fruit shelf-life estimation is a crucial area of research in the food industry, focusing on enhancing food safety, reducing waste, and improving supply chain efficiency. The application of Convolutional Neural Networks (CNNs) has revolutionized shelf-life prediction by enabling automated, real-time monitoring of fruit freshness. While numerous studies have explored the use of deep learning for fruit quality assessment, well-defined research objectives are essential to guide future investigations. A structured literature survey provides insights into key research objectives that drive advancements in CNN-based fruit shelf-life estimation. This study aims to identify critical research goals, including data acquisition, model optimization, accuracy improvement, real-world deployment, and ethical considerations.

2.2.1 Development of High-Quality Datasets

A major research objective in CNN-based fruit shelf-life estimation is the collection and enhancement of high-quality datasets. CNN models require extensive labeled datasets to improve accuracy and generalization. Existing datasets often lack diversity in fruit types, environmental conditions, and ripeness stages. Research must focus on creating comprehensive, publicly available datasets that include images of fruits from various regions, storage conditions, and time intervals.

2.2.2 Optimization of CNN Architectures for Higher Accuracy

Another key research objective is optimizing CNN architectures to improve accuracy while reducing computational complexity. Traditional CNN models require substantial computing power, making them impractical for real-time applications. Research should explore lightweight architectures, hybrid deep learning models, and transfer learning techniques to enhance model performance [14, 15].

2.2.3 Real-World Deployment and Scalability

Despite advancements in CNN-based fruit analysis, real-world deployment remains a challenge. The objective of scalability is to ensure that AI-driven shelf-life estimation can be integrated into existing food supply chain management systems. Future research should address model adaptation for different storage environments, real-time IoT integration, and mobile-friendly AI applications [7].

2.2.4 Ethical AI Implementation and Compliance with Food Safety

Ensuring ethical AI implementation and compliance with food safety regulations is a critical research objective. CNN-based models must adhere to data privacy laws, algorithmic fairness, and transparency in decision-making. Researchers should develop frameworks that prevent biased predictions and ensure responsible AI deployment in food monitoring systems.

3. SYSTEM ANALYSIS

3.1 Existing System

Predicting the remaining shelf-life of food products presents a dichotomy based on their inherent characteristics and handling. For sealed, dried, highly processed, and chemically preserved items, established mathematical models offer reliable predictions due to their relative stability and predictable degradation patterns. Conversely, the shelf-life prediction for unsealed, unprocessed, and biologically active foods, exemplified by fresh fruits, vegetables, milk, fish, and meat transported alongside diverse fresh commodities under dynamic logistic conditions, remains a significant and multifaceted challenge. This complexity does not stem from a scarcity of sophisticated mathematical models designed to address the primary factors contributing to their deterioration. Instead, the practical application of these theoretical frameworks is severely hampered by numerous limitations, including their pronounced specificity to individual food species, the frequent requirement for destructive testing methodologies, and the often-prohibitive time investment associated with necessary pre-testing procedures. Consequently, despite the theoretical advancements, current models often prove unsustainable and impractical for widespread implementation in real-world scenarios, particularly within the complexities of modern food supply chains.

3.2 Proposed System

This paper proposes a system that leverages deep learning for the non-destructive prediction of fruit shelf life based on visual data. The core of the system is a comparative study between two distinct neural network models: a classification model to categorize fruits into discrete ripeness stages, and a regression model to predict the specific number of days of remaining shelf life. To ensure a fair and effective comparison, both models are built upon the highly efficient MobileNetV2 architecture, utilizing transfer learning from weights pre-trained on the ImageNet dataset. This approach allows the system to achieve

high accuracy even with a specialized and limited dataset. A key objective of the proposed system is to be practical for real-world deployment in the supply chain. To this end, the trained models are converted to the TensorFlow Lite (TFLite) format, optimizing them for high-performance inference on edge devices. This enables real-time, on-site analysis without the need for constant cloud connectivity, making the solution scalable and accessible.

4. METHODOLOGY

4.1 Dataset

The foundation of this research is the "Fruit Arrival Time and Shelf-Life Dataset" from Kaggle [16], which provides images of bananas categorized into five distinct stages of ripeness. The dataset contains images of bananas categorized into five classes based on their age: "Banana(1-2)", "Banana(3-4)", "Banana(5-7)", "Banana(8-10)", "BananaExpired". The dataset was split into training and testing sets. To ensure the quality of the training data, the dataset underwent a preprocessing and cleaning phase. This involved programmatically verifying each image to identify and discard any corrupt or invalid files, thereby ensuring the integrity of the dataset before model training. For the regression model, the remaining shelf-life in days was calculated for each image based on its class. Taking 11 as an average number of days until expiry, the formula used was

$$\text{remaining_shelf_life} = 11 - \max(\text{age in class name})$$

4.2 Model Architecture

Two CNN models were developed, both based on the MobileNetV2 architecture for a fair comparison. The MobileNetV2 architecture was selected for this study due to its high efficiency and performance, making it particularly well-suited for deployment on edge devices. As a state-of-the-art model designed for mobile and resource-constrained environments, it provides an excellent balance between accuracy and computational cost. Furthermore, leveraging MobileNetV2 pre-trained on the ImageNet dataset allows for the application of transfer learning. This approach enables the model to utilize a rich set of pre-learned features, which is highly beneficial for achieving strong performance on a specialized task with a limited dataset, such as the one used in this research.

The general architecture for both models in this study consists of the MobileNetV2 base model, with its convolutional layers frozen to retain the pre-trained weights. On top of this frozen base, a custom head was added. This head begins with a GlobalAveragePooling2D layer, which reduces the spatial dimensions of the feature maps. Following this, a Dense layer with 128 neurons and a ReLU activation function is used to learn higher-level features.

For the classification model, the final layer is a Dense layer with 5 neurons, corresponding to the five ripeness classes, and uses a softmax activation function to output class probabilities. For the regression model, the final layer is a single Dense layer with a linear activation function, which outputs a continuous numerical value representing the predicted remaining shelf-life.

4.3 Model Training

Classification Model: The model was trained using the Adam optimizer with a sparse categorical cross-entropy loss function. To improve generalization and prevent overfitting, data augmentation techniques such as rotation, shifting, shearing, zooming, and horizontal flipping were applied to the training

data. A hold-out validation strategy was employed, where 20% of the training data was set aside as a validation set.

Regression Model: The regression model was trained using the Adam optimizer with a mean squared error (MSE) loss function. The primary evaluation metric was mean absolute error (MAE). Data augmentation techniques were also applied to the training data. Similar to the classification model, a 20% validation split was used.

4.4 Edge Deployment and Optimization

Following the training phase, both models were optimized for edge deployment. The trained Keras models were converted into the TensorFlow Lite (TFLite) format. This conversion process creates a highly efficient and portable version of the models, employing techniques such as quantization and operator fusion to significantly reduce the model's file size and computational complexity. The resulting TFLite models are designed for low-latency inference on resource-constrained hardware, making them ideal for real-world applications on edge devices.

5. SYSTEM DESIGN

The system architecture is centered around a transfer learning approach using the pre-trained MobileNetV2 model. The process begins with the dataset of fruit images, which is split into training, validation, and testing sets. The core of the architecture is the MobileNetV2 base model, with its weights frozen to leverage the powerful features learned from the ImageNet dataset.

Two distinct predictive models are constructed on top of this common base. For the classification task, a custom head is added, which consists of a GlobalAveragePooling2D layer, a Dense layer with 128 neurons (ReLU activation), and a final Dense output layer with 5 neurons and a softmax activation function to correspond to the five ripeness classes. For the regression task, a similar head is used, but the final output layer is a single Dense neuron with a linear activation function to predict the continuous value of remaining shelf life.

Both models are trained and validated independently. After training, the final step in the architecture is the conversion of the saved Keras models into the TensorFlow Lite (TFLite) format. This optimization step produces lightweight, efficient models ready for performance evaluation and deployment on edge computing devices, completing the workflow from data input to deployable asset.

6. IMPLEMENTATION

The system was implemented in Python using the TensorFlow and Keras libraries for model development and training. The implementation followed a structured workflow, from data preparation to final model optimization, to create and evaluate both the classification and regression models for shelf-life prediction.

6.1 Data Collection and Preprocessing

The initial step involved acquiring the "Fruit Arrival Time and Shelf Life Dataset" from Kaggle [16]. A data cleaning script was implemented to ensure the integrity of the dataset. This script iterated through all image files and used the Python Imaging Library (PIL) to identify and remove any corrupted or invalid images that would otherwise cause errors during training. For the classification model, the directory names served as the class labels. For the regression model, these class labels were programmatically converted into numerical targets representing the remaining shelf life in days, based on the formula described in the methodology. The dataset was then

split into training and testing sets to ensure a fair evaluation of the models' performance on unseen data.

6.2 Building the Model

The core of the implementation was the construction of the deep learning models using the TensorFlow Keras API. A transfer learning approach was adopted by using the MobileNetV2 architecture with weights pre-trained on ImageNet as the base model. The convolutional layers of this base were frozen to retain the learned low-level features. A custom head was then added for the specific prediction tasks. This head consisted of a Global Average Pooling(2D) layer to reduce the feature map dimensions, followed by a Dense layer with 128 neurons and a ReLU activation function. For the classification model, the final output layer was a Dense layer with 5 units and a softmax activation. For the regression model, the output was a single Dense layer with a linear activation.

6.3 Model Training and Validation

The training process was configured and executed for both models. An *ImageDataGenerator* class was used to apply data augmentation to the training set, including transformations like rotation, shifting, shearing, zooming, and horizontal flipping, to improve model generalization and reduce overfitting. The classification model was compiled with the Adam optimizer and the *sparse_categorical_crossentropy* loss function. The regression model was also compiled with the Adam optimizer but used the *mean_squared_error* (MSE) loss function. Both models were trained with a validation split of 20%, allowing for the monitoring of performance on a hold-out set during training to identify the best-performing epoch.

6.4 Edge Optimization and Evaluation

After training, the best-performing model checkpoints were saved. The final implementation step focused on optimizing these models for edge deployment. A conversion script was used to transform the saved Keras models (.h5 format) into the TensorFlow Lite (.tflite) format. This process creates lightweight, optimized models suitable for low-latency inference on resource-constrained devices. The performance of these TFLite models was then rigorously evaluated on the test dataset to confirm that the conversion process did not degrade accuracy or Mean Absolute Error and to measure the significant reduction in prediction latency.

7. RESULTS

7.1 Classification Model Performance

The performance of the trained classification CNN model, based on the MobileNetV2 architecture, was evaluated on the held-out test dataset. The model achieved a final test accuracy of 74.75%. The training process showed the model accuracy improving over epochs, while the validation accuracy provided a measure of the model's ability to generalize to new data.

The classification report in Table 1 provides a more detailed breakdown of the model's performance for each ripeness class, showing the precision, recall, and F1-score. These metrics offer insights into the model's ability to correctly identify positive instances (precision) and to find all positive instances (recall) for each class. The F1-score is the harmonic mean of precision and recall, providing a single metric that balances both concerns.

Table 1. Classification Report

Class	Precision	Recall	F1-score	Support
Banana(1-2)	0.95	0.90	0.92	20
Banana(3-4)	0.71	0.63	0.67	19
Banana(5-7)	0.66	0.95	0.78	20
Banana(8-10)	0.76	0.65	0.70	20
BananaExpired	0.71	0.60	0.65	20

The training and validation performance is plotted visually using accuracy and loss curves as shown in Fig 1. To further understand the model's performance and identify specific areas of confusion, a confusion matrix was generated as shown in Fig 2. This matrix visually represents the number of correct and incorrect predictions for each class, highlighting which classes are being misclassified.

7.2 Regression Model Performance

The regression model, which also utilizes transfer learning with MobileNetV2, was evaluated on the same test set. The model achieved a Mean Absolute Error (MAE) of 1.44 days on the test set. This indicates that, on average, the model's predictions of the remaining shelf-life are off by approximately 1.44 days. A plot of the training and validation performance is visualized in Fig 3.




Fig 1: Accuracy and loss curves plotted during training and validation for the classification model




Fig 2: Confusion matrix for the classification model over test set.




Fig 3: MAE and loss curves plotted during training and validation for the regression model.

7.3 Edge Model Performance

Both the classification and regression models were converted to the TensorFlow Lite (TFLite) format and evaluated on the test set. The TFLite classification model achieved a test accuracy of 74.75%, and the TFLite regression model achieved a test MAE of 1.44 days. This confirms that the conversion process did not result in a loss of performance for either model.

Crucially, the performance evaluation showed a significant improvement in inference speed. The average prediction latency for the TFLite models was 3.6372 ms. This low latency is critical for real-time applications and demonstrates the effectiveness of TFLite in optimizing models for deployment on edge devices.

7.4 Comparison of Models

The classification model provides a simple and interpretable output, categorizing bananas into distinct ripeness stages. This is useful for quick assessments in a retail environment. The test accuracy of 74.75% shows a good performance for this task.

The regression model, on the other hand, provides a continuous numerical value for the remaining shelf-life, offering more detailed and actionable information for inventory management and supply chain optimization. The MAE of 1.44 days indicates a reasonable level of accuracy for predicting the exact day of shelf life. The choice between a classification and a regression model depends on the specific application. For a high-level overview of fruit ripeness, a classification model may be sufficient. For more precise inventory management and waste reduction, a regression model is more suitable.

8. CONCLUSION

The primary objective of this work was to develop and compare efficient and accurate systems for predicting the shelf life of fruits by leveraging the capabilities of CNNs and Edge Computing. The research successfully demonstrated that this integrated approach provides a powerful solution to the significant challenge of food spoilage in the supply chain.

The developed systems, a classification model and a regression model, both based on the MobileNetV2 architecture, effectively predict the freshness of fruits based on visual data. The classification model achieved a test accuracy of 74.75%, while the regression model achieved a mean absolute error of 1.44 days.

The integration with an edge computing paradigm is a key contribution of this work. By converting the models to TensorFlow Lite, the critical need for real-time, on-site analysis is addressed, reducing latency and dependency on centralized cloud infrastructure. The TFLite models' low average prediction latency of 3.64ms makes the solution practical for deployment in diverse environments across the agricultural supply chain.

In summary, this research validates that the synergy between deep learning-based image analysis and edge computing offers robust, scalable, and effective methods for shelf-life prediction. The successful implementation and validation of both systems underscore their potential to significantly reduce food waste, optimize inventory management, and enhance overall efficiency in the food industry. The choice between the two models depends on the desired level of granularity for the prediction.

9. FUTURE ENHANCEMENTS

While the current systems provide a strong foundation, there are several avenues for future enhancement that could further improve their capabilities and impact:

Expanded Multi-Fruit Dataset: The current models were trained on a dataset of bananas. A significant enhancement would be to expand the dataset to include a wider variety of fruits and vegetables, each with its unique spoilage characteristics. This would improve the models' generalizability and make the system a more versatile tool for the broader agricultural industry.

Integration of Internet of Things (IoT) and Multi-Modal Sensor Data: Future iterations could develop a holistic monitoring system by incorporating data from real-time IoT sensors. This would involve tracking key environmental parameters that affect spoilage, such as ambient temperature, humidity, and ethylene gas concentrations, during transport and storage. Fusing this multi-modal data with the visual information from the CNN could lead to significantly more accurate and nuanced shelf-life predictions.

Development of a Full Decision Support System: The predictive models could become the core of a comprehensive Decision Support System. This would involve creating a user-friendly, web-based dashboard for real-time monitoring and visualization of freshness data and environmental trends. The system could include an automated alert and notification feature to flag batches approaching critical shelf-life thresholds, enabling proactive interventions to reduce waste.

Longitudinal Shelf-Life Prediction: The current regression model predicts the remaining shelf life from a single image. A more advanced system could be developed to analyze a time-series of images and sensor data. By incorporating Recurrent Neural Networks (RNNs) or LSTMs, the model could learn the trajectory of ripening and spoilage, potentially leading to more accurate predictions of the future rate of decay.

Integration with Supply Chain Management Systems: To maximize its practical impact, the system could be integrated via Application Programming Interfaces (APIs) with existing inventory and supply chain management software. This would

allow for automated decision-making, such as dynamically adjusting prices, re-routing shipments based on predicted freshness, or automatically flagging batches of produce for quick sale.

10. REFERENCES

- [1] L. Liu, P. Ouyang, W. Wang, F. Li, and D. Chen, "A review of fruit ripeness detection using computer vision," *Journal of Food Engineering*, vol. 283, p. 110030, 2020.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, 2012, pp. 1097-1105.
- [3] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-Based Learning Applied to Document Recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278-2324, 1998.
- [4] A. Bhargava and A. Bansal, "Fruit recognition using convolutional neural networks," in *2018 8th International Conference on Cloud Computing, Data Science & Engineering (Confluence)*, 2018, pp. 1-5.
- [5] S. S. Pan and Y. C. Wang, "A deep learning approach for fruit quality classification," in *2019 IEEE International Conference on Consumer Electronics-Taiwan (ICCE-TW)*, 2019, pp. 1-2.
- [6] M. Satyanarayanan, "The Emergence of Edge Computing," *Computer*, vol. 50, no. 1, pp. 30-39, 2017.
- [7] M. Z. Uddin, M. T. I. Khan, and M. A. Iqbal, "An IoT-based automated system for monitoring of fruit freshness," in *2019 1st International Conference on Advances in Science, Engineering and Robotics Technology (ICASERT)*, 2019, pp. 1-6.
- [8] M. A. Wazzan, S. A. Al-Ghamdi, and A. M. Al-Ghamdi, "A Real-Time Fruit Quality Classification System Using Deep Learning," *Arabian Journal for Science and Engineering*, vol. 46, pp. 11845-11861, 2021.
- [9] Y. Safari, J. Nakatumba-Nabende, R. Nakasi, and R. Nakibuule, "A Review on Automated Detection and Assessment of Fruit Damage Using Machine Learning," *International Journal of Computer Applications*, vol. 184, no. 1, pp. 1-10, 2022.
- [10] S. S. Shenoy, R. Mishra, and R. Chaturvedi, "Comparative Analysis for Predicting Shelf life of Fruits Using Advanced Deep Learning Approaches," in *2022 International Conference on Applied Artificial Intelligence and Computing (ICAAIC)*, 2022, pp. 586-591.
- [11] S. S. Alahmari and T. Salem, "Food State Recognition Using Deep Learning," *IEEE Access*, vol. 8, pp. 195515-195525, 2020.
- [12] A. Abougharib, M. Awad, and M. Ndiaye, "Remaining Shelf-Life Estimation of Fresh Fruits and Vegetables During Transportation," *IEEE Access*, vol. 8, pp. 175586-175599, 2020.
- [13] A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *2016 IEEE*

- Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.
- [15] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *International Conference on Learning Representations (ICLR)*, 2015.
- [16] Abhinav099802, et al. 2023. Time-fruits Market Availability. Data set. Kaggle. <https://www.kaggle.com/datasets/abhinav099802/time-fruits-market-availability>.
- [17] F. Chollet, "Keras," GitHub, 2015. [Online]. Available: <https://github.com/fchollet/keras>.
- [18] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.
- [19] P. Wes McKinney, "Data Structures for Statistical Computing in Python," in *Proceedings of the 9th Python in Science Conference*, 2010, pp. 51-56.
- [20] H. Zhu and X. Li, "Deep Learning Based Food Category Classification and Recognition," in *2018 International Conference on Sensing, Diagnostics, Prognostics, and Control (SDPC)*, 2018, pp. 634-638.
- [21] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90-95, 2007.
- [22] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 1874-1883.
- [23] M. A. Fadlullah, Z. M., Tang, F., Mao, B., Kato, N., Akashi, O., Inoue, T., & Mizutani, K., "State-of-the-Art Deep Learning: Evolving Machine Intelligence Toward Tomorrow's Intelligent Network Traffic Control Systems," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2432-2455, 2017.
- [24] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779-778.
- [25] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.
- [26] R. Girshick, "Fast R-CNN," in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 1440-1448.
- [27] J. Long, E. Shelhamer, and T. Darrell, "Fully Convolutional Networks for Semantic Segmentation," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3431-3440.
- [28] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, 2015, pp. 91-99.
- [29] V. B. S. Prasath, H. A. F. Al-Dmour, A. H. A. El-Kader, and S. A. H. Al-Hassan, "Fruit image classification using deep learning," in *2017 8th International Conference on Information Technology (ICIT)*, 2017, pp. 810-815.
- [30] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [31] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems," *arXiv preprint arXiv:1603.04467*, 2016.