DeepAgriScan: Intelligent Produce Recognition System for Smart Retail Billing

Dr. K. Siva Krishna[1], J. Mounika[2], M. Jahnavi[3], M. Bhanu Manohar Reddy[4], B. Devi Vara Prasad[5], Dr. K. Siva Krishna[6]

[1]M.Tech.,(Ph.D.,) Associate professor, Information Technology, KHIT
[2,3,4,5]Student, Information Technology, KHIT

Abstract
DeepAgriScan is an intelligent automated system designed to address the limitations of manual fruit identification and billing in modern retail environments. Fruits and vegetables, unlike packaged goods, do not carry barcodes, making automated identification inherently challenging. The proposed system leverages deep learning and computer vision to capture, process, and classify fruit images with high accuracy. A Convolutional Neural Network (CNN) model is trained on a diverse dataset of fruit images and is capable of recognizing fruits based on visual features such as shape, color, and texture. Upon successful identification, the system retrieves the corresponding price from a structured database and automatically generates the billing output, eliminating manual intervention entirely.
The pipeline integrates image preprocessing (resizing, normalization, augmentation), CNN-based feature extraction and classification, price retrieval, and bill generation into a unified end-to-end system. Implemented using Python and TensorFlow, the system achieves an average top-1 classification accuracy of 95.5% across six fruit categories. Experimental results demonstrate that DeepAgriScan significantly reduces checkout time and billing errors compared to traditional manual methods. The system is designed for scalability and can be extended to support additional produce categories, real-time video streams, and integration with weighing machines for weight-based pricing.

Keywords: Deep Learning, Convolutional Neural Network (CNN), Computer Vision, Fruit Recognition, Image Processing, Smart Retail, Automated Billing, TensorFlow, Artificial Intelligence.

1. Introduction

The rapid evolution of Artificial Intelligence (AI) and Machine Learning (ML) has catalyzed transformative changes across diverse sectors, including retail, agriculture, and logistics. In retail environments, particularly supermarkets and grocery stores, the efficiency and accuracy of the checkout process directly determine customer satisfaction and store productivity. While packaged goods rely on standardized barcode scanning for identification, fresh produce such as fruits and vegetables presents a unique challenge: these items are loose, perishable, and devoid of machine-readable labels.
In traditional retail billing, cashiers identify fruits visually and enter codes manually. This process is inherently slow, susceptible to human error, and inconsistent across different operators. Misidentification is particularly common for fruits with similar visual characteristics, such as apples and tomatoes, or oranges and sweet limes. During peak shopping hours, these inefficiencies compound, resu-

lting in long queues, incorrect bills, and reduced customer trust.

To address these challenges, this paper presents DeepAgriScan: Intelligent Produce Recognition for Smart Retail Billing, an end-to-end automated system that employs deep learning and computer vision to identify fruits from camera images and seamlessly generate billing outputs. The system captures a fruit image, applies preprocessing techniques to normalize and enhance image quality, and feeds the processed image into a trained CNN model for classification. Upon identification, the fruit price is retrieved from a structured database and the total bill is computed automatically.

The primary contributions of this work are: (i) a complete end-to-end pipeline from image capture to bill generation; (ii) a CNN-based classification model achieving 95.5% average accuracy across six fruit categories; (iii) integration of recognition and billing into a single deployable system; and (iv) a scalable architecture that supports future expansion to additional produce categories and hardware peripherals.

1.2 Problem Statement

Modern retail billing systems depend on barcode scanning for product identification. However, fresh fruits and vegetables are typically sold as loose, unpackaged items without barcodes, making automated identification inherently difficult. The manual billing approach introduces several operational challenges: increased checkout time, higher likelihood of misidentification, pricing inconsistencies, and elevated labor dependency. Moreover, existing partial automation solutions lack a unified pipeline that integrates recognition, pricing, and bill generation, thereby limiting their practical applicability in retail environments.

The problem is therefore defined as: the development of an intelligent, real-time system capable of automatically recognizing fruits using image processing and deep learning, and seamlessly connecting the recognition output to a billing mechanism, so as to improve speed, accuracy, and overall operational efficiency in retail checkout.

1.3 Objectives

The objectives of this project are as follows:

- To develop a CNN-based fruit recognition model capable of accurately classifying multiple fruit categories from camera images.
- To design and implement an image preprocessing pipeline that enhances image quality for improved model performance.
- To integrate the fruit recognition module with a price database and billing engine into a unified end-to-end system.
- To evaluate the system under real-world conditions including varying lighting, background, and camera quality.
- To design a user-friendly interface suitable for deployment at self-checkout kiosks in retail stores.
- To provide a scalable architecture enabling future expansion to additional produce categories and hardware integrations.

1.4 Scope

DeepAgriScan is scoped for deployment in supermarkets, grocery stores, and self-checkout kiosks handling fresh produce. The system supports recognition of common fruit categories including apples, bananas, oranges, tomatoes, grapes, and mangoes. Image input is accepted via live camera feeds or uploaded images. The current version supports static (per-unit) pricing with provisions for future weight-based pricing integration. The system is developed for desktop and kiosk environments running

Windows or Linux operating systems and is implemented entirely using open-source technologies to minimize deployment cost.

2. Literature Survey

The domain of automated fruit recognition and retail checkout has attracted considerable research interest in recent years. This section reviews key works that inform the design of DeepAgriScan, identifies methodological approaches employed in prior literature, and highlights research gaps that the proposed system addresses.

2.1 Review of Prior Work

Rojas-Aranda et al. (2020) [1] proposed a fruit classification system for retail environments using MobileNetV2 augmented with color histogram features. The system achieved high accuracy even under the presence of plastic bag occlusions and was designed with computational efficiency for edge deployment. However, the system addressed only the classification component and did not integrate pricing or billing functionality.

Xiao et al. (2023) [2] conducted a comprehensive survey of deep learning architectures for fruit detection in agricultural harvesting contexts, covering YOLO, Faster R-CNN, SSD, and ResNet variants. While providing valuable comparative analysis of detection models, the scope was limited to agricultural applications and did not address retail-specific requirements such as billing integration or checkout user interfaces.

Bongulwar and Talbar (2021) [3] proposed a robust CNN model specifically designed to handle intra-class variation and inter-class similarity in fruit images. The model demonstrated improved performance on visually confusable pairs such as apples and tomatoes, but the implementation did not extend to real-time systems or integrate with any billing mechanism.

Additional studies [4, 5] have explored vision-based deep learning approaches for produce recognition in smart retail contexts, demonstrating the feasibility of automated checkout using computer vision. These works confirm that deep learning can reliably automate produce identification but stop short of delivering a complete end-to-end billing solution.

Table 1: Literature Survey Summary

Paper & Year	Authors	Key Contribution	Methodology	Limitations
Fruit Classification for Retail Stores Using Deep Learning (2020)	Jose Luis Rojas-Aranda et al.	Used MobileNetV2 with color features to classify fruits in retail stores, achieving high accuracy	CNN, MobileNetV2, Image Preprocessing	Performance drops under plastic bags or poor lighting
Fruit Detection and Recognition Based on Deep Learning for Automatic Harvesting (2023)	Feng Xiao, Haibin Wang, et al.	Comprehensive review of deep learning architectures for fruit detection and recognition in	YOLO, Faster-RCNN, SSD, ResNet	Focused on harvesting, not retail billing integration

		agriculture		
A Review on Fruit Recognition using CNNs (2023)	Various Authors	CNN-based classification of fruits and vegetables based on visual features; reviewed state-of-the-art methods	CNN, Feature Extraction, Image Processing	No end-to-end billing system integration
Robust CNN Model for Fruit Recognition (2021)	Deepali M. Bongulwar, S.N. Talbar	Proposed a robust CNN to handle intra-class variation and inter-class similarity in fruit images	Deep Learning, CNN	High computational cost; limited dataset diversity
Vision-Based Deep Learning for Smart Retail (2022)	Various Authors	Vision-based deep learning systems to support automated produce recognition in smart retail environments	Deep Learning, Computer Vision	Does not cover complete automated billing pipeline

2.2 Research Gaps

Based on the review of existing literature, the following critical gaps have been identified:

- Lack of End-to-End Automation: Most systems address only fruit classification and do not provide a complete pipeline encompassing identification, price retrieval, bill generation, and output display.
- Absence of Billing Integration: Existing research rarely combines recognition with automated billing, making them impractical for direct retail deployment.
- Poor Real-Time Performance: Many models are evaluated on static datasets under controlled conditions and fail to account for the variability present in live retail environments.
- Inability to Handle Multiple Objects Simultaneously: Most systems are designed for single-object classification and cannot handle multiple fruits in a single transaction.
- High Dependency on Large, Curated Datasets: Deep learning models require substantial labeled training data, and many existing studies use limited or controlled datasets, reducing generalizability.
- Absence of Weight-Based Pricing: Fruits are typically priced by weight in real stores; existing systems assign fixed per-unit prices and do not integrate with weighing hardware.

DeepAgriScan is designed to address gaps (i), (ii), and (iii) directly, while providing an architectural foundation for future resolution of gaps (iv), (v), and (vi).

3. System Analysis

3.1 Existing System

In the existing retail billing paradigm, fruits and vegetables are identified manually by cashiers or store staff. Upon customer arrival at the billing counter, each fruit is visually examined, identified by name or code, and entered into the billing software. Prices are either pre-stored in the system or manually entered. This process is repeated for every item, creating a linear dependency on human effort.

The existing system is functionally adequate for a low-volume, low-diversity environment but is ill-suited for the demands of modern retail. Key disadvantages include: time-consuming manual entry; high error probability during peak-hour operations; inconsistency across operators; inability to differentiate visually similar fruit varieties; and complete absence of AI or image processing capabilities.

3.2 Proposed System: DeepAgriScan

DeepAgriScan replaces the manual identification and entry process with an automated pipeline. A camera captures the fruit image at the checkout point. The image is preprocessed and fed to a trained CNN model, which classifies the fruit. The identified fruit name is used to query a price database, and the total bill is generated automatically. The process is completed within seconds, enabling high-throughput checkout without operator intervention.

Table 2: Comparison of Existing System vs. DeepAgriScan

Feature	Existing System	DeepAgriScan (Proposed)
Identification Method	Manual by store staff	Automated CNN-based deep learning
Barcode Requirement	Mandatory for packaged items	Not required; image-based recognition
Billing Speed	Slow; depends on operator	Fast; automatic within seconds
Error Rate	High due to human errors	Low; model-based classification
Real-Time Processing	Not supported	Fully supported via camera input
Scalability	Limited; manual updates needed	Scalable with additional training data
Self-Checkout Support	Not supported	Supported via kiosk integration

4. System Design and Architecture

4.1 System Architecture Overview

The DeepAgriScan architecture is organized as a six-layer pipeline, illustrated conceptually below. Each layer performs a distinct function, and data flows sequentially from image capture to billing output.

- **Input Layer:** The user places a fruit in front of a camera, or uploads an image through the web interface. The camera captures a high-resolution JPEG/PNG image and passes it to the preprocessing module.
- **Preprocessing Layer:** The captured image undergoes resizing to a standardized 224×224 pixel resolution, normalization of pixel values to the [0, 1] range, and optional augmentation (horizontal flip, brightness jitter) to improve model robustness.
- **Feature Extraction Layer:** The preprocessed image is fed into the CNN model's convolutional and pooling layers, which automatically learn and extract discriminative features related to fruit shape,

color distribution, and texture patterns.

- **Fruit Identification Layer:** The flattened feature vector is passed through fully connected layers with a SoftMax output, producing a probability distribution over the fruit classes. The class with the highest probability is selected as the identified fruit.
- **Price Calculation Layer:** The identified fruit label is used to query the price database. For the current implementation, prices are stored in a structured CSV/MySQL table. The unit price is retrieved and multiplied by the quantity entered or detected.
- **Output Layer:** The system renders the identified fruit name, unit price, quantity, and total bill on the user interface. A printable receipt is generated, completing the checkout transaction.

4.2 CNN Model Architecture

The CNN model is constructed using TensorFlow/Keras. The architecture comprises three convolutional blocks, each consisting of a Conv2D layer with ReLU activation followed by MaxPooling2D, and a Dropout layer to mitigate overfitting. The convolutional blocks are followed by a Flatten layer and two fully connected (Dense) layers, culminating in a SoftMax classification head with neurons equal to the number of fruit categories.

Input images are resized to 224×224×3. The first convolutional block uses 32 filters of size 3×3; the second uses 64 filters; and the third uses 128 filters. Max pooling is applied with a pool size of 2×2 after each block. The final Dense layers contain 512 and 256 neurons respectively, with ReLU activations. The model is trained using the Adam optimizer with a learning rate of 0.001 and categorical cross-entropy loss. Early stopping with a patience of 10 epochs is employed to prevent overfitting.

4.3 Software and Hardware Requirements

The system is implemented using the following software stack:

- Operating System: Windows 10 / Ubuntu 20.04 LTS
- Programming Language: Python 3.8+
- Deep Learning Framework: TensorFlow 2.x / Keras
- Image Processing Library: OpenCV 4.x
- Database: MySQL (fruit price storage); CSV for lightweight deployments
- Development Environment: Visual Studio Code
- Additional Libraries: NumPy, Pandas, Matplotlib, Scikit-learn

Minimum hardware requirements include an Intel Core i3 processor (i5 or above recommended), 4 GB RAM (8 GB recommended), 20 GB available storage, and a standard USB or integrated webcam. GPU acceleration (NVIDIA CUDA-compatible) is recommended for model training but is not required for inference at deployment.

5. Methodology

5.1 Dataset

The model is trained on a dataset comprising approximately 8,000 labeled fruit images across six categories: Apple, Banana, Orange, Tomato, Grapes, and Mango. Images were sourced from publicly available fruit classification datasets, supplemented with images captured in retail-simulated environments. The dataset is split into 80% training, 10% validation, and 10% test subsets using stratified sampling to ensure balanced class representation.

5.2 Image Preprocessing

All images are resized to 224×224 pixels and normalized to the [0, 1] range prior to model input. Data augmentation is applied during training to increase dataset diversity and improve generalization. Augmentation operations include random horizontal and vertical flips, rotation within ±15 degrees, brightness adjustment within ±20%, and random zoom up to 10%. Augmentation is applied exclusively to the training set; validation and test sets receive only resizing and normalization.

5.3 Model Training

The CNN model is trained for up to 100 epochs with a batch size of 32. The Adam optimizer is used with an initial learning rate of 0.001, reduced by a factor of 0.5 if validation loss plateaus for 5 consecutive epochs. Early stopping with a patience of 10 epochs is applied based on validation accuracy. Model checkpointing saves the best-performing weights based on the validation metric. The training process converges within approximately 60-70 epochs, with validation accuracy stabilizing above 94% across all categories.

5.4 Billing Integration

Upon model inference, the predicted fruit class label is passed to the billing module. The billing module queries the price database using the label as a key and retrieves the unit price. The system calculates the total cost as the product of unit price and quantity. The final bill is displayed on the user interface and can be exported as a printed or digital receipt. The price database is managed via a simple administrative interface, enabling store managers to update prices without modifying system code.

6. Results and Discussion

6.1 Classification Performance

The trained CNN model was evaluated on the held-out test set of 1,600 images. Table 3 presents the top-1 accuracy and precision metrics per fruit category.

Table 3: Per-Category Classification Accuracy

Fruit Category	Training Samples	Testing Samples	Top-1 Accuracy (%)	Precision (%)
Apple	1,200	300	96.4	95.8
Banana	1,100	280	97.1	96.5
Orange	1,000	250	94.8	94.2
Tomato	1,150	290	95.2	94.7
Grapes	900	220	93.5	93.1
Mango	1,050	260	96.0	95.4
Overall (avg.)	6,400	1,600	95.5	95.0

The system achieves an average top-1 accuracy of 95.5% and average precision of 95.0% across all six categories. Banana achieved the highest accuracy (97.1%) owing to its distinctive elongated shape and uniform color profile. Grapes exhibited the lowest accuracy (93.5%), likely due to the challenge of distinguishing individual grape clusters from backgrounds with similar color tones. These results are competitive with state-of-the-art lightweight models reported in literature while maintaining the advantage of an integrated billing pipeline.

6.2 System Performance

End-to-end transaction time (from image capture to bill generation) was measured across 100 test transactions in a simulated retail environment. The average transaction time was 2.3 seconds, compared to an average of 8.7 seconds for the manual process, representing a 73.6% reduction in checkout time. The system demonstrated stable performance under varying lighting conditions, with accuracy degrading by less than 2% under low-illumination scenarios when preprocessing was applied.

6.3 Discussion

The results confirm that the CNN-based approach is well-suited for the retail produce recognition task. The integrated billing pipeline addresses the primary research gap identified in the literature review: prior works focused on classification in isolation, whereas DeepAgriScan provides an operational, deployable system. The 95.5% average accuracy, while not perfect, represents a substantial improvement over manual billing accuracy (estimated at 85-90% in high-volume environments based on retail efficiency studies) and significantly reduces variance across operators.

Observed failure cases were primarily concentrated in two scenarios: (i) multiple fruits placed simultaneously in frame without spatial separation, and (ii) severe occlusion by plastic bags under adverse lighting. Both scenarios are targeted in the future work roadmap through multi-object detection and improved preprocessing for low-light conditions.

7. Testing

A multi-phase testing strategy was employed to validate system correctness and performance across all functional and non-functional requirements.

- **Unit Testing:** Each module (image preprocessing, CNN inference, price database query, billing computation) was tested independently using predefined test cases to verify module-level correctness.
- **Integration Testing:** The data flow between adjacent modules was verified: preprocessing output dimensions were confirmed compatible with model input; model output labels were validated against database keys; billing totals were checked against expected values.
- **System Testing:** End-to-end tests were conducted using real fruit images captured in a simulated retail environment. The system was required to correctly identify the fruit and generate the corresponding bill within 3 seconds for each test case.
- **Interface Testing:** The user interface was tested for responsiveness, display accuracy, and error messaging under invalid input conditions (e.g., blurred images, non-fruit objects).
- **Stress Testing:** The system was tested under a load of 50 concurrent image submissions to assess throughput and identify bottlenecks in the inference pipeline.

8. Conclusion

This paper presented DeepAgriScan, an intelligent produce recognition system designed for automated retail billing. By integrating a CNN-based fruit classification model with a structured billing pipeline, the system eliminates manual identification and automated checkout for fresh produce in retail environments. The system achieves an average top-1 classification accuracy of 95.5% across six fruit categories and reduces transaction time by approximately 73% compared to manual processes.

DeepAgriScan addresses a critical gap in existing literature by providing not merely a recognition model but a deployable end-to-end system capable of operating in real retail conditions. The modular architecture ensures maintainability and scalability, enabling straightforward extension to additional produce categories through model retraining with supplemental data.

9. Future Enhancements

- Real-Time Multi-Object Detection: Extending the system to detect and classify multiple fruits simultaneously in a single image frame using object detection architectures such as YOLO or Faster R-CNN.
- Weighing Machine Integration: Connecting the recognition system with digital weighing hardware to enable weight-based pricing, which is the standard pricing model for fresh produce in most retail markets.
- Mobile Application Support: Developing a mobile client to extend the system's reach to smaller retail stores and enable customer-facing self-checkout via smartphone cameras.
- Cloud Deployment and Centralized Price Management: Migrating the price database and model inference to a cloud backend, enabling centralized price updates across multiple store locations.
- Freshness and Quality Assessment: Incorporating ripeness and defect detection capabilities using transfer learning, enabling the system to assess produce quality in addition to identity.
- Expanded Produce Categories: Retraining the model on a broader dataset covering vegetables, exotic fruits, and regional produce to increase system utility across diverse retail contexts.

Acknowledgement

The authors express their sincere gratitude to Dr. K. Siva Krishna, Assistant Professor, Department of Information Technology, Kallam Haranadha Reddy Institute of Technology, for his invaluable guidance and continuous support throughout this project. The authors also thank Dr. KSR Prasad, Professor and Head of the Department, and the management of KHIT, Guntur, for providing the necessary infrastructure and laboratory facilities. Gratitude is also extended to Smart Bridge for providing the internship platform and mentorship that facilitated the development of this work.

References

1. Jose Luis Rojas-Aranda, Jose-Joel Gonzalez-Barbosa, Juan Humberto Sossa-Azuela, Elvia Ruth Palacios-Hernandez, Deni Torres-Roman, "Fruit Classification for Retail Stores Using Deep Learning," Entropy, 2020.
2. Feng Xiao, Haibin Wang, Yueqin Xu, Ruiqing Zhang, "Fruit Detection and Recognition Based on Deep Learning for Automatic Harvesting: An Overview and Review," Agronomy, Vol. 13, 2023.
3. Deepali M. Bongulwar, S. N. Talbar, "Robust CNN Model for Fruit Recognition," International Journal of Image Processing, 2021.
4. Various Authors, "A Review on Fruit Recognition using Convolutional Neural Networks," International Journal of Computer Science and Engineering, 2023.
5. Various Authors, "Vision-Based Deep Learning Systems for Automated Produce Recognition in Smart Retail," Proceedings of the International Conference on Intelligent Systems, 2022.