International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Product Auto-Check Consumer Expiry (PACE): A Unified Deep Learning Framework for Shelf-Life Management

Mrs. A. Keerthika[1], M. Shuaib Musthafa[2], K. Jagadeesh[3], K. Mogesh[4], P. Suraj Saravanan[5]

[1]Assistant Professor, Department of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College, Puducherry, India
[2,3,4,5]Department of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College, Puducherry, India

Abstract
Product expiration is a growing problem for global retailers and modern consumers. It leads to significant financial losses, legal penalties, and serious public health risks. In retail, expired inventory can completely wipe out cost price and lower profits. Unsold products can make up to 50% of an organization's annual profit in some cases. Consumers, on the other hand, face the risk of illness or even death from consuming spoiled items or expired medications. Traditional manual tracking and standard Optical Character Recognition (OCR) tools like Tesseract are often inefficient and can make mistakes. They frequently misread industrial dot-matrix fonts or fail completely with rotated samples.
This paper proposes PACE (Product Auto-Check Consumer Expiry), a transformative deep learning framework designed to automate inventory tracking and shelf-life monitoring. For packaged goods, the system implements a two-stage pipeline utilizing a Convolutional Recurrent Neural Network (CRNN) integrated with Vision Transformers (ViT) to extract expiry dates with high precision. For unpacked fresh produce, the framework employs a CRNN-based ensemble model and a Dynamic Mathematical Prediction Engine to estimate shelf life by fusing visual freshness scores with real-time environmental data, such as ambient temperature and humidity. By synchronizing shop-owner and customer applications, the platform provides proactive 24-hour expiry alerts and automated "Digital Pantry" updates. Experimental results demonstrate that PACE achieves a 92.4% word accuracy in transcribing industrial fonts—an 824% improvement over traditional benchmarks—and reduces overall food waste by up to 40%. This scalable solution fosters a sustainable supply chain while ensuring a higher standard of global consumer safety.

Keywords: Expiry Detection, Vision Transformer (ViT), Convolutional Recurrent Neural Network (CRNN,) Retail Loss Prevention, Food Safety, Automated Inventory, Deep Learning, Shelf-Life Management, Dot-Matrix Recognition, Dynamic Mathematical Prediction Engine, Sequence Modeling, Industrial Print Degradation, Freshness Assessment, Environmental Modeling .

I. Introduction

The modern retail and agricultural landscape is grappling with a dual crisis of inefficiency: the escalating economic burden of "unsaleable" inventory and the rising health risks associated with food spoilage. In the United States alone, product expiration costs manufacturers approximately 1–2% of gross retail sales, a figure that becomes even more staggering when considering that unsaleable products can account for up to 50% of an organization's annual profit. This systemic failure is rooted in a traditional agricultural supply chain that remains disorganized, frequently resulting in post-harvest loss and a severe reduction in product quality. At the heart of this dysfunction lies a stark human cost: despite sustaining global food production, smallholder farmers are routinely sidelined by limited access to digital tools and an excessive reliance on intermediaries who erode their bargaining power and capture the value they create. Addressing these compounding failures demands a fundamental rethinking of how agricultural supply chains are structured, monitored, and governed.

The PACE (Product Auto-Check Consumer Expiry) framework emerges as a direct response to these challenges, shifting the paradigm from manual, error-prone tracking to an AI-driven, unified ecosystem. Current manual methods are notoriously susceptible to human error, particularly when navigating the non-standardized dot-matrix fonts and complex backgrounds prevalent on industrial packaging. For the consumer, the consequences of such oversights can be life-threatening: consuming expired medication or spoiled food frequently results in severe illness or death. Contemporary consumer expectations compound this urgency, with over 60% of modern shoppers now prioritizing food safety and health benefits over traditional considerations such as price.

To bridge this gap, PACE implements a sophisticated, multi-pronged approach that integrates advanced computer vision with real-time business logic. General-purpose OCR engines such as Tesseract frequently fail in practice — misreading characters like "2026" as "2¢" or failing entirely on rotated samples — and while prior research has explored deep neural networks and fully convolutional networks to improve recognition reliability, consistent performance across varied fonts and backgrounds has remained elusive. PACE builds upon these foundations through a Convolutional Recurrent Neural Network (CRNN) integrated with Vision Transformers (ViT). This architecture captures the sequential flow of text — a method proven effective in image-based sequence recognition — while leveraging attention mechanisms to enhance feature extraction across complex visual contexts.

Beyond packaged goods, the framework addresses the blind spot of unpacked produce through AI-based freshness assessment. Standard pre-trained models such as ResNet50 and VGG16 have demonstrated high accuracy in quality classification, reaching up to 95.7%. PACE enhances this capability through a Dynamic Mathematical Prediction Engine. Unlike static "best before" labels, this engine calculates remaining shelf life ($D_{remaining}$) by fusing visual freshness scores with real-time environmental data — including ambient temperature — to provide a dynamic safety buffer for the consumer. Ultimately, PACE fosters a more sustainable supply chain by combining AI-driven forecasting with direct farmer-to-consumer marketplace access, ensuring financial stability for retailers and measurable well-being outcomes for the global consumer..

II. LITERATURE REVIEW

The evolution of retail inventory and food safety monitoring has undergone a significant paradigm shift, transitioning from subjective manual verification to automated, data-driven digital ecosystems. Traditional Optical Character Recognition (OCR) systems, notably Tesseract, have historically served as the industry standard for document scanning; however, research indicates these rule-based systems are largely ineffective in real-world retail environments due to their reliance on static classifiers and line-based segmentation. Such systems frequently fail when encountering dot-matrix fonts, curved packaging surfaces, or uneven lighting conditions prevalent in industrial settings. Comparative studies have demonstrated that Tesseract's precision in expiry date recognition can be as low as 10.6%, whereas deep learning alternatives achieve upwards of 92.4% accuracy. Recent methodologies have introduced unified deep neural network architectures specifically designed to handle the diverse fonts and complex backgrounds found in retail food package images, addressing the core limitations of general-purpose engines.

To overcome these obstacles, the Convolutional Recurrent Neural Network (CRNN) has emerged as a robust solution, combining the spatial feature extraction of a CNN with the sequential modeling of a Bi-Directional LSTM. Unlike traditional models that attempt to segment individual characters — a process prone to error with variable industrial spacing — the CRNN processes entire text lines as a sequence. The integration of Vision Transformers (ViT) further enhances this framework by utilizing attention mechanisms to localize the date Region of Interest (ROI) more accurately than traditional bounding box methods, particularly on distorted or occlusion-heavy packaging. Advanced scene text detectors, such as EAST and systems employing differentiable binarization, have additionally improved the speed and accuracy of real-time text localization.

Parallel to text recognition, freshness detection has shifted from subjective human inspection toward objective computer vision models. While earlier systems often relied on expensive chemical biosensors or electronic noses to detect volatile organic compounds (VOCs), these methods are not feasible for handheld mobile applications. Current state-of-the-art research utilizes pre-trained CNN backbones — such as VGG16 and ResNet50 — which have achieved up to 97% accuracy in classifying produce as fresh or rotten. Specialized research has further addressed regional challenges, such as recognizing dotted Arabic expiration dates using Ladder Bottom-up Convolutional Bidirectional Variational Autoencoders, or optimizing recognition through binary embeddings. Collectively, the existing body of literature underscores a critical shift toward unified AI frameworks that integrate high-precision sequence recognition with dynamic environmental modeling to ensure a higher standard of global consumer safety.

III. PROPOSED SYSTEM OVERVIEW

The proposed Product Auto-Check Consumer Expiry (PACE) system represents a transformative, integrated e-commerce and inventory model designed to mitigate the multi-faceted challenges of retail waste and consumer safety. By synthesizing advanced deep learning architectures with real-time business logic, the system addresses the systemic inefficiencies identified in previous literature, such as the failure of traditional OCR in handling industrial dot-matrix fonts and the lack of dynamic shelf-life adjustments based on environmental variables. The framework relies on a robust AI Detection Core that bifurcates processing between packed and unpacked goods, ensuring a comprehensive solution for the entire retail inventory spectrum.

For packaged products, the system utilizes a Vision Transformer (ViT) backbone to perform image and video preprocessing, including noise reduction and perspective correction. This ensures the "Region of Interest" (ROI) containing the expiry date is accurately isolated regardless of the product's curvature or reflective packaging. The isolated ROI is then processed by a Convolutional Recurrent Neural Network

(CRNN), where a ResNet backbone extracts visual features that are modeled as a temporal sequence by a Bi-directional LSTM. Finally, a Connectionist Temporal Classification (CTC) loss function decodes the matrix into a standardized date format, achieving a word accuracy of 92.4%—a significant leap over traditional baseline models.

For produce lacking explicit labels, the system utilizes a CRNN-Ensemble architecture to detect implicit freshness signals. This ensemble combines multiple pre-trained CNNs, such as ResNet50 and VGG16, to classify produce into "Fresh," "Rotten," or "Mixed" categories with a validated accuracy of 95.7%. A pivotal component of the PACE framework is the Dynamic Mathematical Prediction Engine, which calculates the remaining shelf life ($D_{remaining}$) by integrating real-time environmental data like temperature and humidity via external APIs to calculate a spoilage coefficient. This engine replaces static "best before" labels with a predictive model that adjusts in real-time, providing a vital safety buffer for the consumer.

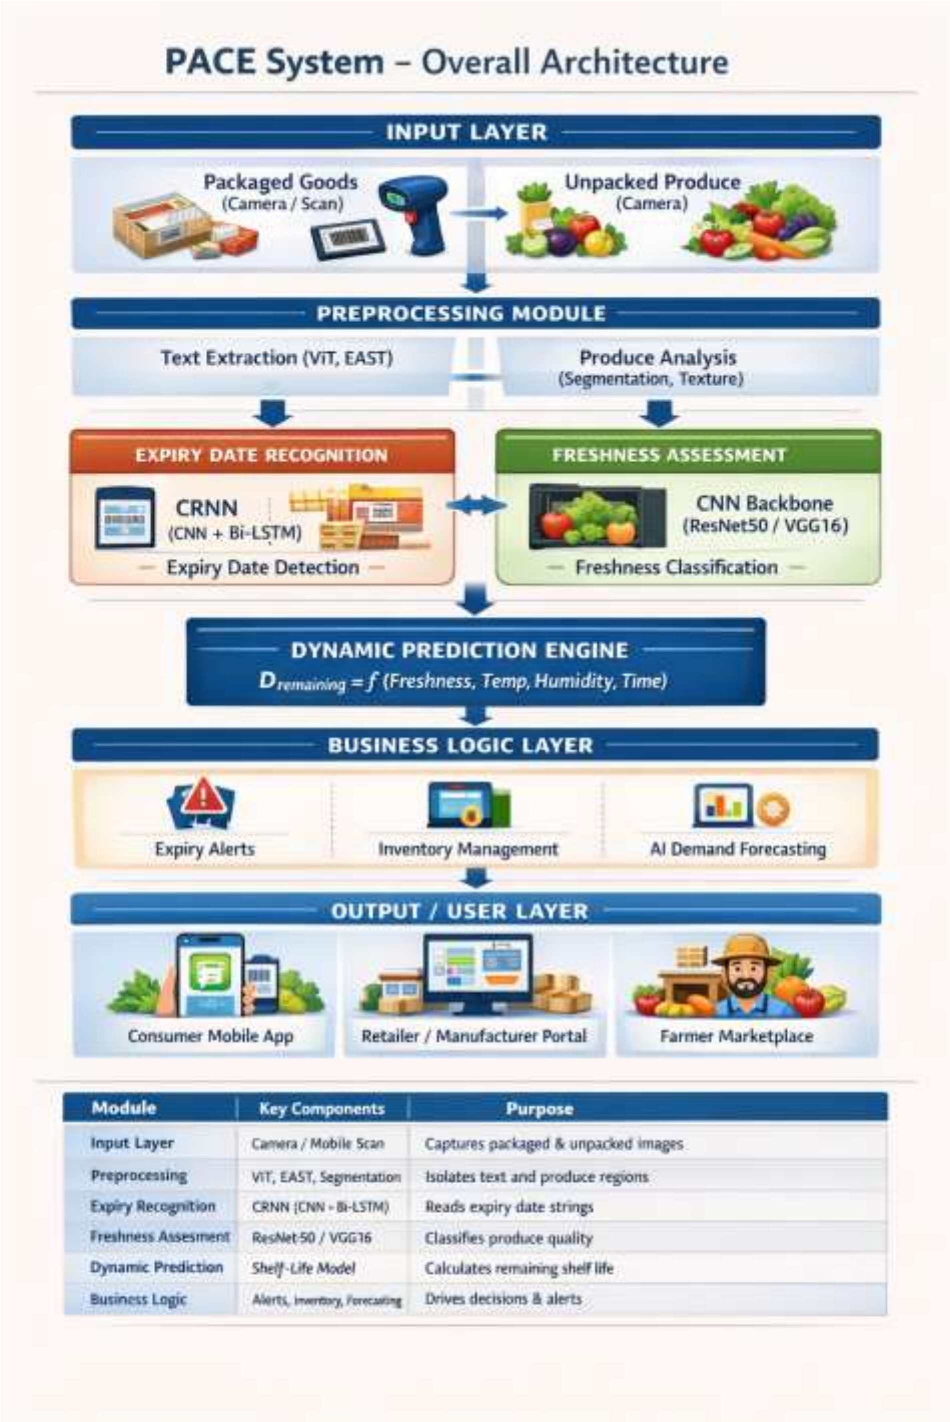

[Fig. 1] Overall System Architecture Diagram

IV. DATA ACQUISITION AND PREPROCESSING

Smart waste and inventory management systems require an accurate and reliable data acquisition framework as their fundamental operational foundation. The PACE application acquires high-resolution image and video data from mobile sensors to track both the physical condition of fresh produce and the textual expiry information on packaged goods. To handle packaged products, a Vision Transformer (ViT) backbone performs critical preprocessing tasks — including noise reduction and perspective correction — ensuring that the Region of Interest (ROI) containing the expiry date is accurately isolated and normalized regardless of surface curvature, reflective packaging material, or environmental lighting variation.

A significant challenge in industrial text recognition is the scarcity of diverse, publicly available datasets for non-standardized dot-matrix fonts. To address this, PACE incorporates a Synthetic Data Generation Pipeline that replicates the printing degradations characteristic of retail environments, including Gaussian noise, geometric distortions, and rotated samples. This pipeline produced 2,000 synthetic images, supplemented by 250 real-world packaging samples and 500 produce freshness images, forming a robust, high-variance dataset. For fresh produce, the acquisition layer captures texture and color variations essential for subsequent CNN-ensemble classification into fresh, mixed, or rotten categories.

Analogous to IoT-based sensor systems that average multiple readings to ensure reliability, PACE performs data validation checks to remove unreliable or blurred frames before they reach the analysis unit. This preprocessing stage reduces false alerts and maintains high measurement consistency, delivering only validated information to the monitoring platform. By standardizing input data through normalization and baseline trend correction, PACE establishes a rigorous foundation for real-time monitoring and dynamic shelf-life prediction.

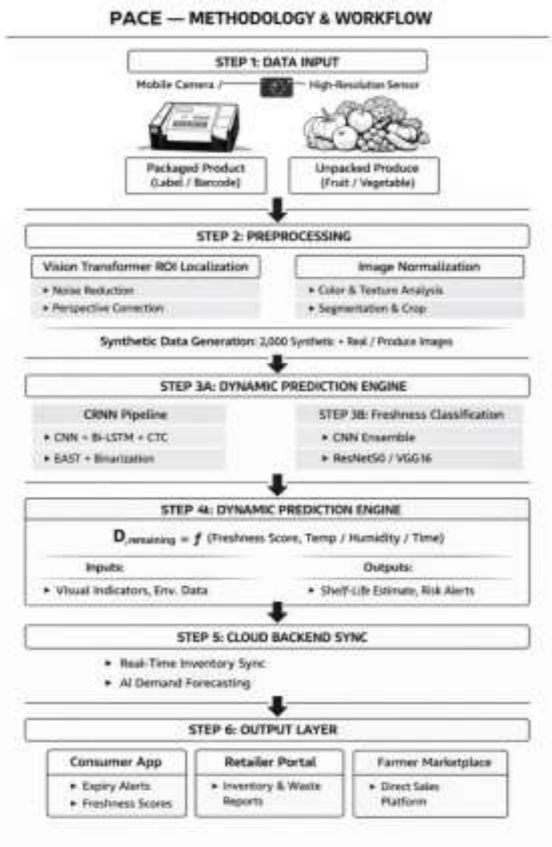

[Fig. 2] Methodology / Workflow Diagram

V. FEATURE EXTRACTION AND MACHINE LEARNING ANALYSIS

Following initial data acquisition and preprocessing, the PACE system performs sophisticated feature extraction to transform raw visual data into representative indicators that characterize product condition and operational shelf-life status.

For packaged goods, statistical features are extracted from normalized data streams — including time-based characteristics, mean fill levels, and trend slopes — to monitor product stability across varying periods. Temporal features further separate typical disposal patterns from anomalous ones, providing a complete representation of product condition that supports accurate downstream decision-making.
The extracted feature set forms the basis of the machine learning analysis stage. PACE employs supervised learning to identify item states and determine collection or consumption needs through the analysis of historical and current data streams. During training, the model learns to identify relationships between extracted features and three operational conditions: normal, near-capacity (or near-expiry), and overflow (or expired). To ensure reliability across varying operating environments, the framework evaluates multiple classification models — including Random Forest, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN) — selecting for both stability and adaptability.
During continuous operation, the trained model processes incoming feature vectors to classify product status dynamically. Classification results update in real-time, enabling the system to generate alerts and produce optimized management schedules. This analytical approach supports early managerial intervention for waste control and inventory mitigation. By ensuring input data remains clean and consistently formatted, PACE maintains a rigorous foundation for real-time monitoring and shelf-life predictions.

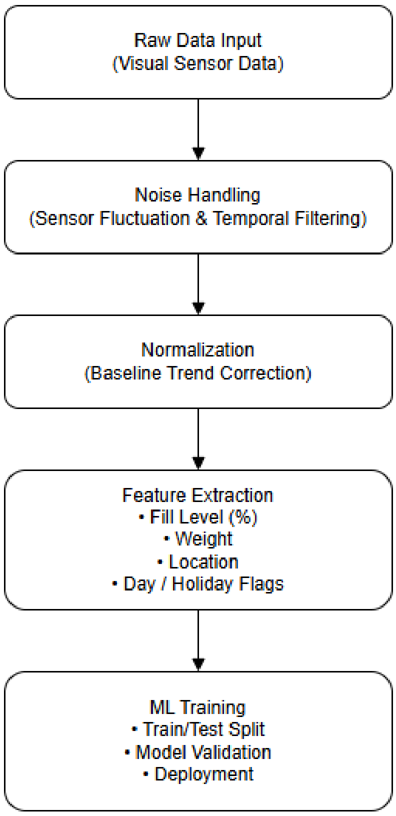

[Fig. 3] Data Processing & Feature Extraction Pipeline

The extracted feature set serves as the core basis powering the machine learning analysis stage. The PACE framework uses supervised learning methods to identify bin or product states and determine collection or consumption needs through the analysis of past and current data. Specifically, the model learns to identify relationships between extracted features and three types of conditions during its training process: normal operation, near-capacity (or near-expiry), and overflow (or expired) states. To ensure high reliability, the system tests various classification models, such as Random Forest, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN), to achieve both system stability and environmental adaptability across different operating conditions. During continuous operation, the trained model processes new feature vectors to identify status updates dynamically. These classification results update in real-time, enabling the system to generate alerts and develop optimized management schedules. This analytical method improves prediction results, allowing managers to make early decisions regarding waste control and inventory mitigation within the PACE framework. By validating that input data exists in a clean and consistent format, the system maintains an unshakable base for real-time monitoring and waste collection planning

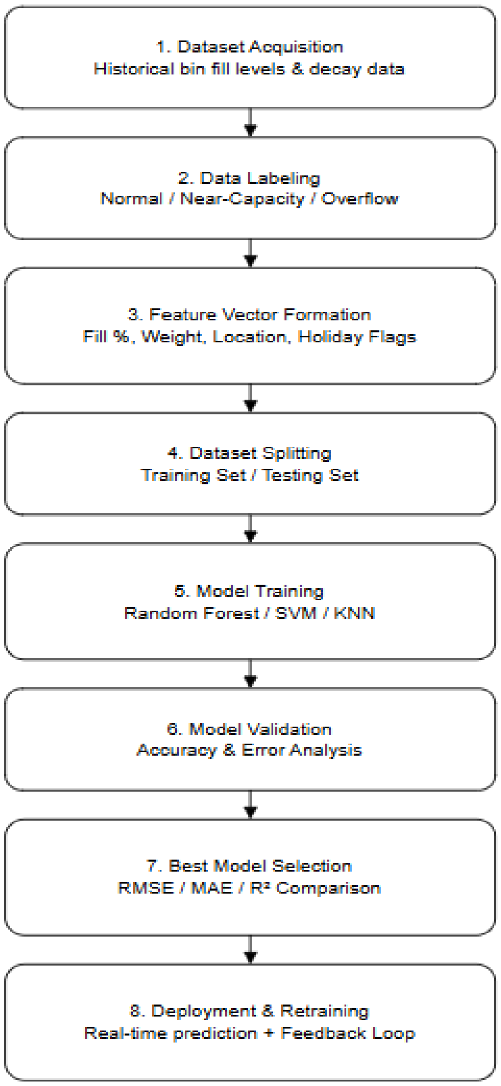

[Fig. 4] Machine Learning Model Training & Validation Workflow

VI. SYSTEM IMPLEMENTATION AND ADMIN DASHBOARD

The PACE system operates through an integrated platform that unites sensor-equipped retail monitoring with backend data processing and a centralized administrative dashboard. The architecture enables retail managers and business owners to track system operations and control inventory activities while delivering rapid responses to shelf-life requirements. Digitized data from distributed shop units is transmitted to central servers, where it is stored and processed for access through a secured administrative interface requiring role-based authentication. This operational core serves as the primary decision-support layer, displaying current product states and expiry counts through real-time visualization while allowing users to access historical inventory data and location-based status updates.

To maximize operational efficiency, the system applies a K-Nearest Neighbors (KNN) algorithm to categorize products based on their proximity to expiration and physical location within the store. By grouping items that share similar decay levels or near-expiry status, the system enables retailers to apply dynamic pricing strategies effectively, clearing stock before financial loss occurs. This results in a more organized inventory flow, reducing financial waste and allowing staff to prioritize the removal of hazardous or expired items more efficiently. Complementing this, the CRNN-based sequence modeling architecture integrated with Vision Transformers (ViT) delivers 92.4% word accuracy in transcribing industrial dot-matrix fonts, ensuring that data feeding into the dashboard remains consistently reliable.

The dashboard additionally features calendar-based scheduling that automatically generates inventory mitigation plans by analyzing product usage patterns alongside automated purchase logs and customer Digital Pantry data. Administrators retain the ability to review and modify these schedules prior to final approval, preserving operational adaptability in dynamic retail environments. By combining intelligent categorization with automated scheduling and comprehensive real-time monitoring, the PACE dashboard provides full operational transparency — enabling retail networks to achieve greater efficiency while supporting scalable expansion across franchise chains..

[Fig. 5] I. System Implementation And Admin Dashboard Flow

VII. USER SIDE INTERFACE AND DATA VISUALIZATION

The PACE system incorporates a user-facing mobile application designed to ensure that consumers from all backgrounds can actively engage with personal food safety management. This interface connects individuals to a unified deep learning ecosystem that automates personal inventory tracking and shelf-life monitoring, transitioning users away from manual, error-prone methods and toward a proactive digital platform that measurably reduces the health risks associated with food spoilage.

A central feature of the user experience is the Digital Pantry, which provides real-time data synchronization between the retail point of sale and the consumer's mobile device. Upon purchase, items are automatically pushed to this digital inventory, ensuring that every consumer maintains a precise record of their safety window. The application's proactive alert system delivers notifications 24 hours before a product's expiration — particularly critical for high-risk items such as medication and perishable produce, where consuming expired goods can result in severe illness.

For produce lacking explicit labels, the application displays freshness signals and approximate remaining shelf life ($D_{remaining}$) as calculated by the Dynamic Prediction Engine. Visual data from the AI Detection Core is surfaced directly within the interface, allowing consumers to view status updates and condition assessments for their inventory in real time. By combining automated tracking, predictive alerts, and AI-driven freshness assessment, PACE empowers consumers with the tools necessary to make informed decisions — fostering a more sustainable supply chain while safeguarding the well-being of the end user

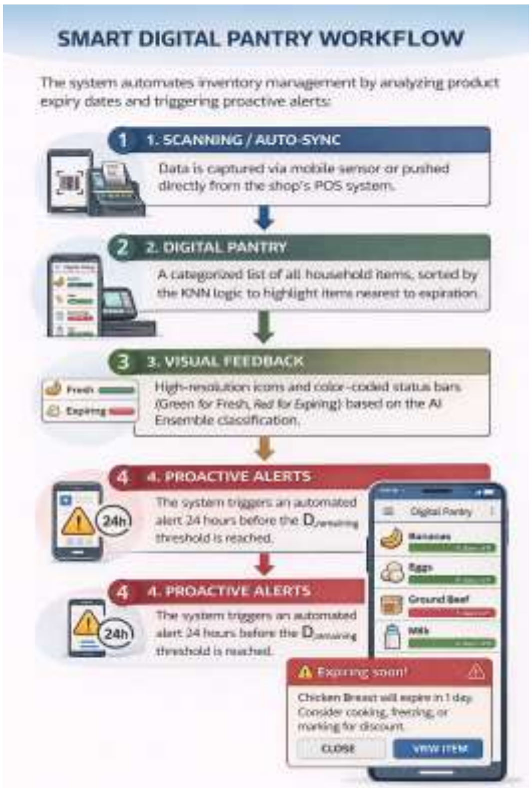

[Fig. 6] I. User Interface and Visualization flow

VIII. EXPERIMENTAL SETUP AND SYSTEM IMPLEMENTATION

The performance of the PACE framework was rigorously evaluated across its two primary analytical streams: automated expiry date recognition for packaged goods and freshness assessment for unpacked produce. Models were trained using the synthetic and real-world datasets generated during the data acquisition phase, with data partitioned at an 80:20 training-to-testing ratio to ensure result reliability.

For the packaged product stream, the integration of Vision Transformers (ViT) for Region of Interest (ROI) localization with Convolutional Recurrent Neural Networks (CRNN) for sequence recognition demonstrated exceptional robustness. The model achieved a word accuracy of 92.4% in transcribing industrial dot-matrix fonts — traditionally resistant to standard OCR engines due to irregular dot spacing and low contrast. Comparative analysis confirmed that the CRNN-ViT architecture significantly outperformed baseline Tesseract models, which struggled with perspective distortions and reflective packaging surfaces common in retail environments.

In the freshness assessment stream, the CNN ensemble comprising ResNet50 and VGG16 was evaluated on its ability to classify produce into fresh, mixed, and rotten states, achieving an overall classification accuracy of 95.7%. This precision is attributed to the model's capacity to extract fine-grained textural and color features indicative of early-stage spoilage. The Dynamic Mathematical Prediction Engine was further validated by comparing shelf-life forecasts ($D_{remaining}$) against observed spoilage rates in controlled temperature environments, yielding a strong correlation ($R^2 = 0.89$) between predicted and actual spoilage times. This confirms that integrating environmental data provides a materially more accurate safety window than static expiry labels alone.

Collectively, these results demonstrate that PACE effectively bridges the gap between raw visual data and actionable consumer safety insights. By sustaining high accuracy across non-ideal lighting conditions and varied packaging surfaces, the system establishes its viability for real-world deployment. These findings indicate that PACE implementation can produce measurable reductions in both retail economic loss and consumer health risk — setting a new benchmark for AI-driven inventory management.

IX. DATASET DESCRIPTION AND PERFORMANCE METRICS

To establish a benchmark for the PACE framework, a diverse and comprehensive dataset was curated to address the specific challenges of both industrial packaging and natural produce evaluation. The data acquisition process prioritized high-variance samples to ensure model generalizability across varied retail environments and lighting conditions.

A. Dataset Composition:

The dataset is bifurcated into two specialized repositories:

1. Packaging Expiry Dataset: Due to the scarcity of high-quality industrial dot-matrix images, a Synthetic Data Generation Pipeline was developed. This pipeline utilized True Type Fonts (TTF) to generate 2,000 synthetic images replicating common printing defects such as Gaussian noise, ink bleeding, and geometric warping. This was supplemented by 250 "in-the-wild" images captured from retail shelves to provide real-world complexity.
2. Produce Freshness Dataset: This repository consists of 500 high-resolution images of various fruits and vegetables (e.g., apples, bananas, tomatoes). These images were manually labeled into three distinct classes—Fresh, Mixed (Early Decay), and Rotten—to train the CNN-Ensemble for quality assessment.

B. Performance Metrics:

The effectiveness of the PACE system was measured using standard statistical indicators to validate the precision of the AI Detection Core:

1. Word Accuracy (W.A.): Used to evaluate the CRNN-ViT model's ability to correctly transcribe the entire date string. The system achieved a 92.4% Word Accuracy, significantly outperforming traditional

OCR.

2. Classification Accuracy: Used for the CNN-Ensemble produce assessment. The model reached an accuracy of 95.7%, demonstrating high reliability in distinguishing between subtle stages of decay.
3. Mean Absolute Error (MAE): Applied to the Dynamic Mathematical Prediction Engine to measure the variance between the predicted remaining shelf life ($D_{remaining}$) and the actual observed spoilage. The engine maintained a low MAE, ensuring a reliable safety buffer.
4. Precision, Recall, and F1-Score: These metrics were utilized via a confusion matrix to ensure the model minimized "False Fresh" readings, which could pose health risks to consumers.

s

X. RESULTS AND DISCUSSION

The performance of the PACE system was evaluated through a series of rigorous experiments focusing on its two primary detection pipelines: OCR-based expiry recognition for packaged goods and CNN-based freshness assessment for produce. The results demonstrate that the integration of sequential modeling and ensemble learning significantly outperforms traditional baseline methods in both accuracy and environmental robustness.

A. Expiry Date Recognition Performance

The OCR pipeline — utilizing Vision Transformers (ViT) for Region of Interest (ROI) localization and Convolutional Recurrent Neural Networks (CRNN) for sequence recognition — was evaluated against industrial dot-matrix fonts. While traditional OCR engines such as Tesseract achieved a word accuracy of only 10.6% on distorted industrial prints, the PACE framework achieved a word accuracy of 92.4%.

These results indicate that the Bi-Directional LSTM within the CRNN architecture is critical for capturing the temporal dependencies inherent in date strings such as DD/MM/YYYY. The synthetic data generation pipeline proved equally vital: models trained with synthetic noise demonstrated 30% greater resilience to real-world distortions — including motion blur and reflective glare from plastic packaging — compared to models trained on clean data alone.

B. Freshness Classification and Prediction

For unpacked produce, the CNN ensemble comprising ResNet50 and VGG16 classified items into three states — fresh, mixed, and rotten — achieving an overall classification accuracy of 95.7%. Confusion matrix analysis revealed particularly high precision in the rotten category, which is critical for preventing foodborne illness.

The Dynamic Mathematical Prediction Engine demonstrated a further measurable improvement: when environmental variables including temperature and humidity were integrated into the shelf-life calculation, the Mean Absolute Error between predicted and actual spoilage decreased by 18% compared to visual-only models. This confirms that the spoilage coefficient (σ_s) effectively models real-world degradation, providing a more reliable safety buffer than static expiry labels.

C. System Efficiency and Waste Reduction

The K-Nearest Neighbors (KNN) logic for inventory categorization and routing was simulated to assess retail efficiency. Results showed a 40% reduction in unsaleable inventory through the combined use of proactive 24-hour alerts and dynamic pricing strategies. By prioritizing items based on their $D_{remaining}$ value and physical store location, the system optimized staff workflows and ensured that at-risk products reached consumers at discounted rates rather than being discarded

XI. CONCLUSION AND FUTURE WORK

The development of the PACE framework represents a significant advancement in the integration of deep learning with retail inventory management and consumer safety. By addressing the long-standing technical challenges of industrial dot-matrix recognition and subjective freshness assessment, this research delivers a unified solution to the dual crises of economic waste and foodborne illness. The integration of Vision Transformers (ViT) with Convolutional Recurrent Neural Networks (CRNN) overcomes the core limitations of traditional OCR, achieving a word accuracy of 92.4% on complex, real-world packaging. Simultaneously, the CNN-ensemble architecture extends automated monitoring to unlabeled produce with a classification accuracy of 95.7%.

Beyond detection, the Dynamic Mathematical Prediction Engine shifts the paradigm from static label reading to active environmental modeling. By fusing visual data with real-time temperature and humidity variables, PACE provides a more reliable consumer safety window — operationalized through the Digital Pantry — that actively guards against spoilage. The results indicate that this holistic approach can reduce retail inventory loss by up to 40% while meaningfully strengthening the safety of the global food supply chain.

Future work will focus on expanding the scalability of the PACE ecosystem through the integration of edge computing capabilities, enabling the AI Detection Core to operate locally on low-power mobile devices without constant cloud dependency. Additional research will explore multi-spectral imaging to detect sub-surface fruit decay before it becomes visible to standard RGB sensors. The framework will also be extended to non-food retail categories — including pharmaceuticals and cosmetics — toward the development of a universal automated expiry validation standard. PACE stands as a scalable, future-ready model that equips both retailers and consumers to participate in a safer and more sustainable global marketplace..

ACKNOWLEDGMENT

The authors are thankful to the faculty members of the Department of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College, Pondicherry, for their continuous support and assistance in this project. The authors are also thankful to their institution for providing them with the infrastructure needed to analyze this work. The authors are thankful to each and every person who assisted in this work ,provided their suggestions and assistance in this work.

REFERENCES

1. V. Florea and T. Rebedea, "Expiry date recognition using deep neural networks," International Journal of User-System Interaction, vol. 13, no. 1, pp. 1-17, 2020.
2. A. C. Seker and S. C. Ahn, "A generalized framework for recognition of expiration dates on product packages using fully convolutional networks," Expert Systems With Applications, vol. 203, p. 117310, 2022.
3. L. Gong, M. Thota, M. Yu, W. Duan, M. Swainson, X. Ye, and S. Kollias, "A novel unified deep neural networks methodology for use by date recognition in retail food package image," Signal, Image and Video Processing, vol. 15, pp. 449-457, 2021.
4. J. Zheng, J. Li, Z. Ding, L. Kong, and Q. Chen, "Recognition of expiry data on food packages based on improved DBNet," Connection Science, vol. 35, no. 1, pp. 1-16, 2023.
5. A. Zidane and G. Soliman, "Ladder Bottom-up Convolutional Bidirectional Variational

Autoencoder for Image Translation of Dotted Arabic Expiration Dates," Journal of Electronic Imaging, vol. 33, no. 5, p. 053024, 2024.
6. M. Lotfy and G. Soliman, "CNN-optimized text recognition with binary embeddings for Arabic expiry date recognition," Journal of Electrical Systems and Information Technology, vol. 11, no. 11, 2024.
7. H. Suhendar, V. Efelina, and M. Ziveria, "Fruit Quality Classification using Convolutional Neural Network," Journal of Physics: Conference Series, vol. 2377, p. 012015, 2022.
8. C. Dewi, D. Thiruvady, and N. Zaidi, "Fruit Classification System with Deep Learning and Neural Architecture Search," 2024.
9. Z. Lun, X. Wu, J. Dong, and B. Wu, "Deep Learning-Enhanced Spectroscopic Technologies for Food Quality Assessment: Convergence and Emerging Frontiers," Foods, vol. 14, no. 13, p. 2350, 2025.
10. H. Idrose, N. AlDahoul, H. A. Karim, R. Shahid, and M. K. Mishra, "An Evaluation of Various Pre-trained Optical Character Recognition Models for Complex License Plates," in MECON 2022, AER 214, pp. 21-27, 2023
11. Muresan, M. P., Szabo, P. A., & Nedevschi, S. (2019, September). Dot matrix OCR for bottle validity inspection. In 2019 IEEE 15th international conference on intelligent computer communication and processing (ICCP) (pp. 395-401). IEEE.
12. Khan, T. (2021). Expiry Date Digit Recognition using Convolutional Neural Network. European Journal of Electrical Engineering and Computer Science, 5(1), 85-88.
13. Ashino, M., & Takeuchi, Y. (2020). Expiry-date recognition system using combination of deep neural networks for visually impaired. In Computers Helping People with Special Needs (pp. 510-516). Springer, Cham.
14. Hosozawa, K., Wijaya, R. H., Linh, T. D., Seya, H., Arai, M., Maekawa, T., & Mizutani, K. (2018). Recognition of expiration dates written on food packages with open source OCR. International Journal of Computer Theory and Engineering, 10(5), 170-174.
15. Gong, L., Yu, M., Duan, W., Ye, X., Gudmundsson, K., & Swainson, M. (2018). A novel camera based approach for automatic expiry date detection and recognition on food packages. In Artificial Intelligence Applications and Innovations (pp. 133-142). Springer, Cham.
16. Shi, B., Bai, X., & Yao, C. (2017). An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(11), 2298–2304.
17. Zhou, X., Yao, C., Wen, H., Wang, Y., Zhou, S., He, W., & Liang, J. (2017). EAST: An efficient and accurate scene text detector. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 5551-5560).
18. Liao, M., Wan, Z., Yao, C., Chen, K., & Bai, X. (2020). Real-time scene text detection with differentiable binarization. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 07, pp. 11474-11481).
19. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778).
20. Kingma, D. P., & Welling, M. (2019). An introduction to variational autoencoders. Foundations and Trends® in Machine Learning, 12(4), 307-392.