IJIREEICE
ISSN (O) 2321-2004, ISSN (P) 2321-5526
International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.021 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 14, Issue 8, August 2026
DOI: 10.17148/IJIREEICE.2026.14806

Computer Vision-Based Stock Quality Analysis and Rotten Quantity Detection for Fresh Produce in Quick-Commerce Warehouses

Vignesh Murali, Sarvesh S, Pranav H, Charulatha R.T

Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India

Abstract: Quick-commerce fulfillment centers (dark stores) operate under strict order-to-delivery timelines, making manual quality assessment of perishable fruits and vegetables highly inefficient and error-prone. This research paper presents a robust, universal Computer Vision (CV) system integrated with a non-linear regression model for automated freshness classification, rot segmentation, and quantitative stock loss estimation. The system utilizes HSV color space analysis and edge gradient filters to segment the produce body, eliminate complex backdrops (including wooden tables and studio backgrounds), and isolate localized rot spots. Furthermore, we introduce a quantitative stock estimation model that maps pixel density ratios to physical weight, allowing operators to automatically determine the exact quantity (in kilograms) of rotten produce within a bulk batch. Tested across tomato, banana, capsicum, and spinach images, the system successfully eliminates false positives from healthy features like green stems and natural yellow/red skin variations, achieving a classification accuracy of 96.8%.

Keywords: Computer Vision, Freshness Detection, Stock Quality Analysis, Rot Spot Masking, Image Segmentation, Quantity Estimation, Quick Commerce.

I. INTRODUCTION

The rapid growth of quick-commerce platforms (such as Zepto and Blinkit) has transformed urban retail, promising grocery delivery within 10 to 15 minutes. To sustain this throughput, dark stores rely on speed and high efficiency. However, fresh produce is highly perishable, and manual inspections at sorting bays are subjective, slow, and susceptible to human error. This leads to defective produce reaching customers, resulting in high refund rates and customer churn.

Computer vision (CV) offers a scalable alternative. However, standard color-thresholding algorithms fail when confronted with the complex visual environments of dark stores, such as wooden crates, cardboard backgrounds, and shadows. Additionally, standard models frequently misclassify natural green stems on fresh tomatoes, or the healthy yellow and red skins of fresh bell peppers, as rot. To address these limitations, this paper proposes an advanced, universal computer vision framework that incorporates background segmentation, green stem filtering, texture edge entropy, and a quantitative stock estimation model.

Beyond classification accuracy, dark store operations demand a system that can be dropped into an existing receiving-bay workflow without additional hardware such as calibrated lighting rigs, depth sensors, or hyperspectral cameras. A single top-down RGB camera feed, already present at most sorting stations for loss-prevention purposes, is sufficient input for the proposed pipeline. This constraint shaped every design decision in the framework: each stage relies only on standard RGB/HSV arithmetic and morphological operations that execute in real time on commodity edge hardware, avoiding the latency and cost overhead of GPU-bound deep networks.

The key contributions of this work are summarized as follows:

1) A universal background and stem segmentation stage that generalizes across wooden crates, cardboard trays, and studio backdrops without per-SKU recalibration.

2) A localized rot-spot isolation model that separates true decay signatures from natural skin-color variance, eliminating the false positives that plague naive thresholding methods.

3) A quantitative, pixel-to-weight stock estimation model that converts a rot-area ratio directly into a kilogram-level spoilage estimate for inventory write-offs.

4) An empirical evaluation across four produce categories — tomato, banana, capsicum, and spinach — demonstrating 96.8% classification accuracy under real dark-store imaging conditions.

DOI: 10.17148/IJIREEICE.2026.14806

II. RELATED WORK

Automated produce quality assessment has been studied extensively in agricultural post-harvest research, though most existing approaches were designed for controlled packhouse or laboratory environments rather than the fast-moving, visually cluttered conditions of a quick-commerce dark store.

A. Color-Threshold and Rule-Based Methods

Early defect-detection systems relied on fixed RGB or HSV thresholds tuned to a single fruit variety under uniform studio lighting. While computationally inexpensive, these rule-based approaches are brittle: a threshold tuned for ripe red tomatoes routinely misfires on green stems, sun-blushed shoulders, or cultivar-to-cultivar color drift, and performance degrades sharply once the background changes from a plain white sheet to a wooden crate or a cardboard liner — both common at dark-store receiving bays.

B. Deep Learning-Based Freshness Classifiers

Convolutional neural network (CNN) classifiers trained on curated produce-image datasets have reported strong accuracy for binary fresh/rotten labeling. However, these models are typically trained on single, centered, well-lit fruit images and act as coarse classifiers rather than localized rot detectors — they cannot indicate where the rot is located on the produce surface, nor can they be trivially converted into a quantitative weight-loss estimate. Their reliance on GPU inference also complicates deployment on the low-cost edge devices typically installed at high-throughput sorting bays.

C. Segmentation-Based Quality Estimation

More recent work has explored pixel-level segmentation to localize surface defects on individual fruits, primarily for high-value export produce such as apples and citrus, where a single unit is scanned in isolation on a conveyor with fixed lighting. These pipelines generally assume single-object framing and do not address the bulk-batch scenario relevant to quick-commerce, where dozens of individual items of mixed size and orientation must be assessed together and translated into an aggregate stock-loss figure in kilograms.

The present work differs from prior art in three respects: (i) it is explicitly designed to tolerate the uncontrolled, multi-texture backgrounds found in dark-store receiving bays; (ii) it separates natural produce features (stems, vines, cultivar-specific skin tones) from true decay signatures using a dedicated exclusion layer rather than a single global threshold; and (iii) it closes the loop from pixel-level detection to a physically meaningful kilogram-level stock-loss estimate that can be consumed directly by warehouse management software.

III. PROPOSED METHODOLOGY

The proposed architecture consists of three core layers: Background & Stem Segmentation, Localized Rot Spot Isolation, and Quantitative Weight Estimation. Figure captures from a top-down camera are first normalized for exposure, then passed sequentially through each layer, as summarized in the pipeline description below.

A. System Overview

At the receiving bay, a fixed-position camera captures an overhead frame of the produce batch as it is emptied onto the inspection tray. The raw frame is converted from RGB to HSV color space, since HSV separates chromatic information (hue, saturation) from illumination (value), making the subsequent segmentation stages far more robust to the uneven lighting typical of a warehouse floor compared to operating purely in RGB space. The frame then passes through the background/stem exclusion mask, the rot-spot isolation filter, and finally the quantitative weight-estimation module, which emits a per-batch spoilage report.

B. Background & Stem Segmentation

To ensure analysis is performed strictly on the produce surface, background pixels (such as white backdrops, studio grey surfaces, and wooden tables) are removed using a multi-band HSV and RGB threshold filter. In addition, green stems are isolated to prevent false rot triggers on fresh produce:

Background Filter: Pixels with RGB values > 220, or uniform wood-tan ratios (R > G > B with center weighting), are segmented out.

Stem & Vine Recognition: Stems are isolated using HSV limits: Hue 65° to 165°, Saturation > 18%, Value > 25%. These regions are excluded from the rot calculation.

After the two filters are applied, a morphological opening operation (3×3 structuring element) removes isolated noise pixels along the produce boundary, and a connected-component analysis discards any residual background islands smaller

Impact Factor 8.021 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 8, August 2026

DOI: 10.17148/IJIREEICE.2026.14806

than 0.5% of the frame area, yielding a clean produce-body mask that forms the denominator (Pbody) used later in Section IV.

C. Localized Rot Spot Isolation

Rotten patches differ from natural skin colors (like red tomatoes or yellow bananas) by exhibiting local blemish drops, dark necrotic zones, and mold. The system flags pixels as rotten if they do not match healthy produce skin ranges and satisfy darkness criteria (Brightness < 50) or muddy discoloration (R < 120, G < 80, B < 60).

Each candidate rot pixel is additionally validated against its 8-connected neighborhood: isolated single-pixel outliers, which typically arise from camera sensor noise or specular highlights, are suppressed unless they belong to a contiguous blemish cluster of at least 15 pixels. This neighborhood check reduces spurious rot flags on otherwise healthy produce without weakening sensitivity to genuine decay spots, which by nature spread across contiguous regions of the surface.

D. Texture Edge Entropy Filtering

Certain healthy surface textures — such as the ridged skin of a bell pepper or the fibrous cut end of a spinach stalk — can locally satisfy the darkness or discoloration criteria used in rot detection. To suppress these cases, a Sobel edge-gradient filter is applied to the candidate rot regions, and the local edge entropy is computed over a sliding 5×5 window. Genuine rot and mold growth exhibit low-entropy, texturally flat regions, whereas natural ridges and creases produce high-entropy, structured gradients. Candidate pixels with edge entropy above an empirically set threshold are reclassified as healthy, which is the mechanism primarily responsible for eliminating false positives on ridged capsicum skins and fibrous spinach stems observed during testing.

E. Non-Linear Regression Calibration

Because different produce types exhibit different pixel-density-to-mass relationships (a tomato and a bunch of spinach of the same pixel footprint do not weigh the same), a per-category non-linear regression model is fitted offline using reference batches of known weight. The model maps a produce type's average pixel body area to an expected mass-per-pixel coefficient, which calibrates the quantitative weight-estimation stage described in Section IV so that the reported kilogram figures remain accurate across the tomato, banana, capsicum, and spinach categories evaluated in this study.

IV. QUANTITATIVE STOCK ROT DETECTION MODEL

A key feature of the proposed system is its ability to translate pixel-level rot area ratios into physical stock weight quantities. When a bulk batch of produce is scanned at the receiving bay, the total stock weight (Wtotal in kg) is logged. The system computes the active produce body pixels (Pbody) and the localized rot pixels (Prot). The weight of the rotten stock (Wrotten) and healthy stock (Whealthy) is estimated as follows:

$$Wrotten = Wtotal \times (Prot / Pbody)$$
$$Whealthy = Wtotal - Wrotten$$

This formulation allows dark store operators to immediately quantify inventory financial losses and write off damaged batches automatically before they are placed on racks.

For example, consider a 20 kg crate of tomatoes logged at receiving. If the system computes Pbody = 159,783 pixels and Prot = 3,355 pixels for the batch, the estimated rotten weight is Wrotten = 20 × (3,355 / 159,783) ≈ 0.42 kg, leaving Whealthy ≈ 19.58 kg available for dispatch. This per-batch figure is written directly to the warehouse management system, allowing automatic partial write-offs instead of rejecting or accepting an entire crate on a binary pass/fail basis.

Because the regression coefficients from Section III-E are produce-specific, the same pixel-ratio formulation generalizes across categories with very different densities — from dense tomatoes to leafy, low-density spinach bunches — without requiring a separate detection model to be trained for each SKU.

V. EXPERIMENTAL RESULTS

A. Dataset and Experimental Setup

The framework was evaluated on a dataset of fresh and rotten produce images spanning four categories — tomato, banana, capsicum, and spinach — captured under dark-store-representative conditions, including wooden crate backgrounds, cardboard trays, and plain studio backdrops, to test the robustness of the background segmentation stage described in Section III-B. Each image was independently labeled as FRESH or ROTTEN by manual inspection to establish ground truth. The classification threshold was set to a combined Rot Area Ratio of 10.0% to decide if the

DOI: 10.17148/IJIREEICE.2026.14806

produce batch is FRESH or ROTTEN, a threshold chosen empirically to balance false rejections of lightly blemished but sellable produce against false acceptances of batches with meaningful spoilage.

B. Classification Results

Table I summarizes representative results across the four sample categories, comparing the computed rot ratio against ground truth and the resulting system decision.

TABLE I. Freshness Classification and Rot-Ratio Results

Sample Image Category	Ground Truth	Body Pixels (Pbody)	Rot Pixels (Prot)	Rot Ratio (%)	System Decision
Fresh Vine Tomatoes	Fresh	159,783	3,355	2.1%	FRESH (Pass)
Fresh Bell Peppers	Fresh	64,967	1,169	1.8%	FRESH (Pass)
Shriveled Red Pepper	Rotten	85,420	12,898	15.1%	ROTTEN (Block)
Blackened Banana	Rotten	98,120	76,926	78.4%	ROTTEN (Block)

Across all four categories, the system correctly separated lightly blemished, sellable produce (Rot Ratio well under the 10.0% threshold) from genuinely spoiled batches (Rot Ratio well above threshold), with a clear margin between the FRESH and ROTTEN classes rather than borderline cases clustering near the decision boundary.

C. Performance Metrics

To quantify classification quality beyond raw accuracy, precision, recall, and F1-score were computed against the manually labeled ground truth across the full test set. Table II reports these metrics alongside the overall classification accuracy of 96.8% achieved by the proposed system.

TABLE II. Overall Classification Performance

Metric	Value
Accuracy	96.8%
Precision (Rotten class)	95.4%
Recall (Rotten class)	97.1%
F1-Score (Rotten class)	96.2%
False Positive Rate (Fresh misclassified)	2.6%

D. Ablation Study

To isolate the contribution of each processing stage, the system was re-evaluated with individual modules disabled. Table III shows the resulting drop in accuracy when the stem-recognition filter and the texture edge-entropy filter (Section III-D) are removed, confirming that both stages meaningfully reduce false rot triggers on healthy produce features.

TABLE III. Ablation Study — Contribution of Each Module

Configuration	Accuracy	False Positive Rate
Full pipeline (proposed)	96.8%	2.6%
Without stem/vine exclusion	89.3%	11.4%
Without texture edge-entropy filter	91.7%	8.9%
Without background segmentation	84.1%	16.8%

The largest single degradation occurs when background segmentation is disabled, confirming that robustness to wooden crates and cardboard trays — the dominant visual clutter in a dark-store receiving bay — is the most critical component for real-world deployment, ahead of even the stem and texture-based refinements.

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering
Impact Factor 8.021 ⁘ Peer-reviewed & Refereed journal ⁘ Vol. 14, Issue 8, August 2026
DOI: 10.17148/IJIREEICE.2026.14806

VI. DISCUSSION

The results in Section V indicate that the largest practical gains over naive color-thresholding come not from the rot-detection logic itself, but from the preceding exclusion stages — background segmentation and stem/vine recognition — that prevent non-rot pixels from ever reaching the classifier. This mirrors the operational reality of a dark store, where visual clutter (crates, trays, shadows) and natural produce variation (stems, ripening color shifts) are far more common sources of error than genuinely ambiguous rot boundaries.

The quantitative weight-estimation model (Section IV) further distinguishes this framework from binary pass/fail classifiers used in prior work. By reporting a kilogram-level spoilage figure per batch rather than a single accept/reject decision, the system supports partial write-offs, which better reflects real inventory economics — a crate of tomatoes with 2% spoilage should not be discarded in the same way as one with 78% spoilage, even though both would be flagged as containing rotten produce under a purely binary scheme.

VII. LIMITATIONS AND FUTURE WORK

The current pixel-to-weight regression model (Section III-E) is calibrated per produce category using reference batches of known weight, which means onboarding a new SKU requires a short calibration pass before the quantitative estimates are reliable. Extending the regression model to a self-calibrating scheme that infers density coefficients from a small number of live receiving-bay samples would reduce this onboarding overhead.

Additionally, the present evaluation was conducted on single top-down frames; produce items that overlap or occlude one another in a densely packed crate can undercount Pbody for the occluded item. Incorporating a lightweight depth cue or a secondary camera angle is a promising direction for improving accuracy on tightly packed, overlapping batches. Future work will also explore extending the categorical coverage beyond tomato, banana, capsicum, and spinach to the broader SKU catalog typical of a quick-commerce dark store, and integrating the reported spoilage figures directly into automated inventory write-off and vendor-scorecard pipelines.

VIII. CONCLUSION

This research paper presented an advanced, real-time Computer Vision system for automated produce freshness classification and rotten weight quantification. By incorporating background segmentation, healthy color protection, green stem isolation, texture edge-entropy filtering, and a quantitative weight estimation model, the system achieved a 96.8% classification accuracy, eliminating the historical problem of false rot triggers on stems and colored skins. The ablation study confirms that robustness to uncontrolled dark-store backgrounds is the single largest contributor to this accuracy gain. This solution enables quick-commerce dark stores to enforce zero-spoilage dispatches, automate partial inventory write-offs at the kilogram level, and reduce overall inventory loss.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of their affiliated institution's Department for providing the computational resources and produce-image datasets used in this study.

REFERENCES

[1] L. Wang et al., "Dynamic Bandwidth and Wavelength Allocation Scheme for Next-Generation Wavelength-Agile EPON", *J. Optical Commun. Networking*, vol. 9, no. 3, pp. 33-42, 2017.
[2] M. P. McGarry, M. Reisslein and M. Maier, "Ethernet Passive Optical Network Architectures and Dynamic Bandwidth Allocation Algorithms", *IEEE Commun. Surveys & Tutorials*, vol. 10, no. 3, pp. 46-60, 2008.
[3] S. Zhang, C. Zhu, J. K. O. Sin, and P. K. T. Mok, "A novel ultrathin elevated channel low-temperature poly-Si TFT", *IEEE Electron Device Lett.*, vol. 20, no. 2, pp. 569–571, 1999.
[4] A. Bhargava and A. Bansal, "Fruits and vegetables quality evaluation using computer vision: A review", *J. King Saud Univ. — Computer and Information Sciences*, vol. 33, no. 3, pp. 243-257, 2021.
[5] S. Naik and B. Patel, "Machine vision based fruit classification and grading — A review", *Int. J. Computer Applications*, vol. 170, no. 9, pp. 22-34, 2017.
[6] R. Kaur and M. Kumar, "HSV color space based image segmentation for fruit defect detection", *Proc. Int. Conf. Signal Processing and Communication*, pp. 112-117, 2019.
[7] P. Mehta and S. Rao, "Deep learning based rotten fruit detection using convolutional neural networks", *Int. J. Engineering Research & Technology*, vol. 9, no. 6, pp. 455-460, 2020.
[8] T. Iqbal and F. Ahmed, "Post-harvest quality monitoring of perishables using image processing: A survey", *J. Food Engineering and Technology*, vol. 8, no. 2, pp. 89-101, 2022.