# Edge-Ready Citrus Disease Screening Using Classical Vision Features: Computational Efficiency, Robustness, and Threshold-Tunable Canker Detection

**Viet-Hung. Pham**
Thu Dau Mot University
**Hoa-Cuc. Nguyen**
Thu Dau Mot University
**Ngoc-Duong. Nguyen**
Thu Dau Mot University
**Bich-Ngoc. Mach**
Thu Dau Mot University
**Trinh Thi Nhu Quynh**
Thu Dau Mot University
**Duong Thi Kim Chi**
chidtk@tdmu.edu.vn

Thu Dau Mot University

**Research Article**

**Keywords:** Edge AI, citrus disease screening, citrus canker, lightweight computer vision, HSV–LBP–HOG features, robustness evaluation, threshold tuning, CPU-only deployment, sustainable agriculture

**Posted Date:** May 11th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-9594591/v1



**Additional Declarations:** No competing interests reported.

# Edge-Ready Citrus Disease Screening Using Classical Vision Features: Computational Efficiency, Robustness, and Threshold-Tunable Canker Detection

Viet-Hung. Pham[1], Hoa-Cuc. Nguyen[1], Ngoc-Duong. Nguyen[1],
Bich-Ngoc. Mach[1], Trinh Thi Nhu Quynh[2], Duong Thi Kim Chi[3,✉]

[1]Faculty of Engineering and Technology, Thu Dau Mot University, No. 06 Tran Van On, Phu Loi Ward, Ho Chi Minh City, Viet Nam

[2]School of Law and Management, Thu Dau Mot University, No. 06 Tran Van On, Phu Loi Ward, Ho Chi Minh City, Viet Nam

[3]Institute of Digital Technology, Thu Dau Mot University, No. 06 Tran Van On, Phu Loi Ward, Ho Chi Minh City, Viet Nam

✉Email: chidtk@tdmu.edu.vn

## Abstract

Edge-based citrus disease screening requires not only accurate recognition but also low latency, modest memory use, robustness to imperfect image acquisition, and flexible decision thresholds for practical field operation. This study evaluated a lightweight classical vision pipeline as an operational screening engine for orange fruit disease detection under CPU-only deployment constraints. RGB images were represented using HSV color histograms, Local Binary Pattern texture descriptors, and Histogram of Oriented Gradients shape features, followed by one-vs-rest Logistic Regression classification. Beyond classification accuracy, the system was assessed through computational profiling, perturbation robustness, descriptor-level accuracy–latency trade-offs, and threshold-tunable canker detection. The pipeline achieved macro-F1 ≈ 0.94 while requiring only ~0.482 ms/image for feature extraction, negligible classification latency, ~106 MB RAM, and ~6% CPU utilization. Robustness analysis showed stable performance under brightness, contrast, crop, rotation, and moderate noise perturbations. These findings support classical feature-based vision as a practical, transparent, and resource-efficient edge-screening strategy for sustainable citrus disease monitoring.

***Keywords:*** *Edge AI; citrus disease screening; citrus canker; lightweight computer vision; HSV–LBP–HOG features; robustness evaluation; threshold tuning; CPU-only deployment; sustainable agriculture*.

## 1. Introduction

Image-based disease recognition has become an important component of precision agriculture, particularly for fruit crops whose market value depends strongly on visible surface quality. In citrus production, diseases such as citrus canker and melanose can reduce commercial grade, increase post-harvest losses, and prompt unnecessary chemical intervention when diagnosis is delayed or uncertain [1]. Although many recent studies have emphasized classification accuracy, practical orchard screening imposes a broader set of requirements. A deployable system must operate with low latency, modest memory consumption, limited energy demand, and stable performance under imperfect image acquisition [2, 3]. In many citrus-growing environments, screening devices are expected to work with commodity cameras, low-cost processors, intermittent connectivity, and limited technical maintenance [4]. Under these conditions, a high-performing model is useful only if it can also be executed reliably at the edge.

Deep learning has substantially advanced plant disease detection and has achieved impressive accuracy in several controlled datasets [5]. However, deep convolutional and transformer-based models often require large annotated datasets, accelerator-class hardware, repeated calibration, and complex model-update procedures [6]. These requirements can limit their use in smallholder orchards, mobile inspection units, or post-harvest sorting settings where CPU-only operation remains common [7]. Moreover, field deployment is frequently affected by changing sunlight, mild blur, sensor noise, framing variation, glare, and fruit orientation [8]. A model that performs well on clean test images may therefore become less useful if its robustness and computational cost are not quantified. This gap has made edge-oriented evaluation as important as model design itself.

The present study addressed this deployment gap by evaluating a lightweight classical vision pipeline as an operational screening engine for citrus disease recognition. Instead of positioning handcrafted features merely as a lower-complexity alternative to deep learning, the pipeline was assessed from the perspective of field readiness [9, 10]. RGB fruit images were processed through deterministic feature extraction using HSV color histograms, Local Binary Pattern texture descriptors, and Histogram of Oriented Gradients edge–shape features, followed by one-vs-rest Logistic Regression classification [11, 12]. This architecture offered several practical advantages: its execution path was compact, its memory footprint was modest, its output probabilities enabled threshold control, and its feature structure remained transparent enough for diagnostic interpretation.

The contribution of this work is therefore not limited to reporting another classification result. It provides an operational evaluation framework that connects disease recognition with real deployment constraints. The analysis quantified CPU-only latency, inference overhead, memory usage, CPU utilization, robustness under common test-time perturbations, threshold-tunable canker screening behavior, and descriptor-level accuracy–latency trade-offs [10, 13]. The results showed that the pipeline maintained strong diagnostic performance while requiring approximately 0.482 ms per image for feature extraction, negligible classification latency, about 106 MB of RAM, and roughly 6% CPU utilization on commodity hardware. These findings support the view that classical feature-based vision can serve as a "just-enough" edge-AI solution for frequent, local, and transparent citrus disease screening, particularly in resource-constrained agricultural workflows where practical deployability is as critical as accuracy.

## 2. Materials and Methods

### 2.1. Edge-oriented screening framework

The proposed system was evaluated as an edge-oriented citrus disease screening framework rather than as a standalone laboratory classifier. Each RGB image of orange fruit was processed through a deterministic vision pipeline composed of image standardization, handcrafted feature extraction, one-vs-rest probabilistic classification, and threshold-based decision support. The screening task involved three diagnostic categories: healthy fruit, citrus canker, and melanose. However, the operational emphasis was placed on rapid and adjustable canker screening, because missed canker cases may delay orchard-level intervention and increase disease-management uncertainty.

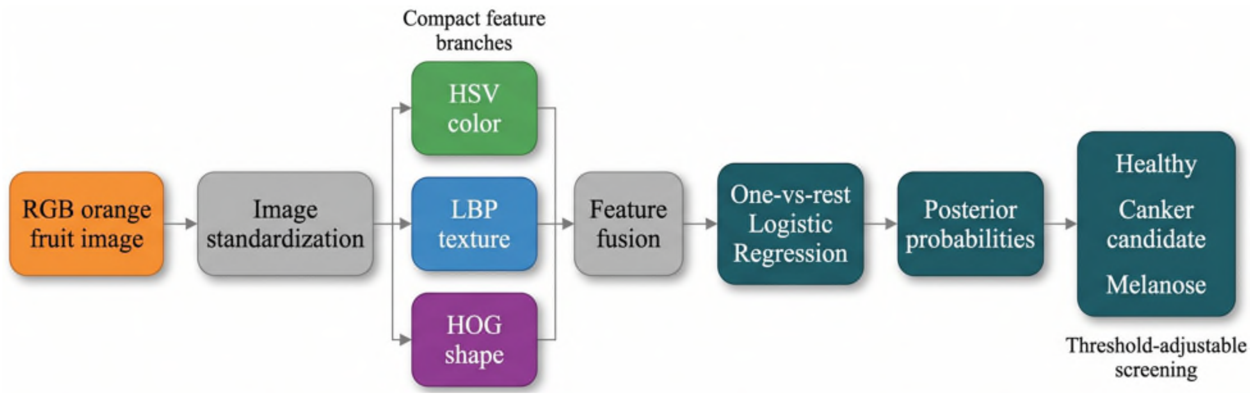


**Figure 1.** *Edge-oriented citrus disease screening framework using lightweight classical vision features*

As illustrated in **Figure 1**, the framework converts each RGB image into a compact descriptor using HSV color histograms, Local Binary Pattern texture features, and Histogram of Oriented Gradients shape information. The concatenated feature vector is then passed to a Logistic Regression classifier that produces class-level posterior probabilities. The final screening output can therefore be interpreted either as a three-class prediction or as a canker-focused risk signal. This structure is particularly suitable for edge deployment because the computationally expensive steps are deterministic, the classifier is linear, and decision thresholds can be adjusted without model retraining. This design directly connects visual disease cues with operational screening requirements reported in the original experimental pipeline .

### 2.2. Hardware and computational profiling protocol

Computational profiling was performed under CPU-only conditions to reflect realistic low-cost deployment scenarios. The system was evaluated on an Intel Core i5-12500H processor with 16 GB RAM, without GPU acceleration. Four operational variables were recorded: feature-extraction latency, classifier inference latency, CPU utilization, and RAM footprint. The profiling protocol separated the descriptor stage from the classification stage because edge deployment depends not only on total runtime but also on identifying the dominant computational bottleneck.

For a test workload containing $N$ images, the mean per-image feature-extraction latency was computed as

$$t_{\text{feat/img}} = \frac{T_{\text{feat}}}{N}, \tag{1}$$

where $T_{\text{feat}}$ is the measured batch feature-extraction time. Effective throughput was then estimated as

$$Q = \frac{N}{T_{\text{feat}}}, \tag{2}$$

expressed as images per second. Classifier overhead was evaluated separately using batch inference time,

$$t_{\mathrm{inf/img}} = \frac{T_{\mathrm{inf}}}{N}. \tag{3}$$

As shown conceptually in **Figure 2**, this decomposition allowed latency to be interpreted at two levels: descriptor computation and decision computation. Such separation is important because the original results showed that feature extraction dominated total runtime, whereas Logistic Regression inference contributed negligible overhead.

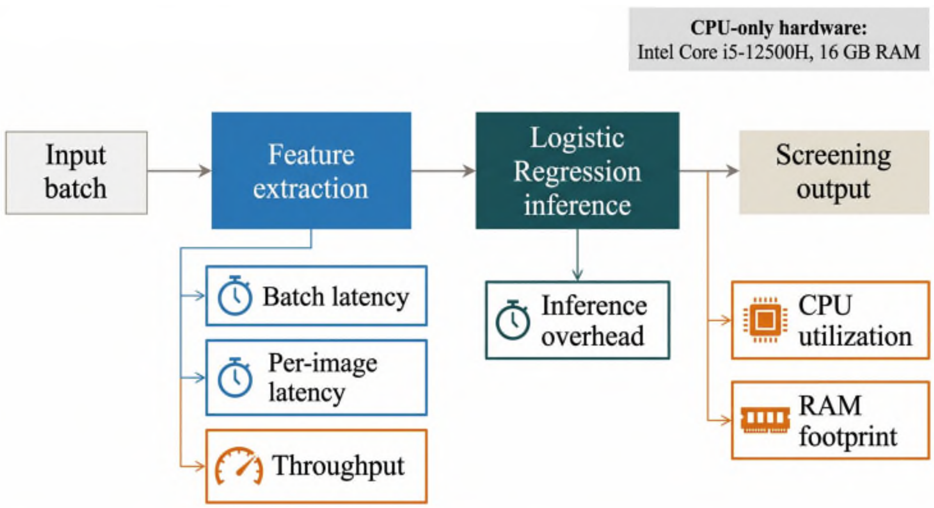


**Figure 2.** *CPU-only computational profiling protocol for edge-based citrus disease screening*

### 2.3. Perturbation-based robustness evaluation

Robustness was evaluated using controlled test-time perturbations that represented common acquisition imperfections in orchard and post-harvest environments. The clean held-out test set was transformed under moderate brightness shifts, contrast shifts, Gaussian noise, random cropping followed by resizing, and small in-plane rotations. These perturbations were selected because citrus fruit images collected outside laboratory conditions are frequently affected by variable sunlight, camera exposure differences, sensor noise, slight framing inconsistency, handheld tilt, and partial scale changes.

Let $f(\cdot)$ denote the trained screening pipeline and $g_k(\cdot)$ denote the $k$-th perturbation applied to an input image. Robustness was assessed by comparing the class-level and macro-averaged F1-scores between the clean input and the perturbed input:

$$\Delta F1_k = F1_{\mathrm{clean}} - F1_{g_k}. \tag{4}$$

A small $\Delta F1_k$ indicated stable screening behavior under that acquisition condition. **Figure 3** summarizes this evaluation logic by linking each perturbation type to its field-level interpretation and downstream metric comparison. This design allowed robustness to be treated as an operational property of the screening system, rather than as a secondary classification statistic. The original robustness results showed that the pipeline remained stable under most moderate perturbations, with the strongest degradation occurring under heavier Gaussian noise.

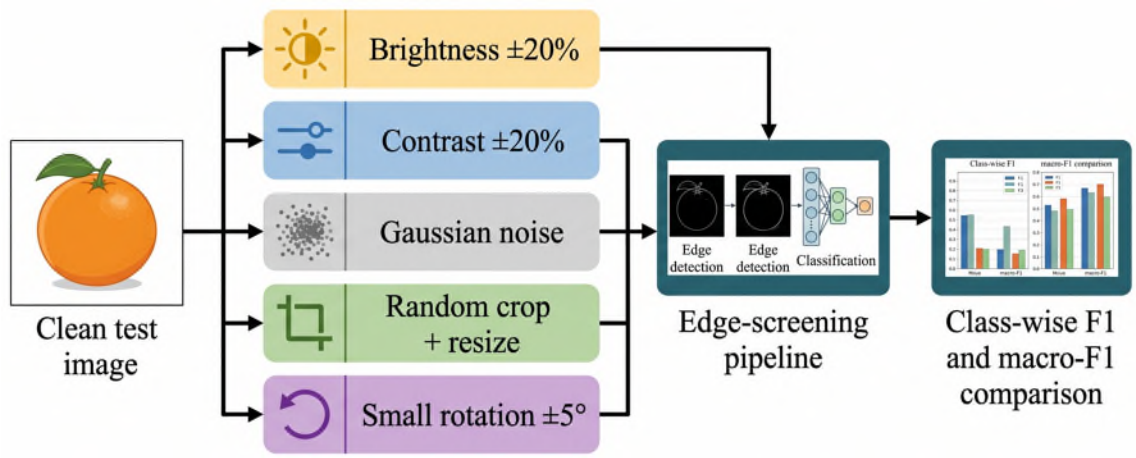


**Figure 3.** *Perturbation-based robustness evaluation under field-like image acquisition conditions*

## 2.4. Threshold-tunable canker screening protocol

The one-vs-rest Logistic Regression classifier produced posterior probabilities for each class, enabling canker-focused threshold tuning without retraining the model. For citrus canker, the decision rule was expressed as

$$\hat{y}_{\text{canker}} = \begin{cases} 1, & p_{\text{canker}} \geq \tau_{\text{canker}}, \\ 0, & p_{\text{canker}} < \tau_{\text{canker}}, \end{cases} \tag{5}$$

where $p_{\text{canker}}$ is the posterior probability of canker and $\tau_{\text{canker}}$ is the operating threshold. Three threshold settings were evaluated to represent distinct deployment modes: a high-sensitivity screening mode, a balanced baseline mode, and a high-precision confirmatory mode.

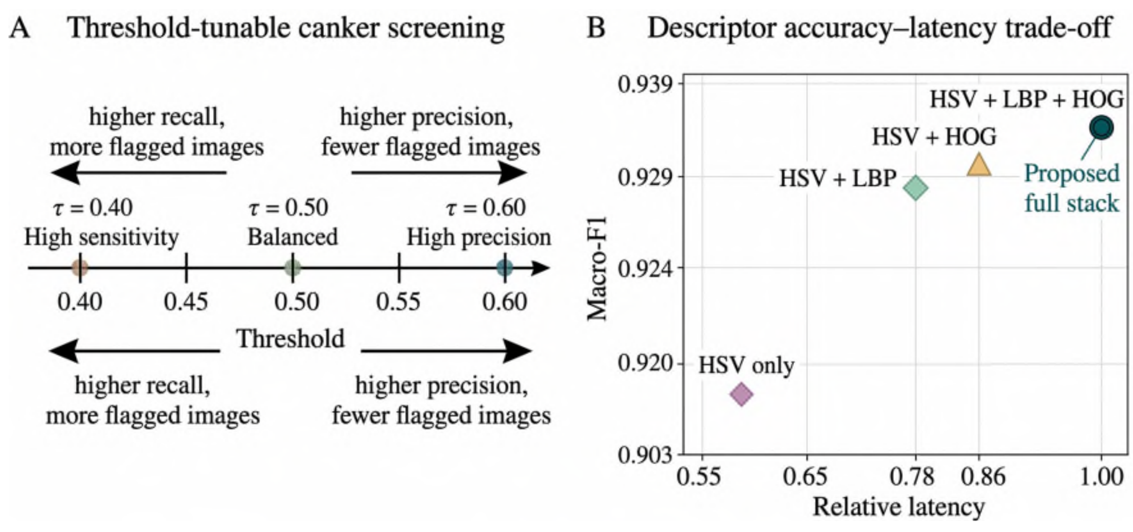


**Figure 4.** *Threshold-tunable canker screening and descriptor-level accuracy–latency trade-off analysis*

As illustrated in **Figure 4A**, lowering the canker threshold prioritizes recall and increases the number of images flagged for follow-up inspection, whereas raising the threshold prioritizes precision and reduces false alarms. This threshold-tunable behavior is operationally important because different agricultural workflows impose different costs for missed disease and unnecessary manual review. In orchard surveillance, early screening may favor sensitivity; in post-harvest grading, confirmatory operation may favor precision. The original results showed that this adjustment changed canker precision–recall behavior while preserving the same underlying model and inference cost .

### 2.5. Accuracy–latency trade-off analysis

Descriptor-level trade-off analysis was conducted to determine whether the full HSV–LBP–HOG representation was justified under edge constraints. Four variants were compared while keeping the Logistic Regression classifier unchanged: HSV only, HSV + LBP, HSV + HOG, and the full HSV + LBP + HOG descriptor. For each variant, macro-F1, overall accuracy, and relative latency were recorded.

Relative latency was defined as

$$L_{\mathrm{rel}} = \frac{t_{\mathrm{variant}}}{t_{\mathrm{full}}}, \tag{6}$$

where $t_{\mathrm{variant}}$ is the feature-extraction time of a given descriptor set and $t_{\mathrm{full}}$ is the latency of the complete HSV–LBP–HOG configuration. As shown in **Figure 4B**, this analysis positioned each descriptor variant within an accuracy–latency space. HSV-only extraction offered the lowest computational cost but produced the weakest performance, whereas the full descriptor achieved the highest macro-F1 at the largest, but still practical, relative latency. This analysis clarified the deployment choice: the full descriptor was retained as the primary edge-screening configuration, while reduced variants remained possible fallback options for more constrained hardware.

## 3. Results

### 3.1. CPU-only computational efficiency

The CPU-only computational profile confirmed that the proposed screening pipeline was suitable for low-cost edge deployment. On the held-out workload of 1,170 orange fruit images, the complete feature-extraction stage required 563.73 ms, corresponding to approximately 0.482 ms per image and an effective throughput of about 2,075 images per second. In contrast, the Logistic Regression inference stage required only 0.29 ms for the full batch, making the classification overhead practically negligible. As summarized in **Table 3**, this runtime structure showed that the computational burden was concentrated almost entirely in the deterministic HSV–LBP–HOG descriptor extraction stage, while the linear classifier added almost no operational delay.

**Table 3.** Measured computational profile on CPU-only hardware

| Component / metric | Measured value | Per-image equivalent | Operational interpretation |
|---|---|---|---|
| Feature extraction time | 563.73 ms / batch | 0.482 ms/image | Dominant processing stage |
| Logistic Regression inference | 0.29 ms / batch | ~0.00025 ms/image | Negligible decision overhead |

| | | | |
|---|---|---|---|
| Effective throughput | ~2,075 images/s | — | High screening capacity |
| RAM footprint | ~106.4 MB | — | Suitable for commodity edge hardware |
| CPU utilization | ~6% | — | Low processing load for continuous operation |

**Figure 5** visualizes this imbalance between descriptor computation and classifier inference. The latency breakdown highlights that optimization efforts should target feature extraction rather than the classification head. This finding is important for edge deployment because it identifies a clear and controllable bottleneck: runtime can be further improved through region-of-interest extraction, coarser HOG grids, integral histogram caching, or hardware-level implementation of handcrafted descriptors without redesigning the model. The measured 106.4 MB RAM footprint and approximately 6% CPU utilization further indicate that the system can run on commodity processors while leaving sufficient resources for image acquisition, storage, user interface, or wireless communication. Overall, the profiling results support the pipeline as a practical CPU-only screening engine rather than a laboratory-only classifier.

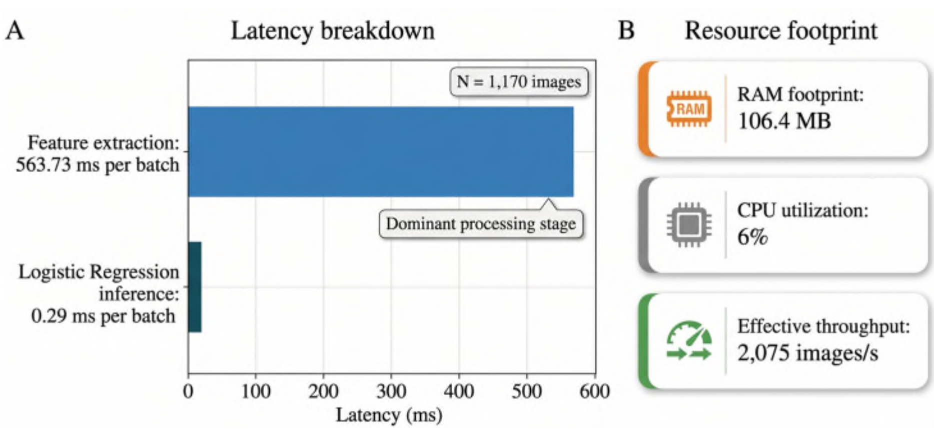


**Figure 5.** *End-to-end latency and resource usage profile*

### 3.2. Robustness under acquisition perturbations

The robustness evaluation showed that the proposed edge-screening pipeline maintained stable performance under several acquisition conditions commonly encountered in orchard and post-harvest imaging. Under the clean test condition, the model achieved a macro-F1 score of 0.939. When moderate brightness variation was introduced, macro-F1 decreased only slightly to 0.931. A similar pattern was observed under contrast shifts, where macro-F1 remained at 0.928. These results indicate that the color–texture–shape descriptor retained discriminative information even when illumination and exposure were not perfectly controlled.

As reported in **Table 4**, the strongest degradation occurred under Gaussian noise with $\sigma = 10$, where macro-F1 decreased to 0.902. Even in this more challenging condition, performance remained above 0.90, suggesting that the pipeline did not collapse under noise but became less confident in fine-grained texture separation. Random crop with resizing produced a macro-F1 of 0.932, while small in-plane rotation retained 0.935, demonstrating that modest framing and handheld orientation changes had limited effect on diagnostic performance. **Figure 6** presents these robustness trends as a degradation profile relative to the clean baseline. The pattern shows that brightness, crop, and rotation perturbations caused only mild performance loss, whereas noise produced the clearest decline. From an operational perspective, this result indicates that routine field variation can be tolerated, but image quality control remains important when sensor noise, compression artefacts, or low-light acquisition become severe.

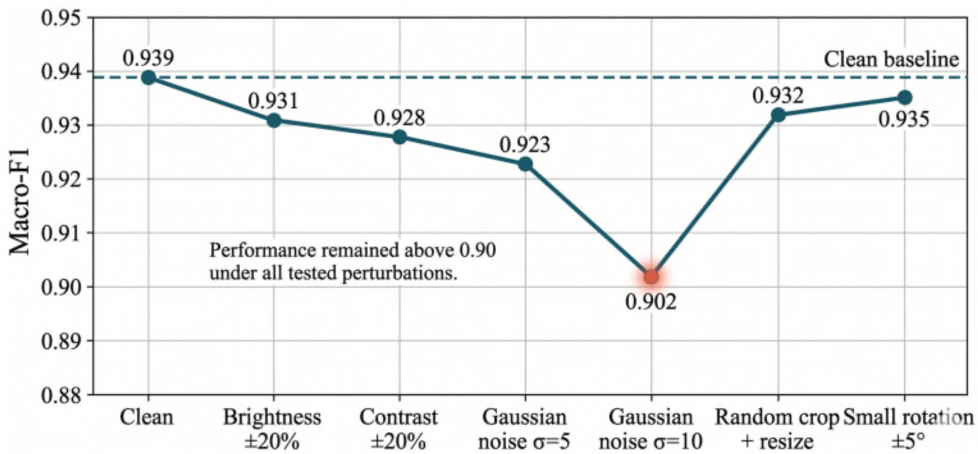


**Figure 6.** *Robustness degradation profile under field-like perturbations*

**Table 4.** Robustness of the edge-screening pipeline under test-time perturbations

| Perturbation scenario | Canker F1 | Healthy F1 | Melanose F1 | Macro-F1 |
|---|---|---|---|---|
| Clean image | 0.927 | 0.933 | 0.958 | 0.939 |
| Brightness ±20% | 0.918 | 0.925 | 0.949 | 0.931 |
| Contrast ±20% | 0.915 | 0.922 | 0.946 | 0.928 |
| Gaussian noise, $\sigma = 5$ | 0.910 | 0.918 | 0.941 | 0.923 |
| Gaussian noise, $\sigma = 10$ | 0.885 | 0.897 | 0.923 | 0.902 |
| Random crop + resize, 90% area | 0.919 | 0.928 | 0.950 | 0.932 |
| Small rotation, ±5° | 0.921 | 0.931 | 0.952 | 0.935 |

### 3.3. Threshold-based canker screening behavior

The threshold analysis demonstrated that the same trained classifier could be adapted to different operational priorities without retraining. Because the one-vs-rest Logistic Regression model produced posterior probabilities for citrus canker, the canker decision threshold could be adjusted to favor sensitivity, balance, or precision. This behavior is central to field screening because the cost of a missed canker case is not always equivalent to the cost of a false alarm. In early orchard surveillance, a high-sensitivity setting is preferable because suspicious fruit can be flagged for follow-up inspection. In post-harvest grading or confirmatory inspection, a high-precision setting may be more useful to avoid excessive manual review.

**Table 5** shows the three evaluated operating modes. At the high-sensitivity threshold of 0.40, canker recall increased to 0.949, while precision decreased to 0.905, and approximately 40% of images were flagged as canker candidates. The balanced baseline at 0.50 achieved precision of 0.942, recall of 0.913, and F1-score of 0.927, while flagging 32.3% of images. When the threshold was raised to 0.60, precision increased to 0.963, whereas recall decreased to 0.887, and the flagged fraction declined to 26%. **Figure 7** illustrates this precision–recall shift, showing a clear movement from screening-oriented sensitivity to confirmatory specificity. Importantly, these operational modes were obtained without changing the feature extractor, classifier, or computational cost. This makes threshold tuning a practical deployment mechanism rather than an additional modeling burden.

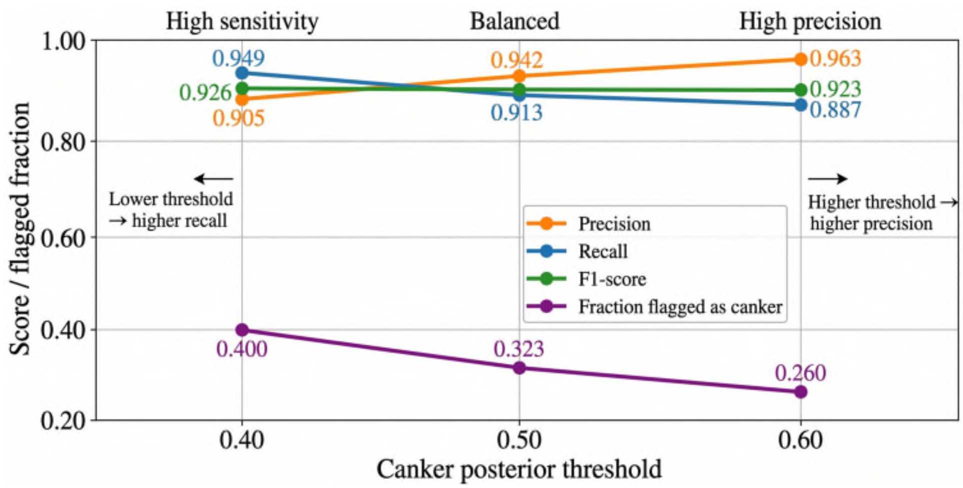


**Figure 7.** *Precision–recall shift across canker posterior thresholds*

**Table 5.** Operating points for canker-focused screening

| Operating mode | Canker threshold | Precision | Recall | F1-score | Fraction flagged as canker |
|---|---|---|---|---|---|
| High-sensitivity screening | 0.40 | 0.905 | 0.949 | 0.926 | 0.400 |
| Balanced baseline | 0.50 | 0.942 | 0.913 | 0.927 | 0.323 |
| High-precision confirmatory | 0.60 | 0.963 | 0.887 | 0.923 | 0.260 |

### 3.4. Descriptor-level accuracy–latency trade-off

The descriptor-level comparison clarified the computational value of each feature group under edge constraints. As shown in **Table 6**, the HSV-only variant achieved the lowest relative latency (0.55) but produced the weakest performance, with macro-F1 of 0.903 and accuracy of 0.900. This result indicates that color information alone was insufficient to fully separate citrus canker, melanose, and healthy fruit, especially when lesion morphology and rind texture contributed important diagnostic cues.

**Table 6.** Accuracy–latency trade-off across descriptor variants

| **Descriptor variant** | **Macro-F1** | **Overall accuracy** | **Relative latency** |
|---|---|---|---|
| HSV only → Logistic Regression | 0.903 | 0.900 | 0.55 |
| HSV + LBP → Logistic Regression | 0.924 | 0.922 | 0.78 |
| HSV + HOG → Logistic Regression | 0.929 | 0.927 | 0.86 |
| HSV + LBP + HOG → Logistic Regression | 0.939 | 0.939 | 1.00 |

Adding LBP texture features improved macro-F1 to 0.924, while the HSV + HOG configuration achieved a slightly higher macro-F1 of 0.929. These intermediate variants confirmed that both micro-texture and edge–shape information contributed meaningful gains beyond color. The full HSV + LBP + HOG descriptor achieved the highest macro-F1 (0.939) and overall accuracy (0.939) at the reference relative latency of 1.00. **Figure 8** positions these variants in an accuracy–latency space and shows that the full descriptor sat near the favorable edge of the trade-off curve. Although it required the largest relative latency, the absolute runtime remained compatible with CPU-only screening, making the performance gain operationally worthwhile. For highly constrained devices, HSV + HOG may serve as a reduced fallback configuration; however, for commodity edge hardware, the full descriptor provided the most reliable balance between diagnostic performance and processing cost.

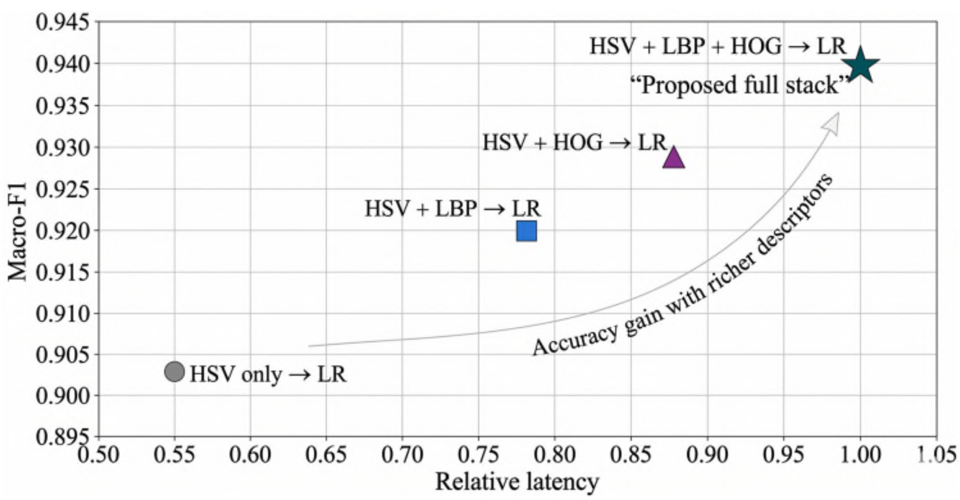


**Figure 8.** *Pareto-style accuracy–latency comparison of descriptor variants*

## 4. Conclusion

This study established a lightweight classical vision pipeline as an operationally viable edge-screening framework for citrus disease recognition. Beyond reporting classification performance, the work quantified the practical deployment properties that determine field usability, including CPU-only latency, inference overhead, memory footprint, robustness to acquisition perturbations, threshold-tunable canker detection, and descriptor-level accuracy–latency trade-offs. The HSV–LBP–HOG + Logistic Regression pipeline achieved macro-F1 ≈ 0.94

while requiring only ~0.482 ms per image for feature extraction, negligible classification time, ~106 MB RAM, and ~6% CPU utilization. Performance remained stable under brightness, contrast, crop, rotation, and moderate noise perturbations, supporting reliable operation under imperfect imaging conditions. The canker threshold analysis further showed that the same trained model could shift between high-sensitivity screening and high-precision confirmation without retraining. These findings position the proposed framework as a transparent, resource-efficient, and deployment-ready solution for frequent citrus disease screening in sustainable agricultural workflows.

## Ethics Statement

This study uses an open, publicly available orange-fruit image dataset; no human/animal data are involved. All usage follows the dataset's terms.

## Data Availability

All experiments reference the open dataset described in the source papers; class balancing to 2,600 images per class is documented therein. Feature code and training scripts can be packaged for release upon acceptance.


## Acknowledgements

We thank the creators and maintainers of the open orange-fruit dataset used in both source papers.


## Competing Interest Declaration:

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References


[1] H. T. Rauf, B. A. Saleem, M. I. U. Lali, M. A. Khan, M. Sharif, and S. A. C. Bukhari, "A citrus fruits and leaves dataset for detection and classification of citrus diseases through machine learning," *Data in brief,* vol. 26, p. 104340, 2019, doi: 10.1016/j.dib.2019.104340.

[2] M. Shoaib, A. Sadeghi-Niaraki, F. Ali, I. Hussain, and S. Khalid, "Leveraging deep learning for plant disease and pest detection: a comprehensive review and future directions," *Frontiers in Plant Science,* vol. 16, p. 1538163, 2025, doi: 10.3389/fpls.2025.1538163.

[3] R. Sujatha, S. Krishnan, and J. M. Chatterjee, "Advancing plant leaf disease detection integrating machine learning and deep learning," *Scientific Reports,* vol. 15, p. 11552, 2025, doi: 10.1038/s41598-024-72197-2.

[4] S. Xing, M. Lee, and K. K. Lee, "Citrus pests and diseases recognition model using weakly dense connected convolution network," *Sensors,* vol. 19, no. 14, p. 3195, 2019, doi: 10.3390/s19143195.

[5] S. Qadri, H. Ahsan ul, M. A. Javed, M. Imran, and M. A. Bashir, "Machine vision approach for classification of citrus leaves using fused features," *International Journal of Food Properties,* vol. 22, no. 1, pp. 2072–2089, 2019, doi: 10.1080/10942912.2019.1703738.
[6] G. Liu, S. Mao, and J. Li, "Image recognition of citrus diseases based on deep learning," *Computers, Materials & Continua,* vol. 66, no. 1, pp. 457-466, 2021, doi: 10.32604/cmc.2020.012165.
[7] S. Kumar, A. K. Pandey, D. Raghav, G. Gupta, and V. Srivastava, "A deep learning approach for multiclass orange disease classification," in *2024 2nd international conference on disruptive technologies (ICDT)*, 2024: IEEE, pp. 184-189.
[8] M. S. Krishna, P. Machado, R. I. Otuka, S. W. Yahaya, F. Neves dos Santos, and I. K. Ihianle, "Plant Leaf Disease Detection Using Deep Learning: A Multi-Dataset Approach," *J,* vol. 8, no. 1, p. 4, 2025, doi: 10.3390/j8010004.
[9] N. Butt, M. M. Iqbal, S. Ramzan, A. Raza, L. Abualigah, N. L. Fitriyani, and I. T. R. Yanto, "Citrus diseases detection using innovative deep learning approach and Hybrid Meta-Heuristic," *PLOS ONE,* vol. 20, no. 1, p. e0316081, 2025, doi: 10.1371/journal.pone.0316081.
[10] M. Islam, M. Habib, and M. K. Hossain, "Lemon and Orange Disease Classification using CNN-Extracted Features and Machine Learning Classifier," *arXiv,* 2024, doi: 10.48550/arXiv.2408.14206.
[11] S. Faisal, K. Javed, S. Ali, A. Alasiry, M. Marzougui, M. A. Khan, and J. H. Cha, "Deep transfer learning based detection and classification of citrus plant diseases," *Computers, Materials & Continua,* vol. 76, no. 1, pp. 895–914, 2023, doi: 10.32604/cmc.2023.039781.
[12] A. Goyal and K. Lakhwani, "Integrating advanced deep learning techniques for enhanced detection and classification of citrus leaf and fruit diseases," *Scientific Reports,* vol. 15, p. 97159, 2025, doi: 10.1038/s41598-025-97159-0.
[13] L. Kujur, V. Gupta, and A. Singhal, "A hybrid multi-optimizer approach using CNN and GB for accurate prediction of citrus fruit diseases," *Discover Applied Sciences,* vol. 7, p. 151, 2025, doi: 10.1007/s42452-025-06593-2.