Rule-Based Banana Ripeness Classification via Mean-RGB Thresholding

A Transparent Statistical Baseline for Edge Deployment

Dwi Ahmad Dzulhijjah and Karina Auliasari*

Institut Teknologi Nasional Malang, Malang, Indonesia

* *Corresponding author*

Algorithm originally developed as Digital Image Processing coursework (2020); web reimplementation and technical note (2026).

musa.dwiahmad.com

Abstract—Banana ripeness directly affects eating quality, shelf life, and market price, yet manual grading remains subjective and inconsistent. This note documents a deliberately simple, training-free classifier that assigns a banana image to one of three ripeness classes (unripe, ripe, overripe) using only the mean colour of the segmented fruit region. The background is removed by grayscale thresholding, the per-channel means of the remaining pixels are computed, and a three-branch decision rule is applied on the ordering of the red, green, and blue means together with their average intensity. The algorithm, conceived in 2020 without reference to prior literature, is reimplemented here as a fully client-side web application that performs all computation in the browser. On a controlled white-background image set the rule reproduces the expected class for every sample; on uncontrolled images it degrades predictably. We present the method, its reference measurements, and an explicit account of its failure modes, and position it as a transparent baseline against which learned and hybrid methods can be compared on resource-constrained edge devices.

Index Terms—banana ripeness, colour features, image segmentation, rule-based classification, edge computing, OpenCV.

I. INTRODUCTION

Colour is among the most direct visual cues of fruit maturity, and colour-based grading of bananas predates the modern deep-learning era by decades [1]. As a banana ripens, chlorophyll in the peel breaks down and is replaced first by yellow carotenoids and later by brown melanoidins, producing a characteristic green to yellow to brown progression. A large body of work exploits this progression using RGB or HSV colour statistics combined with classifiers such as fuzzy logic, k-nearest neighbours, support vector machines, and neural networks [2], [3].

This note documents a minimal instance of that family. The classifier was written in 2020 as an undergraduate Digital Image Processing assignment, without consulting external references, as a one-week exercise in turning a visual intuition directly into executable rules. It uses no machine learning and no training step: it computes the mean colour of the fruit and applies three hand-specified conditions. In 2026 the same logic was re-expressed as a browser-based demonstration that uploads nothing to a server. We document the method here both as a record of the original work and as a controlled baseline for subsequent comparison with learned and hybrid approaches.

II. METHOD

A. Background segmentation. The input image is converted to grayscale using the BT.601 luma weighting and thresholded at a fixed value $T = 190$; pixels brighter than T are treated as background and set to black. An optional elliptical morphological erosion (5×5 kernel) cleans the fruit boundary. The complement of the threshold serves as the fruit mask.

B. Colour measurement. The per-channel means, denoted R, G, and B, are computed over the masked fruit region only, so that the black background does not bias the statistics. The mean intensity is taken as $avg = (R + G + B) / 3$.

C. Decision rule. The class is assigned by the first matching condition:

Condition	Assigned class
$G > R > B$	Unripe (green-dominant)
$R > G > B$ and $R \geq 100$	Ripe (yellow)
$avg < 80$ and $R < 100$	Overripe (dark / brown)

otherwise	Unclassified

Table I. The three-branch decision rule. Conditions are evaluated top to bottom; the first match wins.

We note a documentation discrepancy in the original 2020 material: the accompanying prose described the unripe condition as G > B > R, whereas the source code used G > R > B. The implementation here follows the source code, which is the variant consistent with the measured reference data of Table II. Throughout, R, G, and B denote the mean channel values defined above.

III. REFERENCE MEASUREMENTS AND BEHAVIOUR

Table II reports the mean channel values measured on the original controlled image set, in which each banana is photographed against a white background. The three classes separate cleanly: the unripe class is green-dominant, the ripe class is red-and-green dominant with low blue, and the overripe class collapses to a low overall intensity. On this set every sample is assigned its expected class, and the browser reimplementation reproduces the same assignments.

Class	R	G	B	avg
Unripe	44	162	128	111
Ripe	178	138	31	116
Overripe	73	52	57	61

Table II. Mean RGB and average intensity per class on the controlled 2020 reference set (8-bit channels, 0–255).

IV. LIMITATIONS

The simplicity that makes the method transparent also bounds its validity. (i) *Background assumption.* Fixed thresholding at T = 190 assumes a white studio background; photographs taken at a market, on the tree, or in the hand introduce non-fruit pixels into the mean and corrupt the decision. (ii) *Bright yellow leakage.* Because the threshold operates on luma, very bright yellow peel can exceed T and be discarded as background; a saturated yellow of (R,G,B) = (244,200,66), for example, yields luma ≈ 198 > 190 and is masked out. (iii) *Illumination.* Yellow, neon, or shadowed lighting shifts the channel means systematically. (iv) *Variety.* Thresholds tuned on Cavendish-like fruit do not transfer directly to Kepok, Raja, or Mas bananas, whose natural colour ranges differ. (v) *Spatial averaging.* A partially ripe fruit is reduced to a single mean, discarding spatial detail. (vi) *No confidence and no detection.* The output is a hard label with no confidence estimate, and the method does not verify that the object is a banana at all.

V. CONCLUSION AND FUTURE WORK

A three-line colour rule, derived from intuition alone, classifies banana ripeness correctly under controlled imaging and fails in well-understood ways outside it. Its value today is not accuracy but transparency and negligible cost: it runs in a browser and would run on an 8-bit microcontroller. It therefore serves as an honest lower bound for a planned comparison study. Future work replaces fixed RGB thresholding with HSV or adaptive segmentation, adds classical (SVM, random forest) and compact deep-learning classifiers, introduces a confidence-gated cascade that invokes the learned model only on ambiguous inputs, and benchmarks accuracy against latency, memory, and energy across edge hardware (ESP32-class microcontrollers and single-board computers) using a purpose-built multi-variety Indonesian dataset.

REFERENCES

[1] F. Mendoza and J. M. Aguilera, “Application of image analysis for classification of ripening bananas,” *Journal of Food Science*, vol. 69, no. 9, pp. E471–E477, 2004.

[2] A. H. Kahfi, M. Hasan, and R. L. Hasanah, “Classification of banana ripeness based on color and texture characteristics,” *Journal of Computer Networks, Architecture and High Performance Computing*, vol. 5, no. 1, pp. 10–17, 2023.

[3] L. Kamelia et al., “Ripeness level classification of banana fruit based on HSV color space and k-nearest neighbour,” *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 10, no. 2, pp. 941–947, 2021.

[4] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.

[5] K. Auliasari, M. Orisa, F. Riski, and D. A. Dzulhijjah, *Pengolahan Citra Digital Menggunakan Python dan PyCharm*. Malang: ITN Press, 2025. ISBN 978-634-04-1537-7.

[6] D. A. Dzulhijjah, “Mendeteksi Kematangan Buah Pisang,” course manuscript, Institut Teknologi Nasional Malang, 2020.