A Lightweight Real-Time Tomato Leaf Disease Detection System for Edge-Based Smart Agriculture: Failure Modes and Error Containment

Miles Holland, Patrick Jenkins, Damian Phillips | Review Article | Published 2026-08-27

ABSTRACT

Aggregate performance can conceal concentrated failures, so errors must be classified by cause, consequence, and the controls available to contain them. This structured evidence review evaluates "A Lightweight Real-Time Tomato Leaf Disease Detection System for Edge-Based Smart Agriculture" alongside nine author-disjoint, topically matched publications in edge agricultural vision. It compares construct definitions, evaluation choices, operating assumptions, and reported limitations instead of treating bibliographic similarity as empirical equivalence. Viewed through error taxonomy and failure containment, the map separates claims supported by the available record from questions that still require full-text extraction, replication, or new experiments. The synthesis is interpretive rather than meta-analytic and therefore does not present a pooled effect estimate or a new causal result. The resulting agenda pairs a documented error taxonomy with stress tests, escalation rules, and safeguards for high-consequence failures.

KEYWORDS

edge agricultural vision; error taxonomy and failure containment; evidence synthesis; reproducibility; research evaluation

Research Context

Zhao et al. (2026) [1] provides the focal point for this review. Its topic is consequential because progress in edge agricultural vision depends not only on a proposed method, material, dataset, or workflow, but also on the credibility of the evidence chain that connects design decisions to reported outcomes. The present article therefore asks a deliberately bounded question: how should the focal work be interpreted when the primary lens is error taxonomy and failure containment? This framing supports comparison while avoiding claims that cannot be established from bibliographic records alone.

The review follows a structured evidence-mapping protocol. First, the focal work defines the problem boundary. Second, nine adjacent publications are selected by controlled topical terms. Third, complete author sets are normalized and screened so that no two references in this manuscript share an author. Fourth, each item is assigned an explicit analytical role. This protocol makes the inclusion logic inspectable and prevents a long bibliography from masquerading as a coherent synthesis.

Evidence Map

Evidence node 2 is Saha et al. (2026) [2], which addresses "Edge-Ready Lightweight CNN Architectures for Tomato Leaf Disease Detection Using Transfer Learning". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 3 is Salunke et al. (2025) [3], which addresses "Tomato Leaf Disease Detection using Lightweight Deep Convolution Network". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 4 is Mostafa et al. (2026) [4], which addresses "LeafDet: A lightweight and interpretable deep learning framework for tomato leaf disease detection". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 5 is Ibrahim et al. (2026) [5], which addresses "A Systematic Review of Lightweight and Robust Deep Learning Approaches for Tomato Leaf Disease Detection". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 6 is Gookyi et al. (2024) [6], which addresses "Enabling Intelligence on the Edge: Leveraging Edge Impulse to Deploy Multiple Deep Learning Models on Edge Devices for Tomato Leaf Disease Detection". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 7 is Bellout et al. (2025) [7], which addresses "LT-YOLOv10n: A lightweight IoT-integrated deep learning model for real-time tomato leaf disease detection and management". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 8 is Yang et al. (2026) [8], which addresses "Embedded intelligence in agriculture: A lightweight deep learning network for low-cost tomato leaf disease detection". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 9 is Lu et al. (2026) [9], which addresses "Lightweight Real-Time Detection Transformer for Tomato Leaf Disease Recognition in Complex Agricultural Scenarios". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Evidence node 10 is Roopasinghe et al. (2026) [10], which addresses "A Lightweight Incremental Learning Framework for Tomato Leaf Disease Detection in Dynamic Field Deployments". Its controlled title terms place it directly within edge agricultural vision; in this review it is used to test how the focal argument behaves when viewed through error taxonomy and failure containment.

Taken together, references [1]-[10] form a compact, conflict-free map rather than a vote count. They span the focal contribution, adjacent methods, evaluation settings, and implementation considerations. Their value lies in triangulation: agreement can identify stable design assumptions, while differences reveal where metrics, datasets, populations, hardware, or operating conditions limit direct comparison.

Analytical Framework

The first analytical layer is construct alignment. A useful comparison requires that the outcome named in one study correspond to the outcome measured in another. For manuscript 03-P10-006, this means distinguishing nominally similar terms from operationally equivalent measures. The second layer is evidence strength. Peer-reviewed articles, conference papers, and preprints may all contribute, but their claims should be weighted by methodological transparency, sample or benchmark coverage, and the availability of uncertainty estimates.

The third layer concerns transfer. A result can be methodologically coherent yet fragile outside its original setting. Under the lens of error taxonomy and failure containment, transfer should be tested along at least four axes: data composition, task definition, resource envelope, and decision cost. The fourth layer is failure analysis. Aggregate performance alone cannot show whether errors concentrate in rare classes, difficult operating conditions, minority subgroups, boundary cases, or high-cost decisions. A credible research program therefore reports both central tendency and structured failure slices.

The fifth layer is reproducibility. At minimum, readers need a precise description of inputs, preprocessing, comparison baselines, evaluation metrics, and exclusion rules. Where code or data cannot be released, a procedural specification and sensitivity analysis can still make the work more auditable. This is especially important when the focal contribution is recent or when a formal venue record is still evolving.

Implications for Research and Practice

For researchers, the evidence map recommends designing experiments backward from the decision that the system or scientific claim is intended to support. That approach clarifies which baselines are meaningful, which uncertainty measures matter, and which ablations can attribute gains to specific components. It also discourages benchmark-only conclusions when the application depends on latency, reliability, calibration, manufacturing variation, environmental exposure, or human interpretation.

For practitioners, the key implication is staged adoption. A promising result should first be reproduced in a controlled environment, then stress-tested under shifted conditions, and finally monitored after deployment. Decision thresholds should reflect asymmetric costs rather than headline accuracy alone. Documentation should preserve data provenance, versioned configurations, and known failure conditions so that later changes can be traced.

Limitations and Research Agenda

This article is a structured evidence synthesis, not a new empirical trial. The adjacent works were selected for topical relevance and author independence; those criteria do not make their datasets or outcomes interchangeable. The next step is full-text extraction of study design, sample characteristics, effect estimates, uncertainty intervals, and implementation details. A preregistered comparison matrix would strengthen the analysis further.

Three questions follow. First, does the focal claim remain stable under alternative datasets or populations? Second, which component contributes most when cost and uncertainty are evaluated jointly? Third, what monitoring signal would reveal degradation early enough for intervention? Answering these questions would turn the present evidence map into a cumulative research program centered on error taxonomy and failure containment.

References

1. Zhao, R., Deng, F., Que, H., Liu, M., Yue, X., & Mu, L. (2026). A Lightweight Real-Time Tomato Leaf Disease Detection System for Edge-Based Smart Agriculture. Sensors, 26(11), 3474. https://doi.org/10.3390/s26113474

2. Saha, J., Gunasekaran, H., Rajkumar, S., & B, L.-K. (2026). Edge-Ready Lightweight CNN Architectures for Tomato Leaf Disease Detection Using Transfer Learning. . https://doi.org/10.22541/au.176961329.96736990/v1

3. Salunke, P., & Katti, J. (2025). Tomato Leaf Disease Detection using Lightweight Deep Convolution Network. 2025 Third International Conference on Augmented Intelligence and Sustainable Systems (ICAISS), 1236-1243. https://doi.org/10.1109/icaiss61471.2025.11042133

4. Mostafa, V., Sozol, M.-S.-S., & Mondal, M.-R.-H. (2026). LeafDet: A lightweight and interpretable deep learning framework for tomato leaf disease detection. PLOS One, 21(5), e0349501. https://doi.org/10.1371/journal.pone.0349501

5. Ibrahim, F.-S., El-Saber, N.-A.-S., & Mohamed, I.-S. (2026). A Systematic Review of Lightweight and Robust Deep Learning Approaches for Tomato Leaf Disease Detection. Engineering Research Journal (Shoubra), 55(1), 0-0. https://doi.org/10.21608/erjsh.2026.471579.1504

6. Gookyi, D.-A.-N., Wulnye, F.-A., Wilson, M., Danquah, P., Danso, S.-A., & Gariba, A.-A. (2024). Enabling Intelligence on the Edge: Leveraging Edge Impulse to Deploy Multiple Deep Learning Models on Edge Devices for Tomato Leaf Disease Detection. AgriEngineering, 6(4), 3563-3585. https://doi.org/10.3390/agriengineering6040203

7. Bellout, A., Zarboubi, M., Elhoseny, M., Dliou, A., Latif, R., & Saddik, A. (2025). LT-YOLOv10n: A lightweight IoT-integrated deep learning model for real-time tomato leaf disease detection and management. Internet of Things, 33, 101663. https://doi.org/10.1016/j.iot.2025.101663

8. Yang, Z., Dai, B., Wang, R., An, X., Ma, D., & Yun, Y. (2026). Embedded intelligence in agriculture: A lightweight deep learning network for low-cost tomato leaf disease detection. Engineering Applications of Artificial Intelligence, 175, 114612. https://doi.org/10.1016/j.engappai.2026.114612

9. Lu, J., Liu, Y., Yu, M., & Wei, G. (2026). Lightweight Real-Time Detection Transformer for Tomato Leaf Disease Recognition in Complex Agricultural Scenarios. . https://doi.org/10.21203/rs.3.rs-9832324/v1

10. Roopasinghe, R.-U.-N., Vasanthapriyan, S., & Chathumini, K.-G.-L. (2026). A Lightweight Incremental Learning Framework for Tomato Leaf Disease Detection in Dynamic Field Deployments. 2026 IEEE International Research Conference on Smart Computing and Systems Engineering (SCSE), 1-6. https://doi.org/10.1109/scse70081.2026.11499804