JOURNALS | Journal of Visual Artificial Intelligence

Received: 01 October 2026; Revised: 18 November 2026; Accepted: 24 November 2026; Published Online: 26 November 2026.

*J. Vis. Artif. Intell.*, 2026, 1(3), 26111 | Volume 1 Issue 3 (December 2026) | DOI: https://doi.org/10.64189/vai.26111

© The Author(s) 2026

This article is licensed under Creative Commons Attribution NonCommercial 4.0 International (CC-BY-NC 4.0)

# Deep Learning-Based Maize Leaf Blight Detection and Treatment Recommendation with Comparative Evaluation and Web-Based Deployment of YOLO Architectures

Muhammad Adamu Shuaibu[1,2,*] I. T. Balami,[1] Godiya Yakubu,[1] and Abdulrasheed Ibrahim[1]

[1] *Center for Embedded Artificial Intelligence and Smart Energy Systems, Abubakar Tafawa Balewa University, Bauchi, Nigeria*
*2 Department of Embedded Artificial Intelligence, Abubakar Tafawa Balewa University, Bauchi, Nigeria*
**Email:* Mashuaibu.pg@atbu.edu.ng (Muhammad Adamu Shuaibu)

## Supporting Information

### Part A. Process and per-model diagnostics (Figs. S1-S17)

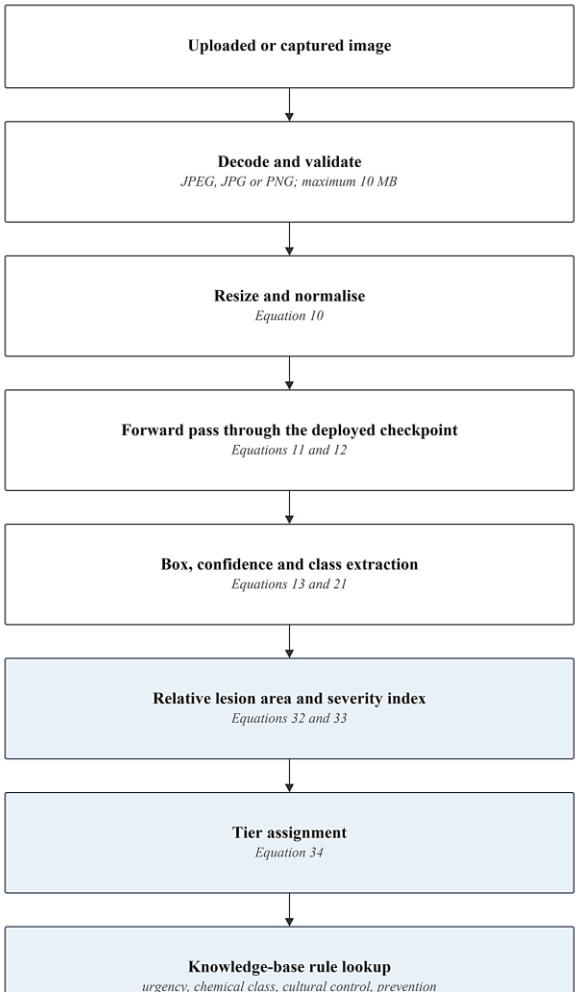



**Fig. S1:** Request-level data flow of a single inference request. The image is decoded and validated against the accepted formats and the 10 MB limit, resized and normalised by Eq. 10, passed once through the deployed checkpoint (Eqs. 11 and 12), reduced to boxes, confidences and classes (Eqs. 13 and 21), converted into a relative lesion area and a severity index (Eqs. 32 and 33), assigned a tier by Eq. 34, and resolved against the knowledge base before the response is rendered. Cited in Section 3.11.

The figure documents the internal sequence of a single request rather than the tier boundaries of the system, and it is provided so that the additive latency decomposition of Eq. 36 can be located on the pipeline. The first request after process start additionally includes the one-time weight-loading cost, which is why the warm and cold latency figures are reported separately in Section 4.11.

**Figs. S2-S5 — YOLOv8n**

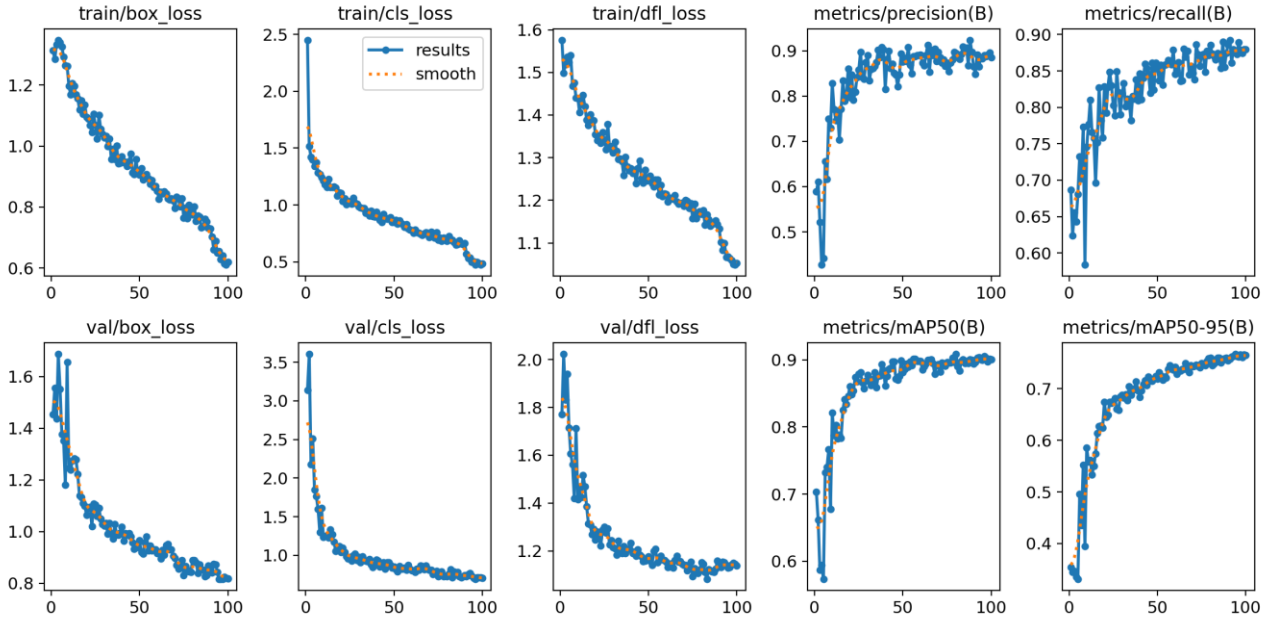



**Fig. S2:** Training and validation curves for YOLOv8n across the hundred epochs of the common protocol, showing the two mean average precision measures in the upper panels and the box, classification, and distribution-focal loss components in the lower panels.

The four panels that follow document the training and evaluation behaviour of YOLOv8n under the protocol of Table 6. They are supplementary to the comparative analysis of Section 4 and are reproduced in full so that the convergence behaviour of each architecture can be inspected independently.

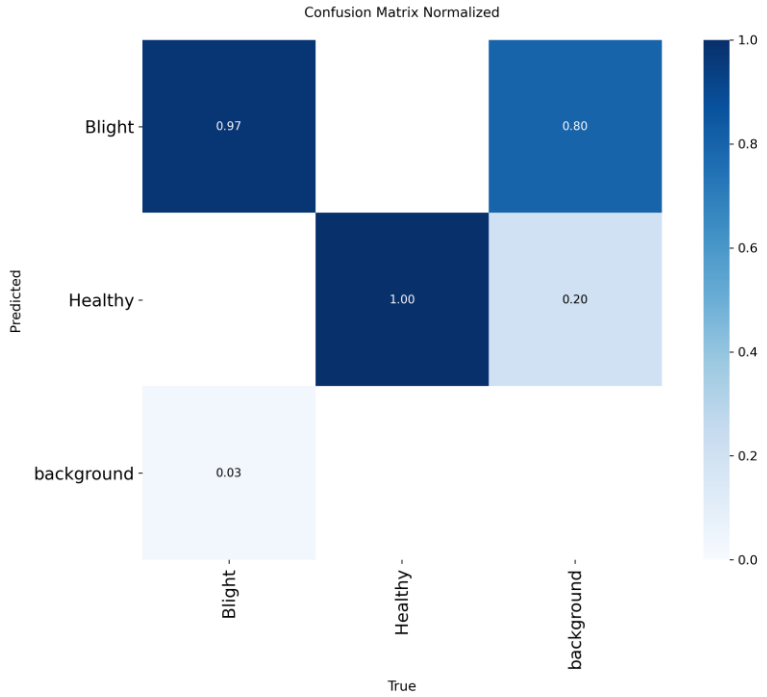



**Fig. S3:** Normalised confusion matrix for YOLOv8n on the 165-image test partition. The background row and column record predictions that were not matched to any annotation, so the matrix is informative about a tendency to fire on non-lesion regions rather than about object-level accuracy.

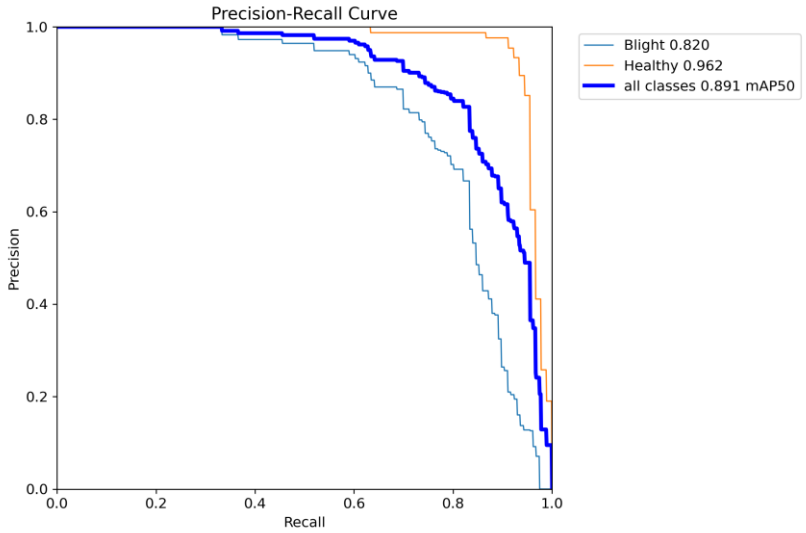



**Fig. S4:** Precision-recall curve for YOLOv8n across the confidence sweep, from which the average precision values of Table 14 are integrated.

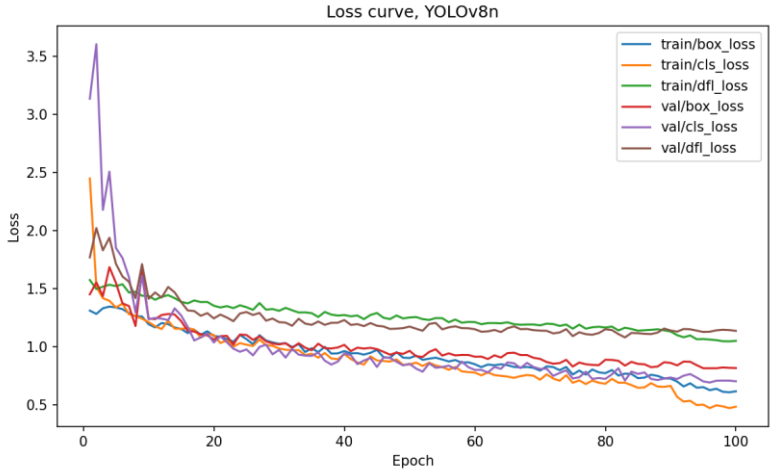



**Fig. S5:** Training and validation loss components for YOLOv8n across the hundred epochs, showing the divergence between the falling training box loss and the plateaued validation metrics discussed in Section 4.8.

### Figs. S6-S9 — YOLOv9t

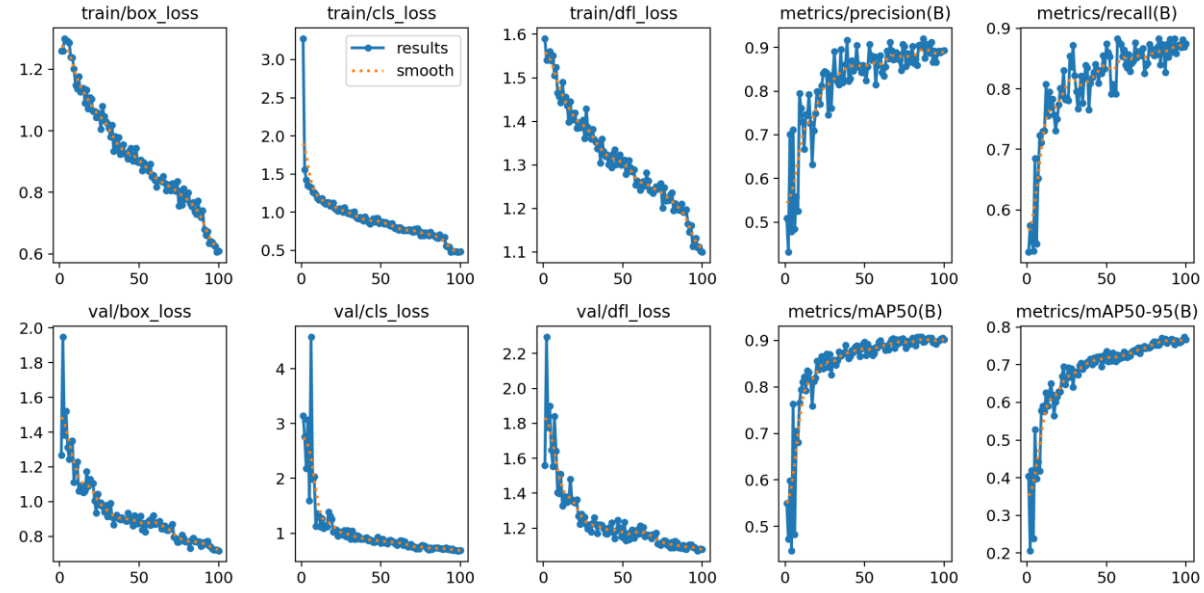



**Fig. S6:** Training and validation curves for YOLOv9t across the hundred epochs of the common protocol, showing the two mean average precision measures in the upper panels and the box, classification, and distribution-focal loss components in the lower panels.

The four panels that follow document the training and evaluation behaviour of YOLOv9t under the protocol of Table 6. They are supplementary to the comparative analysis of Section 4 and are reproduced in full so that the convergence behaviour of each architecture can be inspected independently.

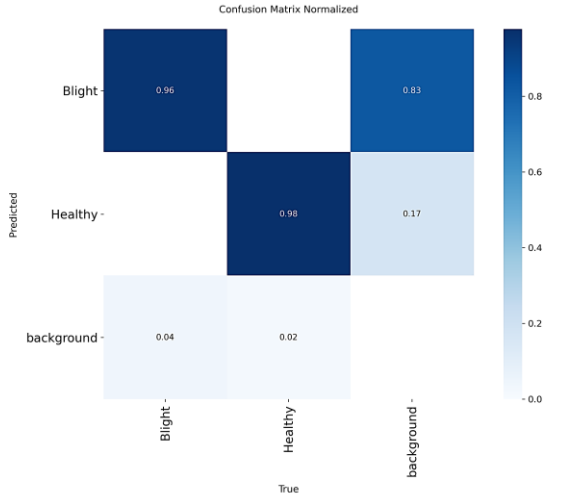



**Fig. S7:** Normalised confusion matrix for YOLOv9t on the 165-image test partition. The background row and column record predictions that were not matched to any annotation, so the matrix is informative about a tendency to fire on non-lesion regions rather than about object-level accuracy.

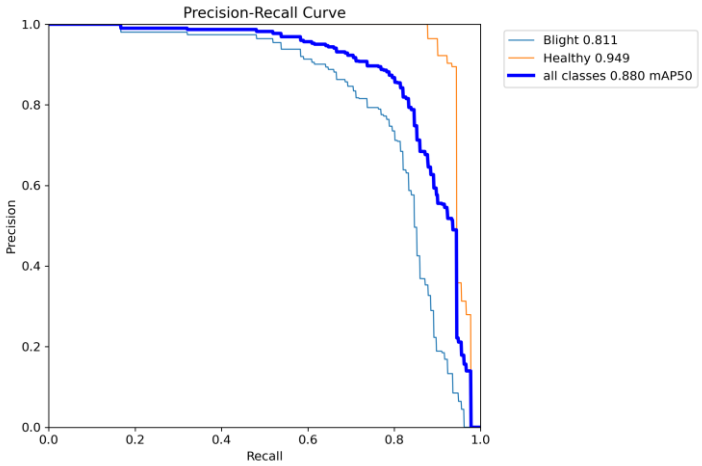



**Fig. S8:** Precision-recall curve for YOLOv9t across the confidence sweep, from which the average precision values of Table 14 are integrated.

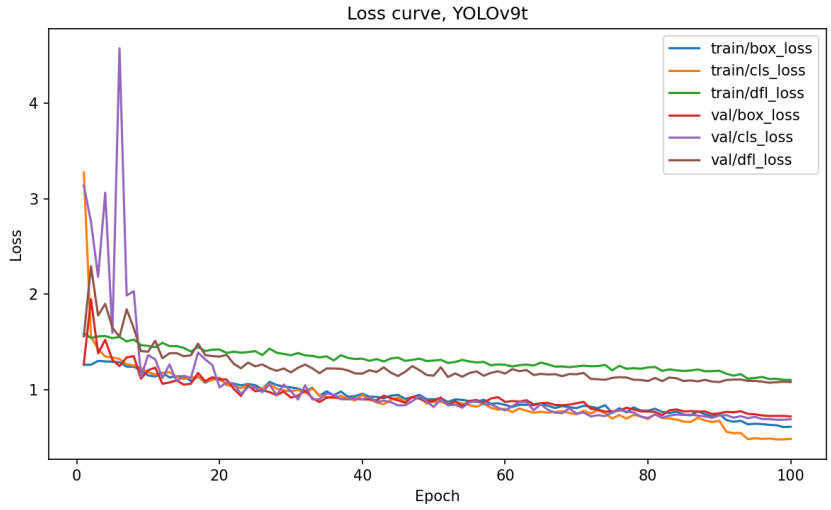



**Fig. S9:** Training and validation loss components for YOLOv9t across the hundred epochs, showing the divergence between the falling training box loss and the plateaued validation metrics discussed in Section 4.8.

**Figs. S10-S13 — YOLOv10n**

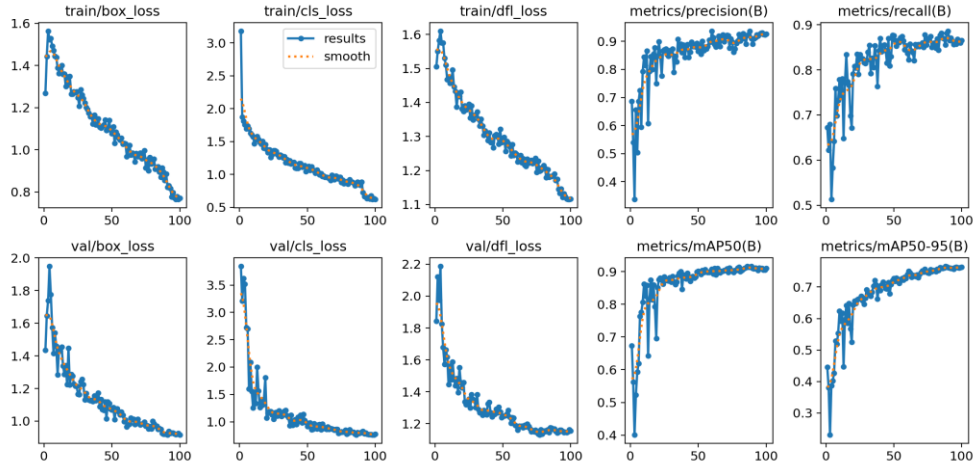



**Fig. S10:** Training and validation curves for YOLOv10n across the hundred epochs of the common protocol, showing the two mean average precision measures in the upper panels and the box, classification, and distribution-focal loss components in the lower panels.

The four panels that follow document the training and evaluation behaviour of YOLOv10n under the protocol of Table 6. They are supplementary to the comparative analysis of Section 4 and are reproduced in full so that the convergence behaviour of each architecture can be inspected independently.




**Fig. S11:** Normalised confusion matrix for YOLOv10n on the 165-image test partition. The background row and column record predictions that were not matched to any annotation, so the matrix is informative about a tendency to fire on non-lesion regions rather than about object-level accuracy.

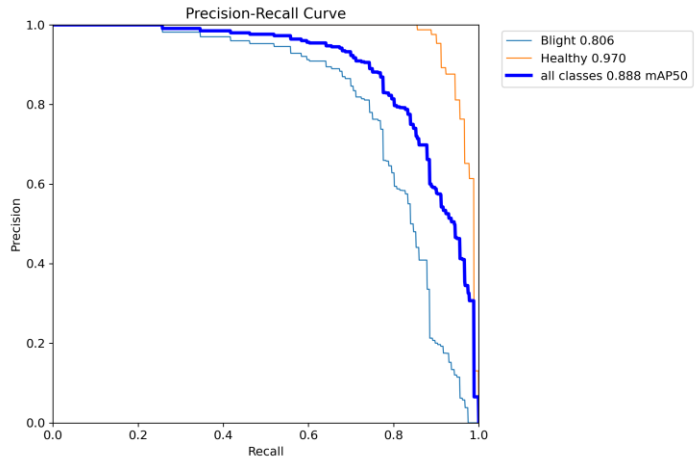



**Fig. S12:** Precision-recall curve for YOLOv10n across the confidence sweep, from which the average precision values of Table 14 are integrated.

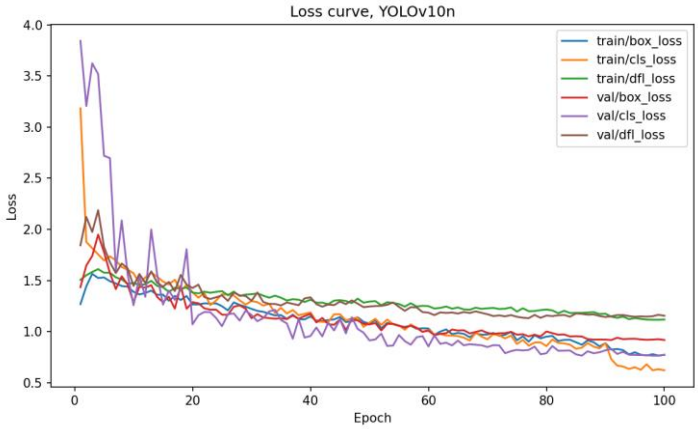



**Fig. S13:** Training and validation loss components for YOLOv10n across the hundred epochs, showing the divergence between the falling training box loss and the plateaued validation metrics discussed in Section 4.8

.

**Figs. S14-S17 — YOLO11n**

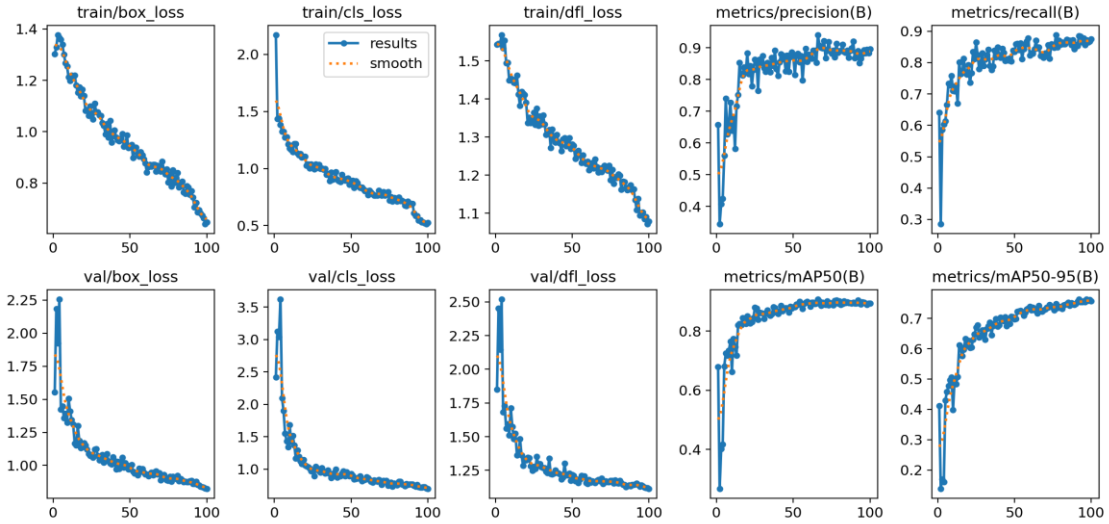



**Fig. S14:** Training and validation curves for YOLO11n across the hundred epochs of the common protocol, showing the two mean average precision measures in the upper panels and the box, classification, and distribution-focal loss components in the lower panels.

The four panels that follow document the training and evaluation behaviour of YOLO11n under the protocol of Table 6. They are supplementary to the comparative analysis of Section 4 and are reproduced in full so that the convergence behaviour of each architecture can be inspected independently.

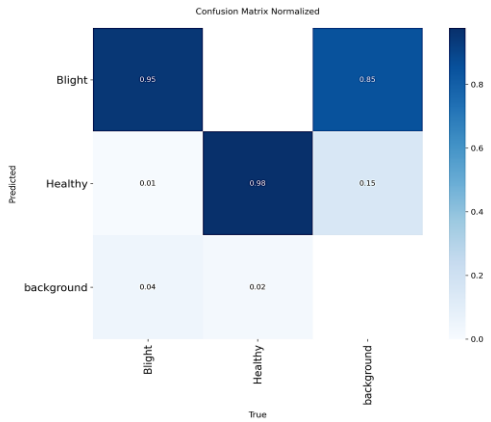



**Fig. S15:** Normalised confusion matrix for YOLO11n on the 165-image test partition. The background row and column record predictions that were not matched to any annotation, so the matrix is informative about a tendency to fire on non-lesion regions rather than about object-level accuracy.

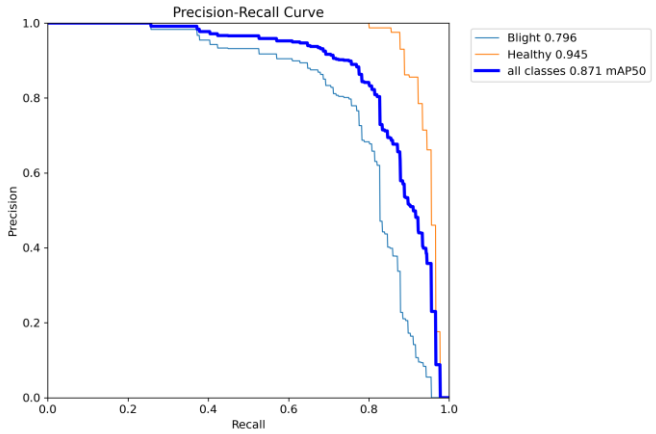



**Fig. S16:** Precision-recall curve for YOLO11n across the confidence sweep, from which the average precision values of Table 14 are integrated.

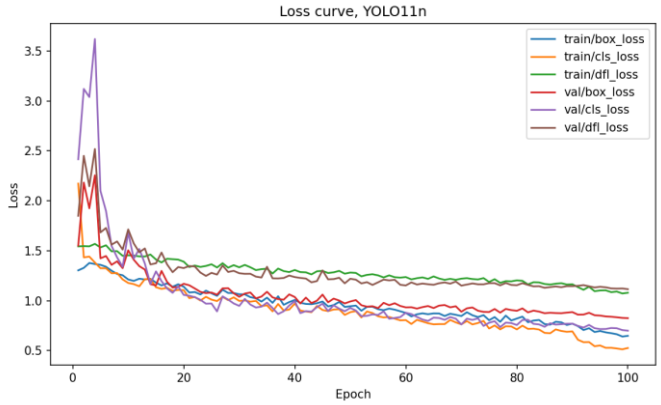



**Fig. S17:** Training and validation loss components for YOLO11n across the hundred epochs, showing the divergence between the falling training box loss and the plateaued validation metrics discussed in Section 4.8.

**Part B. Secondary interface figures (Fig. S18)**

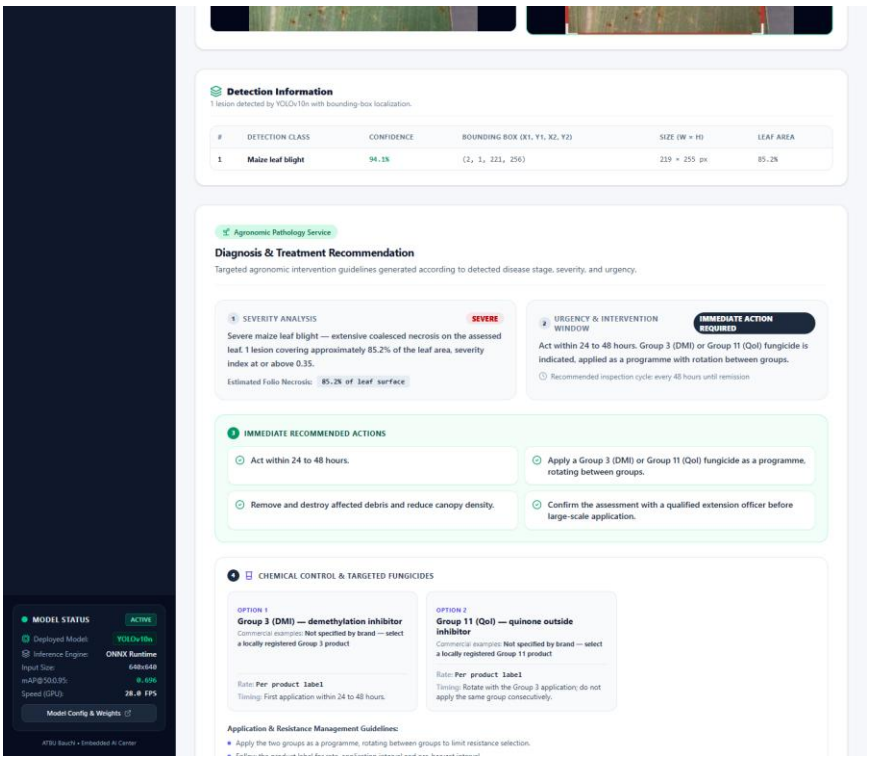



**Fig. S18:** Lower region of the result view for the reference scan, showing the severity analysis, the urgency and intervention window, the immediate recommended actions, and the chemical-control block returned by the knowledge base. Cited in Section 4.11.

The upper region of the same result view is reproduced as Fig. 16 of the main text. Together the two panels document that the interface states the severity index, the tier, the urgency window, the chemical class by FRAC group, the cultural control and prevention entries, and the safety disclaimer as separate fields, which is the presentation consequence of keeping perception and advice apart.

**Part C. Interface figures formerly numbered 14 to 23 (Figs. S19-S27)**

The nine figures below were numbered 14 to 23 in the original submission and have been moved to this file in order to consolidate the main text. Their order is preserved, so that former Fig. 14 corresponds to Fig. S19 and former Fig. 23 to Fig. S27. The former Fig. 18, the scan-history archive, is retained in the main text as Fig. 17.

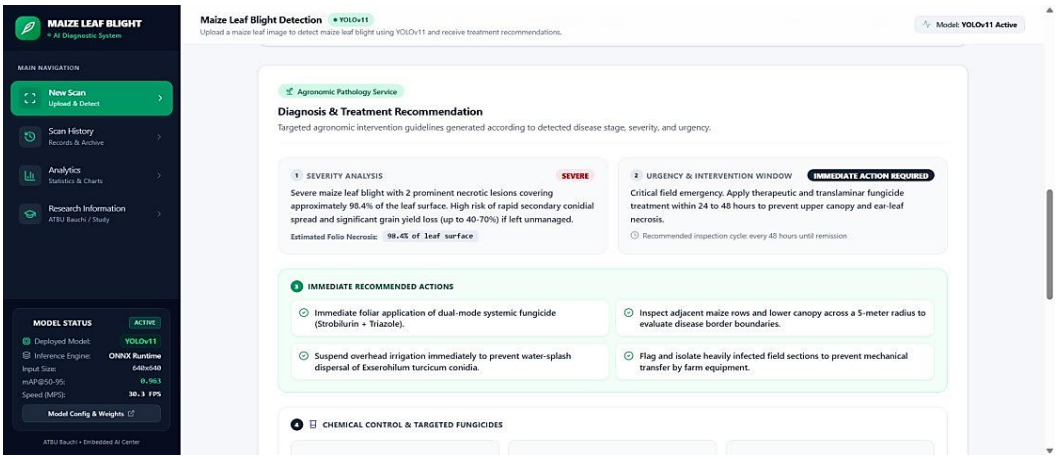



**Fig. S19:** Structured agronomic diagnosis and treatment recommendation returned for a blighted sample (formerly Fig. 14 of the original submission), comprising the diagnosis, the severity statement, the urgency and intervention window, the chemical-control block, the cultural-control block, and the prevention block.

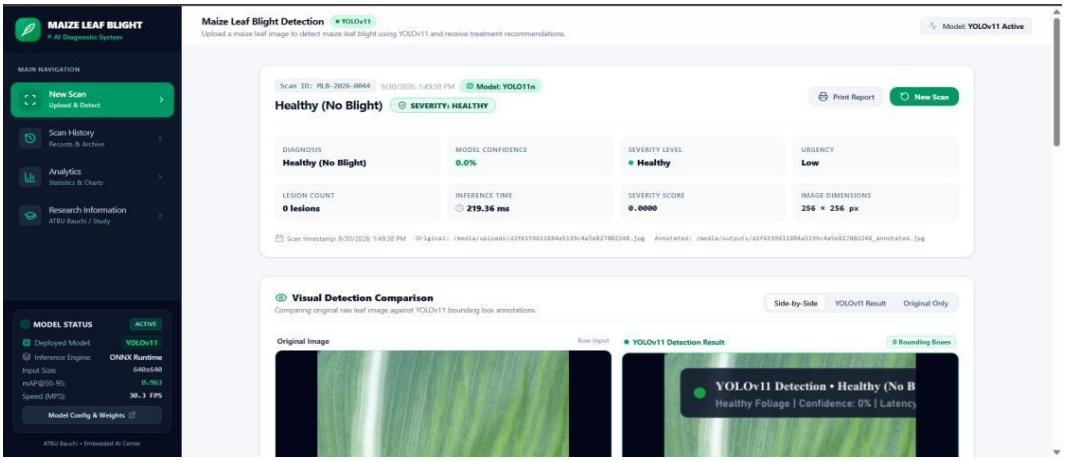



**Fig. S20:** Detection result overview for healthy maize foliage (formerly Fig. 15), showing a single healthy-class detection and the absence of any blight-class detection above the operating threshold.

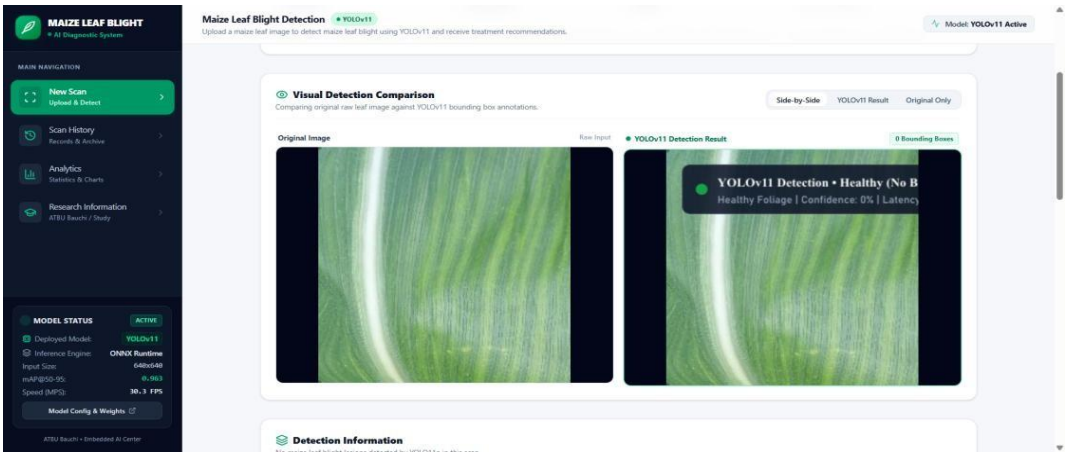



**Fig. S21:** Visual detection comparison and lesion-absence verification for a healthy maize leaf (formerly Fig. 16), with the original and annotated images displayed side by side.

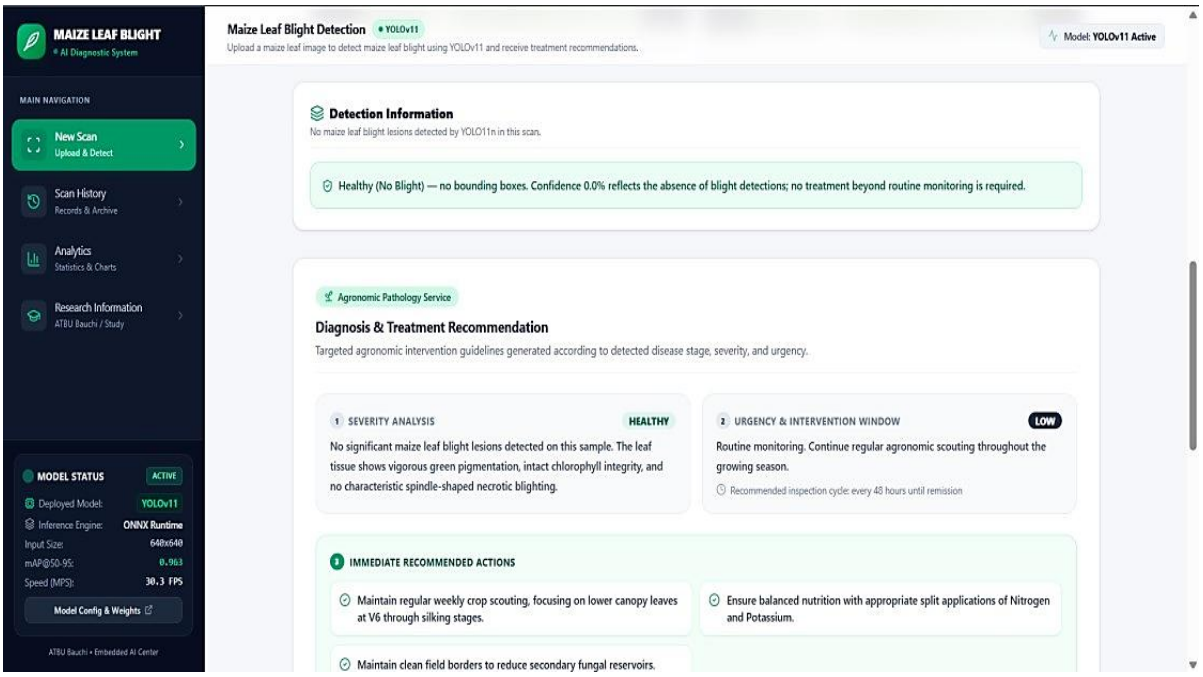



**Fig. S22:** Agronomic monitoring recommendations and routine management returned for the healthy class (formerly Fig. 17), that is the tier-0 record of Table 8.

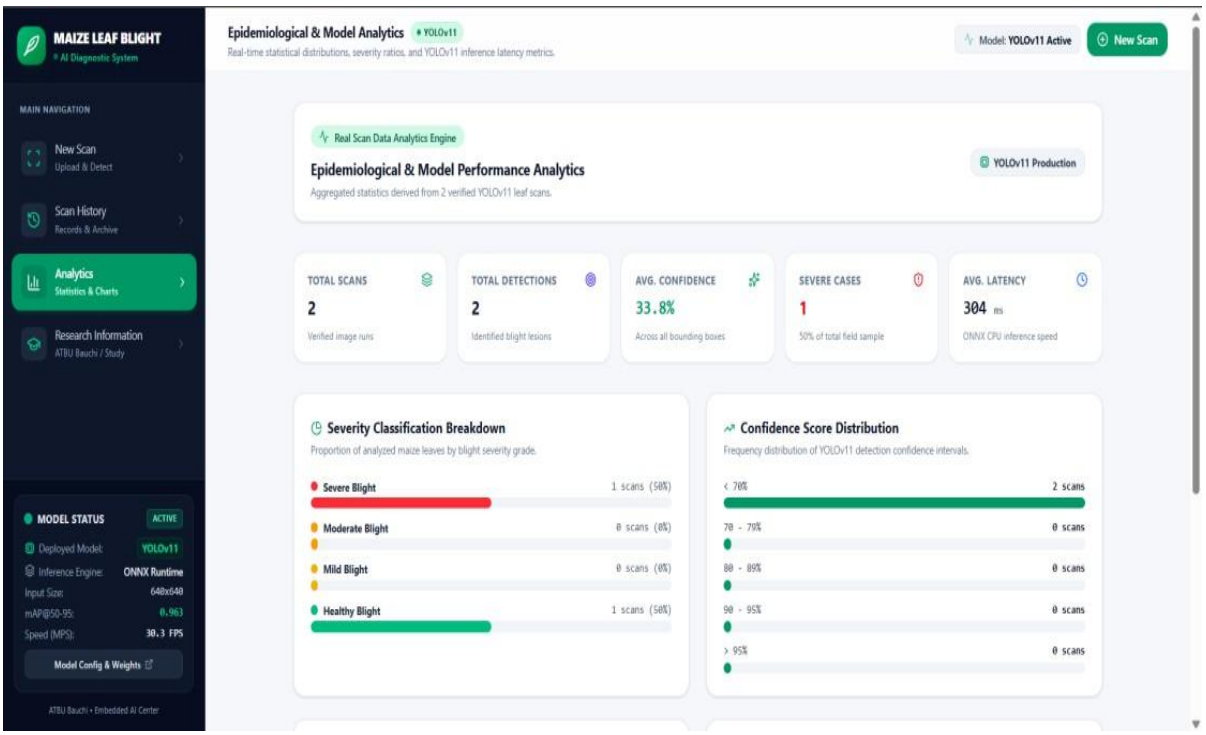



**Fig. S23:** Epidemiological and model performance analytics dashboard (formerly Fig. 19), aggregating the persisted scan records by severity tier and by model.

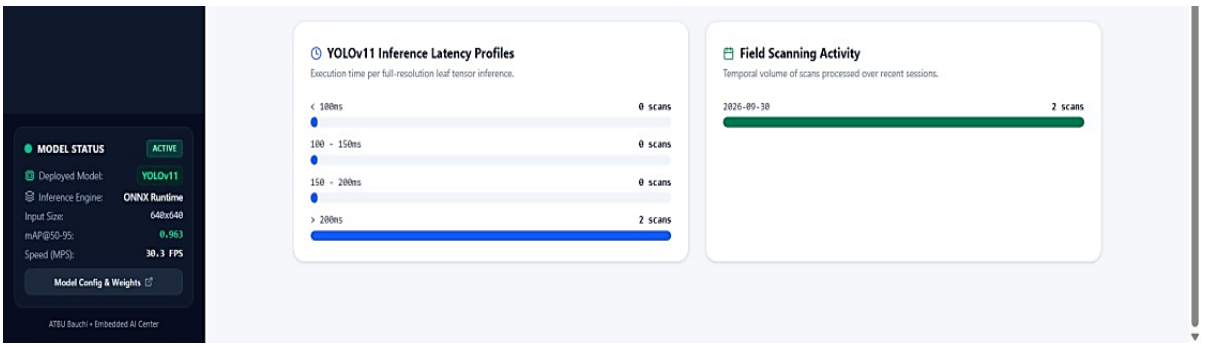



**Fig. S24:** Analytics panel reporting the inference-latency distribution reconstructed from the audit store and the field scanning activity (formerly Fig. 20).

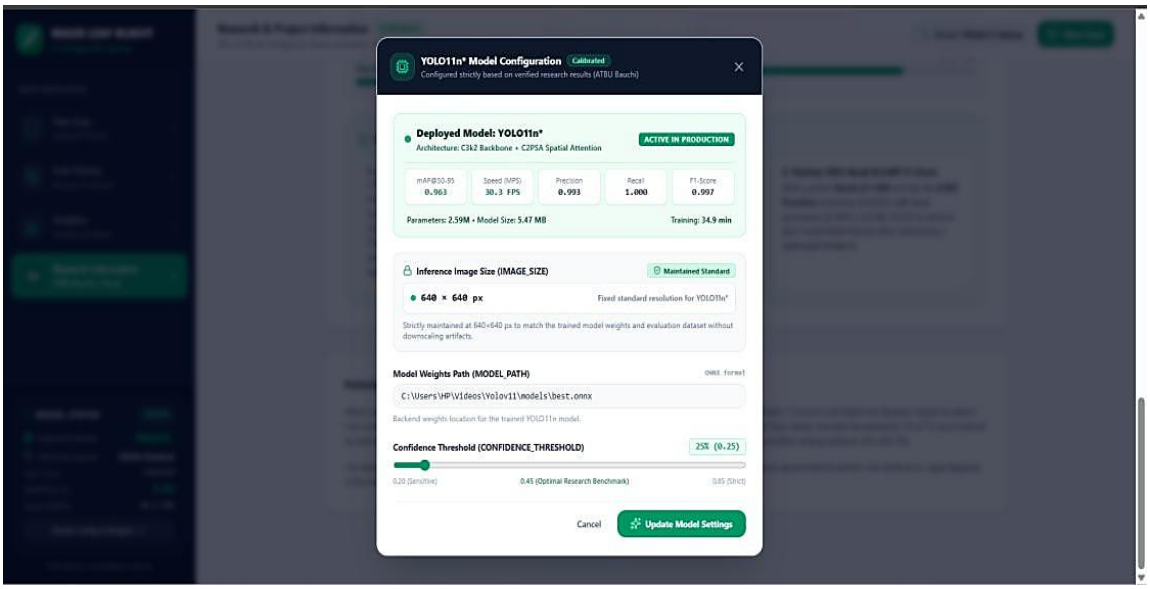



**Fig. S25:** Model runtime configuration and inference parameter modal (formerly Fig. 21), through which the operating confidence threshold and the displayed model identifier are set.

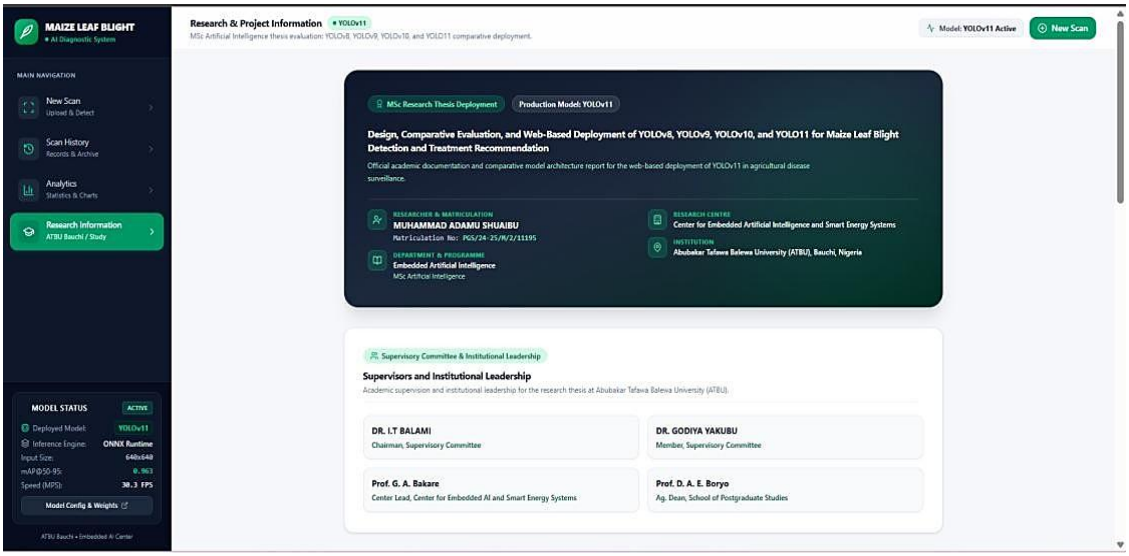



**Fig. S26:** Research project information screen (formerly Fig. 22), documenting the institutional affiliation and the supervisory committee of the study.

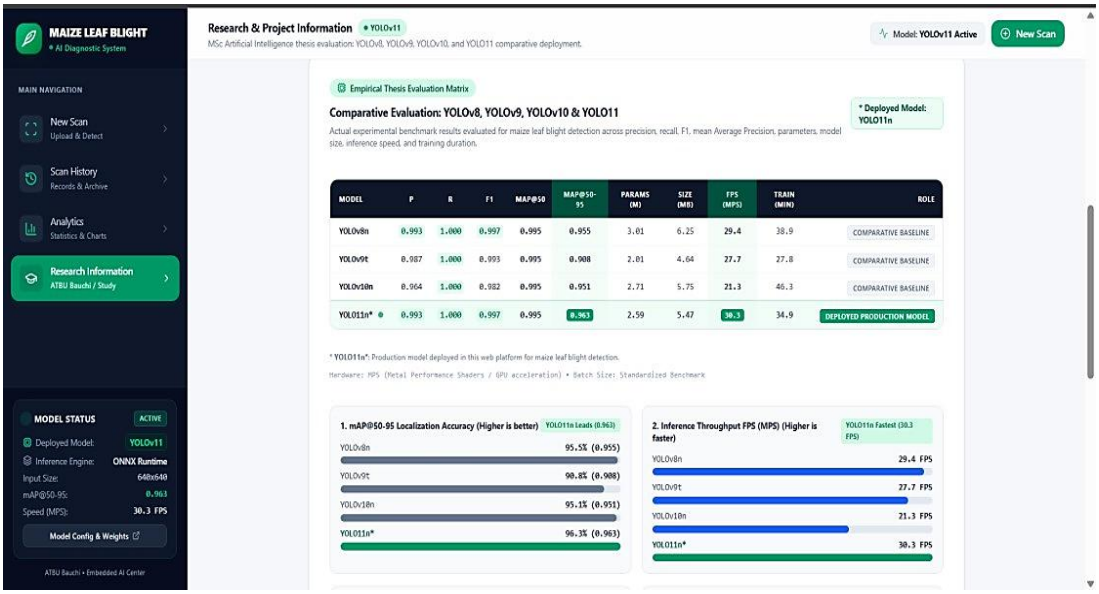



**Fig. S27:** In-system comparative evaluation matrix and performance benchmarks for the four detectors (formerly Fig. 23), as displayed by the running application.

## Part D. Deployment and interface-state figures (Figs. S28-S31)

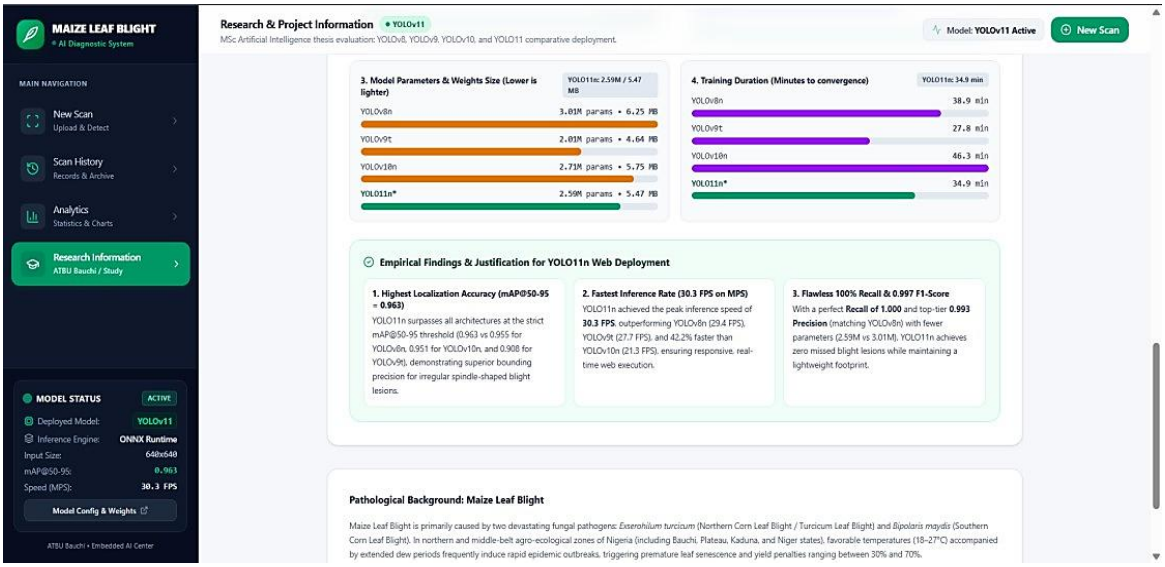



**Fig. S28:** In-system model parameter analysis and deployment justification, reporting the parameter count, weight footprint, and measured throughput of the deployed artefact against the alternatives. Cited in Section 4.11

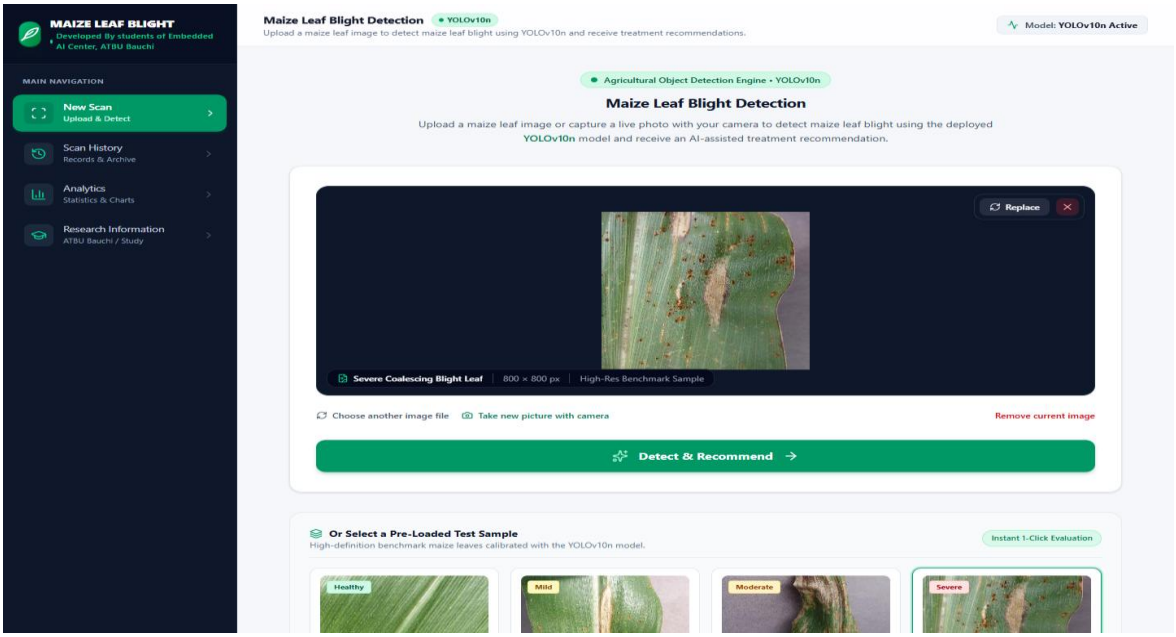



**Fig. S29:** Sample-selection state of the landing view: one of the four documented sample images has been loaded and the detection control has become active. Cited in Section 4.11.

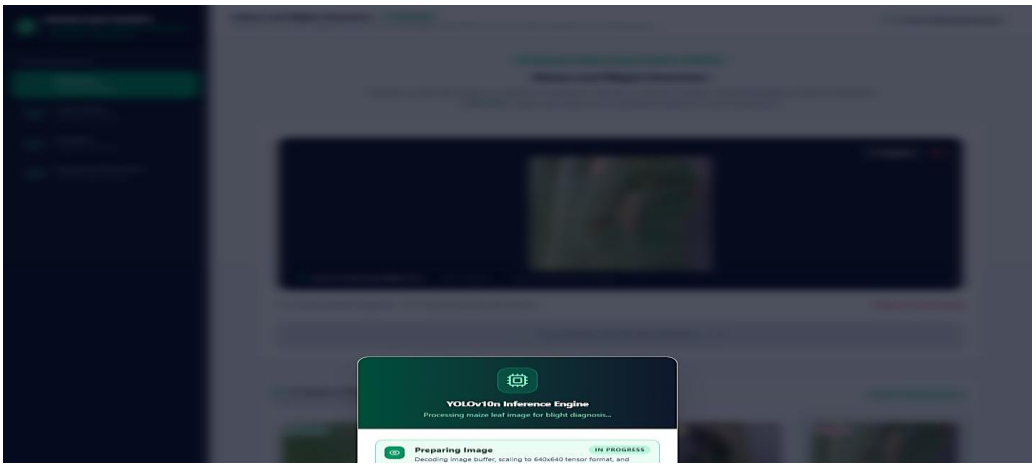



**Fig. S30:** Detection-in-progress state, in which the control is disabled while the request is outstanding. Cited in Section 4.11.

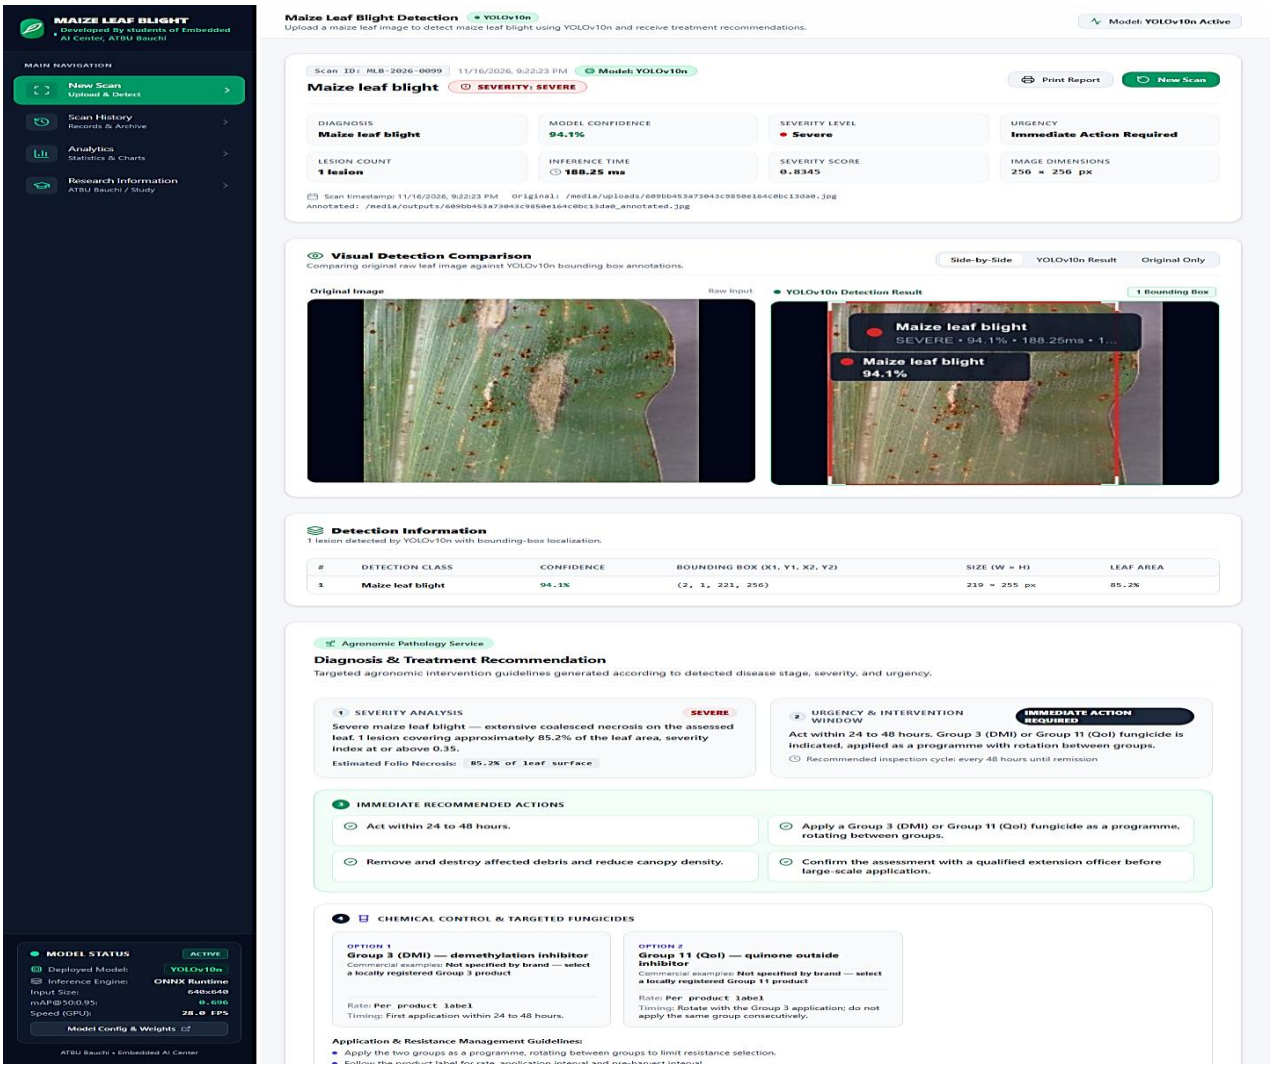



**Fig. S31:** Complete result view on a single page, comprising the scan record, the visual detection comparison, the detection table, and the full advisory record. Cited in Section 4.11.

All interface figures were captured from the running service on the four documented sample images. No figure in this file has been retouched, and no interface element has been added or removed after capture.