ICPA
17th INTERNATIONAL CONFERENCE ON PRECISION AGRICULTURE
ConBAP 2026
CONGRESSO BRASILEIRO DE AGRICULTURA DE PRECISÃO E DIGITAL
13-16 July 2026 | Porto Alegre, Brazil
ISPA
AsBraAP

From Render to Field: Detecting Asian Soybean Rust Using Models Trained Exclusively on Synthetic Imagery

Augusto de Freitas[1,2], Eduardo Farinati Leite[1,2], Giusepe Disconzi Dallegrave[1,2], Lucas Brizola Fontoura[1], Rafael Volpe de Freitas[1], Vinicius Silveira Mello[1,2], João Francisco Valiati[1,2]

[1] Center for Embedded Devices and Research in Digital Agriculture (CEDRA) of SENAI-RS,

[2] SENAI Innovation Institute for Sensing Systems (ISI-SIM), São Leopoldo, Brazil.

A paper from the Proceedings of the
17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture
13-15 July 2026
Porto Alegre, Brazil

Abstract.

Training machine learning models for crop disease detection require large, annotated datasets that are costly and difficult to obtain under variable field conditions. Asian Soybean Rust (ASR) can reduce soybean yield by up to 90% and costs Brazilian producers over US$2 billion per season in fungicide applications and yield losses. Despite this impact, existing machine learning studies on ASR remain scarce with no publicly available RGB datasets annotated at the lesion level. The data problem is compounded by visual symptoms that evolve gradually, vary with environmental conditions, and are difficult to capture consistently in real-world imagery. This work investigates whether models trained exclusively on synthetic data can classify and segment lesions of this disease on real soybean images, with the long-term goal of enabling earlier detection that supports timely intervention and reduces crop losses. Building on prior work that applied procedural synthetic data generation for disease detection, this study extends this approach to a new crop and pathology and introduces instance segmentation as an additional task. A procedural pipeline was developed combining physically based modeling, real image textures, and parametric 3D geometry to generate disease symptoms with control over lesion morphology, distribution, and severity progression. Complete crop scenes were constructed with variable plant geometry, growth stages, canopy structure, and lighting conditions, simulating diverse field scenarios and disease development over time. A YOLO-based model was trained exclusively on this synthetic dataset and evaluated on real ASR leaf photographs across two

The authors are solely responsible for the content of this paper, which is not a refereed publication. Citation of this work should state that it is from the Proceedings of the 17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture. EXAMPLE: Last Name, A. B. & Coauthor, C. D. (2026). Title of paper. In Proceedings of the 17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture (unpaginated, online). Monticello, IL: International Society of Precision Agriculture and Brazilian Congress on Precision and Digital Agriculture.

tasks: the model was first used as a classifier, images in which any lesion was detected were labeled as infected, achieving 97.73% accuracy and 98.56% F1-score. For semantic segmentation of the lesions, the model obtained 0.404 mAP@25%, reliably localizing diseased regions despite limited precision in individual lesion boundaries, a trade-off acceptable for field-level decision-making, where confirming disease presence matters more than delineating lesion boundaries. The evaluation was performed on real leaf photographs from the MH-SoyaHealthVision dataset that contains only image-level annotations, used as a proxy for aerial imagery given the absence of publicly available segmentation-annotated datasets for ASR. A subset of 20 images was manually annotated with pixel-level segmentation masks, as the original dataset contained only classification labels. The results demonstrate that synthetic data alone is sufficient to train models that classify ASR with high accuracy and segment its lesions functionally, circumventing the absence of annotated datasets for this pathology. By removing this barrier, the proposed approach enables the development of detection systems that support timely intervention against ASR, directly reducing fungicide costs and yield losses in soybean production

Keywords.
Asian Soybean Rust, Synthetic Data, Disease Detection, Procedural Generation, Domain Adaptation, Instance Segmentation

Introduction

Soybean is one of the most economically important crops worldwide, serving as a primary source of protein and oil for food, feed, and biofuel production. Major producers such as Brazil, United States, and Argentina rely heavily on soybean yields for both domestic use and export-driven economies (Hamza et al., 2024). However, this productivity is constantly threatened by plant diseases, among which Asian Soybean Rust (ASR), caused by the fungus *Phakopsora pachyrhizi*, stands out as one of the most destructive, capable of spreading rapidly under favorable environmental conditions, causing severe defoliation and yield losses of up to 90% under severe infection, and costing Brazilian producers alone over US$2 billion per season in fungicide applications and yield losses (Godoy et al., 2015; Silva et al., 2020; Murithi et al., 2016; Childs et al., 2018; Langenbach et al., 2016).

Timely detection of ASR is critical for effective management, as fungicide applications are most efficient when applied at early stages of infection. In practice, however, disease monitoring is still largely manual, relying on field scouting performed by local farmers (Khakimov et al., 2022). This process is labor-intensive, time-consuming, and often limited in spatial and temporal coverage. Recent advances in computer vision and deep learning offer a promising alternative, enabling automated disease detection from images captured at different scales, including proximal leaf imagery and aerial data acquired via unmanned aerial vehicles (UAVs) (Li et al., 2023; Adimas, A.K. et al., 2025; Jeong et al., 2025). Such approaches have the potential to support precision agriculture by enabling rapid, large-scale monitoring and targeted intervention.

Despite this potential, the development of robust machine learning models for ASR detection faces a fundamental bottleneck: the lack of large, high-quality annotated datasets, a challenge well-documented across plant disease detection in general (Pacal et al., 2024), but particularly acute for ASR given the gradual evolution of symptoms, variability across environmental conditions, and the high cost of expert labeling. While several studies have explored deep learning for plant disease detection, most rely on classification tasks using curated datasets, and very few address lesion-level segmentation (Zhou et al., 2023; Pacal et al., 2024). Moreover, our review of publicly available datasets yielded no RGB benchmark for ASR with pixel-level annotations, further limiting progress in lesion-level segmentation for this disease.

One emerging strategy to address data scarcity is the use of synthetic data generation. Prior work has shown that procedurally generated datasets can be effective for training models in agricultural contexts, offering full control over scene composition and annotation at virtually no labeling cost (Klein et al., 2024; Cieslak et al., 2024; Farinati Leite et al., 2025). However, the application of synthetic data to ASR remains largely unexplored, particularly for more complex tasks such as instance or semantic segmentation of its lesions.

In this work, we investigate the feasibility of training deep learning models solely on synthetic data for the classification and segmentation of ASR lesions in real soybean imagery. The hypothesis is that the domain gap between synthetic and real imagery can be reduced by controlling lesion morphology, disease progression, and scene variability during procedural generation. To test this, we train a YOLO-based model exclusively on a synthetically generated dataset and evaluate it on real leaf photographs across both tasks. Due to the scarcity of annotated lesion segmentation datasets, evaluation relies in part on classification benchmarks; we argue that performance on these benchmarks nonetheless demonstrates that synthetic data captures enough of the relevant visual structure to support generalization across tasks. By addressing the data scarcity problem, this study aims to advance the development of scalable, annotation-efficient solutions for early ASR detection.

Methodology

This section describes the methodology adopted to address the research question. It covers four components: the procedural pipeline used to generate synthetic training data, the dataset composition and train/test split design, the model architecture selected for detection and segmentation and the evaluation protocol and metrics used to assess performance on real photographs.

Procedural Synthetic Data Generation Pipeline

This work builds on a procedural generation pipeline to produce annotated synthetic images of soybean plants affected by ASR, without requiring any real labeled data. Once the 3D assets are created, the pipeline is fully automated, generating individual samples at will from parameterized distributions. Each image is generated by sampling from these distributions, which control two main aspects of the scene: lesion appearance, modeled to reflect the known visual characteristics of ASR symptoms, and plant structure, varying geometry and growth stage across soybean developmental phases.

Lesion morphology is modeled procedurally based on the known visual characteristics of ASR symptoms. Infected leaves typically present small, angular, water-soaked lesions that evolve into raised pustules (uredinia) ranging from pale yellow to tan and dark brown, depending on the infection stage and the host resistance profile. Lesion coloration, size, and density also vary with environmental conditions and geographic region, with more aggressive phenotypes reported in tropical climates such as Brazil and Colombia (Goellner et al., 2010). To capture this variability, the pipeline samples from parameterized distributions governing lesion size, shape, color gradient, and density, simulating the visual spectrum of ASR symptoms across different infection stages, from early water-soaked lesions to dense pustule clusters. Lesion placement follows a biologically informed spatial distribution on leaf surfaces, reflecting natural patterns of disease spread, typically initiating on lower canopy leaves and progressing upward.

Plant scenes are constructed using parametric 3D geometry that varies leaf arrangement, canopy density, and growth stage across samples. Physically based rendering (PBR) is used in conjunction with real image textures to produce visually plausible leaf surfaces under diverse lighting conditions. This combination of geometric variability and realistic rendering reduces the domain gap between synthetic training data and real field photographs, exposing the model to viewpoint, illumination, and occlusion variation expected at inference time. Figure 1 presents a visual comparison between real ASR infected leaves and synthetic samples generated by the pipeline.

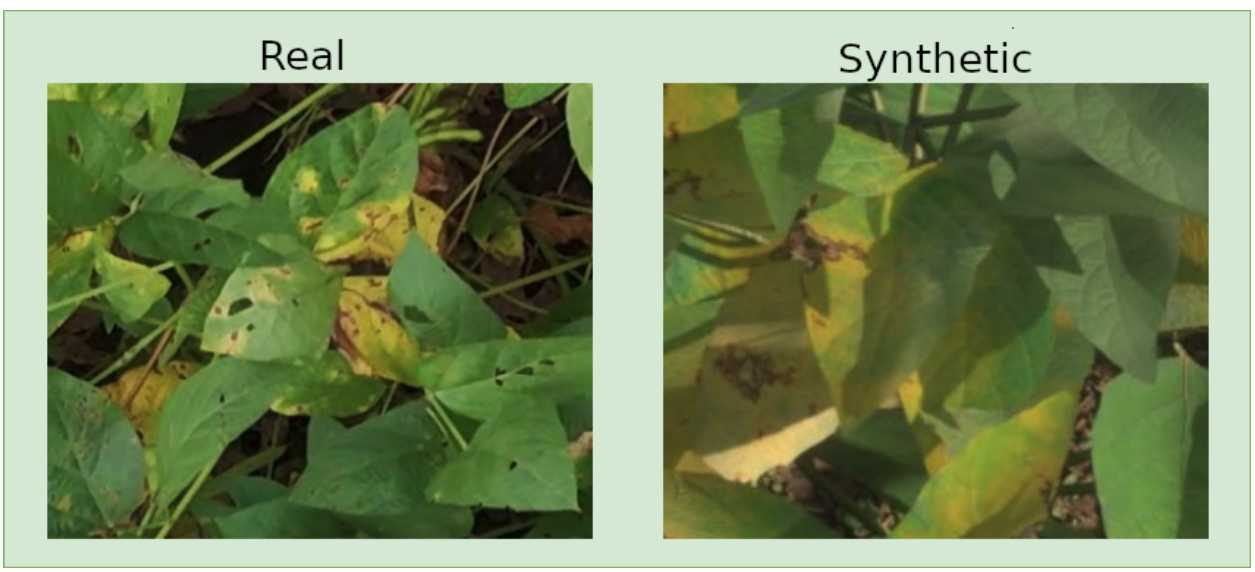

Figure 1: Comparison between real ASR infected leaves and synthetic samples generated by the pipeline.

The pipeline produces both RGB images and paired instance segmentation masks at no additional annotation cost, as ground-truth masks are derived directly from the geometry during rendering.

Dataset Description

The training and validation data consist entirely of synthetically generated images produced by the pipeline described above. The dataset comprises 5,000 images, split into training and validation subsets following an 80/20 ratio, yielding 4,000 training images and 1,000 validation images. Each image is accompanied by instance-level segmentation masks and bounding box annotations in YOLO format.

Evaluation was performed entirely on photographs from the MH-SoyaHealthVision dataset (Shinde et al., 2024), a publicly available collection of real soybean leaf images with image-level classification labels covering multiple diseases and pests, including ASR. Since this dataset provides no lesion-level annotations, a subset of 20 images was manually annotated by the authors with pixel-level instance segmentation masks to enable quantitative evaluation of the segmentation task. The full available test set was used for the classification task, where ground truth is a per-image binary label.

This strict train/test separation is central to the study design, as the primary research question concerns the ability of synthetically trained models to generalize real-world imagery without any real labeled data during training.

Model Architecture

A YOLO-based architecture (Redmon et al., 2015) was selected for its inference speed, native support for instance segmentation, and suitability for deployment in resource-constrained field scenarios such as edge devices mounted on UAVs. Specifically, YOLO26s (Sapkota et al., 2026) and its segmentation variant (yolo26s-seg) were adopted, extending the standard detection head with a lightweight mask prediction branch. This enables the model to produce both bounding box predictions and pixel-level instance masks, supporting both tasks evaluated in this work.

The architecture operates on input images of 1024x1024 pixels. The backbone is a CSP-based feature extractor that produces multi-scale feature maps, which are passed to a path aggregation network (PAN) neck for feature fusion across scales. The detection and segmentation heads share the backbone and neck, with the mask branch producing one segmentation mask per detected instance. This design is particularly suited to detecting small, densely distributed objects such as ASR lesions, as it explicitly models features at multiple spatial resolutions (Hidayatullah and Tubagus, 2026).

Configuration, Evaluation and Metrics

The model was initialized with weights pre-trained on the COCO dataset and fine-tuned for 200 epochs using stochastic gradient descent (SGD) with an initial learning rate of 0.01 and linear decay with momentum of 0.937 and a weight decay of 5×10^{-4}. Standard YOLO data augmentation was applied, including random horizontal flips, mosaic augmentation, and color jitter. Training was conducted on a single Quadro RTX 4000 GPU with a batch size of 4.

A single train/validation/test split was used, with no cross-validation, as the training domain (synthetic) and the test domain (real photographs) are categorically distinct. Cross-validation over the synthetic training set would not provide a meaningful estimate of real-world generalization. The held-out real-image test set was used once, after training was complete, to report final performance.

- For classification, the model's instance predictions are aggregated at the image level: any image in which at least one lesion is detected above a confidence threshold is labeled as infected. Accuracy measures the proportion of images correctly classified as infected or healthy (1). The F1-score is the harmonic mean of precision and recall (2), providing a single measure that balances both error types.

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN) \quad (1)$$

$$F1 = (2 \cdot \text{precision} \cdot \text{recall}) / (\text{precision} + \text{recall}) \quad (2)$$

where TP, TN, FP and FN denote true positives, true negatives, false positives and false negatives, respectively.

- For instance segmentation, performance is reported using mean Average Precision (mAP), computed against the manually annotated pixel-level masks. mAP summarizes detection quality across confidence thresholds by averaging the area under the precision-recall curve for each class (3). A predicted lesion is counted as a true positive only if its overlap with the ground-truth mask exceeds a defined Intersection over Union (IoU) threshold, where IoU measures the ratio of the overlapping area to the total area covered by both masks (4):

$$IoU = \text{Area of Overlap} / \text{Area of Union} \quad (3)$$

$$mAP = 1 / C \cdot \Sigma \int Pc(r)\, dr \quad (4)$$

where C is the number of classes and Pc(r) is the precision at recall r for class c.

In field-level disease monitoring, management decisions rely on identifying the presence and approximate location of infected regions to trigger fungicide application, rather than precisely delineating individual lesion boundaries. Accordingly, mAP@25 is adopted as the primary metric, as it rewards correct spatial localization while tolerating boundary imprecision consistent with this operational requirement. mAP@50 is reported as a secondary metric to characterize boundary delineation quality.

Results

This section evaluates the proposed methodology on real soybean leaf images. Classification metrics are first reported to assess disease identification performance, followed by instance segmentation results that quantify lesion localization accuracy. Visual examples are then presented to illustrate prediction behavior and to provide additional insight into the model's ability to generalize from synthetic training data to real-world images.

Classification Results

Table 1 reports classification performance on the real-image test set, where images with at least one lesion detected above the confidence threshold are labeled as infected.

Table 1. Classification performance on the real-image test set.

Metric	Value
Accuracy	0.9773
Precision	0.9774
Recall	0.9940
F1	0.9856

The model achieves 97.73% accuracy and an F1-score of 98.56%, demonstrating that a YOLO model trained exclusively on synthetic data can classify ASR infection with high reliability on real soybean leaf photographs. The strong recall (99.40%) indicates that the model rarely misses genuinely infected images, which is the operationally critical failure mode in disease monitoring contexts, a false negative delays intervention and allows the disease to spread, whereas a false positive triggers an unnecessary but recoverable fungicide application. The precision of 97.74% reflects an equally low false positive rate, confirming that the model seldom flags healthy tissue as infected and that both error modes remain within ranges acceptable for field-level decision support.

These results should be interpreted in the context of a binary classification task, with a single disease class. This constrained problem formulation reduces inter-class ambiguity and likely contributes to the high performance observed; extending the pipeline to multiple foliar diseases would introduce overlapping symptoms and is expected to increase classification difficulty.

Segmentation Results

Table 2 reports instance segmentation performance on the 20 manually annotated real images. The model achieves mAP@25 of 0.404 and mAP@50 of 0.374. This difference reflects the difficulty of meeting a 50% overlap criterion for ASR lesions, which are small, irregularly shaped, and clustered in high-density regions that make individual instance delineation genuinely ambiguous even for human annotators.

Table 2. Instance segmentation performance on the manually annotated subset.

Metric	Value
mAP@50	0.374
mAP@25	0.404

Figures 2 and 3 present qualitative examples from the evaluation set. In images with well-defined lesions on visible leaf surfaces, predicted masks align closely with ground-truth annotations. Boundary imprecision is more pronounced in high-density lesion clusters. In healthy images, the model produces no detections, indicating that learned features are specific to ASR lesion appearance rather than general leaf texture.

Figure 2. Comparison between prediction and ground truth labels on infected test data.

Figure 3. Comparison between prediction and ground truth labels on healthy test data.

Together, these results indicate that the model localizes diseased regions in the majority of evaluated images, with performance degrading at the boundary level rather than at the level of disease presence detection, which is the operationally relevant distinction for triggering agronomic interventions.

Limitations

The segmentation evaluation relies on 20 manually annotated images, limiting the statistical representativeness of the reported metrics. Since no public pixel-level ground truth exists for ASR, annotation was performed by the authors, introducing subjectivity given the ambiguous nature of lesion boundaries. Additionally, only one architecture was evaluated without comparison to alternative models or real-data baselines, a deliberate scoping decision for a proof-of-concept study; without a real-data baseline, the relative contribution of synthetic data quality versus architectural choices remains an open question for future ablation studies. Finally, the synthetic pipeline does not capture the full morphological diversity across soybean cultivars or the shifting symptom expression driven by climate variability, meaning model performance may degrade in conditions the pipeline was not designed to represent.

Conclusion

This work investigated whether deep learning models trained exclusively on synthetic data can classify and segment ASR lesions in real soybean imagery, addressing a fundamental bottleneck in the field: the absence of publicly available pixel-level annotated datasets for this pathology. A YOLO26s model trained entirely on procedurally generated data achieved 97.73% accuracy and 98.56% F1-score for classification, and mAP@25 of 0.404 for lesion segmentation, without exposure to any real labeled image during training.

The model performs best in conditions well-represented by the synthetic pipeline: images with clearly visible lesions on non-occluded leaf surfaces across a range of lighting conditions. In these scenarios, the domain gap between synthetic and real data does not meaningfully impair performance, supporting the hypothesis that procedural generation can substitute for real-world annotation. Performance degradation occurs primarily at the boundary delineation level rather than at disease presence detection, consistent with the residual domain gap between synthetic and real tissue appearance. This distinction is operationally relevant: classification results support deployment in field monitoring pipelines, while segmentation provides sufficient spatial localization for canopy-level decision-making despite imprecise lesion boundaries.

Several directions for future work follow directly from the current limitations. The manually annotated subset of 20 images released alongside this study represents a first pixel-level benchmark for ASR segmentation; expanding it through community annotation or AI-assisted labeling would yield more statistically robust evaluation. Fine-tuning the synthetically trained model on a small number of real annotated images could close the remaining domain gap at low annotation cost, combining the scalability of synthetic pretraining with the precision of real-data adaptation. Finally, extending evaluation to aerial imagery at UAV altitudes would bring the study closer to the intended deployment scenario and enable a more direct assessment of readiness for field-scale deployment.

Acknowledgments

This work has been fully funded by the project FieldEye supported by the Center for Embedded Devices and Research in Digital Agriculture (CEDRA) of SENAI-RS, with financial resources from the PPI IoT/Manufatura 4.0 / PPI HardwareBR of the MCTI, grant number 053/2023, signed with EMBRAPII.

References

ADIMAS, A.K.; MEKONEN, M.Z.; ASSEGIE, T.A.; SINGH, H.K.; MAZUMDAR, I.; GUPTA, S.K.; SALAU, A.O.; TIN T.T. Soybean leaf disease detection and classification using deep learning approach. Bulletin of Electrical Engineering and Informatics (BEEI). Vol 14, No 4. 2025.

CHILDS, S.P.; BUCK, J.W.; LI, Z. Breeding soybeans with resistance to soybean rust (Phakopsora pachyrhizi). Plant Breed. 2018;137(3):250-261.

CIESLAK, M.; GOVINDARAJAN, U.; GARCIA, A.; CHANDRASHEKAR, A.; HADRICH, T.; MENDOZA-DROSIK, A.; et al. Generating diverse agricultural data for vision-based farming applications. 2024.

FARINATI LEITE, E.; BRIZOLA FONTOURA, L.; DE FREITAS, A.; DISCONZI DALLEGRAVE, G.; VOLPE DE FREITAS, R.; MELLO, V.S.; VALIATI, J.F. 2025. VitiForge: a new procedural pipeline approach for grapevine disease identification under data scarcity. Front. Plant Sci. 16:1706973. doi: 10.3389/fpls.2025.1706973

GODOY, C.V.; BUENO, A.D.F.; GAZZIERO, D.L.P. Brazilian soybean pest management and threats to its sustainability. Outlooks Pest Manag. 26(3): 13-117. 2015.

GOELLNER, K.; LOEHRER, M.; LANGENBACH, C.; CONRATH, U.; KOCH, E.; SCHAFFRATH, U. Phakopsora pachyrhizi, the causal agent of Asian soybean rust. Mol Plant Pathol. 2010 Mar;11(2):169-77. doi: 10.1111/j.1364-3703.2009.00589.x. PMID: 20447267; PMCID: PMC6640291.

HAMZA, M.; BASIT, A.B.; SHEHZADI, I.; TUFAIL, U.; HASSAN, A.; HUSSAIN, T.; SIDDIQUE, M.U.; HAYAT, H.B. Global Impact of Soybean Production: A Review. Asian Journal of Biochemistry, Genetics and Molecular Biology (p. 12-20). 2024.

HIDAYATULLAH, P.; TUBAGUS, R. YOLO26: A Comprehensive Architecture Overview and Key Improvements. arXiv preprint arXiv:2602.14582. 2026.

KALUZNY, J.; SCHERECKENBER, Y.; CYGANIK, K.; ANNIGHOFER, P.; PIRK, S.; MICHELS, D. L.; et al. 2024. Laesi: Leaf area estimation with synthetic imagery. arXiv preprint arXiv:2404.00593. doi: 10.48550/arXiv.2404.00593

KLEIN J.; WALLER R.; PIRK S.; PALUBICKI W.; TESTER M.; MICHELS, D.L. (2024) Synthetic data at scale: a development model to efficiently leverage machine learning in agriculture. Front. Plant Sci. 15:1360113. doi: 10.3389/fpls.2024.1360113

JEONG, M.; PARK, S.; KWON, S.; LIM, K.; JUNG, D.; LEE, H.; KIM, H.J; SHIN, J. Rapid detection of soybean nutrient deficiencies with YOLOv8s for precision agriculture advancement. Sci Rep. 2025.

KHAKIMOV, A.; SALAKHUTDINOV, L.; OMOLIKOV, A.; UTAGANOV, S. Traditional and current-prospective methods of agricultural plant diseases detection: a review. IOP Conf Ser Earth Environ Sci 951(1):012002. 2022.

LANGENBACH, C.; CAMPE, R.; BEYER, S.F.; MUELLER, A.N.; CONRATH, U. 2016. Fighting Asian Soybean Rust. Front. Plant Sci. 7:797. doi: 10.3389/fpls.2016.00797

LI, W.; YU, X.; CHEN, C.; GONG, Q. Identification and localization of grape diseased leaf images captured by uav based on cnn. Comput. Electron. Agric. 214, 108277. 2023.

MURITHI, H.M.; BEED, F.; TUKAMUHABWA, P.; THOMMA, B.P.H.J.; JOOSTEN, M.H.A.J. Soybean production in eastern and southern Africa and threat of yield loss due to soybean rust caused by Phakopsora pachyrhizi. Plant Pathol. 2016, 65, 176-188.

PACAL, I.; KUNDURACIOGLU, I.; ALMA, M.H.; DEVECI, M.; KADRY, S.; NEDOMA, J.; et al. 2024. A systematic review of deep learning techniques for plant diseases. Artif. Intell. Rev.57. doi: 10.1007/s10462-024-10944-7

REDMON, J.; DIVVALA, S.K.; GIRSHICK, R.B.; FARHADI, A. You Only Look Once: Unified, Real-Time Object Dection. CoRR. 2015.

SAPKOTA, R.; CHEPPALLY, R.H.; SHARDA, A.; KARKEE, M. YOLO26: Key Architectural Enhancements and Performance Benchmarking for Real-Time Object Detection. 2026.

SHINDE, S.; ATTAR, V. IDEAS (2024). MH-SoyaHealthVision: An Indian UAV and Leaf Image Dataset for Integrated Crop Health Assessment. Mendeley Data, V1.

SILVA, E.; DA GRAÇA, J.P.; PORTO, C.; et al. Unraveling Asian Soybean Rust metabolomics using mass spectrometry and Molecular Networking approach. Sci Rep 10, 138. 2020. https://doi.org/10.1038/s41598-019-56782-4

ZHOU, L.; XIAO, Q.; TAHA, M.F.; XU, C.; ZHANG, C. Phenotypic Analysis of Diseased Plant Leaves Using Supervised and Weakly Supervised Deep Learning. Plant Phenomics, Volume 5. 2023.