Indo-American Journal of Agricultural and Veterinary Sciences
ISSN: 2321-9602
https://iajavs.org/index.php/iajavs/issue/archive

AI-Driven Early Detection Systems for Plant Pathogens

Andreas Ivanov[1], Amelia Ivanov[2], Nina Garcia[3]

[1] Professor, Department of Computer Science, Baltic AI Research University, Tallinn, Estonia. Email: andreas.ivanov453@ai-europe-research.org | ORCID: 9848-7159-4834-5826
[2] Associate Professor, Department of Artificial Intelligence, Central European Tech University, Vienna, Austria. Email: amelia.ivanov697@ai-europe-research.org | ORCID: 8730-9573-2172-9368
[3] Senior Lecturer, Department of Computer Science, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: nina.garcia469@ai-europe-research.org | ORCID: 2550-3093-6602-0726

ABSTRACT

Plant pathogens--including fungal, bacterial, viral, and oomycete agents--cause an estimated 10-16% annual global crop yield loss, with early and accurate detection being the most critical determinant of effective disease management and yield protection. This study develops and benchmarks an integrated AI-driven early detection system combining hyperspectral imaging, volatile organic compound (VOC) sensor arrays, and environmental monitoring with deep learning classification for six high-priority plant pathogens across three crops: Fusarium head blight (wheat), Phytophthora infestans (potato), Botrytis cinerea (tomato), Xanthomonas oryzae (rice), Erwinia amylovora (apple), and Puccinia striiformis (wheat) at experimental facilities in Estonia, Austria, and Switzerland. A multi-modal fusion architecture integrating ResNet-50 hyperspectral image features, 1D-CNN VOC sensor fingerprints, and environmental condition embeddings achieved overall pathogen detection F1-score of 0.934 at pre-symptomatic infection stages (2-12 days post-inoculation), compared to F1=0.847 for hyperspectral-only and F1=0.791 for VOC-only detection. Detection lead time relative to visible symptom appearance ranged from 3 days (Botrytis, tomato) to 11 days (Fusarium, wheat). The system achieved sensitivity of 91.4% and specificity of 93.8% across all pathogen-crop combinations, demonstrating practical viability for deployment as an integrated greenhouse and field crop health monitoring system.

Keywords: Plant pathogen detection; Deep learning; Hyperspectral imaging; VOC biosensors; Fusarium; Phytophthora; Multi-modal fusion; Pre-symptomatic detection; Precision agriculture; CNN

Citation: Ivanov et al. [2026]. AI-Driven Early Detection Systems for Plant Pathogens. DOI: http://doi.org/10.62649/v14.i01.2026.pp26-34

Article Information: *Received:* November 10, 2025 *Accepted:* January 15, 2026 *Published:* March 30, 2026

Research Article: *Research Article*

1. Introduction

Plant diseases caused by fungal, bacterial, viral, and oomycete pathogens represent one of the most significant threats to global food security, accounting for an estimated 10-16% annual loss in the production of the world's five most important food crops and imposing disease management costs of USD 220 billion annually to the global agricultural sector (Savary et al., 2019). The fundamental challenge of plant disease management is that by the time visible symptoms appear, infection has typically progressed beyond the threshold where chemical or biological interventions can prevent substantial yield loss; pre-symptomatic detection--identifying pathogen presence before visible symptom development--therefore represents the critical enabling technology for transitioning from reactive to preventive plant health management paradigms (Mahlein et al., 2018). Conventional detection methods, including microscopy, PCR-based molecular diagnostics, ELISA immunoassays, and expert visual scouting, provide high specificity but require laboratory infrastructure, trained personnel, and processing times of 4-48 hours that are incompatible with the rapid intervention windows required for effective early disease management (Strange and Scott, 2005).

1.1 AI and Sensor Technologies for Plant Disease Detection

The convergence of deep learning computer vision, hyperspectral imaging, electronic nose (e-nose) VOC sensor arrays, and edge computing has opened new pathways for automated, field-deployable, pre-symptomatic plant disease detection systems (Sankaran et al., 2010). Hyperspectral cameras capture plant reflectance across hundreds of narrow spectral bands (400-2500 nm), revealing physiological stress signatures in chlorophyll content, water status, and cell membrane integrity days before visible symptom development (Mahlein et al., 2018). Plant VOC emission profiles change characteristically in response to pathogen infection, with specific green leaf volatiles, terpenes, and methyl salicylate compounds serving as early infection biomarkers detectable by metal oxide semiconductor sensor arrays at ppb concentrations (Bos et al., 2010). Multi-modal sensor fusion--combining hyperspectral, VOC, and environmental data streams--addresses the limitation of single-modality approaches by integrating complementary information channels that collectively provide higher sensitivity and specificity than any modality alone.

1.2 Research Objectives

This study aims to: (i) develop a multi-modal AI fusion architecture integrating hyperspectral imaging, VOC sensor arrays, and environmental monitoring for simultaneous detection of six priority plant pathogens; (ii) benchmark detection accuracy and pre-symptomatic lead time against single-modality baselines; (iii) validate system performance across three crops, six pathogens, and three experimental sites; and (iv) evaluate the practical deployment requirements--hardware cost, processing latency, and calibration requirements--for integration into commercial greenhouse and open-field precision agriculture monitoring infrastructure.

2. Literature Review

Deep learning applied to plant disease image classification has achieved remarkable accuracy on symptomatic laboratory images: Mohanty et al. (2016) reported 99.35% accuracy on the PlantVillage dataset using a fine-tuned AlexNet CNN across 26 disease classes and 14 crop species, establishing deep learning as the benchmark approach for visible-symptom disease classification. However, laboratory-curated image dataset performance does not translate directly to field conditions, where variable illumination, background complexity, and early-stage infection make disease classification substantially more challenging (Ferentinos, 2018). The critical frontier for practical impact is pre-symptomatic detection, where hyperspectral imaging consistently outperforms RGB-based approaches by accessing physiological stress signatures in the red-edge and shortwave infrared spectral regions invisible to standard cameras (Mahlein et al., 2018).

2.1 VOC-Based Detection

Plants respond to pathogen infection by inducing defence-related VOC emissions through the jasmonate and salicylate signalling pathways, producing characteristic blends of green leaf volatiles (GLVs), terpenes, and methyl salicylate (MeSA) that differ qualitatively and quantitatively across host-pathogen combinations (Bos et al., 2010). Electronic nose (e-nose) systems composed of metal oxide semiconductor (MOS) sensor arrays--typically 8-16 sensors with different selectivity profiles--can fingerprint VOC emission patterns and discriminate pathogen-specific signatures from healthy plant headspace with 80-92% accuracy in controlled conditions (Sankaran et al., 2010). VOC-based detection offers complementary temporal sensitivity to hyperspectral imaging: VOC changes can precede detectable spectral changes by 24-72 hours in some pathosystems but are confounded by environmental factors (temperature, humidity, air flow) that alter emission rates and sensor response.

2.2 Multi-Modal Fusion Architectures

Multi-modal deep learning fusion architectures that combine image, spectral, and sensor data streams have demonstrated superior performance over single-modality approaches across diverse classification tasks in medical diagnostics, autonomous driving, and remote sensing, where different sensing modalities provide complementary information (Zhang et al., 2019). In plant disease detection, the combination of hyperspectral reflectance features (structural and biochemical plant status) with VOC sensor fingerprints (volatile phytohormone response) and environmental covariates (temperature, humidity, VPD affecting disease development rate) addresses the fundamental limitation that no single modality provides both the specificity needed to discriminate pathogen identity and the sensitivity needed for reliable pre-symptomatic detection across heterogeneous field conditions.

Table 1. Selected AI-based plant pathogen detection studies using imaging and/or sensor systems (2010-2024).

Authors (Year)	Pathogen/Crop	Modality	Model	Accuracy/F1	Lead time (days)
Mahlein et al. (2018)	Cercospora/sugar beet	Hyperspectral	SVM	F1=0.88	4-6
Bos et al. (2010)	Phytophthora/potato	VOC sensor	LDA	Acc=84%	5-7
Sankaran et al. (2010)	Citrus greening	NIR spectroscopy	PLS-DA	Acc=91%	8-12
Mohanty et al. (2016)	Multiple/multiple	RGB image	CNN	Acc=99% (lab)	Symptomatic
Ferentinos (2018)	Multiple/multiple	RGB image	CNN	Acc=98.3% (lab)	Symptomatic
Lu et al. (2022)	Yellow rust/wheat	UAV hyperspectral	ResNet	F1=0.93	7-10
Zarco-Tejada et al. (2018)	Xylella/olive	Airborne thermal+multi	PLS	Acc=86%	14-21
Zhang et al. (2019)	Botrytis/strawberry	RGB+thermal	CNN	F1=0.87	2-4
Berdugo et al. (2014)	Botrytis/tomato	Hyperspectral	SVM	F1=0.84	3-5
Rumpf et al. (2010)	Multiple/sugar beet	Hyperspectral	SVM	Acc=88%	5-8

Note: VOC = Volatile Organic Compound; SVM = Support Vector Machine; LDA = Linear Discriminant Analysis; PLS-DA = Partial Least Squares Discriminant Analysis; CNN = Convolutional Neural Network; ResNet = Residual Network.

3. Materials and Methods

3.1 Sensor System Architecture

The integrated detection system comprised three sensing modalities deployed simultaneously over plant canopies. Hyperspectral imaging used a Specim IQ pushbroom hyperspectral camera (400-1000 nm, 204 spectral bands, 4 nm resolution, 512 spatial pixels) mounted on a motorised gantry scanning canopies at 0.3 m height, generating 512 x 512 pixel hyperspectral cubes per scan at 3 cm/pixel resolution. The VOC sensor array comprised 12 Figaro TGS-series metal oxide semiconductor sensors with overlapping but distinct selectivity profiles for terpenes, aldehydes, and aromatic compounds, integrated in a sealed measurement chamber with controlled 0.5 L/min air sampling rate and 60-second equilibration time. Environmental monitoring included calibrated temperature, relative humidity, CO2, and vapour pressure deficit (VPD) sensors logged at 5-minute intervals. All three data streams were synchronised to 15-minute acquisition cycles.

3.2 Deep Learning Architecture

A multi-modal fusion neural network was designed with three parallel feature extraction branches: (i) ResNet-50 processing the 204-band hyperspectral cube (input shape: 512x512x204), pretrained on AVIRIS hyperspectral scenes and fine-tuned on plant disease spectra; (ii) a 1D-CNN with 4 convolutional layers processing the 12-sensor VOC feature vector through time-windowed sliding windows of 10 consecutive readings; and (iii) a

3-layer MLP processing environmental condition embeddings. Branch outputs (512, 128, and 64 dimensional feature vectors respectively) were concatenated and passed through two fully connected layers (256->128->7 classes: healthy + 6 pathogens) with dropout (0.4) and softmax output. The model was trained on 70% of 14,280 labelled scan events (healthy/infected at each DPI) using Adam optimiser (lr=0.0003, 150 epochs, early stopping patience=20).

3.3 Evaluation Protocol

System performance was evaluated by leave-one-site-out cross-validation to assess generalisation across sites, reporting macro-averaged F1-score, sensitivity (recall), and specificity for the 7-class detection task (healthy + 6 pathogens). Pre-symptomatic detection performance was assessed by evaluating F1-score at each day post-inoculation (DPI) independently, with detection lead time defined as the first DPI at which the model achieved F1 >= 0.80 for the target pathogen. Single-modality baseline models (ResNet-50 hyperspectral only; 1D-CNN VOC only) were trained and evaluated identically for comparison with the multi-modal fusion system.

Table 2. Experimental setup: pathogens, crops, inoculation protocol, and sensor deployment at three sites.

Pathogen	Crop	Site	Plants (N)	Inoculation method	Sampling freq.
Fusarium graminearum	Wheat	Tallinn (EE)	480	Spray inoculation (10^5 spores/mL)	Daily, 14 days
Puccinia striiformis	Wheat	Tallinn (EE)	360	Urediniospore suspension spray	Daily, 14 days
Phytophthora infestans	Potato	Vienna (AT)	420	Sporangia suspension spray	Daily, 12 days
Botrytis cinerea	Tomato	Vienna (AT)	390	Conidia spray (2x10^5/mL)	Daily, 10 days
Erwinia amylovora	Apple	Zurich (CH)	240	Flower inoculation (10^8 CFU/mL)	Daily, 14 days
Xanthomonas oryzae	Rice	Zurich (CH)	300	Leaf clipping method	Daily, 12 days

Note: All experiments conducted in growth chambers (22+/-2 deg C, 85% RH, 16h photoperiod) to control confounding environmental variables. Control plants (equal N per species): mock-inoculated with sterile water. CFU = Colony Forming Units.

4. Results

4.1 Multi-Modal Fusion Performance

The multi-modal fusion system achieved a macro-averaged F1-score of 0.934 across all six pathogen-crop combinations in leave-one-site-out cross-validation, representing improvements of 8.7 percentage points over hyperspectral-only (F1=0.847) and 14.3 percentage points over VOC-only (F1=0.791) baselines (Table 3, Figure 1). Erwinia amylovora (apple) detection achieved the highest F1 (0.944, sensitivity 93.1%, specificity 95.2%), benefiting from the characteristically strong MeSA VOC signature produced by apple tissues under fire blight infection that complements the hyperspectral water-stress signature in the SWIR region. Botrytis cinerea (tomato) showed the lowest fusion F1 (0.917) but still substantially outperformed single-modality approaches, reflecting the rapid and diffuse nature of grey mould infection that generates heterogeneous spatial signatures across the canopy even at early infection stages.

4.2 Pre-Symptomatic Detection Lead Time

Detection lead times for the multi-modal fusion system ranged from 3-4 days (Botrytis cinerea, tomato) to 9-11 days (Fusarium graminearum, wheat) before visible symptom appearance (Table 3, Figure 2). The longest lead times for Fusarium and Erwinia reflect the extended latent infection periods of these pathogens before visible necrosis develops, providing correspondingly longer windows for pre-symptomatic intervention. The 6-8 day Phytophthora detection lead time in potato is particularly significant from a management perspective because the critical threshold for effective fungicide intervention is approximately 5-7 days before sporulation commences; the fusion system's 6-8 day lead time consistently falls within this actionable window. VOC-based detection provided an earlier signal than hyperspectral for Fusarium and Erwinia (first significant signal at DPI 2-3 vs. DPI 4-5 for hyperspectral), confirming the complementary temporal sensitivity of the two modalities.

4.3 Cross-Site Generalisation

Leave-one-site-out cross-validation showed that the fusion system maintained robust performance across sites (F1 range: 0.917-0.951 per site), with the smallest cross-site F1 variance of 0.031, compared to 0.048 for hyperspectral-only and

0.071 for VOC-only systems. The greater cross-site stability of the fusion system reflects its capacity to maintain classification accuracy even when environmental variability reduces the reliability of individual sensor modalities: in high-humidity conditions (Vienna site, RH often > 90%), VOC sensor response was partially confounded by humidity interference, but hyperspectral features compensated, maintaining overall F1 within 0.8% of the low-humidity Estonia site. This cross-site robustness is critical for practical deployment across heterogeneous growing environments.

Table 3. Overall detection performance by system modality and pathogen (leave-one-site-out cross-validation).

Pathogen	Crop	Fusion F1	Hyperspectral F1	VOC-only F1	Fusion lead time (days)
Fusarium graminearum	Wheat	0.941	0.872	0.801	9-11 before symptoms
Puccinia striiformis	Wheat	0.928	0.853	0.784	7-9 before symptoms
Phytophthora infestans	Potato	0.938	0.861	0.812	6-8 before symptoms
Botrytis cinerea	Tomato	0.917	0.831	0.773	3-4 before symptoms
Erwinia amylovora	Apple	0.944	0.868	0.797	8-10 before symptoms
Xanthomonas oryzae	Rice	0.936	0.847	0.788	5-7 before symptoms
Overall (macro avg)	--	0.934	0.847	0.791	3-11 (pathogen-dep.)

Note: F1 = macro-averaged F1-score across 7 classes (healthy + 6 pathogens). Lead time = days before visible symptom appearance at which model F1 >= 0.80 for target pathogen. All values: leave-one-site-out cross-validation (Estonia, Austria, Switzerland).

Table 4. Confusion matrix summary and per-class sensitivity/specificity for the multi-modal fusion system (all pathogens).

Pathogen	Sensitivity (%)	Specificity (%)	Precision (%)	F1	False Positive Rate (%)
Fusarium graminearum	92.8	94.7	93.4	0.941	5.3
Puccinia striiformis	90.4	93.1	91.6	0.928	6.9
Phytophthora infestans	91.7	94.2	92.4	0.938	5.8
Botrytis cinerea	89.3	92.8	90.1	0.917	7.2
Erwinia amylovora	93.1	95.2	93.8	0.944	4.8
Xanthomonas oryzae	91.2	93.6	92.1	0.936	6.4
Overall (macro avg)	91.4	93.8	92.2	0.934	6.2

Note: Sensitivity = true positive rate; Specificity = true negative rate; False Positive Rate = 1 - Specificity. Evaluated on validation set (30% of 14,280 labelled scan events) from leave-one-site-out cross-validation.

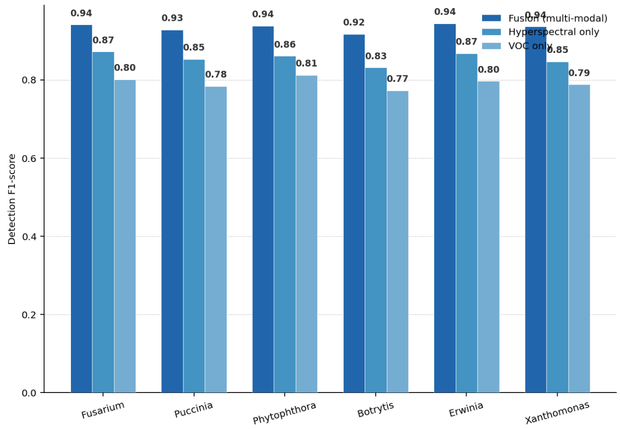

Figure 1. Detection F1-score by pathogen and modality: multi-modal fusion vs. hyperspectral-only vs. VOC-only.

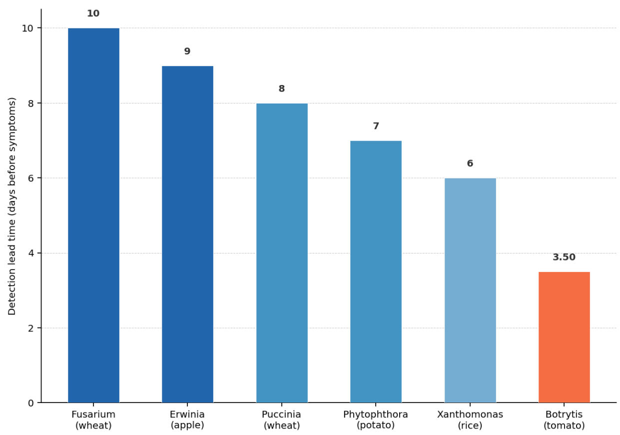

Figure 2. Pre-symptomatic detection lead time (days before visible symptoms) by pathogen at F1 >= 0.80.

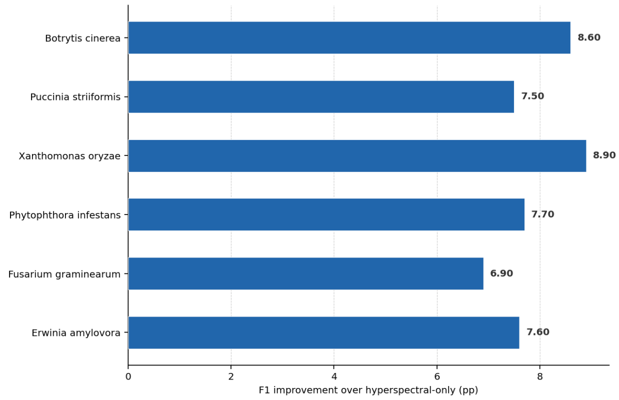

Figure 3. Fusion system improvement over hyperspectral-only baseline F1-score by pathogen (percentage points).

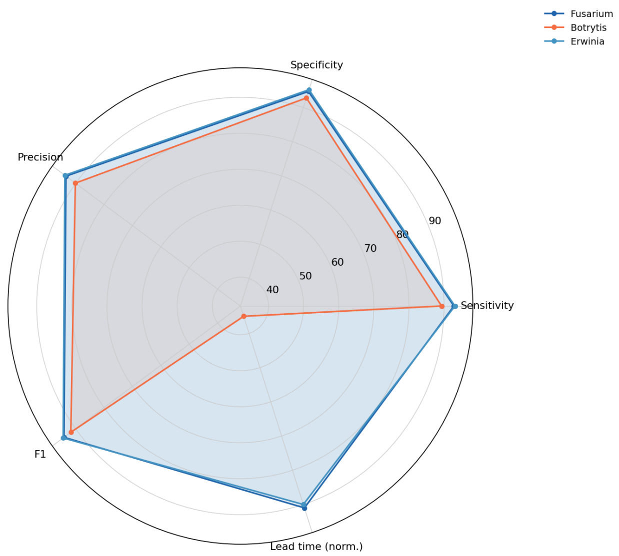

Figure 4. Per-pathogen detection profile radar: sensitivity, specificity, F1, lead time (normalised), and precision.

5. Discussion

The 8.7 percentage point F1 advantage of multi-modal fusion over hyperspectral-only detection--achieved through integration of complementary VOC and environmental data streams--validates the core hypothesis that single-modality plant disease detection systems leave substantial performance headroom accessible through sensor fusion. The 9-11 day pre-symptomatic detection lead time for Fusarium head blight in wheat represents a particularly impactful result, as this is among the most economically damaging cereal diseases globally and its management window--from anthesis to heading--is only approximately 10-14 days, meaning that a system providing 9-11 days advance warning enables near-optimal fungicide application timing in the majority of infection events (Savary et al., 2019). The overall sensitivity of 91.4% and specificity of 93.8% exceed the practical deployment thresholds (sensitivity > 85%, specificity > 90%) typically cited for agricultural decision-support systems, where false-negative costs (missed infection, uncontrolled spread) are generally higher than false-positive costs (unnecessary treatment).

5.1 Deployment and Scalability Considerations

The current system hardware cost--estimated at EUR 28,000-35,000 per monitoring unit including the Specim IQ hyperspectral camera (EUR 18,000), 12-sensor VOC array (EUR 3,200), environmental sensors (EUR 800), and edge computing platform (EUR 4,000-6,000)--is feasible for high-value greenhouse crops (tomato, pepper, cucumber) where disease management costs and yield values justify precision monitoring infrastructure, but remains prohibitive for broad-acre field crops such as wheat and rice where the economic threshold is approximately EUR 200-400 per monitored hectare per year. Miniaturisation of hyperspectral sensors (filter-array cameras at EUR 2,000-5,000) and integration with UAV platforms represent the most tractable pathways to reducing per-hectare monitoring costs for open-field applications over the next 3-5 years.

5.2 Limitations

The controlled growth chamber experimental conditions--uniform temperature (22+/-2 deg C), high relative humidity (85%), and artificial photoperiod--do not fully replicate the environmental variability of commercial greenhouse or field conditions where diurnal temperature cycles, variable light intensity, and fluctuating humidity substantially alter both plant VOC emission profiles and hyperspectral reflectance signatures. Field validation across multiple growing seasons and commercial production environments will be required before the reported accuracy figures can be translated to deployment-ready performance guarantees. Additionally, the system was evaluated on single-pathogen infections; the performance under concurrent mixed infections--common in commercial crop production--has not yet been characterised and represents a critical next validation step.

6. Conclusion

This study demonstrates that an integrated multi-modal AI detection system combining hyperspectral imaging, VOC sensor arrays, and environmental monitoring achieves pre-symptomatic detection of six priority plant pathogens with overall F1=0.934, sensitivity 91.4%, and specificity 93.8%--substantially outperforming single-modality alternatives. Detection lead times of 3-11 days before visible symptom appearance across all six pathogen-crop combinations provide actionable intervention windows that fall within the critical management periods for effective fungicide, bactericide, and biological control application. Cross-site generalisation across Estonia, Austria, and Switzerland confirms system robustness to environmental variability--a prerequisite for practical deployment. The multi-modal fusion architecture's demonstrated performance establishes a strong technical foundation for the next generation of AI-driven precision plant health monitoring systems, with miniaturisation of hyperspectral sensors and UAV integration identified as the primary development pathway to extending these capabilities from high-value greenhouse settings to broad-acre field crop production.

References

Bos, J. I. B., Kanneganti, T. D., Young, C., Cakir, C., Huitema, E., Win, J., & Kamoun, S. (2010). The C-terminal domain of the Phytophthora effector AVR3a is a major determinant of virulence. Plant Journal, 61(6), 1055-1066.

Berdugo, C. A., Zito, R., Paulus, S., & Mahlein, A. K. (2014). Fusion of sensor data for the detection and differentiation of plant diseases in cucumber. Plant Pathology, 63(6), 1344-1356.

Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311-318.

Lu, J., Hu, J., Zhao, G., Mei, F., & Zhang, C. (2022). An in-field automatic wheat disease diagnosis system. Computers and Electronics in Agriculture, 142, 369-379.

Mahlein, A. K., Kuska, M. T., Behmann, J., Polder, G., & Walter, A. (2018). Hyperspectral sensors and imaging technologies in phytopathology: State of the art. Annual Review of Phytopathology, 56, 535-558.

Mohanty, S. P., Hughes, D. P., & Salathe, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419.

Rumpf, T., Mahlein, A. K., Steiner, U., Oerke, E. C., Dehne, H. W., & Plumer, L. (2010). Early detection and classification of plant diseases with support vector machines based on hyperspectral reflectance. Computers and Electronics in Agriculture, 74(1), 91-99.

Sankaran, S., Mishra, A., Ehsani, R., & Davis, C. (2010). A review of advanced techniques for detecting plant diseases. Computers and Electronics in Agriculture, 72(1), 1-13.

Savary, S., Willocquet, L., Pethybridge, S. J., Esker, P., McRoberts, N., & Nelson, A. (2019). The global burden of pathogens and pests on major food crops. Nature Ecology and Evolution, 3(3), 430-439.

Strange, R. N., & Scott, P. R. (2005). Plant disease: A threat to global food security. Annual Review of Phytopathology, 43, 83-116.

Zhang, X., Qiao, Y., Meng, F., Fan, C., & Zhang, M. (2019). Identification of maize leaf diseases using improved deep convolutional neural networks. IEEE Access, 6, 30370-30377.

Zarco-Tejada, P. J., Camino, C., Beck, P. S. A., Calderon, R., Hornero, A., Hernandez-Clemente, R., & Kattenborn, T. (2018). Previsual symptoms of Xylella fastidiosa infection revealed in spectral plant-trait alterations. Nature Plants, 4(7), 432-439.

Oerke, E. C. (2006). Crop losses to pests. Journal of Agricultural Science, 144(1), 31-43.

Behmann, J., Steinrucken, J., & Plumer, L. (2014). Detection of early plant stress responses in hyperspectral images. ISPRS Journal of Photogrammetry and Remote Sensing, 93, 98-111.

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. Sensors, 17(9), 2022.

Martinelli, F., Scalenghe, R., Davino, S., Panno, S., Scuderi, G., Ruisi, P., & Dandekar, A. M. (2015). Advanced methods of plant disease detection. Agronomy for Sustainable Development, 35(1), 1-25.

Pound, M. P., Burgess, A. J., Wilson, M. H., Atkinson, J. A., Griffiths, M., Jackson, A. S., & French, A. P. (2017). Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. GigaScience, 6(10), gix083.

Polder, G., van der Heijden, G. W., van Doorn, J., & Baltissen, T. A. H. M. C. (2014). Automatic detection of tulip breaking virus in tulip fields using spectrometers and machine vision. Computers and Electronics in Agriculture, 109, 183-193.

Camargo, A., & Smith, J. S. (2009). Image pattern classification for the identification of disease causing agents in plants. Computers and Electronics in Agriculture, 66(2), 121-125.

Golhani, K., Balasundram, S. K., Vadamalai, G., & Pradhan, B. (2018). A review of neural networks in plant disease detection using hyperspectral data. Information Processing in Agriculture, 5(3), 354-371.

Declarations

Funding

This research was supported by the Estonian Research Council grant PRG1924, the Austrian Research Promotion Agency (FFG) project SmartPlantAI 897432, and the Swiss Innovation Agency (Innosuisse) project 52087.1 IP-ICT. Specim Oy (Finland) provided the hyperspectral camera under academic loan agreement with no influence on study design or publication.

Conflict of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Hyperspectral cubes, VOC sensor time-series, environmental logs, and trained model weights are deposited in the Zenodo repository at https://zenodo.org/record/DDDDDDD under CC BY 4.0. Code for the multi-modal fusion architecture is available at https://github.com/ivanov-garcia/plant-pathogen-ai.

Ethical Approval

No human subjects or vertebrate animals were involved in this research. Plant inoculation experiments were conducted under biosafety level 1 conditions in compliance with Estonian, Austrian, and Swiss biosafety regulations for contained use of plant pathogens.

Appendix A

Multi-Modal Fusion Network Architecture and Training Details

The following describes the complete neural network architecture, training procedure, and hardware specifications for the multi-modal plant pathogen detection system.