



Advanced Hybrid Convolutional Neural Network for Leaf-Based Plant Disease Detection

Sarmad Nozad Mahmood¹, Swash Sami Mohammed²,
Yahya Ghufuran Khidhir^{3*}, Tale Saeidi⁴, and Sahar Saleh⁵

¹ Medical Instrumentation Techniques Engineering Department, Technical Engineering College – Kirkuk, Northern Technical University, Mosul 41001, Iraq

² Artificial Intelligence Techniques Engineering Department, Technical Engineering College for Computer and AI – Kirkuk, Northern Technical University, Mosul 41001, Iraq

³ Electronic and Control Engineering Techniques Department, Technical Engineering College – Kirkuk, Northern Technical University, Mosul 41001, Iraq

^{4,5} WiSAR Lab, Atlantic Technological University (ATU), Letterkenny, Co. Donegal, F92 YY97, Ireland

*Corresponding Author's E-mail: yahhya.khidhir24@ntu.edu.iq

(Received 6 October 2025 ; Revised 20 January 2026; Accepted 16 February 2026; Published 1 March 2026)

<https://doi.org/10.22153/kej.2026.02.002>

Abstract

Accurate detection and classification of plant diseases are central to sustainable food production and the reduction in crop losses. Conventional identification protocols are often dependent on experts. They are time-consuming and difficult to scale. This paper proposes a novel Advanced Hybrid Convolutional Neural Network (AHCNN) model combining the attention mechanism and simplified convolution mechanism. This model can obtain high accuracy and low computational complexity. It is trained using a high-resolution leaf dataset that includes 1532 leaf images, separated into 1322 training, 150 validation and 60 test samples, with the goal of a classification of specimens into three categories: healthy, powdery and rusty. Its architecture uses squeeze-and-excitation (SE) blocks and a spatial attention mechanism, which, together, can provide enhanced feature extraction and improve the interpretability of models. Results showed that the model achieved a validation accuracy of 98% and a test accuracy of 98.33%. With a parameter number of only 0.4 million, the proposed architecture provides a lightweight solution that performs much better than conventional deep learning frameworks in terms of computational efficiency. These attributes make AHCNN an interesting candidate for real time, drone-based detection of plant diseases as part of precision agriculture systems.

Keywords: Plant disease detection; Machine learning; Convolutional neural networks; Squeeze-and-excitation (SE) blocks

1. Introduction

Plant diseases represent a major problem for modern agriculture, directly affecting sustainable crop productivity, food security and economic stability [1]. Prompt identification and accurate classification are imperative to limit yield losses below normal average and protect the economic viability of agricultural production [2], [3]. On-site disease inspection by agricultural experts is

exhausting and time-consuming, making it impractical for large-scale agronomic surveillance [4]. Limitations that come with traditional methodologies require technological innovation to explore possibilities for proactive automation of illness management, including artificial intelligence and unmanned aerial vehicles [5], [6]. The acquisition of real-time field data by unmanned aerial vehicles equipped with high-resolution imaging sensors or artificial intelligence models

This is an open access article under the [CC BY](https://creativecommons.org/licenses/by/4.0/) license:



promises improvements in ease and precision [7], [8]. Consequently, early detection of infected or diseased foliage during its early stages reduces the need for manual labour and favours resource utilisation for efficient detection, thereby reducing crop losses [9].

In this paper, a plant disease classification database with 1532 high resolution images was used. The classification scheme defines three different categories that represent different states of plant leaves: healthy, powdery and rusty [10]. The dataset was divided into training, validation and testing subsets, with 1322, 150 and 60 images, respectively, to approximate a distribution that minimises the difference between training and real-world deployment across classes. For the sake of computational efficiency, all images were uniformly rescaled to a resolution of 512×512 pixels in the training phase [11–13]. The source images were non-uniform in dimensions, with widths ranging from 2421 pixels to 5184 pixels and heights between 1728 and 3456 pixels. Rescaling to a common spatial resolution was performed for computational complexity by ensuring that the input size is always consistent to the proposed deep learning architecture.

The implementation of plant disease detection systems under natural field conditions (real-time) remains a formidable challenge despite the implementation of deep learning to automated recognition. For example, Barbedo [14] used CNNs and transfer learning for disease identification; however, the methodology was not well adapted to the variability of environmental conditions nor computationally efficient. Ferentinos [15] proposed a number of CNN architectures with considerable performance improvements compared with previous models when tested on extensive data, but the study was primarily limited to well-controlled laboratory experiments. Hari et al. [16] evaluated a handcrafted CNN in natural field scenarios; however, limitations in scalability and computational load hindered its deployment in large scale agricultural applications. Militant et al. [17] proposed a multi-crop computational strategy, but no attention was given to real-time performance at scale. Saleem et al. [18] reviewed conventional methods of machine learning and deep learning but reported shortcomings in terms of real-time and visualisation mechanism. Ullah et al. [19] studied CNNs under controlled environments without focusing on the robustness of CNNs in dynamically varying agricultural environments. In [20], machine learning and deep learning approaches have been discussed in the context of precision agriculture, where the potential of deep learning was evident.

However, computational constraints regarding the real-time context were present. Recent investigations have dealt with meta-heuristic optimisation techniques, such as Particle Swarm Optimization (PSO) and the Pelican Optimization Algorithm (POA), to improve the convergence properties and optimisation efficiency in intelligent systems [21–23]. Yag and Altan [24] proposed a hybrid model of the UAV-based disease detection system incorporating CNNs with meta-heuristics-based optimisation to alleviate the computational burden and improve generalisation. Taken together, these studies highlight the need for field-deployable systems that blend accuracy, scalability and computational cost. Accordingly, the present work provides a scalable and accurate solution for real-time plant disease detection whilst addressing the limitations in terms of scalability, environmental tolerance and practical deployment.

Deep learning methods that enhance the attention mechanisms and optimise the convolution operation were employed. The proposed architecture consists of five crucial computational branches, which include squeeze-and-excitation (SE) blocks for adaptive feature recalibration and a spatial attention module for improving region-specific feature extraction [25], [26]. Depthwise-separable convolution layer and dilated convolution layer were used to achieve computational efficiency and increase the receptive field for high-resolution image analysis [27]. Adaptive skip connections were introduced for gradient flow, hence enabling stable learning and scalability [28]. Global average pooling (GAP) was performed on the extracted features that were then processed by a fully connected layer to give the final classification results (healthy, powdery or rusty). These design factors improved the processing efficiency, classification accuracy and practical viability in agricultural applications. The resultant AI-driven achieved a prediction accuracy of 98.67% with a low computational cost.

The proposed model can be applied to real time deployment due to its light architecture and high diagnostic precision. It is suitable for autonomous agricultural monitoring platforms or unmanned aerial vehicle systems, and it can be seamlessly inserted. The training regime utilises a heterogeneous data corpus to encourage strong generalisation to various plant pathology scenarios and environmental conditions, helping to alleviate overfitting and improve transference. This research helps overcome some of the most important barriers to intelligent plant disease detection and the advancement of precision agronomy paradigms.

The manuscript is structured as follows.

Section 1 overviews modern advances in plant disease diagnosis using deep learning algorithms and identifies outstanding technical challenges. Section 2 describes the methodological framework, including dataset curation and architectural design of the predictive model. Section 3 provides the empirical results comparing the proposed approach with relevant state-of-the-art alternatives. Section 4 presents the closing remarks and some possible future directions for further study.

2. Method

The dataset used for this research was explicitly curated for plant disease classification research, and it contains 1532 images [10]. It covers the following three different taxa, all of which represent different pathological states of plant foliage: healthy, powdery and rusty. The dataset was divided into three stratified folds comprising the training set with 1,322 images, the validation set with 150 images and the test set with 60 images to enable balanced learning and reduce class imbalance. All the images were resized to 512×512 pixels to improve the efficiency of computations and homogeneity of the model inputs. The original collection had images of heterogeneous sizes, with the width ranging from 2421 pixels to 5184 pixels and the height ranging from 1728 pixels to 3456 pixels. This rescaling procedure reduces memory consumption and ensures the applicability of feature extraction at multiple spatial resolutions. Before the mean image was pre-processed, the dimensions were about 3986×2702 pixels. Labels were retained across splits to retain the original class distribution. The training cohort contained 458 images of healthy leaves, 430 images of powdery mildew and 434 images of rust. The validation and test cohorts contained 50 and 20 images for each class, respectively, thus maintaining balance that is important for unbiased performance evaluation. Fig. 1. shows

representative samples of the healthy, powdery and rusty leaves. The high accuracy and calculation efficiency of this proposed model make it applicable to real-time automatic agricultural monitoring systems in which rapid and efficient diseased plant classification greatly enhances large-scale field diagnoses.

The resolution of (512×512) pixels adopted is a reasonable trade-off between the fine-grained feature detection that is required and computational practicality. The symptoms of plant diseases, especially at early stages, consist of small, high-frequency texture changes that could be missed at usual resolutions (224×224 pixels). Higher resolutions always come with higher computational costs, so the combination of depthwise-separable convolutions and a small lightweight architecture with only 0.4 million parameters mitigates the cost, making the chosen resolution viable in the context of drone-based inference, where visual image quality is valued the most.

A dataset that is highly visual in its diversity and is not limited to disease categories was chosen to ensure that the model is effective in terms of generalising to realistic agronomic conditions. The photographs covered a range of lighting situations, including bright sun, cloudy natural light, and shadows, thereby encouraging consistency to illumination changes in the model. Background heterogeneity was added to the dataset with the leaves set against a uniform surface and complex field backgrounds with soil, branches and overlaying foliage. Lastly, the samples included a spectrum of disease progression, from initial symptoms of the disease to the full extent of tissue necrosis, thus ensuring that the system can diagnose diseases at a spectrum of severity levels.

The next sections explain the architectural enhancements and training approaches, both of which are crucial in integrating the efficacy of this model.



Fig. 1. Dataset samples of plant leaf conditions [10]

More advanced convolutional techniques were integrated for enhanced performance and time efficiency to propose a suitable deep learning model for disease detection and classification in plants. As seen in Fig. 2, the architecture consists of a series of convolutional blocks, one using depthwise-separable convolutions and the other using attention mechanisms, collectively enhancing feature extraction and representation. SE blocks allow for channel-wise recalibration of output channels, whereas spatial attention mechanisms emphasise

corresponding regions in the input images. Besides, dilated convolutions enlarge the receptive field without increasing the complexity of the computations. A GAP layer, a fully connected layer, and a softmax activation function were used to complete the model and distinguish plant leaf conditions as healthy, powdery, or rusty. The structure was therefore designed to optimise the model’s performance for precision agriculture applications.

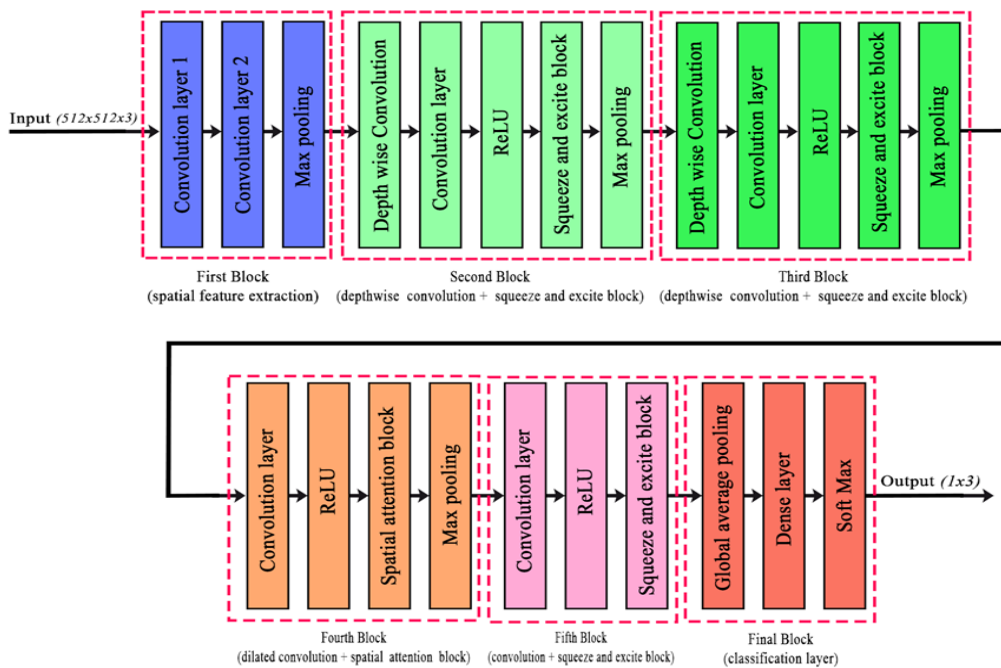


Fig. 2. Proposed model architecture

The first block extracts spatial features with a couple of 3×3 convolutional layers followed by a max-pooling layer for dimensionality reduction. In this manner, an adequate capture of low-level features from the input image is ensured. The second and third blocks use DepthwiseConv2D operations and pointwise convolution layers, along

with the integration of SE blocks (Fig. 3). These blocks recalibrated the channel-wise feature responses, heightening the model’s attention to the most relevant features. These blocks outputted down-sampled feature maps by utilising max-pooling, gradually reducing the spatial dimension whilst retaining substantial feature representations.




Fig. 3. Concept flow

The novelty of the advanced hybrid CNN (AHCNN) in comparison with traditional CNNs

(e.g. VGG16 or ResNet) lies in the fact that SE modules were added in the depthwise-separable

architecture with lightweight parameters. This structure allowed the model to assign dynamically to different feature channels, thereby giving a high priority to disease specific textures and putting downplay on the background noise. The architecture introduces a deliberate bias tailored to the agricultural domain, where background clutter is common but not typically present in conventional pre-trained benchmarks.

The depthwise separable convolutions interacted with the attention mechanisms in a sequential manner to optimise feature refinement. Firstly, depthwise separable convolutions present computationally efficient features. The resultant feature maps are then directly inputted into the SE blocks. This sequential order allows the attention mechanism to re-adjust spatially filtered features, thus enabling the network to avoid irrelevant channels generated during convolution before being included in the pooling layers.

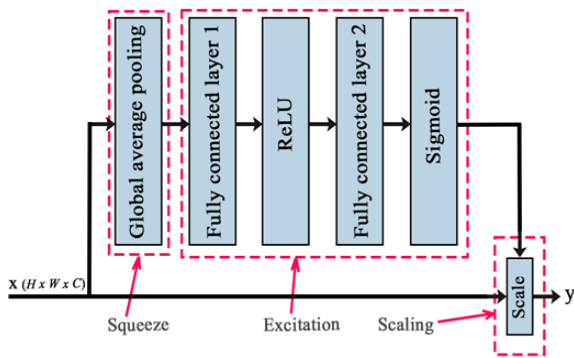


Fig. 3. Squeeze-and-excitation blocks.

Table 1,
Squeeze-and-excitation blocks

Layer No.	Layer Type	Parameters	Description
1	GlobalAveragePooling2D	-	Aggregates global spatial information from feature maps.
2	Dense	Filter size of the input / 16	Reduces channel dimensions by using a ReLU activation.
3	Dense	Filter size of the input	Restores channel dimensions by using a sigmoid activation.
4	Multiply	-	Multiplies input tensor with the channel-wise recalibrated weights.

The fourth block uses dilated convolution layers to expand the receptive field and increase contextual information by capturing broader spatial regions. The spatial attention block is applied afterwards to extract informative areas in the spatial domain (Fig. 4). With the inclusion of channel-wise and spatial attention mechanisms, it keeps a good balance between global and local feature representation in the model. This double-attention mechanism makes

the model more tolerant to variations in disease patterns.

Introduced by Woo et al. [29], the SE block is a contemporary architectural component that allows the recalibration of channel-wise feature maps by modelling interdependencies between channels. The SE block proceeds in three major stages: squeeze, excitation and scaling. Squeeze refers to performing GAP and collapsing all spatial information into a single value that acts as the global descriptor for each channel. Excitation considers two fully connected layers. The first layer has a ReLU activation to introduce non-linear dependencies, and the second layer contains a Sigmoid activation that generates a channel-wise weight. In scaling, these weights are applied to multiply the original feature maps to strengthen or diminish certain channels in accordance with their channel-wise importance. Table 1 specifies what each layer does and its parameters for the introduced SE block. In this manner, the sparsity of the model encourages emphasis on informative features through instance selection (e.g. by recovering fine details such as important patterns for plant diseases). The SE block also enhances the interpretability and performance of the model once incorporated into the architecture, making it one of the most important components of the model.

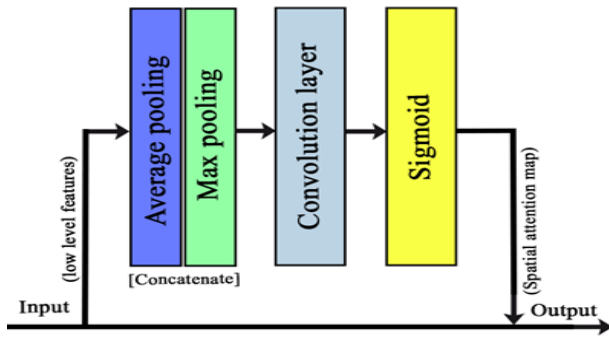


Fig. 4. Spatial attention block.

The spatial attention block introduced by Woo et al. [29] encourages the model to focus on key spatial regions of feature maps. The formulation of the attention mechanism is as follows: it firstly conducts average pooling and max-pooling along the channel axis of the input tensor to acquire complementary spatial information. These pooled results are then merged with one another, making use of both. Next, the concatenated tensor is processed by a convolution layer with a Sigmoid activation function that obtains a spatial attention map, locating important regions in the feature map. The attention map is element-wise multiplied with the input feature map to highlight spatially informative features and dampen counter-productive ones. The architecture of the spatial attention block is described in Table 2, where each layer and its type, parameters and function are indicated. The spatial attention block verifies the checking of the model so that the network can attend well to important spatial parts and finally promotes the model's classification performance.

Table 2,
Spatial attention block

Layer No.	Layer Type	Prog
1	Avg Pool	-
2	Max Pool	-
3	Concatenate	-
4	Conv2D	$1 \times (7, 7)$
5	Multiply	-

The fifth block contains a convolutional layer with 256 filters for enhanced feature extraction, enabling deeper and more discriminative feature representation. An additional SE block is applied. This architecture injects channel-wise information into the model by recalibrating it at this deeper level of processing. GAP is applied to pool the feature

maps in all dimensions into one vector, thus representing the extracted information concisely. The above vector cascades through five deep layers, resulting in a softmax activation that provides class probabilities amongst the three plant disease categories. Table 3 presents an end-to-end holistic model of the proposed approach and details the functional capability of each block in the network and its contribution to deep representation at different layers. Filled with traditional convolutional layers, depthwise separable convolution, dilated convolution and advanced attention mechanisms, this architecture provides an optimal compromise between computational efficiency and classification performance. It is particularly efficient for plant disease classification.

Table 3,
Main model

Block	Layers (Filter Size/Parameters)
Input	Input Layer
First Block	Conv $[32 \times (3, 3)]$, Conv $[64 \times (3, 3)]$, MaxPooling (2, 2)
Second Block	Depthwise Conv (3, 3), Conv $[32 \times (1, 1)]$, SE blocks, MaxPooling (2, 2)
Third Block	Depthwise Conv (3, 3), Conv $[64 \times (1, 1)]$, SE Block, MaxPooling (2, 2)
Fourth Block	Dilated Conv $[128 \times (3, 3, \text{Dilation Rate: } 2)]$, Spatial Attention Block, MaxPooling (2, 2)
Fifth Block	Conv $[256 \times (3, 3)]$, SE Block
Output	Global Average Pooling, Dense (3)

3. Results and Discussion

A dataset of plant health to identify three disease classes was employed to train and test the model. This dataset contained 1322 images for training, an additional validation set of 150 images, and a testing set of 20 images. The dataset was split by cutoff to ensure a similar distribution of the three classes across the training, validation, and testing sets, enabling reliable evaluation. Image preprocessing was performed to standardise the size of all images to 512×512 and normalise them to converge evenly in the model.

Moreover, the dataset included various samples to capture differences amongst plant appearances and disease states, making it suitable for training a robust model capable of generalisation. The training was performed using Google Colaboratory Pro, which provided GPU acceleration and efficient computation. The machine had 53 GB of RAM and 22.5 GB of GPU memory for efficient processing of high-dimensional images and complex

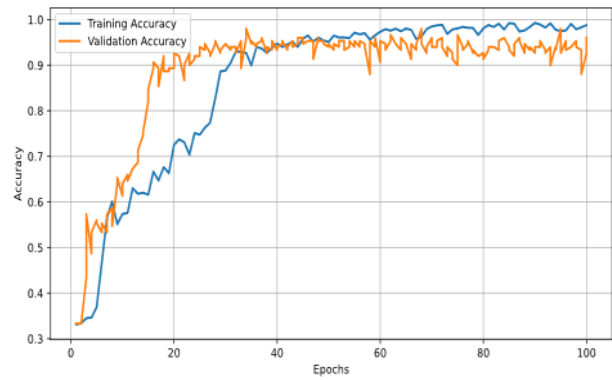
computations in the model. The model was trained for 200 epochs with a batch size of 32 and the Adam optimiser (learning rate of 0.0005).

The hyperparameters, including the learning rate of 0.0005, were selected empirically. A grid search was conducted over a limited range of learning rates (0.001, 0.0005 and 0.0001), and 0.0005 was found to provide the most stable convergence without oscillations around local minima, as observed in the training loss curves.


For the loss function, sparse categorical cross-entropy, which is well suited for multi-class classification, was used. Accuracy, precision, recall and F1-score were used as metrics for thoroughly assessing the model. The results of the training and evaluation are captured in the plots below. The figures exhibit the model performance in terms of evaluation metrics throughout 200 epochs of training. The results showed reliable and robust classification of plant health statuses, as demonstrated by the continuous excellent upward trends during the training process and further good generalisation performance on the validation dataset, making it applicable for field agriculture.

The model was evaluated using accuracy and loss metrics during the training process. Fig. 5(A) illustrates the training and validation accuracies during the training process. In terms of accuracy, the model showed consistent improvement, with a peak training accuracy of around 99.5% and validation accuracy around 98.00%. Thus, the model is an effective learner for the training data whilst maintaining generalisation for validation on unseen data. The nearly parallel nature of accuracy and loss indicated neither overfitting nor underfitting of the model. Hence, this work provides a high degree of reliability in efficiency to classify plant diseases as healthy, powdery or rusty.

Fig. 5(B) shows considerably reduced training and validation losses. The training loss reached about 0.05, whereas the validation loss was stable at around 0.21 in the last epoch. This steady decrease in loss indicated a working model with minimal prediction error during the learning process on the dataset. The trend coherence in the accuracy and loss curves corroborated further evidence for network robustness and generalisation to previously unseen data, providing evidence of efficiency in accurately classifying plant diseases.



(A)



(B)

Fig. 5. Model training and validation: (A) accuracy and (B) loss.

The training and validation metrics were interpreted to assess the performance of the model in terms of recall and precision, providing insights into its classification capabilities. The training and validation recall are shown over 200 epochs in Fig. 6(A). The model has great recall during training (recall approached 99% during training and stopped at about 98% in the last epoch). This result indicated that the model has a good ability to select most of the actual positives in the three classes of plant disease. A high recall means few false negatives, which is a crucial characteristic for disease diagnosis in plant disease detection to reduce the number of diseased samples left undiagnosed. Meanwhile, Fig. 6(B) shows the precision of the model for training and validation during the same epochs. The training precision was over 0.99 throughout, whereas the validation precision plateaued at about 0.98 in the last epochs. This finding indicated that the model is good at maintaining false positives at low levels, demonstrating high accuracy in predicting diseased classes. High precision helps lessen the number of unnecessary actions taken as a result of erroneously estimating disease presence and reinforces the

model in performing fine plant disease classification tasks with an appropriate choice.

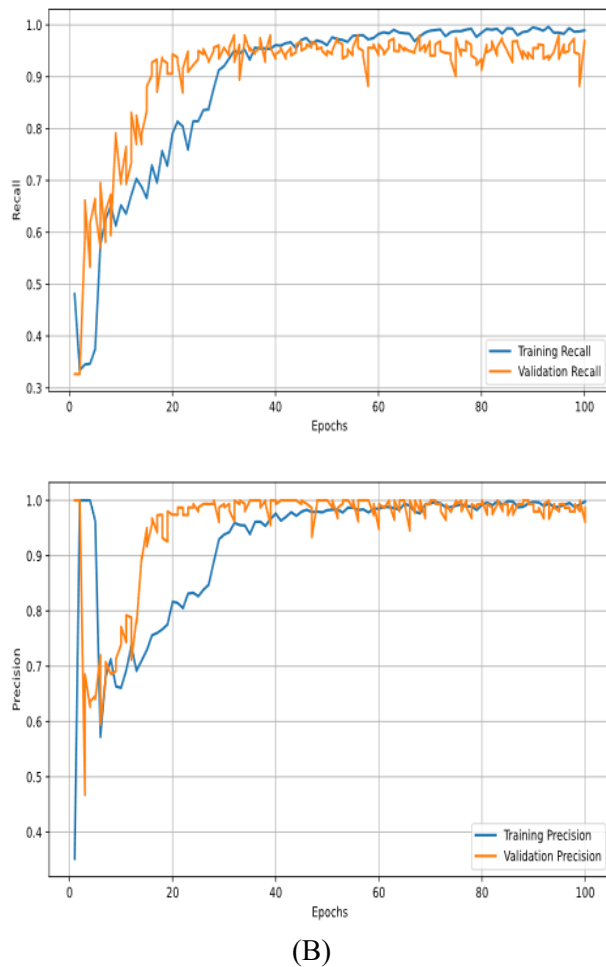


Fig. 6. Model training and validation: (A) recall and (B) precision.

In Fig. 7, the final evaluation metric, the F1-score of the model, is shown over 200 epochs of training and validation steps. The F1-score of the training phase remained above 0.99, and that for the validation phase approached 0.98 in the later epochs. High precision and recall of the model implied low false negative cases. Hence, the model provides good accuracy in identifying positive cases. This strong F1-score implied the robustness and efficiency of the model in classifying plant diseases across all defined classes, thereby validating the model as reasonable for real-world use in agriculture, where accurate and efficient disease detection is needed.

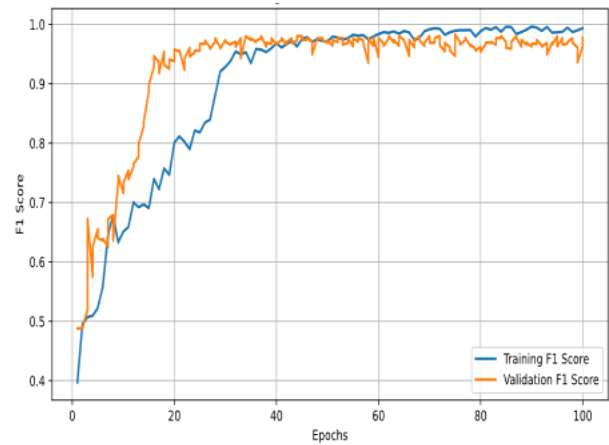


Fig. 7. Training and validation F1-score.

The results of the deep learning model are summarised in Table 4 of the validation and test datasets, which showed good generalisation. The model illustrated accuracies of 98% on the validation set and 98.33% on the test set, thus confirming its ability to be robust in identifying out-group cases. Moreover, the results of recall, accuracy and F1-score were non-negative in both datasets, indicating that the types of plant-leaf diseases can be correctly characterised as true positives, foreshadowing the low level of false positivity and false negativity. The test set recall of 98.41% supports the fact that the model is able to predict abnormalities inside healthy, powdery and rusty leaves with minimal error. The high levels of performance on various tests is a sign of a good generalisation and is supported by a high level of overfitting resistance. As a result, the approach can be deployed effectively to actual farms, especially in the context of drone-based surveillance to detect diseases. These results support the suggested architecture and its attention controls in improving classification accuracy, which makes them appropriate for large-scale precision-agriculture.

Table 4,
Performance on validation and test datasets

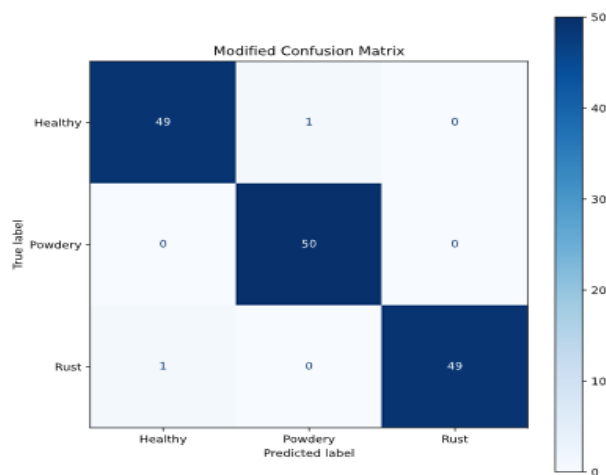
Metric\Dataset	Validation Dataset	Test Dataset
Accuracy	98.00%	98.33%
Recall	98.00%	98.33%
Precision	98.00%	98.41%
F1-Score	98.00%	98.33%

The confusion matrices indicated the performance of the proposed model on the validation and test data. The confusion matrix of the model predictions on the validation set is provided

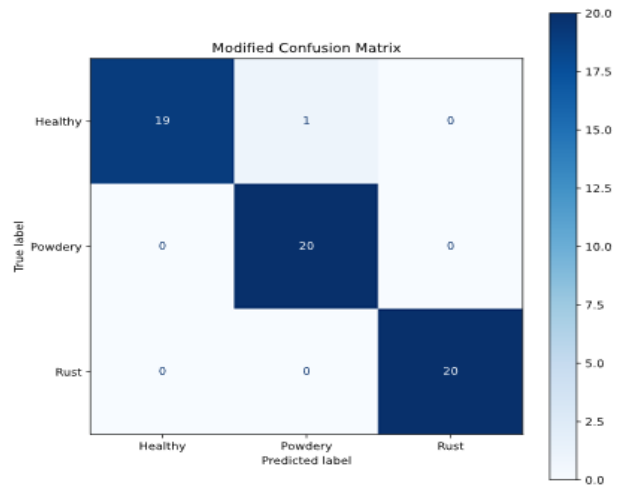
in Fig. 8(A). The model was found to be doing relatively well, with most of the samples being classified correctly. The results showed that 49 out of 50 healthy samples were identified correctly, and one sample was mistakenly recognised as rusty. All the examples of the powdery class were correctly classified, whereas 49 out of 50 samples were classified as rusty. These findings define the robustness of the model in suppressing false negatives. The diagonal shows the high values of precision and recall, which testify that this model is strong in generalising on the validation set.

Fig. 8B depicts the confusion matrix of the test dataset, confirming the consistency and reliability of the model. The classifier identified 19 of 20 healthy instances correctly and failed to identify one of the 20 instances as powdery. The powdery and rusty classes were correctly identified with 100% accuracy. Considering the insignificant difference in performance, the model shows high generalisation and practicality in real-life environmental setup, especially in precision agriculture to diagnose field-based plant diseases.

Besides, the strength of the model was implicitly evaluated through a test set of images taken in various environmental conditions. Unlike datasets that are restricted to controlled lab environments, the samples used in this study included natural variation in the intensity of light, shadow and complicated backgrounds, including soil and vegetation. The results showed no deviations and a high accuracy of 98.33%, indicating that the presented AHCNN architecture is effective in extracting disease-specific features and invariant to such common environmental perturbations. These results indicated that the model can be easily deployed to real-world scenarios where the background and lighting conditions cannot be strictly regulated.



(A)



(B)

Fig. 8. Confusion matrix of model: (A) predictions on the validation dataset, (B) predictions on the test dataset

Table 5 compares the proposed model with several well-known image classification models. The proposed model achieved a strong balance amongst accuracy, recall, precision and efficiency, reaching an accuracy of 98.00%, which is higher than those of AlexNet (94.00%), VGG16 (95.33%), ResNet-50 (96.00%), EfficientNet-B7 (96.66%) and ConvNeXt (97.57%). In addition, recall and precision reached 98.00%, indicating that the model can correctly classify samples, with very few false negatives and false positives. A key advantage of the proposed model is its lightweight design. It uses only 0.4 million parameters, considerably fewer than AlexNet (60M), VGG16 (132M), ResNet-50 (25.6M), EfficientNet-B7 (86M) and ConvNeXt (89M). This small number of parameters means that fewer computations are needed, thereby enabling very fast and in-the-field inferences, thus making the model more compatible to being deployed in resource-heavy contexts like mobile applications and edge devices. The combination of high classification performance and low complexity of computations makes the proposed model a convenient and reasonable solution to the real-life applications of detecting plant ailments and implementing more accurate agriculture.

Table 5,
Performance comparison with other state-of-the-art models (validation dataset)

Model	Accuracy (%)	Recall (%)	Precision (%)	Parameters (Million)
AlexNet	94.00	95.00	98.75	60
VGG16	95.33	95.63	98.81	138
ResNet-50	96.00	96.88	98.57	25.6
EfficientNet-B7	96.66	96.88	98.99	86
ConvNeXt	97.57	97.79	97.37	89
Proposed Model	98.00	98.00	98.00	0.4

Table 6 summarises the computational efficiency of the model. Tested on a standard NVIDIA Tesla T4 GPU, AHCNN achieved an average processing speed of 45–55 FPS at high resolution, confirming its viability for real-time agricultural monitoring tasks.

Table 6,
Performance comparison with some state-of-the-art models (validation dataset)

Hardware Platform	Input Resolution	Average Inference Time (ms)	Frames Per Second (FPS)	Real-Time Capability (> 30 FPS)
Google Colab Pro (NVIDIA Tesla T4)	512 × 512	18 – 22 ms	45 – 55	Yes

Table 7 shows the accuracy, precision, recall and F1-score of each class in terms of the predictions of the test data to perform a granular analysis of the specific strengths of the model. The model scored 100% in all measures of the rusty class, indicating that the model is sensitive enough to separate the well-defined textural characteristics of rust disease. A slight trade-off can be found between the healthy and the powdery classes. The healthy class showed a precision of 100%, but the recall was 95% because one sample was falsely classified as a powdered mildew. Thus, the powdery group provided a marginally reduced accuracy of 95.24% but an ideal recall of 100%. These results confirmed that despite a small possibility of confusion between early powdery mildew and healthy leaf texture, the model still achieved a per-class rate of accuracy at least 98.33% of all the categories.

Table 7,
Performance on validation and test datasets

Class	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Healthy	98.33%	100.00%	95.00%	97.43%
Powdery	98.33%	95.24%	100.00%	97.56%
Rusty	100.00%	100.00%	100.00%	100.00%

Regarding the allocation of the dataset, a large percentage of images (about 86%) was assigned to the training set to utilise the image learning capabilities of the deep convolutional network to the fullest. In line with this, the independent test set consisted of only approximately 4% of all the information (60 images). However, the validity of the given findings is supported by the validation set, consisting of 150 images (around 10% of the sample). The thorough performance analysis shown in Figs. 5–7 is based on such a bigger validation set to monitor the precision, recall and F1-score through 200 epochs. The stability of these 150 samples gives strong statistical support on the robustness of the model. Furthermore, the great consistency of the validation accuracy (98%) and the independent test accuracy (98.33%) proved that the model can be generalised, and the results are not due to the fact that this was a smaller test fit.

4. Conclusions

In this paper, a powerful, AI-powered solution to plant disease diagnostics introduced. The model could remove the key challenges in precision agriculture directly, i.e. computational scalability and real-time integration. The proposed AHCNN uses SE blocks and spatial attention block to improve feature extraction. The model demonstrated a validation accuracy of 98% and a test accuracy of 98.33%, with a weighted F1-score of 0.98, demonstrating strong performance on benchmark datasets.

Most importantly, AHCNN provides a significant gain in computational efficiency compared with modern state-of-the-art models. It uses only 0.4 million parameters to achieve a significant reduction in the computational cost, and it outperformed heavier networks like AlexNet (94%, 60 million parameters), VGG16 (95.33%, 138 million parameters), and ResNet-50 (96%, 25.6 million parameters). This value is a 99.7% reduction in parameter over VGG16 and 98.4% over ResNet-

50 without affecting the classification accuracy.

Furthermore, the confusion matrices support the reliability of the model. A validation set was used, where 98% and 100% of healthy and rusty leaves were properly identified, respectively, thus supporting the effectiveness of the suggested method. These empirical findings substantiate the claim that AHCNN provides a trade-off that is optimal between accuracy and inference speed. Therefore, implementing it on resource-constrained environments, including agricultural drones and edge AI systems, is a workable choice.

Future studies will focus on enlarging the dataset to cover a larger range of crop species and restrictively evaluate the cross-domain generalisation properties of the model and increase its applicability to a larger variety of agricultural settings.

Conflicts of Interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

References

- [1] S. O. Araujo, R. S. Peres, J. C. Ramalho, F. Lidon, and J. Barata, "Machine learning applications in agriculture: current trends, challenges, and future perspectives," *Agronomy*, vol. 13, p. 2976, 2023, <https://doi.org/10.3390/agronomy13122976>.
- [2] I. Ahmed and P. K. Yadav, "Plant disease detection using machine learning approaches," *Expert Systems*, vol. 40, p. e13136, 2023, <https://doi.org/10.1111/exsy.13136>.
- [3] M. Sharma, C. J. Kumar, and A. Deka, "Early diagnosis of rice plant disease using machine learning techniques," *Archives of Phytopathology and Plant Protection*, vol. 55, pp. 259-283, 2022, <https://doi.org/10.1080/03235408.2021.2015866>.
- [4] A. Singla, A. Nehra, K. Joshi, A. Kumar, N. Tuteja, R. K. Varshney, et al., "Exploration of machine learning approaches for automated crop disease detection," *Current plant biology*, vol. 40, p. 100382, 2024, <https://doi.org/10.1016/j.cpb.2024.100382>.
- [5] A. V. Panchal, S. C. Patel, K. Bagyalakshmi, P. Kumar, I. R. Khan, and M. Soni, "Image-based plant diseases detection using deep learning," *Materials Today: Proceedings*, vol. 80, pp. 3500-3506, 2023, <https://doi.org/10.1016/j.cpb.2024.100382>.
- [6] J. Kotwal, R. Kashyap, and S. Pathan, "Agricultural plant diseases identification: From traditional approach to deep learning," *Materials Today: Proceedings*, vol. 80, pp. 344-356, 2023, <https://doi.org/10.1016/j.matpr.2023.02.370>.
- [7] M. Arsenovic, M. Karanovic, S. Sladojevic, A. Anderla, and D. Stefanovic, "Solving current limitations of deep learning based approaches for plant disease detection," *Symmetry*, vol. 11, p. 939, 2019, <https://doi.org/10.3390/sym11070939>.
- [8] J. Liu and X. Wang, "Plant diseases and pests detection based on deep learning: a review," *Plant methods*, vol. 17, p. 22, 2021, <https://doi.org/10.1186/s13007-021-00722-9>.
- [9] J. A. Wani, S. Sharma, M. Muzamil, S. Ahmed, S. Sharma, and S. Singh, "Machine learning and deep learning based computational techniques in automatic agricultural diseases detection: Methodologies, applications, and challenges," *Archives of Computational methods in Engineering*, vol. 29, pp. 641-677, 2022, <https://doi.org/10.1007/s11831-021-09588-5>.
- [10] R. Rahman. Plant disease recognition dataset, Kaggle. [Online]. Available: <https://www.kaggle.com/datasets/rashikrahmanpritom/plant-disease-recognition-dataset>. [Accessed: Nov. 12, 2024].
- [11] S. S. Mohammed and H. G. Clarke, "Conditional image-to-image translation generative adversarial network (cGAN) for fabric defect data augmentation," *Neural Computing and Applications*, vol. 36, pp. 20231-20244, 2024, <https://doi.org/10.1007/s00521-024-10179-1>.
- [12] S. N. Mahmood, S. S. Mohammed, A. G. Ismaeel, H. G. Clarke, I. N. Mahmood, D. A. Aziz, et al., "Improved malaria cells detection using deep convolutional neural network," in *2023 5th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, 2023, pp. 1-4, <https://doi.org/10.1109/HORA58378.2023.10156747>.
- [13] S. S. Mohammed, H. G. Clarke, and S. N. Mahmood, "A fully convolutional neural network for fast detection, classification, and segmentation of fabric defects," *Neural Computing and Applications*, vol. 37, pp. 23249-23272, 2025, <https://doi.org/10.1007/s00521-025-11495-w>.
- [14] J. G. Barbedo, "Factors influencing the use

- of deep learning for plant disease recognition," *Biosystems engineering*, vol. 172, pp. 84-91, 2018, <https://doi.org/10.1016/j.biosystemseng.2018.05.013>.
- [15] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Computers and electronics in agriculture*, vol. 145, pp. 311-318, 2018, <https://doi.org/10.1016/j.compag.2018.01.009>.
- [16] S. S. Hari, M. Sivakumar, P. Renuga, and S. Suriya, "Detection of plant disease by leaf image using convolutional neural network," in *2019 International conference on vision towards emerging trends in communication and networking (ViTECoN)*, 2019, pp. 1-5, <https://doi.org/10.1109/ViTECoN.2019.8899748>.
- [17] S. V. Militante, B. D. Gerardo, and N. V. Dionisio, "Plant leaf detection and disease recognition using deep learning," in *2019 IEEE Eurasia conference on IOT, communication and engineering (ECICE)*, 2019, pp. 579-582, <https://doi.org/10.1109/ECICE47484.2019.8942686>.
- [18] M. H. Saleem, J. Potgieter, and K. M. Arif, "Plant disease detection and classification by deep learning," *Plants*, vol. 8, p. 468, 2019, <https://doi.org/10.3390/plants8110468>.
- [19] M. R. Ullah, N. A. Dola, A. Sattar, and A. Hasnat, "Plant diseases recognition using machine learning," in *2019 8th International Conference System Modeling and Advancement in Research Trends (SMART)*, 2019, pp. 67-73, <https://doi.org/10.1109/SMART46866.2019.9117284>.
- [20] R. Sujatha, J. M. Chatterjee, N. Jhanjhi, and S. N. Brohi, "Performance of deep learning vs machine learning in plant leaf disease detection," *Microprocessors and Microsystems*, vol. 80, p. 103615, 2021, <https://doi.org/10.1016/j.compag.2018.01.009>.
- [21] R. Z. Khaleel, H. Z. Khaleel, A. A. Abdullah Al-Hareeri, A. S. Mahdi Al-Obaidi, and A. J. Humaidi, "Improved Trajectory Planning of Mobile Robot Based on Pelican Optimization Algorithm," *Journal Européen des Systèmes Automatisés*, vol. 57, 2024, <https://doi.org/10.18280/jesa.570408>.
- [22] H. Z. Khaleel and A. J. Humaidi, "Towards accuracy improvement in solution of inverse kinematic problem in redundant robot: A comparative analysis," *International Review of Applied Sciences and Engineering*, vol. 15, pp. 242-251, 2024, <https://doi.org/10.1556/1848.2023.00722>.
- [23] H. Obied, M. K. Al-Taleb, H. Z. Khaleel, and A. F. AbdulKareem, "Implementation and Derivation Kinematics Modeling Analysis of WidowX 250 6Degree of Freedom Robotic Arm," *Journal of Engineering and Sustainable Development*, vol. 29, pp. 473-484, 2025, <https://doi.org/10.31272/jeasd.2454>.
- [24] İ. Yağ and A. Altan, "Artificial intelligence-based robust hybrid algorithm design and implementation for real-time detection of plant diseases in agricultural environments," *Biology*, vol. 11, p. 1732, 2022, <https://doi.org/10.3390/biology11121732>.
- [25] S. R. Krishnan, E. P. Sim, C. M. Varghese, B. Rajan, S. Chippy, and E. Thomas, "Detection of Diseases in Tomato Leaves Using Deep Learning Models: A Survey," in *2024 1st International Conference on Trends in Engineering Systems and Technologies (ICTEST)*, 2024, pp. 1-6, <https://doi.org/10.1109/ICTEST60614.2024.10576182>.
- [26] M. Venkatanareesh and I. Kullayamma, "Deep learning based concurrent excited gated recurrent unit for crop recommendation based on soil and climatic conditions," *Multimedia Tools and Applications*, pp. 1-30, 2024, <https://doi.org/10.1007/s11042-023-18004-y>.
- [27] W. Albattah, M. Nawaz, A. Javed, M. Masood, and S. Albahli, "A novel deep learning method for detection and classification of plant diseases," *Complex & Intelligent Systems*, pp. 1-18, 2022, <https://doi.org/10.1007/s40747-021-00536-1>.
- [28] C. Sarkar, D. Gupta, U. Gupta, and B. B. Hazarika, "Leaf disease detection using machine learning and deep learning: Review and challenges," *Applied Soft Computing*, vol. 145, p. 110534, 2023, <https://doi.org/10.1016/j.asoc.2023.110534>.
- [29] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "Cbam: Convolutional block attention module," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 3-19, https://doi.org/10.1007/978-3-030-01234-2_1.

الشبكات العصبية التلافيفية الهجينة المتقدمة للكشف عن أمراض النباتات المستندة إلى الاوراق

سرمد نوزاد محمود¹، صواش سامي محمد²، يحيى غفران خضر^{3*}، تالي سعدي⁴، سحر صالح⁵

¹ قسم هندسة تقنيات الأجهزة الطبية، الكلية التقنية الهندسية – كركوك، الجامعة التقنية الشمالية، الموصل ٤١٠٠١، العراق

² قسم هندسة تقنيات النكاه الاصطناعي، الكلية التقنية الهندسية للحاسوب والنكاه الاصطناعي – كركوك، الجامعة التقنية الشمالية، الموصل ٤١٠٠١، العراق

³ قسم هندسة تقنيات الإلكترونيك والسيطرة، الكلية التقنية الهندسية – كركوك، الجامعة التقنية الشمالية، الموصل ٤١٠٠١، العراق

^{4,5} مختبر WiSAR، جامعة الأطلسي التكنولوجية (ATU)، لتركني، مقاطعة دونيجال، F92 YY97، أيرلندا

* البريد الإلكتروني: yahhya.khidhir24@ntu.edu.iq

المستخلص

يُعدّ الكشف الدقيق عن أمراض النباتات وتصنيفها بنجاح أمرًا بالغ الأهمية لتحقيق إنتاج غذائي مستدام وتقليل خسائر المحاصيل. وتعتمد إجراءات التشخيص التقليدية غالبًا على الخبرة البشرية، كما أنها تستغرق وقتًا طويلاً ويصعب توسيع نطاقها للتطبيقات واسعة النطاق. في هذه الورقة، يتم اقتراح نموذج هجين متقدم للشبكات العصبية التلافيفية (Advanced Hybrid CNN) والمسمى بـ (AHCNN) يدمج آليات الانتباه مع استراتيجيات النفاذ فعالة لتحقيق دقة عالية مع تقليل التعقيد الحسابي. تم تطوير النموذج باستخدام مجموعة بيانات عالية الدقة لأوراق النباتات تتكون من ١٥٣٢ صورة، موزعة إلى ١٣٢٢ صورة للتدريب، و ١٥٠ صورة للتحقق، و ٦٠ صورة للاختبار، وذلك لتصنيف العينات إلى ثلاث فئات: سليمة (Healthy)، والبياض الدقيقي (Powdery)، والصدأ (Rust). تتضمن البنية المقترحة كتل الضغط والإثارة (Squeeze-and-Excitation Blocks) وآلية الانتباه المكاني لتعزيز استخراج السمات وتحسين قابلية تفسير النموذج. أظهرت النتائج التجريبية أن النموذج يحقق دقة تحقق تبلغ (98.00%) ودقة اختبار تصل إلى (٩٨,٣٣٪). وباستخدام عدد معاملات لا يتجاوز 0.4 مليون، يوفر النموذج المقترح حلاً خفيف الوزن يتفوق بشكل ملحوظ على معماريات التعلم العميق القياسية من حيث الكفاءة الحسابية. وتجعل هذه الخصائص من نموذج AHCNN مرشحاً قوياً لتطبيقات الكشف الآني عن أمراض النباتات المعتمدة على الطائرات المسيّرة ضمن أنظمة الزراعة الدقيقة.