Soft Computing (2024) 28:12929–12941
https://doi.org/10.1007/s00500-024-10348-3

APPLICATION OF SOFT COMPUTING

Real-time disease detection on bean leaves from a small image dataset using data augmentation and deep learning methods

Emmanouil Karantoumanis[1] · Vasileios Balafas[1] · Malamati Louta[1] · Nikolaos Ploskas[1]

Accepted: 11 June 2024 / Published online: 18 November 2024

Abstract

Disease detection in agricultural crops plays a pivotal role in ensuring food security and sustainable farming practices. Deep learning models, known for their ability in image analysis, often demand extensive image datasets and annotations to achieve high accuracy. However, in the case of bean crops, the absence of a publicly available dataset has posed a significant challenge for applying deep learning algorithms to accurately predict diseases. Additionally, the manual annotation of images demands substantial time and resources. This paper introduces an innovative approach to tackle these issues. We introduce a solution for real-time disease detection on bean leaves, despite the lack of bean-specific image data. Initially, we generate a small dataset from real images and annotate them. Then, we utilize images from the existing dataset PlantDoc (Singh et al. in: Proceedings of the 7th ACM IKDD CoDS and 25th COMAD, Association for Computing Machinery, pp 249–253, 2020) from leaves of other plant species. Moreover, to compensate for the limitations of a small image dataset, we employ advanced data augmentation techniques, enriching the training data and enhancing the model's ability to generalize. Our experimental study shows that data augmentation techniques can improve the accuracy of deep learning methods by up to 37%.

Keywords Disease detection · CNN · Bean crop · Data augmentation · Deep learning

1 Introduction

In modern agriculture, early detection and accurate prediction of crop diseases have become crucial to ensure food security, sustainable agricultural practices, and optimal crop yields. One of the innovative approaches applied for early detection of plant diseases is the use of remote sensing technologies, such as Unmanned Aerial Vehicles (UAVs) and Unmanned Ground Vehicles (UGVs), combined with Convolutional Neural Network (CNN)-based algorithms for object detection. This combination not only improves the accuracy of disease detection but also reduces the time and resources required for crop monitoring.

Disease detection in crops holds pivotal importance in modern agriculture for several reasons (Alguliyev et al. 2021). Firstly, it directly impacts crop yield and quality. Detecting diseases early helps farmers take steps to prevent them, reducing crop losses. Secondly, it helps in the optimization of resource utilization. By detecting areas affected by diseases, farmers can efficiently allocate water, fertilizers, and pesticides, reducing unnecessary expenditures and environmental impact. Furthermore, early disease detection contributes to the sustainability of agriculture by reducing the need for excessive chemical treatments, which can harm the environment and human health.

In recent years, CNN-based algorithms have emerged as one of the best solutions for object detection in various domains, including agriculture. Their capability to learn complex patterns and features from images has greatly transformed the field of computer vision, enabling the development of highly accurate and efficient disease detection systems. These algorithms can process large amounts of

Emmanouil Karantoumanis, Vasileios Balafas and Malamati Louta have contributed equally to this work.

✉ Nikolaos Ploskas
nploskas@uowm.gr

Emmanouil Karantoumanis
e.karantoumanis@uowm.gr

Vasileios Balafas
v.balafas@uowm.gr

Malamati Louta
louta@uowm.gr

[1] Department of Electrical and Computer Engineering, University of Western Macedonia, Campus ZEP, 50100 Kozani, Greece

image data and provide real-time information about the health of crops, making them essential tools for modern precision agriculture.

While CNN-based algorithms have demonstrated their value in disease detection, their effectiveness is heavily based on the availability of extensive and diverse training datasets. Publicly available datasets play a pivotal role in training these algorithms and ensuring their robustness. However, no public data exists for images in specific areas. Especially, when it comes to bean crops, a significant gap exists in the availability of such datasets. Despite the economic importance of beans as a basic food crop, there is a notable scarcity of comprehensive datasets for disease detection and prediction in bean plants.

This lack of bean-specific datasets poses a considerable challenge, as deep learning algorithms, including CNNs, require substantial images and annotations to predict diseases accurately. Taking these images is not only time-consuming but also a resource-intensive process. To utilize a valuable set of data, not only the collection of images, which is a costly process by itself but also the annotation of the classes in the images, which is a time-consuming process that experts in the field of agriculture must do. This paper addresses this limitation by proposing a novel data augmentation approach specifically tailored to under-represented crops like beans. Our ultimate goal is to utilize a relatively small image dataset specific to beans and complement it with existing datasets from the leaves of other plants. By applying advanced data augmentation techniques, we aim to bridge the gap between the scarcity of bean-specific data and the requirements of deep learning algorithms, thus facilitating more accurate and efficient disease detection and prediction for bean crops in the Prespa Lakes region of Northern Greece.

The structure of the paper is as follows. Section 2 reviews the field of real-time disease detection using small image datasets and how researchers have adeptly leveraged data augmentation. In Sect. 3, we present the proposed method of disease and pest detection system and the data augmentation techniques. Section 4 presents the computational experiments and an extensive discussion based on the results of the experiments. Finally, Sect. 5 summarizes the work of this paper.

2 Literature review

Disease detection, particularly in the domain of agriculture, represents a challenge that has garnered significant attention over the past few decades. The consequences of undetected or late-detected plant diseases can be vast, leading to diminished yields and economic losses. As such, it remains a well-studied problem across the disciplines of plant pathology, computer vision, and artificial intelligence. Several methods have been proposed to address this concern. Traditional machine learning methods, such as Support Vector Machines and Random Forests, have been employed with notable success (Das et al. 2020; Ramesh et al. 2018; Singh et al. 2022; Govardhan and Veena 2019; Saha and Ahsan 2021). In recent years, deep learning techniques, especially CNNs, have emerged as a dominant force to detect diseases in plants (Ferentinos 2018; Boulent et al. 2019; Shrestha et al. 2020). Furthermore, hybrid models combining features extracted from both traditional and deep learning models have been introduced, adding another dimension to the solutions available (Bedi and Gole 2021).

However, one of the major challenges in employing these sophisticated models is the requirement of large, labeled datasets. In many real-world scenarios, acquiring a substantial dataset is a difficult task, both in terms of time and resources. In such cases where data is sparse, data augmentation techniques come into play (Chug et al. 2023). These techniques artificially expand the dataset by applying a range of transformations, such as rotations, scaling, and cropping, thereby providing the model with a more diverse set of examples to learn from.

In this paper, our focus narrows down to exploring novel methods in real-time disease detection using small image datasets. Specifically, we delve deep into how researchers have leveraged data augmentation in tandem with deep learning techniques to improve the accuracy and efficiency of their models in such constrained data environments.

2.1 Disease detection

Jiang et al. (2019) proposed a deep learning method with CNNs for the real-time detection of apple leaf diseases. They used data augmentation and image annotation technologies to construct a comprehensive dataset, attaining 78.80% mAP (mean Average Precision). Islam et al. (2023) introduced a method for predicting crop diseases using deep learning. They created a web tool, DeepCrop, to aid farmers in identifying plant diseases. During their tests, they assessed multiple deep learning architectures such as CNN, VGG-16 (Simonyan and Zisserman 2014), VGG-19 (Simonyan and Zisserman 2014), and ResNet50 (He et al. 2016) on the PlantVillage (Kaggle 2018) image dataset to identify crop diseases. The ResNet50 model emerged as the best, boasting 98.98% accuracy. Mahum et al. (2023) developed a new framework for detecting potato leaf diseases using a refined deep-learning model. Their method classifies potato leaves into five categories, utilizing the PlantVillage dataset. To overcome data imbalance, they applied a reweighted cross-entropy loss function and integrated the Efficient DenseNet model from DenseNet-201 (Huang et al. 2017), alongside regularization techniques, for efficient disease classification. Their framework reached a 97.2% accuracy.

In the subsequent section, we will focus on existing works that use data augmentation techniques in conjunction with deep learning for effective disease detection, especially when constrained by the dataset size.

2.2 Data augmentation

Data augmentation is a technique widely adopted in the realm of machine learning and deep learning to mitigate the challenges posed by limited data. By artificially expanding the dataset through a series of transformations, data augmentation not only addresses the problem of overfitting but also aids in enhancing the model's capability to generalize across diverse scenarios. Several studies have corroborated the efficacy of data augmentation in improving disease detection rates. Still, many studies do not mention the influence of data augmentations on their outcomes even though they use such techniques, as observed in prior research, e.g., Enkvetchakul and Surinta (2022), Abayomi-Alli et al. (2021), Kaushik et al. (2020), Dai et al. (2023) and Li et al. (2023). In this section, we exclusively examine research works that explicitly discuss the effects of data augmentation on their detection results.

Zhang et al. (2023) introduced a high-quality image augmentation method for enhancing the quality of rice leaf disease samples using a dual GAN framework. By processing the generated pseudo-data through an Optimized-Real-ESRGAN, they achieved high-quality images for disease classification. Their method boosted recognition accuracy by 4.57% on ResNet18 (He et al. 2016) and 4.1% on VGG11 (Simonyan and Zisserman 2014) compared to the original dataset. Compared to solely using WGAN-GP for augmentation, they observed increases of 3.08% for ResNet18 and 3.55% for VGG11. Their method proves especially beneficial for situations with limited training datasets.

Haruna et al. (2023) proposed a synthetic data augmentation method using Style-Generative Adversarial Network Adaptive Discriminator Augmentation (SG2-ADA) to address the challenge of limited and uneven rice leaf disease datasets. By using the variance of the Laplacian filter, they enhanced the performance of Faster-RCNN (Ren et al. 2015) and SSD (Liu et al. 2016) models. After training SG2-ADA for 250 epochs and filtering out low-quality images, they augmented these models for disease detection, achieving a mAP of 93% for Faster-RCNN and 91% for SSD.

Cap et al. (2022) introduced LeafGAN, a system designed to enhance data augmentation in plant disease diagnosis by transforming healthy images into diseased ones, focusing on relevant image areas with its attention mechanism. This made the training data more diverse. In tests on a five-class cucumber disease classification, while CycleGAN (Zhu et al. 2017) improved diagnostic performance by only 0.7% from the baseline, LeafGAN boosted it by a significant 7.4%.

Zeng et al. (2020) delved into the application of deep learning models for detecting the severity of Huanglongbing citrus infections. Using a dataset of 5, 406 HLB-infected citrus leaf images, the InceptionV3 (Szegedy et al. 2016) model emerged as the most efficient, achieving a detection accuracy of 74.38%. Notably, when the team employed deep convolutional generative adversarial networks for data augmentation, doubling the original dataset, the accuracy of the InceptionV3 model significantly increased to 92.60%. This significant improvement underscores the pivotal role of GANs-based data augmentation in enhancing model performance.

Diana Andrushia et al. (2023) used a convolutional capsulet network for detecting diseases on grape leaves. They used an original dataset of 11300 images, where 4000 were images from healthy leaves, while the rest were images with diseased leaves. They used data augmentation techniques, such as rotation, scaling, gamma correction, flipping, and color augmentation to enhance the dataset, reaching a total of 28,000 images. The validation accuracy for the augmented data was 99.12% compared to 92.13% for the original dataset.

In the face of limited data, data augmentation emerges as an indispensable tool, bridging the gap between scarcity and the need for diversity, and propelling models to achieve high accuracy and robustness in disease detection. In this work, we utilize a small dataset of bean crops, data augmentation techniques to enhance the number and diversity of training images, and datasets on other plant species to improve the accuracy of CNN-based disease detection models.

3 Proposed methodology

In Fig. 1, we present a visual representation of the disease and pest detection system utilized within bean cultivation. This system is connected to three distinct subsystems UAV subsystem, the UGV subsystem, and the pest damage detection engine subsystem.

The disease and pest detection system serves as a pivotal tool for data analysis and model management, aiming at the early detection of potential diseases in plants. The process initiates with the activation of the UAV subsystem to gather aerial images. Subsequently, these images are relayed to the pest damage detection engine subsystem, where the best CNN model is chosen to anticipate the diseases that may affect the crops. Figure 2 visually illustrates an instance where the UAV subsystem provides an image to the pest damage detection engine, depicting a scenario marked by the presence of a disease in the field.

The system then evaluates the likelihood of disease occurrence by assessing the confidence level associated with each classified disease. A disease is predicted to be present when the confidence level for any identified disease on any leaf

Fig. 1 Flow chart of the proposed methodology

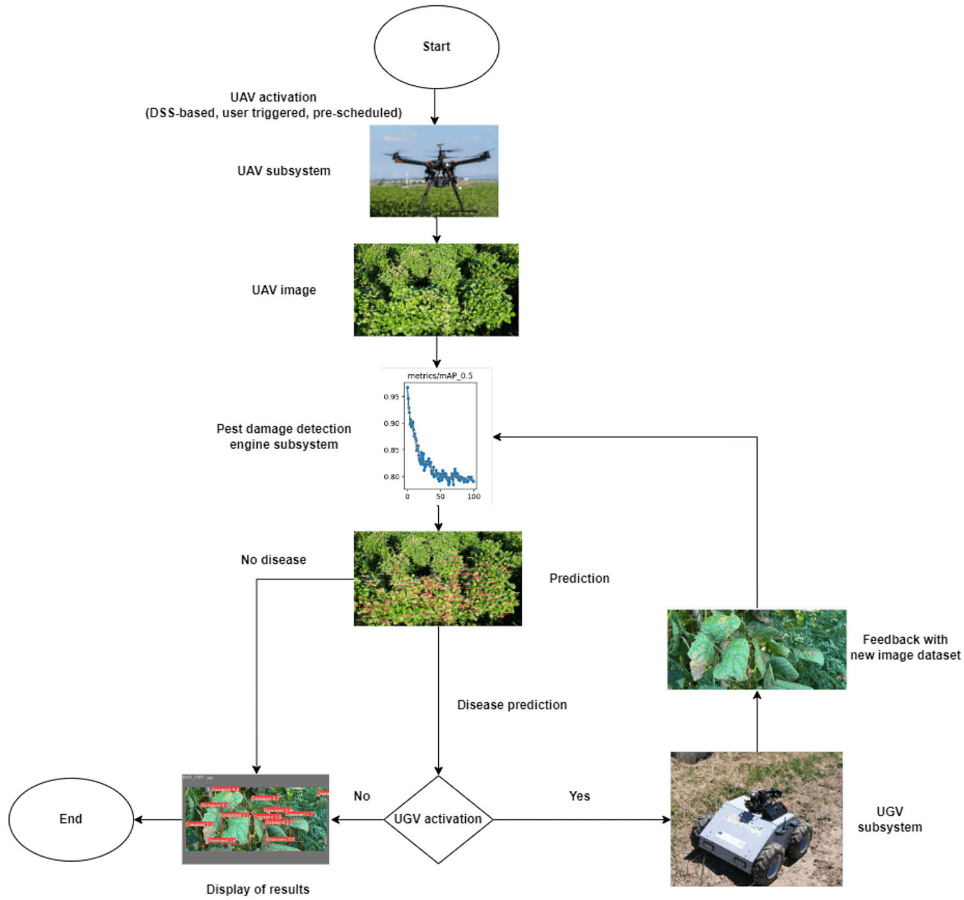

Fig. 2 Real UAV image of the crop

Fig. 3 Model disease predictions on UAV images

exceeds 60%. Figure 3 shows the image after the model prediction, where it is obvious that diseases are present.

Conversely, if no disease is detected or if the confidence level falls below 60%, the disease detection system signals the presence of healthy plants in the image, bringing the process to a halt. This scenario is illustrated in Figs. 4 and 5, featuring an actual UAV image of a healthy portion of the field alongside the model's annotations.

In cases where diseases are detected in the UAV images, the UGV subsystem is called by providing it with the coordinates of the detected disease. The UGV then proceeds to the specified location, capturing images from a closer vantage point. These images have greater sharpness and detail, offering more detailed information about the disease. Figure 6 shows an image of a UGV containing plant leaf diseases.

The new images obtained by the UGV serve the purpose of updating the disease detection model, providing additional data to refine the model's capabilities. This update helps improve the accuracy and performance of the model, allowing it to make more accurate and reliable predictions of plant

Fig. 4 Real UAV image of a healthy crop

Fig. 5 Model healthy predictions on UAV images

Fig. 6 Real image of UGV with crop diseases

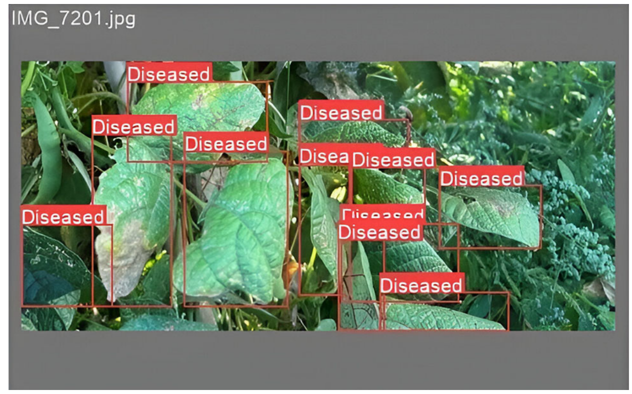

Fig. 7 Image with agronomist annotations of diseases

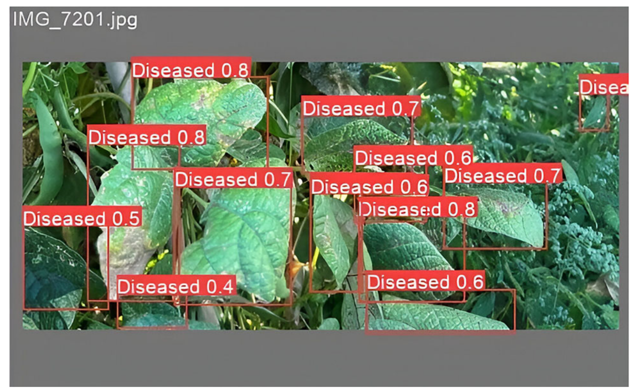

Fig. 8 Image with model predictions of diseases

Fig. 9 Real image of healthy plants

diseases. In this way, the cooperation between the UAV and UGV subsystems allows more accurate data to be collected, the model to be updated and the ability to detect and predict plant diseases to be improved. This process continues as analyses and predictions are repeated, allowing for continuous improvement of the pest damage detection engine subsystem.

Figure 7 exhibits an image from the UGV with disease annotations by the agronomist, while Fig. 8 presents the predicted diseases. As clearly depicted in Fig. 8, leaves are identified as diseased with a confidence level exceeding 60%, prompting a notification of disease presence to the user.

In contrast, Fig. 9 showcases an image from a UGV mission devoid of plant leaf diseases, while Fig. 10 reveals the UGV-captured image featuring agronomist-verified healthy leaves. Figure 11 displays the corresponding predictions, showing that no diseased leaves are detected with a confidence level exceeding 60%, affirming the absence of diseases in the crop, as conveyed to the user.

This work places particular emphasis on the pivotal role of images obtained from the UGV and the subsequent predictions derived from these UGV-captured images. To train the model to make accurate predictions, we employed a two-

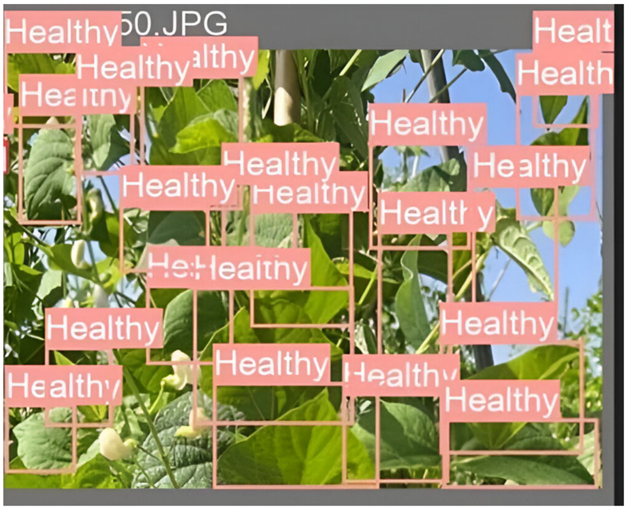

Fig. 10 Image with agronomist annotations of healthy plants

Fig. 11 Image with model predictions of healthy plants

fold strategy. Firstly, we harnessed the power of CNN-based algorithms, pre-trained on richer datasets, to further enhance the accuracy of our predictions. Secondly, we utilized some of the UGV images in which disease and healthy plants were labeled in detail. Moreover, significant emphasis is placed on the outcomes achieved through the application of data augmentation techniques, even when working with only a limited number of crop images.

3.1 Data augmentation techniques

In this subsection, we will briefly describe the data augmentation techniques we experimented with (interested readers may refer to Kaur et al. (2021), Chlap et al. (2021) and Shorten and Khoshgoftaar (2019) (for more details on these techniques):

- Laplacian filtering: Laplacian filtering enhances edges in an image, making them more prominent. It is often used for edge detection and feature extraction.
- Sobel operator: The Sobel operator is used for edge detection by convolving the image with a pair of 3×3 kernels to calculate the gradient of the image in the x and y directions.
- Grayscale: Conversion to grayscale reduces the image to a single channel, simplifying further processing and analysis by removing color information.
- HSV: Hue, Saturation, and Value color space separates color information from intensity, making it easier to work with color-based features in an image.
- LAB: LAB color space represents an image in terms of lightness (L), green-red color (A), and blue-yellow color (B), which can be useful for color-based segmentation and analysis.
- Canny edge detection: Canny edge detection is a multi-stage algorithm that detects edges in an image by finding areas of rapid intensity change. It's widely used for edge extraction.
- Thresholding: Thresholding converts an image into a binary format, where pixels above a certain threshold are set to one value, and those below are set to another. It is used for image segmentation.
- Translation: Translation shifts an image's position horizontally and vertically, which can be useful for alignment or data augmentation.
- Scale: Scaling changes the size of an image, making it larger or smaller. It's often used for resizing images to a standard size.
- Rotation: Image rotation involves rotating the image by a specified angle. It's used for correcting image orientation or creating variations for training data.
- Resize: Resizing adjusts the dimensions of an image while maintaining its aspect ratio. It is commonly used to prepare images for specific input sizes in machine learning models.
- Crop: Cropping removes a portion of an image, isolating a specific region of interest. It is useful for focusing on specific details or reducing image size.
- Gaussian blur: Gaussian blur applies a blur filter to an image, smoothing out noise and reducing fine details. It is used to reduce noise and prepare images for further processing.
- Median blur: Median blur replaces each pixel's value with the median value of the pixels in its neighborhood. It is effective at reducing salt-and-pepper noise.
- Bilateral blur: Bilateral blur is a filtering technique that smooths an image while preserving edges. It is useful for noise reduction while retaining important image features.

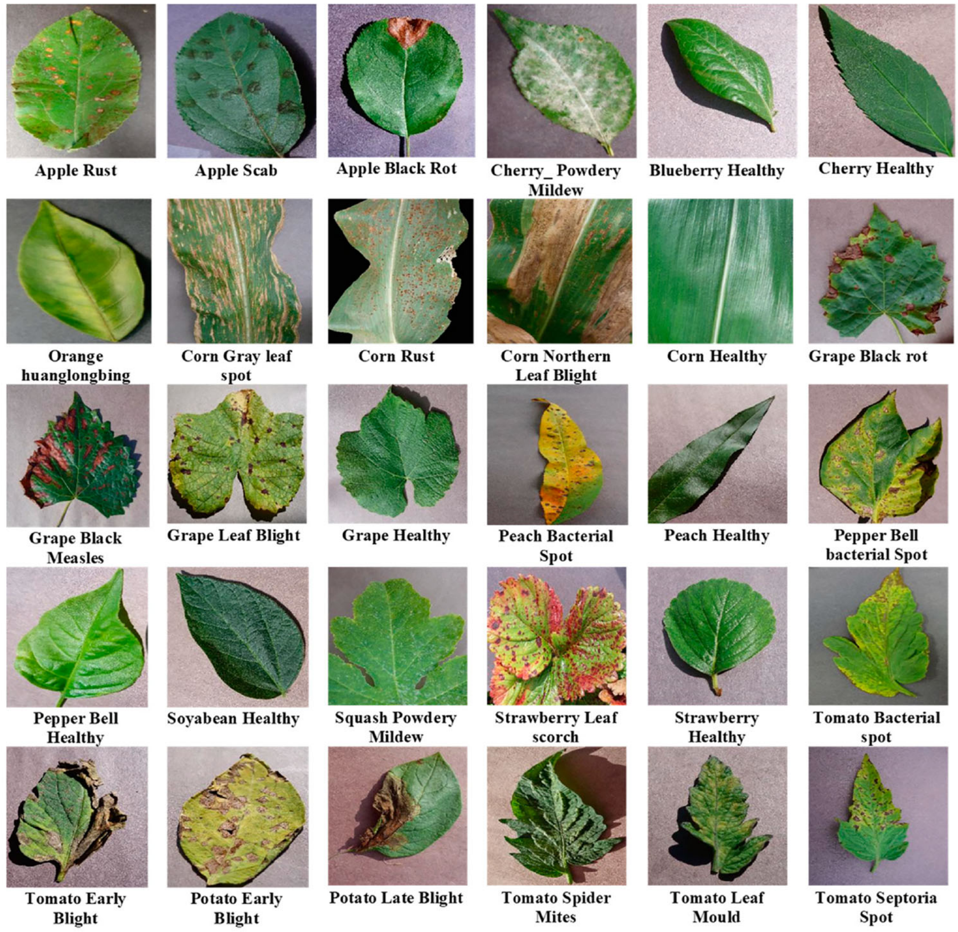

Fig. 12 Sample of PlantDoc dataset

4 Computational experiments

In this section, we detail the computational experiments carried out to assess the effectiveness of the YOLOv5 (Jocher 2022) model coupled with data augmentation techniques in detecting diseases on bean leaves.

4.1 Datasets

We trained our models using two distinct datasets. The first dataset is the widely recognized PlantDoc (Singh et al. 2020), which contains a substantial collection of plant images. The PlantDoc dataset holds significant importance in the realm of plant disease identification and diagnosis, leveraging machine learning and image processing techniques. The second dataset is a custom one that we generated from images of bean crops in the Prespa Lakes region in Northern Greece.

In Fig. 12, we provide a sample of the PlantDoc dataset. This valuable collection of images within the PlantDoc dataset showcases various plant species exposed to diverse factors affecting plant health, including diseases, pest attacks, and other influencing parameters. Comprising 2,569 images, this dataset encompasses 13 distinct plant species and 30 classes, covering both diseased and healthy states. It serves as a crucial resource for tasks such as image classification and object detection. Data for the PlantDoc dataset is sourced from multiple environments, including fields, farms, nurseries, and greenhouses, capturing images in different seasons and conditions, thereby ensuring its diversity and realism. The PlantDoc dataset includes multiple plant species including trees, shrubs, flowers, vegetables, and more. For each plant species, the dataset includes different diseases and pathogens that can affect them. For example, it can include images of plants that have been affected by fungi, bacteria, viruses, and more. Each image in the PlantDoc dataset is accompanied by detailed labels describing the plant species, the specific disease or pest attack observed, and other important information. These labels are critical for training machine learning models. Applications include automatic plant disease detection, real-time plant health monitoring, and increasing productivity in agriculture.

Our initial model training utilized the PlantDoc dataset to establish the initial weights of its variables. While this initial training allowed the model to grasp general structures and patterns related to plant diseases, it fell short of achieving high accuracy in real-world disease detection sce-

Fig. 13 Sample of custom dataset with annotations

narios. To enhance the model's performance, we created our own dataset. This dataset comprises 75 new plant images captured exclusively by a UGV, characterized by a diverse set of attributes, along with approximately 4000 annotations. From the 75 images, we kept 15 for validation and the other 60 to train the model. Since the images are taken from a real field, each image contains a lot of information about the bean leaves. Note here that the PlantDoc dataset has about three with four annotations per image while ours has about 53 annotations per image. We used the real dataset to focus on a wider range of bean pathologies. Our objective was to fine-tune the model, augmenting its sensitivity to intricate details and substantially elevating its disease detection accuracy. This process, involving validation and training with our custom dataset, markedly improved the model's capacity to accurately and reliably detect diseases in various plant species. In Fig. 13, we offer a snapshot as a sample of the dataset, captured by the UGV and meticulously annotated. More specifically, the dataset includes 75 UGV images of bean crops, four different classes (rust, spot leaf, both diseases, and healthy), and 4000 annotations (1806 for the rust class, 686 for the spot leaf class, 779 for both diseases, and 1215 for the healthy class).

4.2 Data augmentation

By augmenting the data, the model is trained on a more diverse set of examples, which helps in improving the generalization of the model, making it more robust to variations in the input data. This is particularly beneficial in plant disease detection, where the appearance of symptoms can vary greatly depending on a variety of factors including the stage of the disease, lighting conditions, and the angle of the image capture.

In our study, we leveraged data augmentation to artificially expand our dataset, thereby enabling our model to learn more diverse features and be better prepared for real-world deployment. After an extensive computational study of data augmentation techniques, we identified the top 10 methods that yielded the most substantial increase in the model's accuracy. The selected techniques are as follows:

- rotation 180 degrees
- rotation 90 degrees left
- rotation 90 degrees right
- grayscale
- HSV
- LAB
- translation
- resize
- random crop
- Gaussian blur

For every one of the 60 original images, we generated one additional image, for each of the 10 aforementioned techniques. This multiplied the initial set of 60 images into a total of 600 images. By incorporating these 600 newly generated images, we significantly augmented the depth of information available during the model training process, resulting in a substantial boost to its overall accuracy. Data augmentation techniques in the context of bean leaf disease detection proved critical to improving model accuracy since original image data alone was not sufficient. In Fig. 14 we present the original image as an initial reference point, followed by a sequence demonstrating the image's transformation after each applied technique.

4.3 CNN-based object detection models

In recent years, various CNN-based object detection models have been developed and utilized in different fields including agriculture, healthcare, and autonomous vehicles, among others. In our study, we opted to use the YOLOv5 model for disease detection on bean leaves from a small image dataset leveraging data augmentation. This choice was influenced by the superior performance of YOLOv5 in our computational study on previous works (Karantoumanis et al. 2022; Balafas et al. 2023), as well as its good performance reported in various studies (Zhao et al. 2018; Li et al. 2022). YOLOv5, the fifth iteration of the YOLO series, stands as a state-of-the-art, real-time object detection system. The architecture of YOLOv5 is built upon several innovative components that work synergistically to enhance its performance. The backbone of the architecture is CSPDarknet53, a variant of the Darknet architecture, which is known for its efficiency and speed. This backbone is coupled with PANet, a feature pyramid network that enhances information flow, and SAM block, a spatial attention module that helps the network focus on the most informative regions of the input image. YOLOv5 employs a multi-scale prediction strategy, where it makes predictions at three different scales, allowing it to

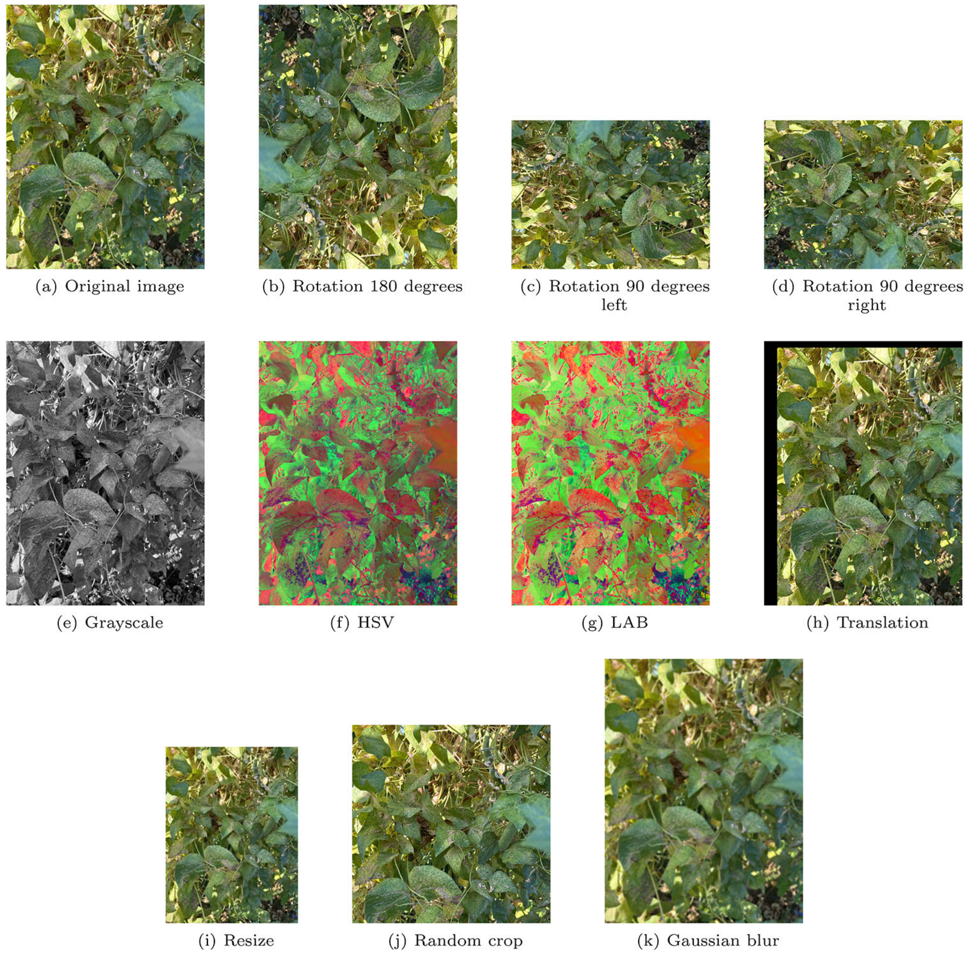

Fig. 14 Visualization of the original image and the modified one after each applied technique

detect objects of various sizes effectively. This is particularly beneficial in plant disease detection where symptoms can vary significantly in size. The network utilizes anchor boxes optimized for the dataset, improving detection accuracy for objects with different aspect ratios. The YOLOv5 model offers improved performance in terms of both speed and accuracy compared to earlier versions. It is capable of detecting objects with high precision and is optimized to function in real-time, making it suitable for applications necessitating instantaneous responses. Designed with user-friendliness in mind, YOLOv5 facilitates easy training and inference procedures. It supports automatic optimization of hyperparameters, simplifying the training process. Unlike its predecessors which were implemented in Darknet, YOLOv5 is implemented in PyTorch, enhancing its accessibility and ease of integration with other deep learning tools and libraries. The model maintains the YOLO tradition of making predictions at multiple scales, effectively detecting objects of various sizes.

4.4 Model training and validation

Utilizing the PyTorch framework, we trained the YOLOv5 model on the augmented dataset. The training process involved optimizing the model's hyperparameters to achieve the best performance. We used YOLOv5's default hyperparameters to fine-tune the model. We validated the model using a separate set of unseen data to ensure its robustness and reliability in disease detection. In Fig. 15, we present the training results of the two-class disease detection system. The training losses are represented by *train/box_loss*, *train/obj_loss*, and *train/cls_loss*, which measure the bounding box predictions, objectness confidence, and classification accuracy, respectively. The metrics include *metrics/precision*, *metrics/recall*, *metrics/mAP_0.5*, and *metrics/mAP_0.5:0.95*, which provide insights into the model's predictive performance and robustness across various Intersection over Union thresholds. Additionally, the validation losses are captured by *val/box_loss*, *val/obj_loss*, and *val/cls_loss*. The learning rates for different parts of the model are represented by *x/lr0*, *x/lr1*, and *x/lr2*, which seem to be decreasing over epochs, a typical strategy in deep learning to ensure convergence.

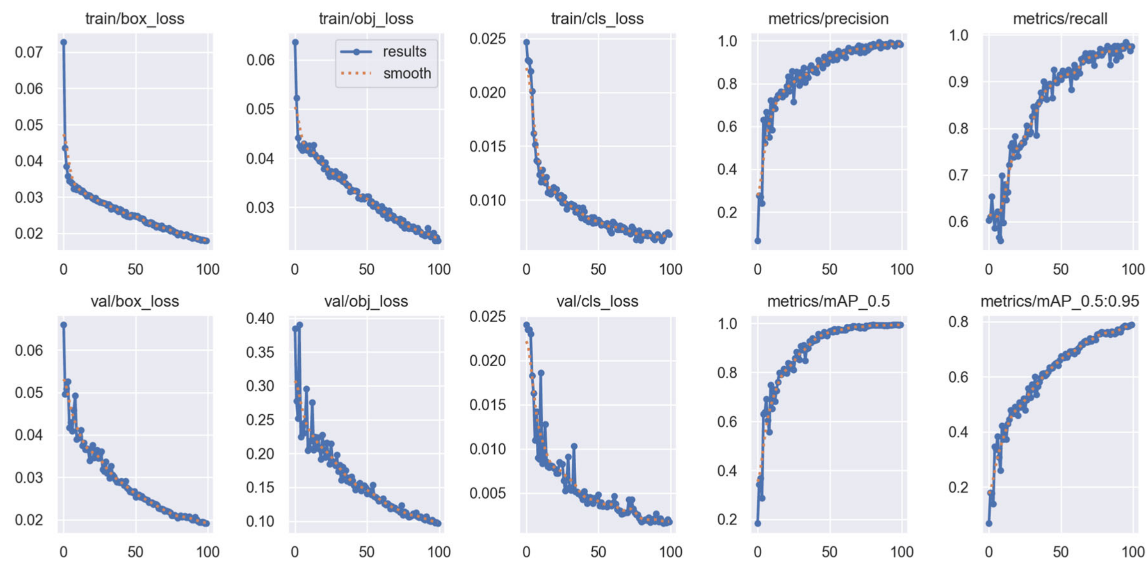

Fig. 15 Training results of the two-class disease detection system

4.5 Results and discussion

The YOLOv5 algorithm and the dataset PlantDoc, our custom dataset, and the combination of these two were used for the experimental procedure. First, we used the PlantDoc dataset to do a first training of the model to generate the basic weights of the variables. This first training helped the model understand the general structures and patterns associated with plant diseases. We then kept a sample of 15 images out of 75 in the custom dataset to evaluate the performance of the models. We divided the experiments into two general categories, two-class and four-class experiments. In two-class experiments, the algorithm's prediction is limited to only a healthy plant or a diseased plant. In contrast, in the four-class experiments, the algorithm's prediction includes the leaf spot disease, rust, the probability that both leaf spot and rust have appeared on a leaf together, and healthy plants. Three sets of datasets were used for both classes of experiments. From the first dataset, we extracted the healthy plant images from PlantDoc and plant images with all the diseases it contained. The models were evaluated on the 15 real images of the custom dataset. The second dataset contained 60 real images of beans and their diseases from the custom dataset, and the evaluation of the models was done on its remaining 15 images. Finally, because we noticed a lack of data in the images of healthy plants since the rates were quite low, we added images from the custom dataset along with an extra 316 images of healthy plants from the PlantDoc dataset.

To evaluate the performance of the model we used the metric mAP. The mAP metric is a popular way to evaluate the performance of an object recognition model in terms of its prediction accuracy. The mAP is calculated based on the prediction results produced by an object recognition model. First, the category and context are predicted for each object in an image. The predictions are then evaluated against the actual state of the objects in the images. mAP is an average of accuracies calculated for each object class. This means taking into account the accuracy for each class separately and then averaging those accuracies.

Table 1 shows the performance of the system with two classes. The classes predicted by the model are healthy plants and diseased plants. More specifically, the results of the PlantDoc dataset show that it achieves the best results with 78% prediction accuracy in the general set, 81% accuracy in diseased plants, and 76% accuracy in healthy plants. On the other hand, the custom dataset has an accuracy of 47% in the general category but has a weakness in healthy plants with an accuracy of 25%. However, this accuracy improves significantly with the addition of the PlantDoc dataset and reaches 64%. However, it does not outperform the results of the PlantDoc dataset. The improvement in results presented by the PlantDoc dataset is due to the fact that it contains many more images of both healthy and diseased plants. This allows it to better predict whether a leaf is healthy or not.

Table 2 shows the performance of the system with four classes. The classes predicted by the model are spot leaf, rust, spot leaf, and rust together on one leaf and healthy plants. This time in the PlantDoc dataset the healthy plant images and only the rust classes are used since this particular dataset does not contain any images of beans or of the four leaf spot diseases. In particular, we notice that the results of the PlantDoc dataset are very low, with the best accuracy achieved in cases where this class is present with a percentage of only 24% and in the general set 16%. The custom dataset has an accuracy of 46% on rust and 27% on the overall set and with the addition of the healthy plants from PlantDoc the accu-

Table 1 Results with two classes

	PlantDoc	Custom dataset	PlantDoc + custom dataset
All categories	0.788	0.447	0.711
Diseased	0.810	0.687	0.768
Healthy	0.765	0.253	0.645

Table 2 Results with four classes

	PlantDoc	Custom dataset	PlantDoc + custom dataset
All categories	0.162	0.277	0.474
Spot leaf	0.133	0.25	0.359
Rust	0.243	0.467	0.649
Both diseases	0.165	0.150	0.284
Healthy	0.105	0.242	0.605

racy rises to 47% on the overall set, 60% on healthy leaves, and 64% on the rust. In the custom dataset, the spot leaf was not seen as much so the percentage for its prediction is lower since a large percentage of the annotations in the images were from rust.

The PlantDoc dataset alone cannot correctly predict the classes, so the combined use of the two datasets is necessary. The information added by the PlantDoc dataset offers more knowledge about plants and their characteristics. This allows the model to more accurately predict the plant class under consideration. Combining the two datasets also allows the model to become familiar with a greater variety of plants and develop a more comprehensive understanding of their characteristics. This can lead to improved model accuracy and performance in classifying plants into multiple classes.

Tables 1 and 2 present the results without the use of data augmentation techniques. In order to achieve greater accuracy in the results, we performed four more experiments. Of these, two experiments involved the two-class category and the other two experiments involved the four-class category. In this case, we used the mentioned techniques on the custom dataset with the aim of achieving more accuracy in the results.

In Table 3 we present the performance of the system with two classes and data augmentation techniques on the custom dataset. For the category with the two classes, the percentage of the custom dataset with the additional images rises by 34.2% and reaches 78%. By adding the healthy images from the PlantDoc dataset as well, the percentage increases to 79.2%, thus providing accurate predictions. Adding these images to the dataset significantly improved the performance of our model and enhanced the reliability of our predictions. In fact, the percentage of diseased plants reaches 82.9%.

In Table 4 we present the performance of the system with four classes and data augmentation techniques on the custom dataset. Regarding the four-class category, the percentage of predictions in the custom dataset increases by 25.1%, reaching 52.8%. The proposed model performs much better than the model that used the PlantDoc dataset (36.6% increase). By inserting the healthy images from the PlantDoc dataset, the prediction rate increases to 59.2%, which is an encouraging increase in accuracy, especially for the four classes. The addition of these images to the dataset significantly improved the performance of our model, enhancing the reliability of our bean crop and disease predictions.

Table 3 Results with two classes and data augmentation techniques

	Custom dataset	PlantDoc + custom dataset
All categories	0.789	0.792
Diseased	0.826	0.829
Healthy	0.744	0.757

Table 4 Results with four classes and data augmentation techniques

	Custom dataset	PlantDoc + custom dataset
All categories	0.528	0.592
Spot leaf	0.334	0.509
Rust	0.698	0.704
Both diseases	0.353	0.412
Healthy	0.72	0.751

Our experiments demonstrated that the YOLOv5 model, complemented by data augmentation techniques, could effectively identify diseases on bean leaves with high precision. The model exhibited a remarkable ability to focus on the most informative regions of the input images, a feature that proved invaluable in detecting subtle and localized symptoms of plant diseases. However, it is crucial to emphasize that the significance of these results was greatly enhanced through the incorporation of a specialized dataset tailored to the unique characteristics of bean leaf diseases. This underscores the importance of using domain-specific datasets to train machine learning models effectively. Furthermore, the multi-scale prediction strategy of YOLOv5 allowed it to detect symptoms of various sizes effectively, showcasing its versatility and applicability in real-world scenarios. The real-time detection capability of the model facilitated swift responses, a critical factor in agriculture where timely intervention can prevent widespread damage.

5 Conclusions

In conclusion, this study presents an innovative approach to tackle the challenge of accurately predicting diseases in bean crops despite the absence of a publicly available dataset. By leveraging the PlantDoc dataset with images from leaves of other plant species and incorporating new real-world images captured by a UGV from the bean crop, we were able to properly train the model with the YOLOv5 algorithm to adapt to the unique characteristics of the specific crop. In this way, we were able to achieve real-time disease detection on bean leaves with greater accuracy than state-of-the-art CNN models. Our approach demonstrates the potential of deep learning methods to overcome the limitations of small image datasets and achieve accurate disease detection in real time.

We conducted disease prediction experiments in two categories, encompassing two and four classes, both with and without the application of data augmentation techniques. Additionally, we employed the PlantDoc dataset, a custom dataset, and a combination of both datasets in separate trials. Our computational experiments yielded particularly promising outcomes, especially when utilizing the combined dataset. Specifically, for the 2-class category, we achieved an accuracy rate of 79%, while for the 4-class category, it reached 59%, with the application of data augmentation techniques. In contrast, without the use of data augmentation techniques, the accuracy was 71% for the 2-class category and 47% for the 4-class category.

Hence, we conclude that the most effective approach for achieving high accuracy in scenarios where public datasets lack sufficient images in certain categories involves two main features: utilizing an existing dataset to establish the model's initial features and incorporating new imagery into the final dataset, as well as the addition of a small expanded dataset from real images using data augmentation techniques to further specialize the model.

Acknowledgements This work has been co-financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation (project code MIS 5047196).

Author Contributions EK: Methodology, Formal analysis and investigation, Writing - original draft preparation; VB: Methodology, Formal analysis and investigation, Writing - original draft preparation; LT: Conceptualization, Writing - review and editing, Funding acquisition; NP: Conceptualization, Methodology, Writing - review and editing, Resources, Supervision.

Funding This work has been co-financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation (project code MIS 5047196).

Data availability Enquiries about data availability should be directed to the authors.

Declarations

Conflict of interest The authors have no relevant financial or nonfinancial interests to disclose.

Ethical approval Ethical approval was not required for this research.

Human participants and/or animals This article does not contain any studies with human participants or animals performed by any of the authors.

References

Abayomi-Alli OO, Damaševičius R, Misra S, Maskeliūnas R (2021) Cassava disease recognition from low-quality images using enhanced data augmentation model and deep learning. Expert Syst 38(7):1–21

Alguliyev R, Imamverdiyev Y, Sukhostat L, Bayramov R (2021) Plant disease detection based on a deep model. Soft Comput 25(21):13229–13242

Balafas V, Karantoumanis E, Louta M, Ploskas N (2023) Machine learning and deep learning for plant disease classification and detection. IEEE Access 11:114352–114377

Bedi P, Gole P (2021) Plant disease detection using hybrid model based on convolutional autoencoder and convolutional neural network. Artif Intell Agric 5:90–101

Boulent J, Foucher S, Théau J, St-Charles PL (2019) Convolutional neural networks for the automatic identification of plant diseases. Front Plant Sci 10:941

Cap QH, Uga H, Kagiwada S, Iyatomi H (2022) Leafgan: an effective data augmentation method for practical plant disease diagnosis. IEEE Trans Autom Sci Eng 19(2):1258–1267. https://doi.org/10.1109/TASE.2020.3041499

Chlap P, Min H, Vandenberg N, Dowling J, Holloway L, Haworth A (2021) A review of medical image data augmentation techniques for deep learning applications. J Med Imaging Radiat Oncol 65(5):545–563

Chug A, Bhatia A, Singh AP, Singh D (2023) A novel framework for image-based plant disease detection using hybrid deep learning approach. Soft Comput 27(18):13613–13638

Dai G, Fan J, Tian Z, Wang C (2023) PPLC-Net: neural network-based plant disease identification model supported by weather data augmentation and multi-level attention mechanism. J King Saud Univ-Comput Inf Sci 35(5):101555

Das D, Singh M, Mohanty SS, Chakravarty S (2020) Leaf disease detection using support vector machine. In: 2020 International Conference on Communication and Signal Processing (ICCSP). IEEE, pp 1036–1040

Diana Andrushia A, Mary Neebha T, Trephena Patricia A, Umadevi S, Anand N, Varshney A (2023) Image-based disease classification

in grape leaves using convolutional capsule network. Soft Comput 27(3):1457–1470
Enkvetchakul P, Surinta O (2022) Effective data augmentation and training techniques for improving deep learning in plant leaf disease recognition. Appl Sci Eng Prog 15(3):3810
Ferentinos KP (2018) Deep learning models for plant disease detection and diagnosis. Comput Electron Agric 145:311–318
Govardhan M, Veena M (2019) Diagnosis of tomato plant diseases using random forest. In: 2019 Global Conference for Advancement in Technology (GCAT). IEEE, pp 1–5
Haruna Y, Qin S, Mbyamm Kiki MJ (2023) An improved approach to detection of rice leaf disease with GAN-based data augmentation pipeline. Appl Sci 13(3):1346
He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 770–778
Huang G, Liu Z, Van Der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp 4700–4708
Islam MM, Adil MAA, Talukder MA, Ahamed MKU, Uddin MA, Hasan MK, Sharmin S, Rahman MM, Debnath SK (2023) Deepcrop: deep learning-based crop disease prediction with web application. J Agric Food Res 14:100764
Jiang P, Chen Y, Liu B, He D, Liang C (2019) Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. IEEE Access 7:77096–77107. https://doi.org/10.1109/ACCESS.2019.2914929
Jocher G (2022) ultralytics/yolov5: v3.1—bug fixes and performance improvements. https://github.com/ultralytics/yolov5. Accessed 8 Oct 2023
Kaggle (2018) Plantvillage dataset. https://www.kaggle.com/datasets/emmarex/plantdisease. Accessed 8 Oct 2023
Karantoumanis E, Balafas V, Louta M, Ploskas N (2022) Computational comparison of image preprocessing techniques for plant diseases detection. In: 2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM). IEEE, pp 1–5
Kaur P, Khehra BS, Mavi EBS (2021) Data augmentation for object detection: a review. In: 2021 IEEE International Midwest Symposium on Circuits and Systems (MWSCAS). IEEE, pp 537–543
Kaushik M, Prakash P, Ajay R, Veni S et al (2020) Tomato leaf disease detection using convolutional neural network with data augmentation. In: 2020 5th International Conference on Communication and Electronics Systems (ICCES). IEEE, pp 1125–1132
Li J, Qiao Y, Liu S, Zhang J, Yang Z, Wang M (2022) An improved YOLOv5-based vegetable disease detection method. Comput Electron Agric 202:107345
Li E, Wang L, Xie Q, Gao R, Su Z, Li Y (2023) A novel deep learning method for maize disease identification based on small sample-size and complex background datasets. Ecol Inform 75:102011
Liu W, Anguelov D, Erhan D, Szegedy C, Reed S, Fu C et al (2016) SSD: single shot multibox detector. In: European Conference on Computer Vision. Springer, pp 21–37
Mahum R, Munir H, Mughal ZUN, Awais M, Sher Khan F, Saqlain M, Mahamad S, Tlili I (2023) A novel framework for potato leaf disease detection using an efficient deep learning model. Hum Ecol Risk Assess Int J 29(2):303–326
Ramesh S, Hebbar R, Niveditha M, Pooja R, Shashank N, Vinod P et al (2018) Plant disease detection using machine learning. In: 2018 International Conference on Design Innovations for 3Cs Compute Communicate Control (ICDI3C). IEEE, pp 41–45
Ren S, He K, Girshick R, Sun J (2015) Faster r-cnn: towards real-time object detection with region proposal networks. In: Cortes C, Lawrence N, Lee D, Sugiyama M, Garnett R (eds) Advances in neural information processing systems, vol 28. Curran Associates, Inc
Saha S, Ahsan SMM (2021) Rice disease detection using intensity moments and random forest. In: 2021 International Conference on Information and Communication Technology for Sustainable Development (ICICT4SD). IEEE, pp 166–170
Shorten C, Khoshgoftaar TM (2019) A survey on image data augmentation for deep learning. J Big Data 6(1):1–48
Shrestha G, Das M, Dey N et al (2020) Plant disease detection using CNN. In: 2020 IEEE Applied Signal Processing Conference (ASPCON). IEEE, pp 109–113
Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556
Singh D, Jain N, Jain P, Kayal P, Kumawat S, Batra N (2020) PlantDoc: a dataset for visual plant disease detection. In: Proceedings of the 7th ACM IKDD CoDS and 25th COMAD. Association for Computing Machinery, pp 249–253
Singh AK, Sreenivasu S, Mahalaxmi U, Sharma H, Patil DD, Asenso E (2022) Hybrid feature-based disease detection in plant leaf using convolutional neural network, Bayesian optimized SVM, and random forest classifier. J Food Qual 2022:1–16
Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z (2016) Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 2818–2826
Zeng Q, Ma X, Cheng B, Zhou E, Pang W (2020) GANs-based data augmentation for citrus disease severity detection using deep learning. IEEE Access 8:177883–177895. https://doi.org/10.1109/ACCESS.2020.3025196
Zhang Z, Gao Q, Liu L, He Y (2023) A high-quality rice leaf disease image data augmentation method based on a dual GAN. IEEE Access 11:21176–21191
Zhao ZQ, Zheng P, Xu S, Wu X (2018) Object detection with deep learning: a review. IEEE Trans Neural Netw Learn Syst 30(11):3212–3232. https://doi.org/10.1109/TNNLS.2018.2869696
Zhu JY, Park T, Isola P, Efros AA (2017) Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision, pp 2223–2232