www.editada.org

International Journal of Combinatorial Optimization Problems and Informatics, 16(2), April 2025, 84-97. ISSN: 2007-1558.
https://doi.org/10.61467/2007.1558.2025.v16i2.870

Precision Agriculture and Object Detection: Deep Learning Models for Crop Disease Management

Eduardo A. Huerta-Mora[1], *Abraham E. Rodriguez-Mata*[2], *Jesús A. Medrano-Hermosillo*[2], *Leonel E. Amabilis-Sosa*[3], *Victor A. Gonzalez-Huitron*[4], *Héctor Rodríguez Rangel*[5]
[1] Facultad de Ingeniería y Tecnología de Mazatlán, Universidad Autónoma de Sinaloa
[2] División de Estudios de Posgrado e Investigación, Tecnológico Nacional de México-Instituto Tecnológico de Chihuahua
[3] División de Estudios de Posgrado e Investigación, SECIHTI-Tecnológico Nacional de México-Instituto Tecnológico de Culiacán
[4] Departamento de Sistemas y Computación, Tecnológico Nacional de México-Instituto Tecnológico de Querétaro
[5] División de Estudios de Posgrado e Investigación, Tecnológico Nacional de México-Instituto Tecnológico de Culiacán

25716.huerta@ms.uas.edu.mx; abraham.rm@chihuahua.tecnm.mx; jesus.mh@chihuahua.tecnm.mx; leonel.as@culiacan.tecnm.mx; victor.gh@queretaro.tecnm.mx; hector.rr@culiacan.tecnm.mx

Abstract. This study presents a comprehensive dataset designed for the visual detection of crop diseases, comprising 43,267 images of 12 crop species across 15 disease classes. The dataset was developed over 14 months of dedicated human effort. To evaluate its effectiveness, several plant disease detection and classification algorithms were tested. The models generated by these algorithms were deployed on mobile devices and specialized hardware, enabling practical applications ranging from drones to Android smartphones, with on-device detection capabilities.
The results highlight the performance of deep learning techniques, with the YOLOv4 algorithm achieving a mean average precision (mAP) of 71.04%, while the VGG model attained 92% precision and 90% accuracy. These findings demonstrate the potential of deep learning in enhancing crop monitoring, offering significant support for pest and disease control in vegetable crops. This work underscores the role of advanced technologies in promoting sustainable agricultural practices.

Keywords: Precision agriculture; Deep learning; Image Classification; Object Detection

Article Info
Received February 8, 2025
Accepted March 3, 2025

1 Introduction

From an economic and social perspective, the agricultural sector plays a crucial role in the Mexican region of Sinaloa, where it represents a significant activity. In 2019, Sinaloa had 1,058,758 hectares under cultivation, contributing 9.1% to the national total agricultural production value. This underscores the importance of agriculture in Sinaloa as a fundamental region of the country and a key driver of economic output.

Horticultural pests and diseases pose a significant challenge to agricultural activities in Sinaloa and across Mexico, as they can severely threaten crop productivity. According to (Velusamy et al., 2021), these issues have a profound impact on the agriculture industry, requiring the development and implementation of innovative solutions to safeguard crop quality and yield. In this context, precision agriculture has emerged as a promising approach that leverages technology to enhance agricultural practices (Memon et al., 2023). One notable example is the use of unmanned aerial vehicles (UAVs) for monitoring and early detection of crop diseases, enabling timely intervention (Velusamy et al., 2021; Mogili & Deepak, 2018). Additionally, high-resolution satellite sensors and other remote sensing technologies further enhance disease detection capabilities, highlighting the potential of technological advancements in addressing agricultural challenges (Zhao & Li, 2018).

The growing adoption of precision agriculture, particularly in developed countries, is evidenced by the increasing use of Global Navigation Satellite System-based technologies (Nowak, 2021). However, further technological advancements are deemed essential to maximize crop yields and mitigate the threats posed by pests and diseases effectively (Buddhi & Joshi, 2022; Brenes & Raventos, 2020). By enhancing productivity and sustainability, precision agriculture techniques have the potential to transform Sinaloa's agricultural landscape. This underscores the pivotal role of technology in securing the future of agriculture in regions heavily dependent on this sector, such as Sinaloa, and calls for a concerted effort toward the development and adoption of innovative agricultural technologies.

In 2022, Sinaloa cultivated a total of 1,029,978 hectares, contributing 7.9% to the national agricultural production value. The Mexican states with the highest contributions were Michoacán (12.7%), Jalisco (12.0%), and Sinaloa (7.9%) (Amarillas, 2022). This highlights the economic and social significance of agricultural activity in the region. However, the presence of horticultural pests and diseases can cause substantial damage if not addressed promptly, making the timely and accurate detection of crop diseases a cornerstone of precision agriculture.

Plants are increasingly threatened by dangerous diseases that significantly reduce the quality and quantity of agricultural products. Early detection and prevention of plant diseases are, therefore, crucial to mitigating these impacts. Traditionally, plant diagnostics rely on visual inspection by experts, with biological examinations used as a secondary option when needed. However, these methods are often costly and time-consuming (Cap et al., 2018).

Advances in technology have introduced various approaches to plant disease detection, including models based on artificial intelligence and other techniques (Mohanty et al., 2016; Liakos et al., 2018, González Huitrón et al., 2025). The frequency of pest and plant disease outbreaks is rising, posing a serious threat to food security. These outbreaks have far-reaching economic, social, and environmental consequences, jeopardizing the stability of the agricultural sector. Despite these challenges, farmers often struggle to detect plant diseases in a timely manner. Beyond consulting fellow farmers or utilizing resources like the Kisan hotline, their options for immediate intervention are limited.

Accurate identification of plant diseases requires specialized expertise, as well as laboratory infrastructure in many cases, to confirm diagnoses and identify diseased leaves effectively (Singh et al., 2020). This underscores the urgent need for accessible and efficient disease detection methods to support farmers and enhance food security.

With the progressive advancements in deep learning and its increasing adoption for detection and classification tasks, it is now possible to design a system that integrates deep learning techniques to process images of horticultural crops from the generated dataset. By combining these advances with specialized hardware and mobile devices, the system aims to enhance early detection of pests and diseases, contributing to improved agricultural practices.

The primary objective of this study is to develop a multidimensional horticultural dataset that will serve as a valuable resource for detecting diseases in diverse crops, greenhouses, and even urban gardens. Deep learning algorithms are evaluated, and the performance of popular models that can be implemented on low-spec hardware and mobile devices with low power consumption is compared. Consequently, a secondary objective of this study is to enable detection on mobile devices, such as Android smartphones, and specialized hardware, including the Raspberry Pi 4 microcomputer.

This manuscript presents the development of a deep learning-based dataset aimed at improving crop disease detection. Section 1 discusses the importance of agriculture in Mexico and the impact of diseases on crop yields. Section 2 reviews related research on the use of deep learning for plant disease diagnosis. Section 3 outlines the creation of a comprehensive horticultural dataset and the training and deployment of detection and classification algorithms on mobile devices and specialized hardware for practical applications. Section 4 evaluates the performance of the models using comprehensive metrics. Finally, Section 5 highlights the study's contributions to precision agriculture and provides recommendations for future research, emphasizing the potential of deep learning in crop disease management.

2. Related work

Deep learning has proven to be a powerful tool for diagnosing plant infections, as evidenced by several studies. This highlights the significant progress achieved in agricultural technology and disease control. Research, such as that by (Jakjoud et al., 2019;

Nagaraju & Chawla, 2020), has demonstrated the considerable potential of convolutional neural networks (CNNs) in this field. For instance, (Jakjoud et al., 2019) showcased the adaptability and efficiency of deep learning models through the implementation of the VGGnet16 architecture. Similarly, the efforts of (Militante et al., 2019; Ahmed & Reddy, 2021) have enabled the identification of various diseases across different plant species. Ahmed's (Ahmed & Reddy, 2021) mobile-based diagnostic system, in particular, has demonstrated practical applicability in real-world scenarios, achieving an impressive accuracy rate of 94%.

Within this research context, the work relevant to this study can be categorized into two primary areas: the compilation of datasets for plant disease detection and the exploration of techniques for disease diagnosis. The former involves the meticulous processes of data collection, curation, and annotation to produce comprehensive datasets. These datasets form the foundation for training and evaluating machine learning models, enabling the development of algorithms capable of identifying a wide range of plant diseases. Such datasets are critical for enhancing the robustness and reliability of these models, ultimately improving their practical utility in addressing agricultural challenges.

The high accuracy rates achieved in diagnosing and classifying plant diseases, as demonstrated by (Jasim & AL-Tuwaijari, 2020; Ferentinos, 2018), confirm the effectiveness of deep learning techniques. However, (Pradhan & Patil, 2019) emphasizes the need for continued research to enhance accuracy and efficiency due to the inherent complexity of plant disease detection. Additionally, (Bora et al., 2022) highlights the untapped potential of advanced equipment and methodologies to revolutionize disease detection processes.

This section examines the techniques and technologies employed in diagnosing and categorizing plant diseases, including the use of convolutional neural networks (CNNs), the development and optimization of deep learning algorithms, and the integration of mobile-based systems for field diagnostics. Each method contributes uniquely to the overarching goal of achieving high accuracy and efficiency in disease detection, underscoring the multidisciplinary nature of the field. The continuous refinement of these approaches, combined with advancements in computing power and improved data accessibility, holds significant promise for enhancing plant health management and ensuring the sustainability of agriculture.

2.1 Datasets focused on plant disease detection

In (Hughes & Salathé, 2015), the "PlantVillage Dataset" is introduced as an open repository containing 54,309 images of 14 crop species and 38 types of plant diseases. At the time of its release in 2015, it was one of the few publicly available datasets specifically designed for plant disease detection.

The dataset includes images of crops such as apple, blueberry, cherry, corn, grape, orange, peach, pepper, potato, raspberry, soybean, pumpkin, strawberry, and tomato. It documents 17 fungal diseases, 4 bacterial diseases, 2 mold diseases (oomycetes), 2 viral diseases, and 1 mite disease. Additionally, 12 of the crop species include images of healthy leaves that show no visible signs of disease. Figure 1 illustrates examples of the dataset's contents. This dataset has played a pivotal role in advancing research on plant disease detection and classification, serving as a foundational resource for training and evaluating deep learning models in agricultural technology.

Fig. 1. Sample images from the PlantVillage dataset.

In (Singh et al., 2020), the "PlantDoc Dataset" is introduced as a dataset designed for the visual detection of plant diseases. It consists of 2,598 images covering 13 plant species and 17 disease classes, yielding a total of 27 classes (17 disease and 10 healthy). The dataset was developed through approximately 300 hours of human labor, involving the annotation of images sourced from the Internet.

To demonstrate the dataset's effectiveness, 3 models were implemented for plant disease classification. The results revealed that using the PlantDoc dataset improved classification accuracy by up to 31%, showcasing its potential to advance plant disease detection. The authors suggest that this dataset can lower the barriers to adopting computer vision techniques in agricultural applications.

Compared to the PlantVillage dataset, which contains images captured in controlled environments, the PlantDoc dataset uses real-world images of healthy and diseased plants. PlantVillage's controlled settings, with uniform backgrounds and lighting, may limit its applicability in real-world scenarios where plant images often include multiple leaves, varying backgrounds, and inconsistent lighting. In contrast, the PlantDoc dataset incorporates such variations, making it more practical for real-world disease detection tasks. Figure 2 shows a comparison between the PlantVillage and PlantDoc datasets.

Fig. 2. Samples of various classes from the dataset PlantDoc show the difference between laboratory-controlled and real-life images.

In (Picon et al., 2019), a dataset was developed using images captured with a mobile phone under real-world conditions. This dataset includes over 100,000 images across crops such as wheat, barley, corn, rice, and rapeseed, encompassing 17 different types of diseases. Similarly, (Tani et al., 2018) describes the creation of a specialized dataset for cucumber leaves, which includes cases of multiple infections. The dataset comprises 48,311 images, including 38,821 leaves infected with one of 11 disease types, 1,814 leaves with multiple infections, and 7,676 healthy leaves. On the other hand, (Liu et al., 2019) developed a smaller dataset consisting of 247 images for experimental purposes. To train a support vector machine (SVM) classifier, 100 images were selected for training, while 72 images were used for validation and 75 for testing. From the training set, 207 tomato samples and 621 background samples were manually cropped to create a comprehensive dataset. To enhance the training data, augmentation techniques, including random rotations from 0° to 360°, were applied. Examples of these datasets are illustrated in Figure 3.

Fig. 3. Examples of the dataset. Top row: tomato samples; bottom row: background with leaves, stems and other objects.

2.2 Techniques for plant disease detection

In the paper by (Sankaran et al., 2010), the use of reliable sensors for monitoring plant health and disease under field conditions was proposed. However, the adoption of sensor-based disease detection remains limited due to the high cost of hardware and the lack of expertise in using such technology, which may restrict its usage. In contrast, (Patil & Bodhe, 2011) focused on extracting shape features for sugarcane leaf disease detection, achieving an impressive average accuracy of 98.60%. More recently, (Grinblat et al., 2016) explored the use of neural networks for identifying three different legume species by analyzing the morphological patterns of leaf veins.

3 Methodology

This section describes the methodology developed for this research, including the strategies used to create a multidimensional dataset derived from horticultural sources. This dataset will be valuable for the agricultural sector and for future research focused on detecting diseases in various crops and classifying leaves using deep learning methods, a subset of artificial intelligence. Additionally, the design of an application incorporating the models trained on this dataset is discussed. An overview of the research methodology is presented in Figure 4.

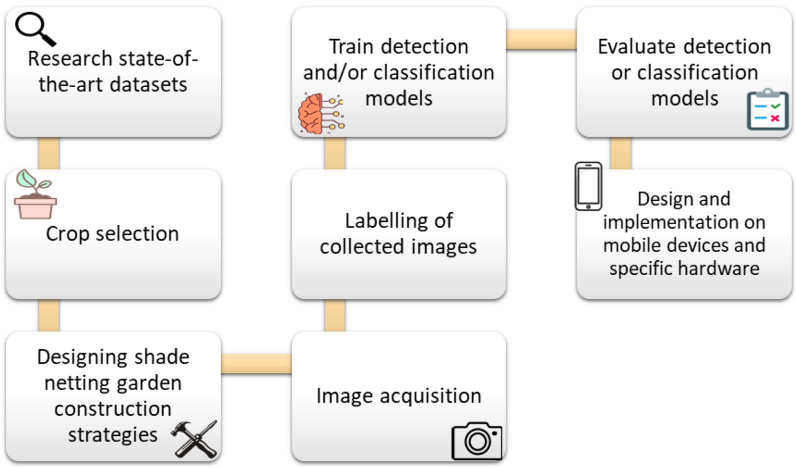

Fig. 4. Methodology.

3.1 Generation of the horticultural dataset

The process of obtaining an effective model for detection and/or classification depends on the data used during model training. This process begins with a preliminary review of existing datasets for foliar disease detection algorithms, such as the widely referenced PlantVillage dataset. Following this, the selection of crops for the dataset is made, focusing on vegetable crops that are commonly grown in the state of Sinaloa, Mexico.

3.1.1 Number of images obtained in the dataset

The dataset generated consists of 43,267 images across 27 different classes. The selected crops include 12 species: cucumber, radish, watermelon, melon, basil, pineapple, pumpkin, chiltepin chili, serrano chili, spinach, sweet potato, and tomato. The dataset includes images of 11 fungal diseases, 1 viral disease, 3 bacterial diseases, and healthy leaves or leaves without visible disease damage for each crop species. Table 1 presents the dataset with selected cultivars and the number of images collected for each class. An example of each class is shown in Figure 5.

Table 1. Healthy and diseased crops in the dataset

Crop	Number of Images
Basil blotch leafminer	2189
Basil downy mildew	1147
Basil healthy	3038

Chiltepin pepper early bacterial blight	919
Chiltepin pepper healthy	2706
Chiltepin pepper whitefly	1026
Cucumber angular leaf spot	854
Cucumber gummy stem blight	1140
Cucumber healthy	2928
Cucumber mosaic virus	656
Cucumber whitefly	1068
Melon angular leaf spot	901
Melon gummy stem blight	1037
Melon healthy	2786
Melon whitefly	1174
Pineapple healthy	2016
Pumpkin healthy	2511
Pumpkin leaf blight	1218
Radish fusarium rot	1089
Radish healthy	3037
Radish whitefly	623
Serrano pepper healthy	1050
Spinach healthy	1002
Sweet potato healthy	2848
Tomato healthy	314
Tomato whitefly	1029
Watermelon healthy	2961
Total	43267

Fig. 5. Sample images for each class in the horticultural dataset.

3.1.2 Equipment used to obtain images of the dataset

The image acquisition process was performed using various devices, as shown in Table 2. The images were captured at different times to ensure the dataset included sunny, darker, images with a controlled background, and images with a noisy background. This approach was intended to create a dataset that functions effectively in the agricultural sector and yields reliable results.

Table 2. Equipment used.

Device	Resolution (MegaPixels)
Smartphone Xiaomi Mi A3	32
Smartphone Huawei P20	24
Camera Canon EOS 250D	24.1

3.1.3 Labeling

Manual image labeling for object recognition is a common task, as it is essential for many supervised learning approaches. Consequently, various labeling tools have been developed (Fiedler et al., 2018).

To train an object detection algorithm, each image in the dataset must be labeled. For this purpose, annotations were created in both the YOLO and PASCAL VOC formats. Several tools are available for labeling images and generating annotation files in TXT format for YOLO and XML format for PASCAL VOC. Some of the most used tools include *LabelImg*, *makesense.ai*, and *Draw Bounding Box*, among others. A sample of image labeling using the *LabelImg* tool is shown in Figure 6.

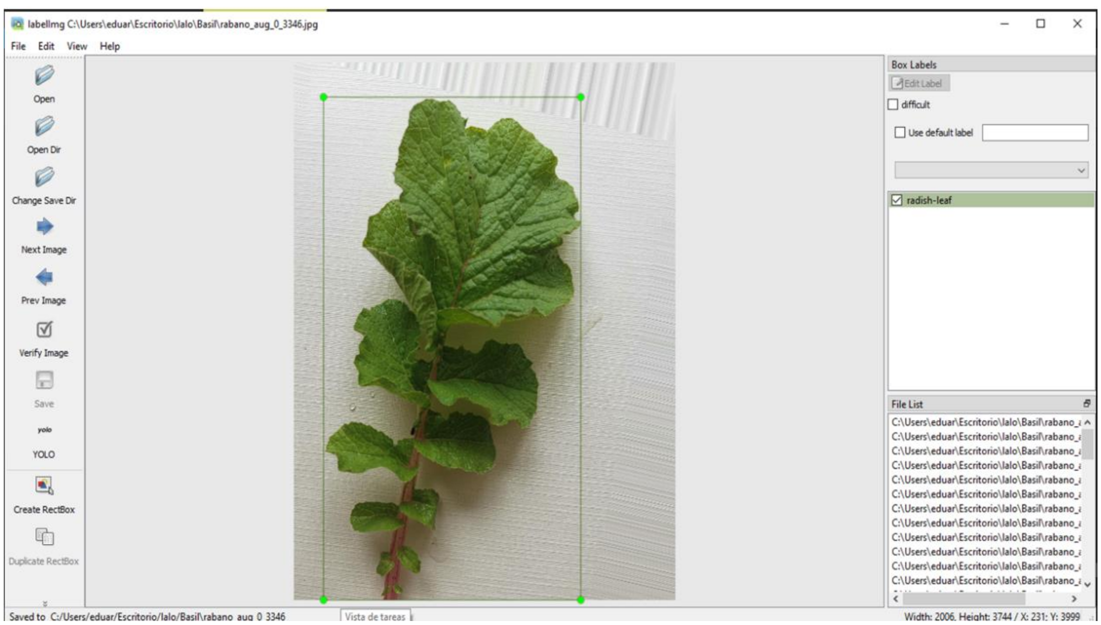

Fig. 6. Screenshot from annotation tool used for labeling

3.1.4 Dataset training

For the detection and classification algorithms, training and validation were performed on a computer equipped with an Intel i5 processor and an NVIDIA GTX 1050 Ti graphics card with 8GB of RAM.

The training process for the classification algorithms was conducted using Python, Keras v2.2.4 (Chollet, 2015), and TensorFlow v1.15.0 (Abadi et al., 2015), which provide a framework for designing and implementing convolutional neural networks (CNNs).

Applications and graphical visualizations were developed using the *matplotlib* library to facilitate the visualization of network activations and monitor the training progress.

The dataset, which contains images of both healthy and diseased leaves, was used for this classification model. It was divided into 80% for training and 20% for validation. The classification task was performed using convolutional neural networks, testing five different architectures. The characteristics of the architecture used are detailed in Table 3.

Table 3. Trained architectures summary

Network Architecture	Depth	Parameters (millions)	Dimensions input size
VGG16	16	15.6	224x224
InceptionV3	48	23.7	299x299
MobileNetV2	19	2.3	224x224
ResNet50V2	50	27.4	224x224
Xception	71	20.9	299x299

The detection algorithm is highly versatile, making it suitable for implementation on mobile devices across different performance ranges. This is due to its low power consumption and the fact that the number of parameters during training depends on the number of convolutional layers in each architecture. A lower number of parameters results in faster real-time detection but decreases the likelihood of successful detection. Various frameworks facilitate the deployment of detection algorithms on mobile devices, with TensorFlow Lite being the most commonly used for Android.

3.1.5 Implementations

The trained detection model was implemented in a mobile application for the Android platform, developed using the *Android Studio* desktop application. The mobile implementation enables farmers and agronomy experts to perform detections more efficiently and accurately using their smartphones.

Additionally, a Raspberry Pi microcomputer was used—a low-cost, single-board computer (SBC) designed to broaden access to computer technology and support computer science education in schools. For classification and image processing tasks, the Raspberry Pi Model 4 was selected. allowing low consumption implementation to image capture using either the official Raspberry Pi Camera Module (Pi-Camera) or a general-purpose USB webcam.

The Raspberry Pi operates on a GNU/Linux-based system, which includes a vast collection of software, most of which is distributed under free or open-source licenses (Pastor, 2019).

4 Results

This section presents the results obtained from the detection application and its implementation on a microcomputer using the developed dataset. Evaluation metrics used to assess the performance of the trained detection algorithms are provided, along with a comparison of the training results, which will help ensure a proper evaluation of the dataset. Additionally, the results of convolutional neural network architectures for classification are presented, using evaluation metrics derived from the confusion matrix.

4.1 Classification Model

Figure 7 shows a flowchart of the real-time classification process using the micro-computer Raspberry Pi. A brief introduction of the classification process using the 4 processes consisting of image acquisition, image pre-processing, training, and evaluation is shown in the image.

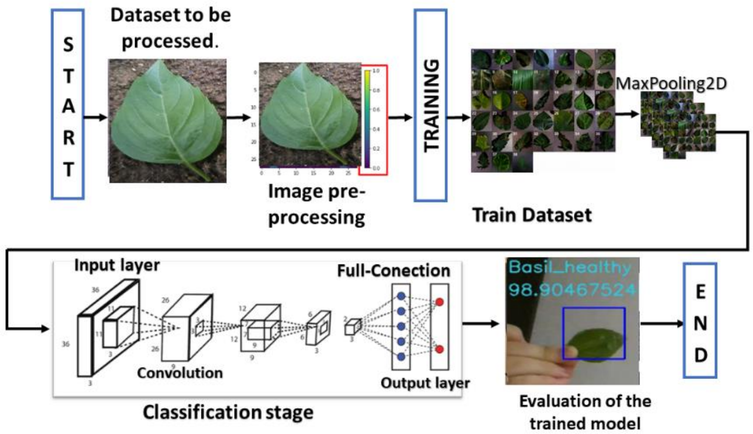

Fig. 7. Flowchart for the Raspberry Pi 4 implementation

A variety of evaluation metrics exist to measure a model's performance, each assessing different aspects of its effectiveness. Performance was evaluated by comparing pre-trained models using various metrics. Typically, the performance of pre-trained models is assessed using a test dataset.

Metrics such as Precision, Accuracy, Recall, and the F1-score were used to evaluate these models.
Precision, also known as "positive predictive value," is defined in Equation (1). It measures the degree of certainty in the model's predictions by indicating the proportion of predicted positives that are actually positive.

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive} \tag{1}$$

Recall, also known as "Sensitivity," measures how effectively the model identifies positive instances. In other words, it evaluates the model's ability to correctly predict positive examples while also quantifying its ability to minimize false negatives.
To calculate recall, the number of true positives is divided by the sum of true positives and false negatives, as shown in Equation (2).

$$Recall = \frac{True\ Positive}{True\ Positive + False\ Negative} \tag{2}$$

The F1-score, the measure of the efficiency of the trained model, is determined as the harmonic mean of the accuracy and recall. It focuses on the analysis of positive classes. A high value of this metric indicates that the model performs best in the positive class. It is defined in the equation (3).

$$F_1 = 2 * \frac{Precision * Recall}{Precision + Recall} \tag{3}$$

Typically, accuracy is the most commonly used metric to evaluate classification performance. This metric calculates the percentage of samples that are correctly classified and is represented in the equation (4).

$$Accuracy = \frac{True\ Positive + True\ Negative}{Total\ Samples} \tag{4}$$

At this stage, an evaluation of the pre-trained models for the plant disease classification task was performed on the different classes of the developed dataset. The results and training times are shown in the table 4.

Table 4. Testing performance (%) for each pre-trained model

Network Architecture	**Recall**	**Precision**
VGG16	85	92
MobileNetV2	62	77
InceptionV3	67	80
ResNet50V2	79	87
Xception	83	94
	Accuracy	F-Score
VGG16	90	90
MobileNetV2	61	64
InceptionV3	65	60
ResNet50V2	78	79
Xception	89	86
	Time (mins)	
VGG16	252	
MobileNetV2	134	
InceptionV3	208	
ResNet50V2	225	
Xception	240	

The VGG16 model has the highest recall value, indicating a low presence of false negatives across all classes. Conversely, a low recall value suggests a higher number of false negatives. Precision measures the proportion of correctly predicted positive cases, while recall quantifies the model's ability to detect actual positives. When these metrics are applied to multiple classes, the values obtained for each class are averaged in a weighted manner. The F1-score provides a balanced measure that combines precision and recall. When the dataset contains an equal number of images per class, precision and recall contribute equally to the F1-score. Accuracy represents the ratio of correct predictions to the total number of examples, expressed as a percentage, and reflects the model’s overall performance in a classification task.

Among the tested architectures, VGG16 achieved the highest performance across all metrics. Its superior accuracy suggests that it is the most efficient architecture for this specific task. The difference lies in the number of parameters, as VGG16 has a lower depth compared to other architectures. However, when considering processing time, MobileNetV2 demonstrated the best efficiency, as it required the shortest training time due to its lower number of parameters. A lower parameter count improves computational efficiency compared to other architectures, with only a minor trade-off in accuracy. This trade-off is acceptable in scenarios where higher processing efficiency is prioritized, particularly for real-time applications with limited hardware resources. For low consumption device implementation on the Raspberry Pi, the trained model was exported as an H5 file, enabling portability to the microcomputer. Figure 8 illustrates the real-time classification process using the Pi-Camera. Each transfer learning architecture was deployed on the Raspberry Pi, with classification results summarized in Table 5. To validate the training of these convolutional neural network architectures, random tests were conducted using various classes.

Table 5. Inference times on Raspberry Pi 4 B

Network Architecture	**Time (seconds)**
VGG16	22
InceptionV3	30
MobileNetV2	12
ResNet50V2	37
Xception	27

Fig. 8. Classification example result for Transfer Learning models

4.2 Detection model

The evaluation of the foliar disease detection model was conducted using metrics that assess the performance of different detection algorithms. For this task, the YOLOv4 architecture was selected and trained using the object detection framework *Darknet* provided by YOLO. The training process lasted approximately 238 hours and was carried out on hardware consisting of an NVIDIA 1050 Ti GPU, an Intel i5 8-core processor with a clock speed of 3.6 GHz, and 8 GB of RAM. The trained detection model was tested using images from the dataset, as well as images not included in the training process. This approach ensured an unbiased evaluation, preventing results that could artificially inflate accuracy. After training, the model was converted to a TFLITE format for deployment on an Android mobile device.

4.2.1 Mean Average Precision metric

The mean Average Precision (mAP) or sometimes simply referred to as AP is a metric used to measure the performance of object detection models. It is defined by equation (5):

$$mAP = \frac{\sum_{q=1}^{Q} AverageP(q)}{Q} \tag{5}$$

Where Q is the number of queries in the set and AverageP(q) is the average precision (AP) for a given query, q.

The mean Average Precision (mAP) for object detection is the average of the AP calculated for all classes. The mAP for a set of detections is the average of the interpolated Average Precision (AP) for each class, where the AP for each class is determined by the area under the precision/recall (PR) curve for the detections. The PR curve is constructed by assigning each detection to its most overlapping ground truth object instance (Henderson & Ferrari, 2017). AP (Average Precision) is a widely used metric for measuring the accuracy of target detectors. Mean accuracy calculates the average accuracy value for retrieval, ranging from 0 to 1. A comparison of mAP results for different detection models is presented in Table 6, and a comparison of training times for these models is presented in Table 7.

Table 6. Average Precision performance per class and detection model (%)

Crop	YOLOv4	YOLOv4-tiny	YOLOv3	YOLOv3-tiny
Basil blotch leaf miner	98.3	88.33	83.73	69.04
Basil downy mildew	98.97	89.83	86.78	75.84
Basil healthy	98.6	82.93	84.72	63.76
Chiltepin pepper early bacterial blight	97.99	0.00	0.00	0.00
Chiltepin pepper healthy	0.00	64.90	83.95	0.00
Chiltepin pepper whitefly	53.17	90.78	83.17	84.50
Cucumber angular leaf spot	89.92	98.39	96.96	96.37
Cucumber gummy stem blight	94.09	88.08	82.16	68.34
Cucumber healthy	88.45	92.46	90.70	87.94
Cucumber mosaic virus	98.42	91.09	87.06	88.22

Cucumber whitefly	0.00	98.09	98.36	97.44
Melon angular leaf spot	99.24	93.94	91.34	87.26
Melon gummy stem blight	95.00	2.13	0.00	0.00
Melon healthy	35.63	40.25	38.83	27.22
Melon whitefly	100.00	54.74	57.36	41.33
Pineapple healthy	46.90	30.73	24.15	18.79
Pumpkin healthy	86.10	87.34	83.06	81.11
Pumpkin leaf blight	54.10	82.14	74.67	71.59
Radish fusarium rot	1.47	0.00	0.00	0.00
Radish healthy	72.99	70.83	70.10	56.49
Serrano pepper healthy	27.10	89.45	86.70	74.26
Spinach healthy	93.89	82.45	74.16	66.62
Sweet potato healthy	93.83	71.07	56.12	46.62
Watermelon healthy	77.64	25.12	17.68	19.15
mAP	**70.90**	**67.29**	**64.66**	**55.08**

Table 7. Training time for detection models

Network Architecture	Time (hours)
YOLOv4	248
YOLOv4-tiny	55
YOLOv3	172
YOLOV3-tiny	38

4.2.2 IOU metric

The Intersection over Union (IoU) evaluation metric is commonly used to measure the overlap between two bounding boxes or masks, as well as to assess the accuracy of an object detector on a given dataset. Any algorithm that produces predicted bounding boxes as output can be evaluated using the IoU metric (Cowton et al., 2019).

In practice, it is highly unlikely that the (x, y) coordinates of the predicted bounding box will exactly match those of the ground truth bounding box. Due to various factors such as the image pyramid scale, sliding window size, and feature extraction methods, a perfect match between the predicted and ground truth bounding boxes is unrealistic.

Therefore, we need an evaluation metric that rewards predicted bounding boxes that overlap significantly with the ground truth. Figure 9 presents examples of both good and poor intersections over the union scores (Rosebrock, 2016).

Fig. 9. Different results and visual behavior for Intersection over Unions.

As shown, bounding boxes that are predicted and overlap more with the ground truth bounding boxes have higher scores than those with less overlap. This makes Intersection over Union (IoU) an excellent metric for evaluating custom object detectors.
A comparison of IoU results using different detection models is presented in Table 8.

Table 8. IOU for detection models

Network Architecture	IOU (%)
YOLOv4	62.23
YOLOv4-tiny	52.43
YOLOv3	64.53
YOLOV3-tiny	49.35

4.3 Mobile application

The process of obtaining a prediction for crop status begins with an image captured by the camera, as shown in Figure 10. The detection model then identifies whether the crop leaves are diseased or healthy and calculates the probability of detection for the identified image. A key limitation in the development of this research lies in the dataset, which was created under sanitary restrictions. The work was conducted using an urban garden with shade netting, but there was insufficient image collection from larger rural areas to capture better images under actual agricultural conditions.

In addition, to train very accurate models for disease detection, we might need a dataset with a larger number of images in each class and a larger number of crops. But, due to the lack of availability of public datasets and the absence of real scenarios for field work, our approach offers a feasible direction to address the current problem of disease detection.

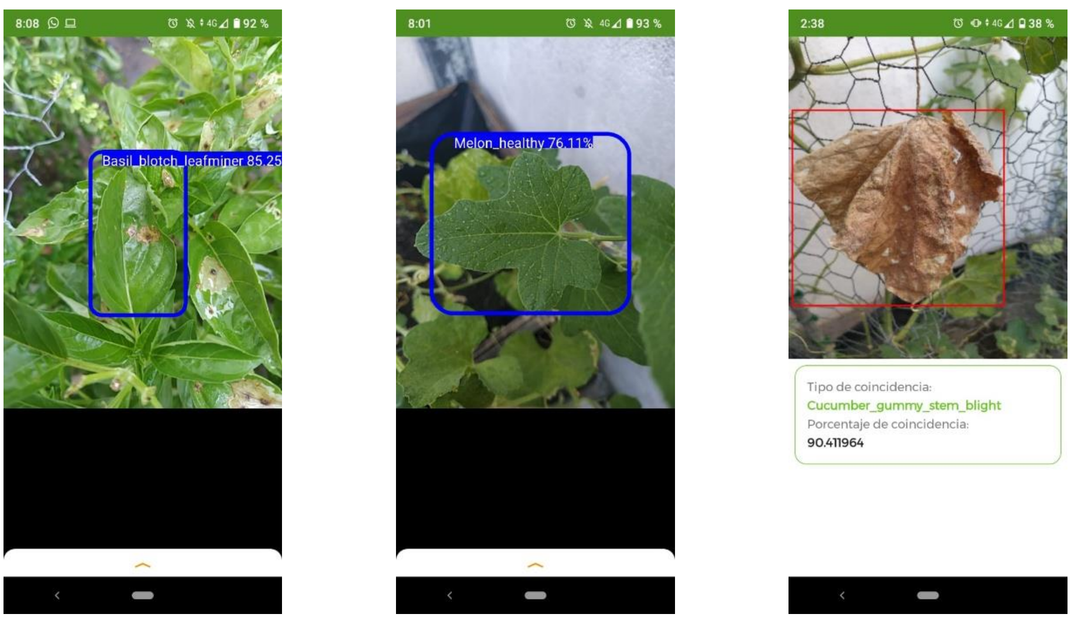

(a) (b) (c)

Fig. 10. Different results for object detection and classification from the implemented Android App

5 Conclusions

Precision agriculture is a critical sector for feeding the growing global population, which demands increasing amounts of food. Therefore, it is essential to grow healthy crops to meet this demand. Farmers in many regions still rely on rudimentary techniques to achieve sustainable harvests. However, these methods are not entirely reliable, as various factors—such as severe droughts, pests, viruses, bacteria, and burns—can significantly impact crops, leading to considerable economic losses.

This dataset was developed to support the agricultural sector and is available to anyone interested in creating systems capable of detecting foliar diseases. Although datasets developed in the last five years cover a limited range of crops, they are insufficient to capture the vast diversity of crops that have not yet been represented. Moreover, many of these datasets include images taken under suboptimal conditions, which could affect the accuracy of foliar disease detection.

This research presents a methodology for developing a dataset that will aid both the agricultural sector and home-based urban gardens. The methodology includes the implementation of an app capable of detecting foliar diseases in crops such as cucumber, radish, melon, basil, and chili chiltepin, among others, using deep learning algorithms.

An architecture based on You Only Look Once (YOLO) version 4 was used to detect plant leaves, and models were trained with various convolutional neural networks to identify diseases in the detected leaves. These models were integrated into an Android application for real-time detection on mobile devices.

Evaluations were conducted using metrics like mean average precision (mAP), which proved useful in demonstrating the efficiency of the trained models. The results were favorable, achieving an mAP of 70.90%. This percentage reflects the average of accuracies across all trained crops and is critical for ensuring accurate detection of healthy and diseased leaves. These results will be valuable in improving detection systems for crop health.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... & Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. Retrieved from https://www.tensorflow.org/
2. Ahmed, A. A., & Reddy, G. H. (2021). A mobile-based system for detecting plant leaf diseases using deep learning. *AgriEngineering, 3*(3), 478–493. Retrieved from https://doi.org/10.3390/agriengineering3030030
3. Amarillas, C. (2022, 6 de junio). Sinaloa en números (CODESIN). Retrieved from https://sinaloaennumeros.codesin.mx/agricultura-en-sinaloa-al-2022/
4. Bora, R., Parasar, D., & Charhate, S. (2022). Plant leaf disease detection using deep learning: A review. *International Journal of Engineering Research & Technology (IJERT), 11*(04), 1–8. Retrieved from https://www.ijert.org/research/plant-leaf-disease-detection-using-deep-learning-a-review-IJERTV11IS040001.pdf
5. Brenes, J. A., & Raventos, G. M. (2020). Research on optimization techniques applied to precision agricultural processes supported by technology in Latin America: Opportunities and a route to follow.
6. Buddhi, D., & Joshi, A. (2022). Computer and electronics in precision agriculture.
7. Cap, H. Q., Suwa, K., Fujita, E., Kagiwada, S., Uga, H., & Iyatomi, H. (2018). A deep learning approach for on-site plant leaf detection. *Proceedings of the 2018 IEEE 14th International Colloquium on Signal Processing & Its Applications (CSPA)*, 150–154. https://doi.org/10.1109/CSPA.2018.8368697
8. Chollet, F. (2015). *Keras*. GitHub. Retrieved from https://github.com/keras-team/keras
9. Cowton, J., Kyriazakis, I., & Bacardit, J. (2019). Automated individual pig localisation, tracking and behaviour metric extraction using deep learning. *IEEE Access, 7*, 108049–108060. https://doi.org/10.1109/ACCESS.2019.2933060
10. Fiedler, N., Bestmann, M., & Hendrich, N. (2018). ImageTagger: An open-source online platform for collaborative image labeling. En *Proceedings of the 17th International Conference on Computer Aided Systems Theory (EUROCAST 2019)* (pp. 162–169). Springer. https://doi.org/10.1007/978-3-030-27544-0_13
11. Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture, 145*, 311–318. https://doi.org/10.1016/j.compag.2018.01.009
12. González Huitrón, V. A., Medrano Hermosillo, J. A., Rodríguez Mata, A. E., Almaráz Damián, J. A., & Ramírez, P. C. (2025). Detección de enfermedades foliares en maíz mediante aprendizaje profundo. *Komputer Sapiens*, *1*, 40–44.
13. Grinblat, G. L., Uzal, L. C., Larese, M. G., & Granitto, P. M. (2016). Deep learning for plant identification using vein morphological patterns. *Computers and Electronics in Agriculture, 127*, 418–424. https://doi.org/10.1016/j.compag.2016.07.003
14. Henderson, P., & Ferrari, V. (2017). End-to-end training of object class detectors for mean average precision. En *Proceedings of the 14th European Conference on Computer Vision (ECCV 2016)* (pp. 198–213). Springer. https://doi.org/10.1007/978-3-319-54193-8_13
15. Hughes, D. P., & Salathé, M. (2015). An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*. Recuperado de https://arxiv.org/abs/1511.08060
16. Jakjoud, F., Hatim, A., & Bouaaddi, A. (2019). Deep learning application for plant diseases detection.
17. Jasim, M. A., & AL-Tuwaijari, J. M. (2020). Plant leaf diseases detection and classification using image processing and deep learning techniques.
18. Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. *Sensors, 18*(8), 2674. https://doi.org/10.3390/s18082674
19. Liu, G., Mao, S., & Kim, J. H. (2019). A mature-tomato detection algorithm using machine learning and color analysis. *Sensors, 19*(9), 2023. https://doi.org/10.3390/s19092023
20. Militante, S. V., Gerardo, B. D., & Dionisio, N. V. (2019). Plant leaf detection and disease recognition using deep learning.
21. Memon, K., Umrani, F. A., Baqai, A., & Syed, Z. S. (2023). A review based on comparative analysis of techniques used in precision agriculture.