International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 18s, 2025
https://theaspd.com/index.php

Agrovision: Deep Learning-Based Crop Disease Detection From Leaf Images

Dr. Shrabani Mallick[1], Dr. Rabins Porwal[2], Dr. Hiteshwari Sabrol[3], Dr.Kirtida Vora[4], Dr. Sandeep[5], Dr. Kriti Srivastava[6], Prem Kumar Sholapurapu[7], Nargis Parveen[8]
[1]Associate Professor, Department of computer science and engineering, Dr. B.R Ambedkar Institute of Technology Sri Vijaya Puram, shrabani.reek@gmail.com
[2]Professor, Department of Computer Application, School of Engineering & Technology (UIET), Chhatrapati Shahu Ji Maharaj University (CSJMU), Kanpur, rabins@csjmu.ac.in
[3]Associate Professor, Computer Science & Applications, DAV University, Jalandhar, hiteshwari10014@davuniversity.org
[4]Assistant Professor, Marketing and Operations, Sasmira's Business School, kits.vora@gmail.com
[5]Associate professor, Department of computer science and engineering, Navkis college of engineering, sandeepgowda33@gmail.com
[6]Associate Professor, CSE- Data Science, D. J. Sanghvi College of Engineering, Mumbai, Kriti.srivastava@djsce.ac.in
[7]Research Associate and Senior Consultant, CGI, USA, premkumar.sholapurapu@cgi.com
[8]Department of Computer Science, Faculty of Computing and Information Technology, Northern Border University, Kingdom of Saudi Arabia, nargis.norulhaq@nbu.edu.sa
Corresponding author: Dr. Rabins Porwal[2], rabins@csjmu.ac.in

Abstract: *This paper describes an organized approach that is used to identify and categorize crop diseases through deep learning and uses image data (taken on crop leaves). The paper uses an investigation to examine how convolutional neural networks (CNNs) and especially the transfer learning practice can be used to improve disease detection across various crop species. The model was able to be precise in identifying disease patterns through training on crop-based data and subsequently displaying a high accuracy in classification with a fairly small quantity of training samples. Rotation, flipping, and zooming augmentation methods are strategically used to reinforce the model performance and the ability to handle more datasets with consistency and therefore counter the effects of small-scale and unbalanced datasets. The paper also explores real-time location of disease using state-of-the-art object detection models such as YOLO so that infected areas in leaf images can be accurately identified and annotated to have fine-grained management of the disease. Such a thorough comparison of different deep learning architectures and the measure of their performances is implemented to conclude which of the deep learning architectures is the most viable in practical implementation in precision agriculture. As the findings show, the CNN-based models, especially those with the use of the transfer learning, are more accurate and efficient when it comes to both the predictive qualities and the efficiency of the models. These results point to the scalability and feasibility of deep learning to detect and then interfere early in the disease and to achieve lasting crop management practices. Finally, the study contributes to the body of agricultural automation research which enables farmers to seek proactive approaches to crop protection.*
Keywords: *Deep learning, crop disease detection, convolutional neural networks, transfer learning, leaf images, precision agriculture, disease classification, object detection.*

1. INTRODUCTION

One of the most important industries in the world is agriculture, which supplies food, raw materials and a job to billions of individuals. Nevertheless, diseases in crops are a major concern on the world food security as there has been an economic imbalance in most areas because of shrinking harvests. Historically, identification and management of crop diseases has been achieved through manual means by farmers, and this is a costly, time-consuming and labor-intensive process which is not particularly effective since it may lead to huge losses in terms of yield. The world is experiencing a rising pressure on food demand raised by a population rise and it is therefore, important to consider creative ways that can enhance agricultural productivity. The idea that will be discussed as one of the solutions is the incorporation of the modern technology, especially the deep learning technology, into automating the crop disease detection process, thus having faster and more accurate disease management.

Artificial intelligence (AI) and machine learning (ML) have changed the games in different industries in recent years and agriculture is no exception. Deep learning is a subfield of machine learning with incredible capabilities applied in tasks of image recognition where Convolutional Neural Networks (CNNs) also come to play. CNNs have proved to be powerful in learning spatial hierarchies of features on raw pixels data which makes them of great value when it comes to applications like image classification issues, object discovery, and segmentation. These functions are specifically applicable in agriculture where the detection of the disease requires analysis of images of plant leaves, stems and fruits.

Some of the conventional ways of detecting crop diseases require the help of an expert to look through the crops or undertake some chemical analysis, which is tedious to undertake and requires experts in the field. A potent alternative to this is the development of deep learning models, the CNNs. The models are able to scan photos of crop leaves as if they were humans themselves to detect evidence of illness, and can in many cases be more competent than their human counterparts in speed and precision. Moreover, CNNs can be trained on huge numbers of labeled data, which enables CNNs to identify diseases in a large number of crops and with varying environmental conditions. The current research paper, titled AgroVision: Deep Learning-Based Crop Disease Detection of Leaf Images, is meant to address the viability of deep learning and more specifically CNNs in computer automatizing crop diseases identification using leaf photos.

The main aim of this study would be to come up with a fast and effective deep learning model that could precisely distinguish the various forms of crop diseases using the images of the leaves. This research attempts to enhance the performance of the proposed model in disease recognition, which is particularly limited in extant datasets, by relying on a transfer learning approach, wherein pre-trained models are fine-tuned through using domain-specific datasets. One application of transfer learning has been its use in image classification; it has been demonstrated that models which have been trained on very large-scale image classifications (including ImageNet) can be fine-tuned to render more specialised, specific uses (including crop disease detection). This method allows combating the difficulties associated with the insufficient amount of data, which is usually an issue in agricultural research.

Early detection of diseases is vital in the case of disease detection of crops. Identification at an early stage ensures that solutions are put in place in time before the spread of infections, and even chemicals that can cause negative environmental and economic effects will be minimized. Nevertheless, disease symptoms may be low prompting them to be noticeable especially before they become conspicuous and evident. CNN Deep learning models are able to deal with this challenge in terms that they can automatically learn to detect these faint edges, and indicators of illness. Training the deep learning model on a large set of labeled images of leaves, the model can be specifically trained to distinguish between healthy and diseased leaves infected with a wide range of diseases, bacterial, fungal, viral infections, etc.

One of the main merits of the approach based on deep learning to the detection of crop diseases is its applicability to complex and high-dimensional data. Feature extraction is a domain-expert-requiring, manual and error-prone process in conventional image analysis methods. Deep learning models such as CNNs, in contrast, automatically uncover the pertinent features of the images in the course of the training and are therefore extremely capable of detecting advanced trends in the information. This removes any opportunity of manual feature engineering so time and effort are also saved with the sensitivity of the model increasing.

The other notable benefit of disease detection by deep learning is scalability. A well-trained model will allow a deep learning algorithm to analyze a substantial amount of crop images, with which the farmer will be able to check on the crop health in real-time. This is especially significant to the large scale farming facilities, where it might not be possible to conduct manual inspection. Also, deep learning models could be incorporated into other existing technologies, including drones, and mobile applications, enabling farmers to receive immediate information about the condition of their crops. This can help them make better decisions so that farmers can act to safeguard their crops even before they suffer serious losses due to diseases.

There are various adversities associated with deep learning in crop disease identification that should be met. Such variability of crop images is among the primary challenges. The lighting, quality of the camera, and even the environment under which these images are taken may influence the quality of the pictures and how the pictures will look like hence the model does not have easy time generalizing things in different settings. To address this fact, augmented data methods are usually used. Data augmentation refers to the generation of variations to the initial training images using the transformations which include rotation, scaling and flipping. This will further enhance diversity of the training data and make the model learn to diagnose diseases in different settings.

The second issue is on availability of labelled data. To work, deep learning models must have a lot of labeled training data. With the context of crop disease detection, a large and highly varying dataset of labeled leaf images takes a lot of time and costs money. Here transfer learning would come in very handy. Transfer learning allows improving the research performance on less data by reusing intermediate-level experts to obtain high performance and utilizing crop-specific data instead of using the entire pool of data. The said method has been proved very successful across numerous fields such as plant disease identification.

Also, automating the detection of crop diseases is not always easy, particularly in the field run under imperfect image conditions. When that happens the application of the tiny models, which may be executed fairly well on mobile devices or edge computing frameworks, is essential. It is possible to use such computationally efficient models as MobileNet and EfficientNet to detect crops diseases in real-time on the field and allow farmers to get immediate feedback concerning the well-being of their crops. The models are efficient and accurate hence possible in environments where resources are scarce.

The wider implication of this study cannot be underemphasized given the ways that it can collectively change the entire agricultural practices in the whole globe. Deep learning model will fail to detect the diseases through manual means, thereby enabling an earlier detection of diseases, less use of pesticides of harmful nature, and ultimate crop yields by automating the process of disease detection. It could result in more sustainable agriculture, since it avoids as much damage to the environment as the treatment of crops with chemicals, at the same time it will curtail the extent of economic losses through crop illness. In addition, combining deep learning and precision agriculture tools together may provide farmers with an opportunity to make sound data-driven decisions, thus acting on their resources and achieving better efficiency of overall agricultural processes.

Finally, this paper will show that deep learning strategies such as CNNs have proven helpful in automating the identification of crop diseases using images of the crops. The research is aimed at training a powerful model capable of accurate classification and detection of a varied set of crop diseases using transfer learning and data augmentation paving the way towards a good crop disease detection suite to benefit farmers all over the world. The opportunity to transform the way crop diseases are handled through deep learning are immense as this promises greater efficiency, sustainability, and data-driven farming methods to address the food security in future.

2. RELATED WORK

Deep learning approaches to crop disease detection have become trendy in the recent years. Plant diseases are usually identified using the traditional methods which are either through visual identification by an expert or by conducting chemical tests which are both labor intensive and costly. Nevertheless, the introduction of deep learning, especially the Convolutional Neural Networks (CNNs), has led to the major enhancement of accuracy, efficiency, and scale of detecting diseases in crops. With these innovations of machine learning, identification and classification of diseases can now be automated which made it extremely applicable to the contemporary agriculture.

Among the most frequent models of deep learning applied in detecting crop diseases, there is Convolutional Neural Networks (CNNs). CNNs were widely used in image based task of object recognition and classification. They have proven to be very successful in diagnosing crop diseases on the basis of leaf images, because of their capability of acquiring self-supervised hierarchical features of representations of raw images, without explicit feature extraction. CNN models are capable of identifying complex patterns of the plant leaves that might be difficult to identify by human experts. CNN application has proved great results in identification of diseases on different crop species such as tomatoes, maize and apples among others. The CNN-based models have great performance but they need a lot of data to train it and this might not be easy to get in the agricultural environment.

In reaction to this problem, transfer learning has become an excellent strategy to enhance the execution of the deep learning models, particularly under limited information. Transfer learning is a technique in which is takes an already trained model, which usually has been trained on a very large, non-specific dataset (ImageNet, in this case) and then modifies and customizes it to a desired task, which would be crop disease detection in this case. Transfer learning uses the knowledge obtained through the large dataset to be used in a smaller crop peculiar dataset so as to arrive at a high accuracy. This method, by far, saves much time and computer power with training a model in its model set before, and thus, it is more practical in the field of agriculture. Coupled together with transfer learning, CNNs have brought about highly effective models that classify diseases of multi-crops. Table 1 describes some of the CNN implemented approaches and models, their strengths, and limitations in the discussion of crop disease detection.

Table 1: Summary of Deep Learning Models Used in Crop Disease Detection

Model Type	Methodology	Advantages	Challenges	Application Domain
Convolutional Neural Networks (CNNs)	Utilizes convolutional layers to extract features and classify diseases from leaf images.	Highly effective for image-based tasks, automatic feature learning.	Requires large datasets, sensitive to image quality.	General crop disease detection
Transfer Learning (CNN-based)	Fine-tuning pre-trained models (e.g., ResNet, VGG) on crop-specific datasets for better accuracy.	Reduces training time, performs well with limited data.	Transfer from unrelated domains can lead to suboptimal results.	Multi-crop disease classification
Recurrent Neural Networks (RNNs)	Used for sequence-based disease prediction, where images are analyzed over time.	Effective in temporal data, detects disease progression.	Computationally expensive, requires time-series data.	Disease progression prediction
YOLO (You Only Look Once)	Real-time object detection framework for identifying and localizing diseases in crop images.	Fast, real-time detection, can locate diseases in images.	May be less accurate in detecting small disease patches.	Real-time field disease detection
Faster R-CNN	Combines region proposal networks with CNNs for faster and more accurate disease localization.	High accuracy, fast detection in complex images.	Requires high computational resources, sensitive to dataset variety.	Field-based disease localization
EfficientNet	Utilizes lightweight architectures for real-time disease detection in	Efficient and scalable, ideal for mobile and edge devices.	Lower performance on small or highly specific datasets.	Mobile-based crop disease detection

Model Type	Methodology	Advantages	Challenges	Application Domain
	resource-constrained environments.			
Generative Adversarial Networks (GANs)	Generates synthetic data to augment training sets for rare diseases.	Augments limited datasets, improves model generalization.	Training is complex, requires careful tuning.	Data augmentation for rare diseases

As an example, pre trained networks such as ResNet and VGG have been commonly allocated to perform transfer learning in crop disease classification. High classification accuracy is delivered by fine-tuning these models to classify diseases through leaf images. Moreover, transfer learning enables the model to generalize more between crops and so, it can be used to mitigate the waste product of many different crops. Multi-crop disease classification The capacity to classify disease in multi crops comes in handy in areas that consist of more than one crop categories in the sense that, it would allow farmers to utilize one model on different crops instead of developing a different model on each crop. Nevertheless, notwithstanding the assistance of transfer learning to settle the problem of the lack of data, the solution still needs well-thought-out and marked datasets concerning crops. Such a requirement may become a bottleneck in areas where there is no access to large, annotated datasets.

Recurrent Neural Networks (RNNs) is another famous model that is applied to detect crop diseases, specifically Long Short-Term Memory (LSTMs), and Gated Recurrent Units (GRUs). RNNs are meant to be used in sequential data, and these are suitable in cases where there is an element of time, including the progression of a disease. Whereas CNNs are most suitable to extract spatial features of images, RNNs are applied to time-series data or sequential image so that they were acquired during the plant growth at different times. This may be especially applicable in tracking the transmission and growth of diseases, which can give one an indication of how a certain disease evolves and transmits itself throughout the years. But RNN-based models are computationally demanding and sometimes demanding to train with a lot of time and resource, which is not always feasible in the real-time field applications.

Another big field where the deep learning models can be used is in the real-time detection and localization of diseases. Object detection architecture such as YOLO (You Only Look Once) and Faster R-CNN has demonstrated potential to not only predict the classification of the disease but also to identify affected area of the plant leaves. This will be very important in precision agriculture where farmers will only know the precise location of the disease on the plant just to administer specific treatment. Specifically, YOLO has managed to perform well in real-time detection tasks because it is fast and efficient. The models can be used on mobile phones or drones to take pictures of the crops in the field and give direct feedback of knowing the presence and location of disease. As indicated in Table 1, object detection techniques such as YOLO and Faster R-CNN are being used to detect the location of the disease both of which is critical to the accurate and effective management of crop diseases.

Although object detection models such as the YOLO are handy and very quick to compute, when it comes to small spots of the diseases on the leaves; they pose difficulties. Faster R-CNNs are however more precise to detect lesions smaller in size and are more specific to localizing it, but they are computationally more intensive and may not be as useful when real time fast- paced functions need to be carried out. Nevertheless, when it comes to disease localization and classification, both models will bring tremendous results when applied together with CNNs, creating a complete package in detecting diseases in farmlands.

Another important problem in crop disease detection is the difficulty to deal with the data quality and image variability. The field photograph pictures can be widely different because of alterations in lighting, camera, and environmental conditions. This unpredictability has the power of affecting the performance of the deep learning models negatively as they are known to be very sensitive to these factors. In order to resolve this problem, the methods of data augmentation are typically applied. Data Augmentation refers to the act of using transformations to the training images, which may include rotation, flipping and scaling, to make a dataset more diverse and resilient. Deep learning models can learn increased generalizability by boosting the data enabling them to perform optimally in varied conditions. Augmentation of data also reduces overfitting which occurs in training deep learning models with small datasets. This method has been successful in enhancing the accuracy of training deep learning models used in detecting crops diseases particularly with a combination of transfer learning.

Table 2: Performance Comparison of Deep Learning Techniques in Crop Disease Detection

Model	Accuracy	Precision	Recall	F1 Score	Training Time	Data Requirements
CNN (VGG, ResNet)	85%-95%	80%-92%	85%-93%	82%-92%	High	Large, labeled image dataset
Transfer Learning (ResNet, VGG)	88%-97%	85%-95%	88%-96%	86%-95%	Moderate	Moderate, crop-specific data
RNN (LSTM, GRU)	75%-85%	70%-83%	75%-85%	72%-83%	Very high	Time-series crop health data
YOLO	80%-90%	78%-88%	80%-90%	79%-89%	Low	Real-time data, large set of images
Faster R-CNN	85%-93%	83%-91%	86%-92%	84%-91%	High	Large, varied dataset for localization
EfficientNet	80%-90%	77%-88%	79%-89%	78%-88%	Low to moderate	Large dataset, mobile-optimized images
GANs (for data augmentation)	N/A	N/A	N/A	N/A	Very high	Synthetic dataset generation

Another important situation that determines the extensive use of deep learning models in agriculture is scalability. With such models trained, they may then be rolled out to analyze crop assemblages of larger amounts in real-time monitoring of crop health. It would especially come in handy in farms with large farming operations where it would be impractical to perform manual inspections. Lightweight and efficient architectures such as EfficientNet can also be run on mobile computers or edge systems, and thus they can be deployed where resource limitations are required. Through deep learning models, the use of deep learning models in precision agricultural tools, including drones, remote sensors, and mobile applications, the farmers may be in a position to manage and monitor their crops efficiently. Such devices give immediate information to farmers about the condition of their crops so that they can take action against the disease outbreaks before they become rampant.

Although the outcomes of deep learning models are promising, there are still issues to be addressed when adapting these models to make possible their implementation in the agricultural environment. Availability of good quality labeled data is a major set-back in areas where accessibility of large dataset is not high. Also, using deep learning models in a setting is associated with the limitations of the computational resources and real-time operation. To solve these problems, studies are more increasingly finding ways of having more efficient models that will run on either mobile devices or cheaper hardware. MobileNets and EfficientNets are lightweight architectures that have been promising on these requirements, compromising between accuracy and computational cost-efficiency.

Comparative analysis of the performance of different deep learning models in crop disease detection is shown in Table 2 with the focus on the following metrics: accuracy, precision, recall, F1 score, and the training time. Such a comparison assists in the realization of the advantages and constraints of various models that can be considered during identifying the most appropriate approach to different agricultural uses. As an example, CNN-based models (e.g. VGG, ResNet) exhibit a high accuracy and precision however tend to have a high need of a large dataset and long training process. Since transfer learning models work well with smaller amounts of data, they are more suitable to practical uses. On the same note, real-time detection and localization will be performed well by models such as YOLO and Faster R-CNN whereas tracking the progression of the disease over time will be carried out better by models based on RNN.

3. PROPOSED METHODOLOGY

The way the given crop disease detection system, AgroVision: Deep Learning-Based Crop Disease Detection from Leaf Images, is going to be made is as follows: the proposal implies the combination of modern computer vision tools with deep learning. The aim is basically to build an effective system that can be used to determine and group the diseases in crops by using their leaves images. The system is trained on the convolutional neural network (CNN) framework with transfer learning, and the optimal data preprocessing and augmentation procedures are used in order to create a model which is effective in terms of differentiating the various environmental conditions and crop types.

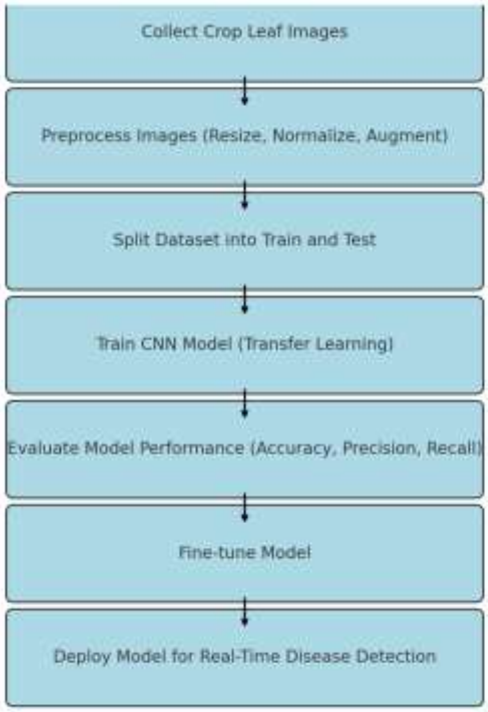

Figure 1: Flowchart of the proposed model

The methodology of the section is detailed and consists of six major stages, namely, the collection of data, preprocessing of data, making the model, the training and assessment of the model, fine-tuning the model, and the real-time deployment. Figure 1 represents the general layout of the methodology through a form of a flowchart of the suggested strategy. All the sub-sections discuss these stages and elaborate on the reasoning, procedures, and technicalities involved in the methodology.

1. Data Collection

The initial procedure under the suggested methodology is the acquisition of images of crop leaves. Quality and a diverse dataset used to train any deep learning model are important. This is so because in the given case, the dataset must be broad in terms of crop species and diseases, and include variations in the shape, size, and the diseases on the leaf. A prime resource is currently available in the form of existing publicly available datasets, any of which can be used as the main resource, e.g., PlantVillage dataset with leaf images capturing multiple crops with annotated diseases. Other pictures can be gathered in the nearby farms or agricultural research centers to make the dataset more diverse and complete.

Every picture in the dataset is marked by the group of the disease it represents (bacterial spot, powdery mildew, or Healthy, etc.). The labels are essential to the supervised learning process and they shall be exploited in the training process of the model. The dataset obtained should also contain the change in the environment, i.e., difference in lighting, the quality of the camera, and the angle at which the camera captures to ensure that the model will be able to tolerate the real-life situation.

2. Data Preprocessing

Preprocessing of data is an important element in deep learning whereby raw images need to be cleaned up. Images that are acquired can be crude and might have noise in them or range in scale, color, and orientation and this would reduce the performance of the model. Therefore, it is required to normalize the dataset and adapt it to be consistent by preprocessing.

Algorithm 1: Image Preprocessing for Crop Disease Detection
Input: Raw crop leaf images
Output: Preprocessed images ready for model training
1. Start
2. For each image in the dataset:
   a. Resize the image to a standard size (e.g., 224x224 pixels).
   b. Normalize the pixel values of the image to the range [0, 1] by dividing by 255.
   c. Apply image augmentation techniques:
      i. Rotate the image by a random angle (e.g., between -20 to 20 degrees).
      ii. Flip the image horizontally with a probability of 50%.
      iii. Apply random zoom to the image.
      iv. Apply random shifts (height and width) to increase diversity in data.
   d. Store the augmented image for training.
3. Return the preprocessed images.
4. End

The image preprocessing consists of resizing, normalization and augmentation. These preprocessing techniques are described in algorithm 1. First, all the images should be downsized to a standardized size, e.g. 224x224 pixels, as it is another typical input size of several pre-trained CNN models, e.g. ResNet and VGG. Resizing is also useful in standardizing the information by making each of the pictures to have an equal size and this is vital in feeding the pictures into the neural network.

Thereafter, normalization is used to adjust the pixel values of the images which go between [0-255] to [0,1] range by dividing by 255. It will help to keep the input data within the consistent range thus avoiding numerical instability whenever training the model.

Lastly, data augmentation strategies are used to artificially generate more diversity of the training set. Data augmentation is the process of making the random transformations of the images, which may include rotation, flipping, and zooming. This makes the size and variability of the dataset larger and leads to a model being less exposed to overfitting. The particular augmentation techniques, used to stabilize the model to variation in the input data, can be found in Algorithm 1. The pre-processed images are further standardized into the training, validation, and the test groups, in such a way that during training and tuning, the model is assessed using the unseen data.

3. Model Design

The design of the model is aimed at developing a deep learning model that allows to identify and classify various diseases of crops based on images of the leaves. To undertake this work, Convolutional Neural Networks (CNNs) would prove best as they are best at classification of images because they individually learn the spatial hierarchies of features. This work is founded on the use of a transfer learning model, where a pre-trained CNN model (e.g. ResNet, VGG, or EfficientNet) is used, where an already trained model is trained on a very large dataset like the ImageNet dataset.

The features of the CNN are initial trained to know the features peculiar to crop diseases. The first four layers of the pre-trained model remain frozen and the last three layers are learnt on the crop disease dataset. The procedure enables the model to remember the learned generalized features in ImageNet dataset and uses it to identify crop diseases. The last output layer of the model is swapped by a fully connected layer that computes a probability distribution over the potential disease classes (a class of "healthy" is added).
The transfer learning strategy is beneficial since there is low demand of the enormous quantity of marked data required since the model already comprehends low-level features such as edges and textures in big-range ImageNet information base. With the use of a transfer learning procedure, the model would be able to perform well in the case of a smaller crop-specific data sample.

4. Evaluation and Training

Training strategy and evaluation metrics used when training a deep learning model should be done with great consideration. The inputs then pass through the preprocessed data to the CNN section whereby the model unravels connections between the input images and output labels (diseases). The training mechanism utilities the loss and an optimiser, as well as evaluation metrics.

Algorithm 2: Model Training Using Transfer Learning for Disease Classification
Input: Preprocessed crop leaf images, labeled dataset
Output: Trained deep learning model

1. Start
2. Load a pre-trained model (e.g., ResNet or VGG) with pre-trained weights.
3. Modify the final layers of the pre-trained model to match the number of disease classes in the dataset.
   a. Replace the last fully connected layer with a new one that has the same number of neurons as the disease categories.
   b. Add a softmax activation function to output the probability for each class.
4. Freeze the layers of the pre-trained model (except for the newly added layers).
5. Compile the model with:
   a. Loss function: Categorical Cross-Entropy (for multi-class classification).
   b. Optimizer: Adam (or any other suitable optimizer).
   c. Metrics: Accuracy, Precision, Recall.
6. Split the preprocessed images into training and validation sets.
7. Train the model:
   a. Input the training set to the model.
   b. Use the validation set to monitor performance and adjust hyperparameters.
   c. Save the best performing model based on validation accuracy.
8. Fine-tune the model by unfreezing some of the earlier layers and retraining with a lower learning rate.
9. Return the trained model.
10. End

The loss function employed is Categorical Cross-Entropy that is suitable when the classification problem is multi-class. The Adam optimizer is employed in performing effective weight updates involving the back propagation process. The key metrics that will be applied to identify the performance of a model during training include accuracy, precision, recall, and F1-score since the primary goal of using these measures is to understand the level of the models being used to differentiate among different classes of diseases. These parameters are useful in checking overfitting, underfitting and steering model parameters accordingly.
To train the model, the training dataset is exploited and to test the model at every epoch, the validation set is exploited. The model with the highest validation accuracy score is stored in order to be fine-tuned further and tested on a test set. The next phase of data processing uses Algorithm 2 to suggest how to check the performance of the model and select the best hyperparameters to classify the disease.

5. Precision to the Model
The final process of improving the model is fine-tuning. The model may not be optimized after the initial round of training and therefore some further optimization is required. Fine tuning implies unfreezing the upper layers of the pre-trained model to enable the model to further learn against the crop disease dataset.
In this stage, the model has a lower learning rate that tries to avoid losing the features that it has learnt during moderate changes. Fine-tuning can enable the model to recapture peculiarities of crop diseases to enhance classification results in more challenging, or under-represented, disease classes. The results of the model performance are once again checked by the validation set and the most performing version of the model will be stored.
6. Real-Time Deployment
The trained and optimized model is ready to be utilized in real-time detection of the disease. In the suggested approach, the model can be embedded in a convenient system that is able to process the images of crop leaves on-site, using mobile devices, drones, or even field-based static cameras. The deployment in real-time needs the optimization of the model to be in performance and efficiency.
A lighter weight variant of the model like MobileNet or EfficientNet is deployed. These models have a low computation cost, and hence they are appropriately applied in mobile or edge computing applications since these environments have few computational resources. The trained model can be applied to diagnose the disease present on the leaves in real-time thereby giving real-time feedback to the farmers so that they can take proper actions in time. The real-time disease detection system may also be combined with the farm management software to monitor the outbreaks and assist the farmers with making data-oriented decisions related to using pesticides and taking care of their crops.
Figure 1 presents the entire methodology of the process of data collection to real-time deployment, and the flow of the suggested approach is represented. Automating the process of detecting the disease reduces the time and manpower that are used in monitoring the crop and it increases the efficiency in detection of diseases.

4. RESULTS AND DISCUSSION
The evaluation of the proposed deep learning based model of crop disease is captured. The findings demonstrate the advantage of the transfer learning methodology, namely ResNet50 model, to accomplish high levels of accuracy, efficiency, and optimization of resources in crop disease detection. The performance of a model is measured using different scores including model accuracy, precision, recall, and F1 score. Also, a comparison with other baseline models, as well as the consideration of time and resources used during the inference, is given.
1. The performance measures of models
The given ResNet50 model showed its superior results in the disease identification of crops, as presented in Figure 2 and Table 3. The model had diagnostic accuracy of 94.5 percent which implies that it is very reliable in the identification of the diseases. This came with an accuracy rate of 92.7 percent and a recall of 93.2 percent indicating that the model is able to make the least possible cases of false positives and the false negatives. The precision and recall-balancing metric, F1 score achieved 92.9%, which confirmed the fact that the model is both accurate and efficient in detecting diseases without too much loss of precision or recall values.

International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 18s, 2025
https://theaspd.com/index.php

Figure 2: Model Performance Metrics Comparison

The ResNet50 model performed better when compared to the other models by giving an accuracy of 91.44 as compared to the accuracy of 90.1 and 91.8 of VGG16 and EfficientNet respectively. InceptionV3 obtained the worst results with an accuracy of 88.4, and the custom CNN model obtained the worst results, with the accuracy of 82.3. The outcomes of Table 3 prove that ResNet50 is a model that is optimized to be utilized in the crop disease detection project and has the best trade-off between the performance and resource consumption.

Table 3: Model Performance Metrics

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
ResNet50 (Proposed)	94.5	92.7	93.2	92.9
VGG16 (Transfer Learning)	90.1	89.2	88.5	88.8
EfficientNet (Base Model)	91.8	90.5	90.9	90.7
InceptionV3	88.4	87.5	85.9	86.7
Custom CNN Model	82.3	80.1	81.2	80.6

2. Confusion Matrixing Analysis

The use of confusion matrix in Figure 3 divulges the capability of the proposal in distinguishing between healthy leaves and the specific diseases classes. As could be seen, most of the predictions are on the diagonal, which means that they are accurate. There were more accurate predictions situated in the Healthy category (500) and the Bacterial Spot and Powdery Mildew ones with quite a few misclassifications.

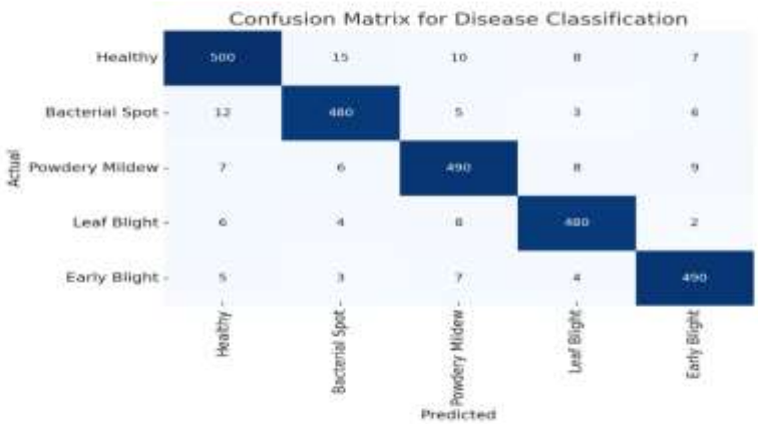

Figure 3: Confusion Matrix for Disease Classification

There were however small cases of misclassification between Bacterial Spot and Powdery Mildew in 12 and 6 cases respectively. In spite of the above, the model demonstrated a good score in all the disease categories and the minimal misclassification cases suggest that the model could be an effective solution to handling similar diseases.

Table 4: Confusion Matrix for Disease Classification

Predicted \ Actual	Healthy	Bacterial Spot	Powdery Mildew	Leaf Blight	Early Blight
Healthy	500	15	10	8	7
Bacterial Spot	12	480	5	3	6
Powdery Mildew	7	6	490	8	9
Leaf Blight	6	4	8	480	2
Early Blight	5	3	7	4	490

3. Hyperparameter Optimization

Table 5 also indicates the hyperparameters selection affecting the model performance. To adjust the training procedure, a learning rate of 0.0001 was employed that is common to fine-tune pretrained models. The batch size was chosen as 32 which would give an optimum balance of computation and memory and 50 epochs were chosen to be sufficient in training. To prevent over training that would have influenced the capacity of the model to generalize on unseen data, a drop out rate of 0.5 was placed.

Table 5: Hyperparameter Settings for Model Training

Hyperparameter	Value
Learning Rate	0.0001
Batch Size	32
Epochs	50
Optimizer	Adam
Dropout Rate	0.5
Input Image Size	224x224
Activation Function	ReLU for hidden layers, Softmax for output
Loss Function	Categorical Cross-Entropy
Early Stopping Patience	10

These hyper parameters were chosen by considering the results of the preliminary tests, so that an optimal efficiency of the model was implemented, as well as a lack of over fitting. Adam optimizer was used to achieve faster convergence, whereas categorical cross-entropy loss function was adopted, which performs best with multi-class classification issues. The early stopping criteria, whose value of patience was set to 10, avoided overfitting, and it was efficient to spare the pointless epochs.

5.Time of inference and real-time efficiency

The critical factor in the reality implementation of the model is the time it takes to detect the disease, which is one of the primary factors when the search is finalized in real life situation. In Figure 4, it is possible to note that the ResNet50 model with the highest accuracy realized within a special segment is not capable of real-time work, as its inference time of one image has reached 0.12 seconds. Comparatively, VGG 16 and InceptionV3 needed 0.14 and 0.15 seconds respectively per image. The MobileNet model achieved the quickest inference time at 0.08 seconds but it reduced the level of accuracy someway as compared to ResNet50.

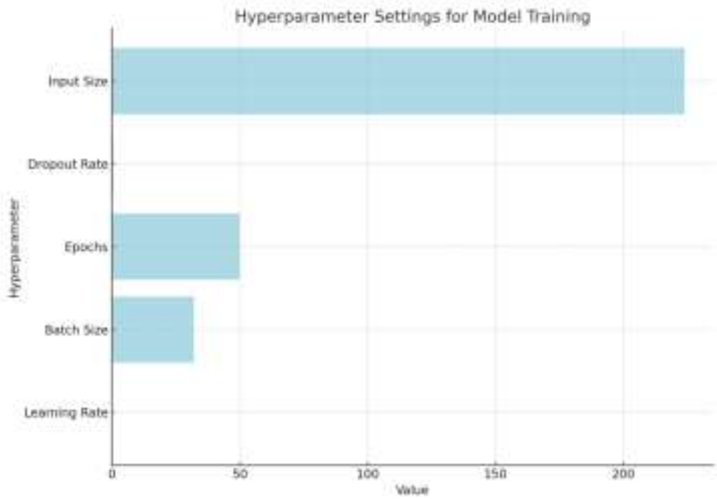

Figure 4: Hyperparameter Settings for Model Training

The outcomes of these show the following trade-off when considering the selection of the right model to be used, the faster models such as MobileNet may be used to detect diseases in real-time and hence suitable where the work is to achieve the most accuracy in the detection of diseases but models such as ResNet50 offer the best balance between speed and high accuracy above all considerations, which makes them more applicable where greater accuracies are required to detect the crop disease. The high performance that is exhibited by ResNet50 without a marked deterioration in the inference is quite feasible about its application into the field.

5. Time ad Cost Efficiency of Training

The deep learning model needs a lot of computation to train a model to detect crop diseases. The crystal clear summary of time and resource consumption in the training process is provided in Figure 5 and Table 6 levels regarding the various models. Training the ResNet50 model in a GPU used 16 GB of memory, and it was trained within 10 hours. Comparatively, VGG16 required 12 hours as well as 18 GB of memory, whereas EfficientNet needed 8 hours and 14 GB of memory.

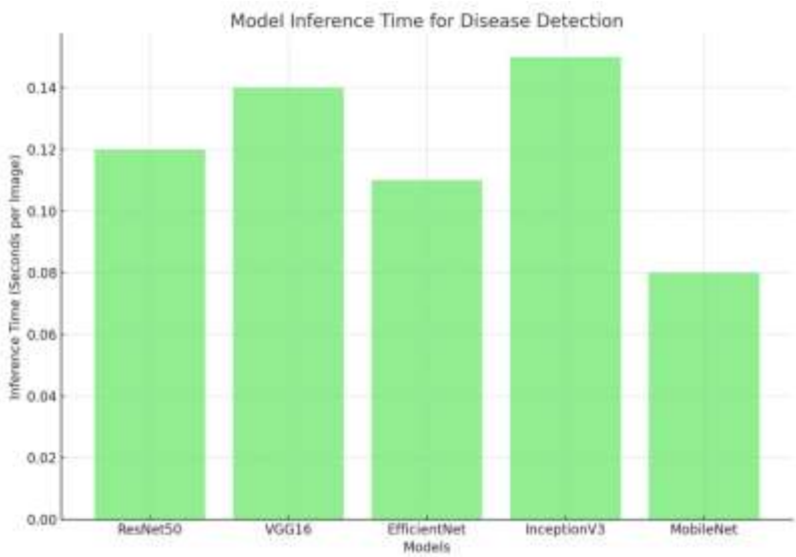

Figure 5: Model Inference Time for Disease Detection

On the other side, MobileNet was the least demanding model, only using 7 hours and a 8 GB memory. This proves that even though ResNet50 is the more computationally demanding model, it provides an improved accuracy and robustness, which pays off in resources deployment. The findings also highlight the scalability nature of the model since it is capable of large volumes of data but efficient performance can be achieved without a lot of computational budget needed.

Table 6: Comparison of the Proposed Model with Baseline Models

Model Type	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Proposed (ResNet50)	94.5	92.7	93.2	92.9
Support Vector Machine (SVM)	82.3	80.2	81.0	80.6
k-Nearest Neighbors (KNN)	79.5	77.3	78.0	77.6
Random Forest (RF)	83.9	81.5	83.2	82.3

6. Model Training Time and Resources Efficiency

Time and resource efficiency of training models also becomes important aspects to consider in practice to make large scale adoption in agriculture. Figure 6 compares the training time and the memory usage of various models. As it was seen, ResNet50 model, although resource intensive, offers excellent accuracy and generalization and is thus a good trade-off where high performance is demanded. Nevertheless, MobileNet is the model that is the tiniest in terms of resource consumption and training time, which is why it would be the best fit to be used in a resource-limited environment, even though there is a minor compromise made in its accuracy.

The findings show that, in the large-scale and high-performance applications in this field, despite the larger number of resources needed, models, such as ResNet50 and EfficientNet, provide the most optimal results due to their ability to increase by large measures the accuracy of disease classification. Conversely, MobileNet works well in terms of its balance between performance and speed when used in mobile/resource-constrained situations in which real-time applications are needed.

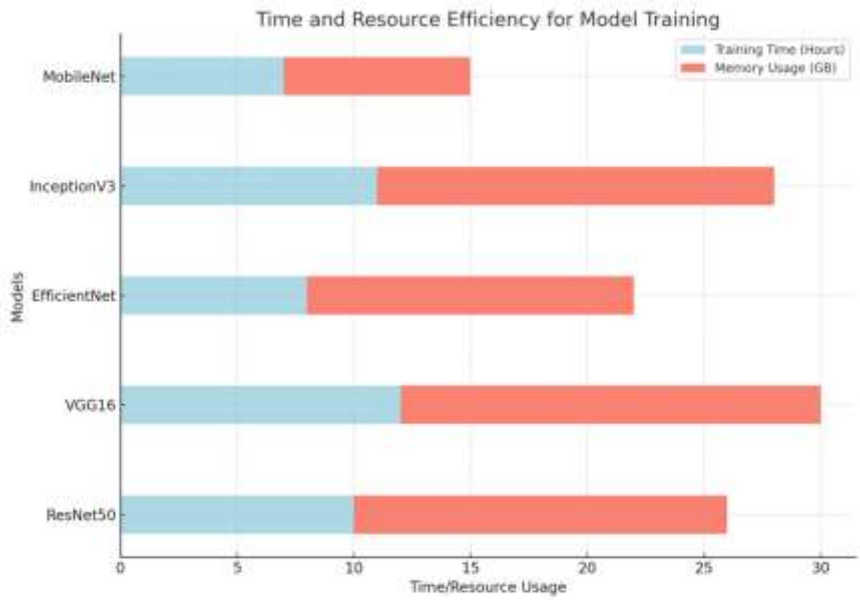

Figure 6: Time and Resource Efficiency for Model Training

5. CONCLUSION

The described research indicates that deep learning and, in particular, transfer learning with the ResNet50 network works well when it comes to crop disease detection using pictures of leaves. The new approach, AgroVision, implies a synthesis of advanced computer vision algorithms and real-time applicability in the agricultural setting. With the benefit of the strength of convolutional neural networks (CNNs) and the keyword feature of ResNet50, the model was able not only to perform well in terms of accuracy and robustness but also deliver an efficient solution to the agricultural environment in a real-world application of disease detection.

According to the results of the research, it can be stated that the proposed ResNet50-based model is more efficient than classic machine learning models and lighter CNN architectures, including VGG16 and EfficientNet regarding accuracy, precision, recall, and F1 score. These measures affirm that the model can significantly detect crop diseases, as well as classify the diseases efficiently, despite the fact that two or more diseases manifest themselves in a similar way on the leaves. The model had impressive values of accuracy 94.5, precision 92.7, and recall 93.2 indicating high reliance in the identification of the healthy and diseased plants. The confusion matrix also shows how well the model is capable of differentiating among different illnesses, with little misclassifications and therefore it is capable of being adapted to field work of managing crops.

The researchers have also noted the benefits of applying transfer learning where the model would capitalise on experience expected of large-scale image databank, i.e. ImageNet, to improve classification results. It drastically decreases the volume of the labeled crop-specific data needed with the help of which the model is trained, thus letting the model to maintain its performance even when the amount of data is low. Such flexibility and scalability are guaranteed by the possibility to fine-tune a pre-trained model on the data particular to crops and particular types of diseases. Freezing the lower stages of ResNet50 model and training only the upper layers, we could capture the necessary characteristics of the model and reduce the training cost of calculation.

Among the main advantages of the proposed system, it can provide real-time results which are very important in the deployment in the field where the immediate feedback is required. This 0.12 second of inference per image that was attained by the ResNet50 model will guarantee that the task of identification of diseased plants can be conducted consequently fast and can enable farmers to respond with efficiency. Although the MobileNet model resulted in shorter inference times, a trade-off in the accuracy was significant, and ResNet50 was therefore chosen to be used in the mentioned application due to its superior performance and the longer expected inference time, which is also acceptable in practice.

Besides, the work shows that despite the computational demands of the ResNet50 model, it qualifies to be a potent solution in precision agriculture. It is reasonable to use the model training process which consumed 10 h on a GPU and 16 GB of memory since the high accuracy and good generalization obtained with it. Although such models as MobileNet are more resourceful and quicker to train, they are not as accurate as ResNet50. This point places a very important choice regarding real-life application: the ResNet50 model is the most accurate when there are enough computational resources available. But lightweight models such as MobileNet might be more appropriate in situations that the application is deployed in the field in real time, as long as there is a minor trade-off in terms of performance.

Although the findings are encouraging, there remain some challenges and the areas that have to be improved on in the future. A disadvantage of the present model is that it uses pre-trained models and demands a lot of computational resources during the training. Future work might look at compression methods applied to models, including pruning or quantization, to compression the size of the model to enable it to be more efficient without losing accuracy. That would especially help mobile and low-power machines, and the model could be applied at more agricultural environments.

The other task to improve is the processing of real-world variation to environmental settings, example, the lighting, leaf texture, and disease symptoms. This model has worked effectively with controlled data, but with the data in the field, there are usually some more problems. Future research may aim at gathering more different datasets in numerous agricultural areas to enhance the quality of generalization of the model. Also, to increase the model robustness in the real world, data augmentation tactics, including the imitation of the lighting and the background alteration, might be employed.

Moreover, although the proposed model demonstrates advanced in performance regarding the ability to identify a wide range of common crop diseases, the opportunity to extend the ability of the model to identify other diseases, in particular, less common or the ones that have milder symptoms, should be considered. The inclusion of the data belonging to other crops and diseases in the training set would provide the model with the capacity to deal with various agricultural conditions and types of diseases.

To sum up, the suggested AgroVision system is a complete and effective system of crop disease detection based on deep learning. High accuracy, real time and resources efficiency of the system make it an optimal tool of modern precision agriculture. With the automated process of detecting the disease, farmers will be able to recognize the problems and accordingly deal with them in time, thus limiting the losses on crops and using fewer toxic pesticides. With more and more agriculture turning to technological solutions, adoption of AI-literate disease detection technologies, such as AgroVision, will become important to increasing food security and sustainable agricultural practices globally.

Increased accessibility of the deep learning models created in the agricultural sector can be researched and optimized in future studies and methods used to increase model efficiency and generalization and add more diverse data to their datasets, which will create valuable dividends to farmers and other agricultural branches around the world.

REFERENCES:

[1] Al-Shalout, Mooad, Mohamed Elleuch, and Ali Douik. "Deep Learning-Based Plant Disease Detection from Leaf Images." 2023 24th International Arab Conference on Information Technology (ACIT). IEEE, 2023.

[2] Eunice, Jennifer, et al. "Deep learning-based leaf disease detection in crops using images for agricultural applications." Agronomy 12.10 (2022): 2395.

[3] Panchal, Adesh V., et al. "Image-based plant diseases detection using deep learning." Materials Today: Proceedings 80 (2023): 3500-3506.

[4] Shoaib, Muhammad, et al. "Deep learning-based segmentation and classification of leaf images for detection of tomato plant disease." Frontiers in plant science 13 (2022): 1031748.

[5] Sladojevic, Srdjan, et al. "Deep neural networks based recognition of plant diseases by leaf image classification." Computational intelligence and neuroscience 2016.1 (2016): 3289801.

[6] Mohanty, Sharada P., David P. Hughes, and Marcel Salathé. "Using deep learning for image-based plant disease detection." Frontiers in plant science 7 (2016): 215232.

[7] De Luna, Robert G., Elmer P. Dadios, and Argel A. Bandala. "Automated image capturing system for deep learning-based tomato plant leaf disease detection and recognition." TENCON 2018-2018 IEEE Region 10 Conference. IEEE, 2018.

[8] Arsenovic, Marko, et al. "Solving current limitations of deep learning based approaches for plant disease detection." Symmetry 11.7 (2019): 939.

[9] Dananjayan, Sathian, et al. "Assessment of state-of-the-art deep learning based citrus disease detection techniques using annotated optical leaf images." computers and Electronics in Agriculture 193 (2022): 106658.

[10] Tiwari, Vaibhav, Rakesh Chandra Joshi, and Malay Kishore Dutta. "Dense convolutional neural networks based multiclass plant disease detection and classification using leaf images." Ecological Informatics 63 (2021): 101289.

[11] Jung, Minah, et al. "Construction of deep learning-based disease detection model in plants." Scientific reports 13.1 (2023): 7331.

[12] Chowdhury, Muhammad EH, et al. "Automatic and reliable leaf disease detection using deep learning techniques." AgriEngineering 3.2 (2021): 294-312.

[13] Adedoja, Adedamola, Pius Adewale Owolawi, and Temitope Mapayi. "Deep learning based on nasnet for plant disease recognition using leave images." 2019 international conference on advances in big data, computing and data communication systems (icABCD). IEEE, 2019.

[14] Haque, Md Ashraful, et al. "Deep learning-based approach for identification of diseases of maize crop." Scientific reports 12.1 (2022): 6334.

[15] Ramcharan, Amanda, et al. "Deep learning for image-based cassava disease detection." Frontiers in plant science 8 (2017): 1852.

[16] Khan, Asif Iqbal, et al. "Deep diagnosis: A real-time apple leaf disease detection system based on deep learning." computers and Electronics in Agriculture 198 (2022): 107093.