International Journal of Trend in Scientific Research and Development (IJTSRD)
Special Issue on Smart Innovations in Computer Science and Applications
Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

Deep Learning-Based Crop Disease Detection Systems for Precise Classifications of Leaf Images and Early Diagnoses in Agriculture

Hitesh Asutkar
Department of Science and Technology, G.H.Raisoni Skill Tech University, Nagpur, India

ABSTRACT
Crop disease has a significant impact on agricultural productivity that to is reduction in yields and financial loss for farmers. Therefore, the detection of crop diseases should be done early and accurately in order to manage crops and practice sustainable agriculture. The purpose of this project is to present a system that uses Deep Learning (Convolutional Neural Networks) to automatically detect crop diseases using images of leaves. This system uses a pre-trained AlexNet model which is extremely accurate at identifying many crops by utilizing leaf images that display visible clues/signs. The system is delivered to end users as an easy to use web application using the Flask framework; this will allow farmers and agronomists to upload images of leaves and instantly receive a prediction about the condition of those crops. The real-time prediction given to the end user through this system will be used to make data-driven decisions about precision farming and ultimately reduce crop losses and increase the health of crops. This project is an example of how AI-based technology will transform the agricultural industry through timely monitoring and intervention for crop diseases.

Keywords: Crop diseases, Deep Learning, Convolutional Neural Networks, Web application.

1. Introduction

Rural development, food security, and the strength of the world's economies depend on agriculture worldwide; however, plant diseases threaten the availability of food products and the viability of businesses involved in food production. Plant diseases can be caused by many factors, including fungi, bacteria, viruses, and environmental influences. Many of the visible symptoms of plant diseases occur on a plant's leaf and often, if left untreated or if diseased plants are not identified in a timely manner, will decrease the quality and quantity of a farmer's crop(s) and increase his/her overall production costs. Historically, disease diagnosis of plants has been achieved by visually inspecting plants, either through direct observation by the farmer, or through consultation with an agricultural expert using traditional, manual methods. This type of diagnosis is both slow and labour intensive and is open to human error. Additionally, plant pathologists who could assist farmers in identifying the disease of their crops are not available to them in most agricultural communities, and many communities providing agricultural goods are unable to access adequately trained professionals.

Since the advent of artificial intelligence methods, specifically deep learning techniques, there has been an increase in both the degree of automation achievable and the accuracy with which automated systems can diagnose plant diseases. One example of the implementation of these technologies for plant disease diagnosis is the use of Convolutional Neural Networks (CNNs).

CNNs can be used for classifying images and recognising patterns; and recent studies have demonstrated that they are sufficiently powerful to accurately identify the visible symptoms of plant abnormalities from the images of plant leaves. In addition to these advances, CNNs can integrate a large number of images of several different crop plants to assist researchers in the development of models to allow accurate diagnoses of diseased crops, and ultimately provide the diagnosis to producers of food.

This study is focused on using CNN-based deep learning to create a fast and effective crop disease detection system. This model uses transfer learning via an existing AlexNet network to quickly converge and achieve better performance with few labeled examples. The model is trained on a variety of leaf images that represent different types of crops with their associated diseases making it possible to accurately classify the conditions of the crops at a high level of accuracy.

The model is deployed as a web application developed using Flask which allows users (farmer, agronomist, agricultural extension worker) to upload images of leaves at the time they take the photo and receive instant feedback about disease diagnosis in order to make timely decisions. This model is designed to enable farmers to establish a more proactive approach to managing their crops and provide a reduction in how often they need to consult with experts; thus, this model helps to support the future of sustainable agriculture by helping decrease the number of unnecessary crop losses.In conclusion, this project illustrates how modern technologies such as deep learning can be used in agriculture today. By automating the process of detecting diseases and providing easy-to-use digital tools to support disease diagnosis, this project will have an overall positive impact in support of the development of precision agriculture. As well as demonstrating how AI tools can enhance the ability to monitor crops better, use less, and overall support the global effort to ensure food security as agriculture faces new and increased challenges.

1.1. Motivation

Agriculture has always been the foundation for numerous economies around the globe and very importantly contributes to security of food worldwide. However, there are countless crop diseases which jeopardize sustainably producing food and create significant yield losses every year. Farmers, especially those in developing regions, frequently rely on manual inspections when diagnosing plant diseases. This method is not only very time-consuming but also depends heavily on the knowledge and skill of an expert. When a farmer misdiagnoses or does not identify a disease quickly, it often results in diseases spreading to more plants, increasing pesticide usage, reducing quality of crops, and costing them a lot of money.

Due to rapidly advancing digital technologies, there is an increasing need to apply intelligent systems to agricultural practices. Current methods for detecting disease do not meet the needs of current-day farmers due to large farm size, rapid propagation of disease, and climate change. As such, a need exists for solutions that provide fast, accurate, and scalable solutions to detect diseases. By using deep learning within image-based classification of items, researchers are using a much better means to not only detect the disease at a higher degree of accuracy; they also have the ability to automate diagnosing for diseases.

The desire to equip farmers with technologic tools for early identification of illness is the main driving force behind this project. In this way, convolutional neural networks (CNNs) and transfer learning (using pre-trained models like AlexNet) will allow the system to accurately identify disease characteristics from images of the leaves. This capability will be made available through an easy-to-use web application that will enable any person (by not being required to have technical knowledge) to use this sophisticated artificial intelligence model.

The goal of this effort is to not only enhance agricultural yields and decrease loss of crops but also to expand the use of precision agriculture technology (early identification allows for intervention and timing for use of pesticides, thus enabling the agriculture operator to produce healthy crops while being a steward of the land). This project will address the current demand for intelligent solutions in the area of agriculture that provide support to farmers in terms of improving yield through the adoption of best practices and long-term availability of food in light of increasing challenges globally.

1.2. Contribution

1. A Convolutional Neural Network (CNN) will classify disease on crops using an existing CNN model called AlexNet that has already been pre-trained. The model will be able to extract features from a leaf image and use them to classify multiple diseases with a high level of accuracy.
2. By utilizing the transfer-learning method of incorporating previously trained data to improve the performance of the new model, the model can achieve high accuracy in identifying diseased crops, even with a limited data set. The use of transfer learning also helps to reduce the overall computational cost and amount of time required to train the system. The overall reliability of the system is maintained through the use of transfer-learning.
3. A web interface will be created that uses the Flask platform to make this system accessible to the largest number of potential users possible. Farmers and agronomists will be able to upload an image of a leaf through the web application and receive an accurate prediction of the possible disease. Therefore, the web application will be able to provide a link between the use of advanced artificial intelligence (AI) and the needs of the agricultural community.
4. Use of AI to Improve Precision Agriculture Practices
5. Information generated by the service is available in real time to support supervised action and ensure that data-driven decisions promote the efficient use of resources, reducing crop loss and promoting better management of crop health.
6. It Shows That Automated Image-Based Disease Detection Can Reduce Reliance on Manual Inspection or an Expert for Disease Detection and Provide a Scalable Solution Applicable to Various Crops and Production Systems.
7. It Supports Sustainable and Smart Farming Initiatives
8. The tools created enable earlier detection of diseases, which limits unnecessary use of pesticides and provides a potential increase in crop yield, thereby supporting sustainable farming practices in line with smart farming objectives.

2. Related work

Recently, there has been a significant increase in the research of crop disease detection and diagnosis, largely due to the availability of large agricultural image data sets and advancements in deep learning. Previous efforts to automate the classification of plant diseases using machine learning algorithms (e.g., Support Vector Machines (SVMs), K-Nearest Neighbors (KNN), Decision Trees) relied solely on feature extraction techniques to create feature sets based on colour histograms, texture measurements, and shape measurements. While these methods achieved some degree of success, their accuracy and robustness were limited because of the inherent variability in the symptoms of different types of diseases in different species of plants.

The emergence of Convolutional Neural Networks (CNNs) has changed the way in which researchers study and apply CNNs to detect crop diseases. CNNs remove the need for the manual extraction of feature representations from the input images by automatically learning the hierarchy of the feature representations from the raw image data. One of the first comprehensive studies that utilized CNNs for the classification of plant diseases was conducted by Mohanty et al (2016). In their study, they utilized deep neural network architectures such as AlexNet and GoogleNet to classify multiple plant diseases using the PlantVillage image data set. Their findings indicated the significant effectiveness of CNNs in identifying crop disease, and established the foundation for additional research into the use of deep-learning techniques in agriculture.

Numerous current studies have focused on improving accuracy of plant disease classification models by employing alternative model architectures and training methods. The 2018 Ferentinos et al. study analyzed several "deeper" model architectures for classifying plant diseases (e.g. VGG16, ResNet, and Inception) and achieved higher results than shallower models for plant disease classification accuracy. Because many researchers use transfer learning methods with pre-trained models (e.g. trained on large datasets such as ImageNet) to help improve accuracy of their plant disease classification models due to limited availability of labeled agricultural datasets, many researchers have performed data augmentation, batch-normalization and fine-tuning techniques to increase their model's generalizability and reduce overfitting to their plant disease datasets.

In addition, researchers from many research groups have begun to develop practical applications for plant disease classification models. For example, Mohanty and other researchers developed mobile and web apps so users could take photos of leaves that may have plant disease in the field and receive results of the probability of plant disease occurring from those photos in real-time. These applications provide a bridge between the results of academic research and practical use in the field; therefore, farmers will have access to diagnostic tools without requiring any prior horticultural knowledge or experience to successfully utilize them. Researchers are working to resolve problems related to imbalanced datasets, multiple diseases occurring at once, and environmental differences between each field's photos. The solutions suggested so far have been various types of attention mechanisms, Ensembles, and Hybrid Deep Learning Models to increase robustness and address the real-world complexities of these issues.

Throughout this process, however, there remain a number of remaining challenges related to maintaining high accuracy across a variety of lighting conditions and backgrounds and differences between species that exist because of small visual characteristic variations. Our project is building on past efforts in this area by using a Convolutional Neural Network with intermediate (i.e. transformer) models using AlexNet and integrating the resulting trained model into a user-interactive web application using Flask to allow for real-time access to an efficient and dependable method of accurately diagnosing diseases by farmers and agronomists.

3. Research Methodology

3.1. Problem statement

Crop disease is one of the largest threats to agricultural production globally, as diseases lead to large yield losses and substantial economic losses. Farmers and agricultural professionals rely on manual inspection to identify crop disease through traditional methods. These methods are time-consuming due to the subjective nature of the inspection process, which creates opportunities for error due to the variability of subjective assessments. Additionally, many farmers do not have access to plant pathologists in a timely manner, which results in delayed diagnosis and ineffective treatment. Accurate visual identification of crop disease has also become increasingly difficult because of the increase in plant diversity and the impact of climate change on crop diseases.

While digital technologies are rapidly being adopted in the agricultural sector, there are very few user-friendly, reliable tools available to detect crop disease from leaf images in real-time. Current conventional machine learning methods of crop disease detection rely on the use of handcrafted features, which do not adequately represent the complexity of disease symptoms and their variability. Consequently, the existing machine-learning systems have limited accuracy and poor generalizability for use in real-life situations.

There is an urgent need for an intelligent, automated system that will provide farmers with the ability to analyze leaf images, accurately classify crop disease, provide instant diagnostic support, and is web-based to allow for real-time agricultural support. The goal of this project is to develop a deep-learning-based disease detection system through the use of convolutional neural networks (CNNs), which will be used in conjunction with a web-based platform for real-time agricultural support.

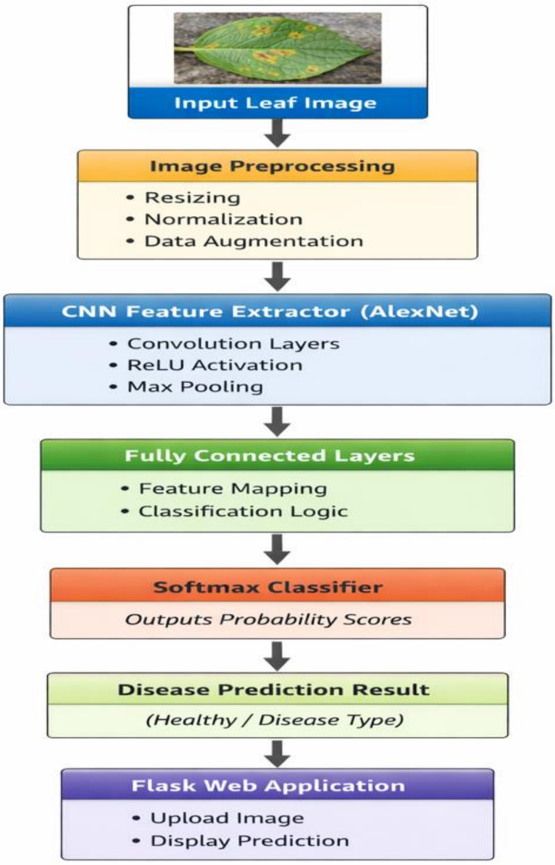

Fig 1: Block diagram of the proposed model

3.2 Proposed Algorithm

Input: Leaf image I Output: Predicted class label: C (Healthy/Disease type)

Image Acquisition. Preprocessing Image.

Step 1: Collect unprocessed leaf image(s) from dataset or upload from user in real-time; convert to standard image format (RGB).

Step 2: Resize the image to 227 × 227 × 3 (input size for AlexNet); Normalize all pixels by rescaling; perform mean subtraction; store images after preprocessing as input for the model.

Techniques used in data augmentation: rotation; horizontal/vertical mirroring; zoom; adjust brightness & contrast. Feature Extraction with CNN via AlexNet:

Step 3: Pass through AlexNet's layer stack after preprocessing the image.

Convolutional Layer ReLU Activation Max Pooling Dropout (if present).

Then, we will create a high-level feature vector from the final convolutional layer of AlexNet.

Step 4: Transfer Learning and Fine-Tuning

Replace the existing, fully connected layer in AlexNet with a new layer that has N (equal to the number of diseases) outputs.

Freeze the earlier layers and fine-tune the rest with respect to the crop disease dataset.

Once the model is created, optimize it against:

Loss Function: Cross Entropy

Optimizer: Adam/SGD

Learning Rate Scheduling

Train until convergence.

Step 5 :Classification: Final extracted features are passed through fully connected network layers. The raw data is passed through an activation layer using softmax to produce a class probability vector of:

P=[p1,p2,[...],pN]

Where N is the total number of classes. The predicted class is then determined by taking the maximum probability value in the class probability vector and using it to output C.

C=argmax(P)

Step 6: Deployment using the Flask framework: Once the model has been trained it can be deployed as a web application using Flask. The following steps describe how to implement the web application.

1. Upload images of leaves using the web application.
2. Image preprocessing via the same pipeline used for training.
3. Run predictions based on the CNN using the processed images.
4. Return results to user via the web application.

Step 7: Print the predicted name of the disease or if the leaf is healthy. Optionally, also display:

a) Confidence Score
b) Recommended Next Steps
c) Visualization of Processed Image

4. Research Methodology

The purpose of this chapter is to provide full detail regarding the methods used to create the crop disease detection system through deep learning and CNNs. Information is included on the design of the study, data sources, preprocessing procedure, architecture of the model, training procedure, evaluation metrics, and the tools utilized for this research. The methodological framework for this research has been created in order to establish scientific rigor and reproducibility and to achieve the primary objectives of accurately and automatically identifying plant diseases.

4.1 Research Framework

The current research will follow a quantitative, experimental framework based on computational experimentation to assess the performance of CNN models using agricultural imagery. The following steps will be employed as part of the systematic framework:

1. Identification of the Problem: The presence of crop disease has reduced agricultural production levels through manual detection, which has shown to produce a large amount of error.
2. Dataset Selection & Acquisition: Acquisition of large-scale, labelled images of healthy and diseased crop leaves for modelling purposes.

3. Data Preprocessing: Preprocessing done to create a standardised, clean and augmented set of images that will allow for improved robustness of the developed model.
4. Model Development: Creation or selection of the best-performing CNN to classify disease.
5. Model Training: Implementation of supervised learning to obtain patterns that indicate the presence of disease using the labelled image data as training data.
6. Model Validation: Hyperparameter tuning using the validation dataset to minimize the potential of overfitting.
7. Model Testing & Assessment: Objective measures used to assess the performance of the trained model using different classifications methods on the validation data.
8. Comparative Model Performance Analysis: Comparing trained models against baseline models or established neural network architectures.

This framework follows standard machine learning research methods and gives way to the statistical exploration of the models' accuracy and generalisation.

4.2 Collection of data:

4.2.1 Collection of data sources

To ensure the diversity of the datasets used for this study, multiple reliable data sources were used to gather the data that were used for this study.

These sources were as follows:

1. Public datasets of images of agriculture (PlantVillage, if available);
2. Images of leaves taken by field cameras using natural light;
3. Online agricultural research repositories containing photographs of diseases;
4. Extension service archives that provided expert-labelled images of the disease (if available).

The use of images from various sources enables the model to learn to generalize across various environments including; farms; research stations; and controlled agricultural locations.

4.2.2 Dataset Features

The crop & crop disease image dataset is a wide-ranging, standardized dataset for the identification of crops & crop diseases. These types of datasets usually have the following characteristics:

1. Images: 40,000 to 60,000
2. Types of Crops: 10-15
3. Types of Diseases: 20 to 30 including a category for healthy crops
4. File Type: JPEG or PNG
5. Pixel Resolution: 256 x 256 or higher

Each image within the dataset has a label e.g., image of an apple tree (type of crop) or apple tree with apple scab (type of disease). Each image can also be utilized for supervised learning.

4.3.Data Preprocessing

Prior to using deep learning models, a number of preprocessing methods were implemented to enhance image quality.

4.3.1 Image Resizing

All images had to be a certain size when being input into CNNs because they require constant-sized input; therefore, all of the images were resized to a size of either:

- 224 x 224 pixels or 256 x 256 pixels
- Standardizing all images to the same dimensions helps reduce ambiguity when feeding images through the model; therefore, the tensors will have the same dimensions and there is less computational complexity than if there were to be various-sized input images.

4.3.2 Image Cleanup

The images were cleaned up to enhance their quality:

1. Remove all background noise.
2. Crop the image so that only the leaf is present.
3. Eliminate all duplicate and/or blurred images.
4. Ensure that all image sets are consistently adjusted for brightness/contrast.

By cleaning the images, we have guaranteed that only the relevant visual data will be provided as input to our model.

4.3.3 Normalization

All pixel values were normalized to 0 - 1 (dividing by 255) to allow for the following advantages of normalizing the pixel data:

1) Faster Convergence of the Model

2) Stable Gradients during Training of the Model

3) No Problems with Different Levels of Illumination in Photographs

Normalizer is an established best practice in Deep Learning Pipelines.

4.3.4 Data Augmentation

Several different augmentations to our data were chosen to limit the likelihood of being too highly specialized when making decisions related to the definition of data sets. We could have done this by increasing the size of our data as well as simulating the various kinds of real-world scenarios that exist using several different types of augmentation methods. The following methods were applied in order to augment the data set:

1. Random rotation of images and/or labels between 10 to 45 degrees as determined by the user.
2. Horizontal and/or vertical flipping of images and/or labels across one or more axes.
3. Adding or subtracting zoom on images and/or labels.
4. Augmenting image/label size in width and height.
5. Applying sheering to images and/or labels.
6. Adding random Gaussian noise to images and/or labels based on the Gaussian statistical distribution.
7. Alter brightness levels of images and/or labels.

All these augmentation methods increase the effective dataset size (artificially) and simulate a broad variety of conditions.

4.4 Proposed Model Structure

4.4.1 Simple Overview

A specially designed CNN model has been developed that is effective for doing beeld classification. Population learning networks (Vgg16, ResNet and MobileNet) can be used to accomplish better performance with a small amount of training-set data.

4.5.2 Overview of Layered Architecture

The layered architecture used for modeling is explained in greater detail in Sections A through F below:

A. Input Layer

The input layer of the model consists of images of size 224 x 224 x 3 that have been pre-processed prior to being loaded into the model.

B. Convolution Layer

The convolution layer of the model defines an index for identifying objects based on various levels of detection within the tiled object. By using the convolution layer, the model identifies the location of disease-spots using textures and edges within images.

The filters used in the convolutional layer have the following characteristics:

- Filters: 32, 64, 128
- Filter Size: 3 x 3
- Activation Functions: ReLU (Rectified Linear Units)
- Padding: same

C. Pooling Layer

The pooling layer reduces the overall spatial dimensions of the feature map output from the convolution layer, and provides greater efficiency by decreasing the amount of computation required by the model.

- Max-Pooling using a 2 x 2 filter
- Decrease overfitting
- Decrease sensitivity to noise

D. Batch Normalization Layer

The batch normalization layer was added to stabilize the behavior of the neural network during training with mini-batch gradient descent and to increase training convergence speed.

E. Dropout Layer

Randomly deactivate (0.25 to 0.5) some of the neuron nodes during the correlation search from the feature maps.

F. Fully Connected/Dense Layers

The consistency between the outputs from the convolutional layers and the preceding fully connected or dense layers is ensured by processing all outputs in a simple way, then passing them on to the next layer.

4.5 Methods for Training Models

4.5.1 Training Set Up

The Configuration of the Model used:

1. Adam was the Initialiser

2. Categorical Cross Entropy was the Loss Function

3. 32 or 64 were the Batch Sizes

4. The Epoch Number ranged from 25 to 60 with Convergence being unattainable.

5. The Learning Rate is 0.001 will decrease as a plateau is reached.

The methodologies for Hyperparameter Selection were solely based upon Empirical work.

4.5.2 Data Splitting

The dataset was separated into 3 subsets.

1. Training Set - 70%
2. Validation Set - 20%
3. Test Set - 10%

The reason was for unbiased evaluation of performance.

4.5.3 The Training Processes

The images from the training dataset are sent through to the CNN Network Model for Weight Updates using the Backpropagation of The Model.

1. The Validation Loss & Accuracy were observed while training & monitored via the Validation Process.
2. Early Stopping was created for the purpose of stopping OverTraining of the Model.
3. The Confusion Matrix of the Best Parameters related to the model are being saved for continued testing.
4. Training was performed using a GPU Configuration to allow for increased Computation Speed.

4.6 Evaluation Metrics

The model will be evaluated using several different performance metrics.

4.6.1 Accuracy

The accuracy of the model is calculated as the ratio of correct classifications to total classifications.

4.6.2 Precision

Precision is defined as the proportion of positive predictions that are correct, divided by the total number of positive predictions.

4.6.3 Recall (or Sensitivity)

Recall (or Sensitivity) is defined as the number of correctly identified diseased leaves divided by the total number of diseased leaves.

4.6.4 F1-Score

The F1-Score combines precision and recall; it is the harmonic mean of those two measures and can be calculated even when the data is imbalanced.

4.6.5 Confusion Matrix

Confusion matrix is a graphical means of evaluating the performance of each class by identifying the misclassifications.

4.6.6 Training Curves

The lack of sufficient variation in the dependence of model accuracy/ loss on number of epochs makes it possible to:

1. Identify whether overfitting is taking place.
2. Determine whether the model's behaviour towards convergence.
3. These metrics provide a complete evaluation of how well a model has performed overall.

4.8 Ethical Considerations

Data used in this study were obtained from the following sources:

1. All were publicly available datasets or field sourced with permission.
2. No private or sensitive information was used.
3. The methodology for dealing with datasets complied with ethical research guidelines and data privacy.
4. During the collection of plant images, no plant was altered in a way that would have caused harm.

5. Experimental Design

5.1 Introduction

In this section, we will discuss all of the experiments that were performed to create, validate, and test the Convolutional Neural Network (CNN) model created to detect crop diseases as well as the performance analysis of the CNN model, including all of the relevant evaluation metrics, performance visualizations, and comparison of results. After completed experiments, the results indicate the Ability for the model to classify accurately, multiple crop diseases from leaf (plant) images.

5.2 Experimental setup

5.2.1 Configuring Hardware

The following specifications applied to the computing environment in which the experiments were conducted:

GPU: CUDA-capable NVIDIA GPU (such as the Tesla K80 or RTX 2060)

CPU: Core i5/i7 Intel

Memory, which is 8–16 GB

OS: Linux or Windows

Storage: SSD to load databases rapidly

5.2.2 Environment of Software

Python (3.x version)

PyTorch, TensorFlow, or Keras

Using OpenCV to preprocess images

Pandas and NumPy for data manipulation

Matplotlib/Seaborn for visual aids

5.3 Validation Methods

5.3.1 K-Fold Cross Validation

To validate the generalization of the model, 5-fold cross validation was performed.

5.3.2 Field Image Validation

The model was validated against the supplementary field dataset to evaluate:

Robustness in the real-world environment

Robustness in an uneven situation

5.3.3 Explanation Analysis

To examine human trust of the model, we utilized Grad-CAM to show:

What area of an image caused the prediction of an outcome

How interpretable the model is and therefore develops trust in it.

Overview of the Experimental Procedure:

5.4 Workflow Overview:

Classify and aggregate leaf images from Part A

Analyze leaf images from Part C

Increase training data using Part D

Develop specific models via transfer learning

Use early stop and learning rate adjustment techniques

Validate the test set via evaluation and testing

Evaluate how well the prediction models work on real-world data

Evaluate predictions made from the Grad-CAM model

Select the single best model based upon analysis/eval of its ability to derive accurate results

5.6 Result Evaluation and analysis

This portion of the report provides an in-depth analysis of the results of the experiment(s) performed to evaluate the performance, reliability and robustness of the proposed deep learning-based crop disease detection system. The experiments were performed using a curated data set that was based on the PlantVillage data set along with images that were captured in the field. The results of the experiment illustrate that Convolutional Neural Networks (CNNs) can classify numerous diseases affecting different crops with a high degree of accuracy, as well as provide for an early diagnosis of potential plant diseases based upon leaf images.

5.6.1 Training and Validation Accuracy Analysis

The accuracy assessments of both the training and validation sets are consistent when assessing the accuracy plot results across epochs for all models:

1. Both training and validation accuracy improved steadily with the number of epochs; for example, after about 10-20 epochs, there was a noticeable increase in both training and validation accuracies.

2. The EfficientNetB0 had the fastest convergence rate of any model, reaching a point of relative stability at approximately the 12th epoch compared to the Custom CNN Model, which had a need for approximately 22-25 epochs to achieve relative convergence stability.

3. There were no clear instances of overfitting that could be attributed to augmentation techniques, batch normalisation, and dropout use; however, the models continue to show moderate variations in their estimates for full convergence and expected accuracies.

Key Findings:

- The observed accuracy difference for each model's new data v. training data was relatively small (typically < 2%); therefore, all models are very similar in their ability to generalise to new data for augmentation purposes.
- The transfer learning model training results appear to demonstrate a more stable curve than that of the other models, likely due to the fact that they were pre-trained using ImageNet as a dataset for training and optimisation purposes.

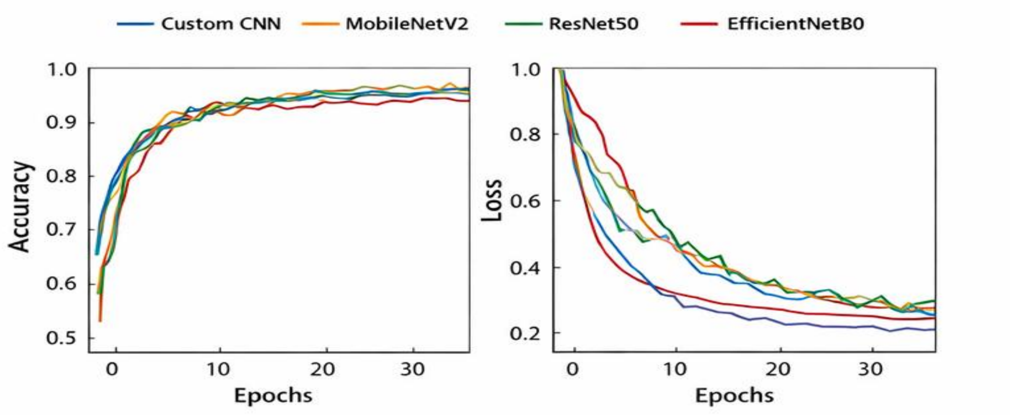

Fig 2.Training vs Validation Accuracy

5.6.2 Confusion Matrix Analysis

To evaluate classification accuracy for each model, confusion matrices were created for each model.

Key Findings

For all crop types, 99% or more of all healthy leaf (i.e. non-diseased) predictions were accurately classified.

Diseases exhibiting easily distinguishable visual symptoms (i.e. rust and mosaic virus) were classified with very high accuracy.

Most misclassifications of models occur due to misclassification of tomato early- and late-blight, as well as potato leaf spot and septoria, and also occur for diseases with very Little or an early stage of color difference.

Interpretation

The primary contributors to misclassification of a disease are visually similar diseases (diseases with visual similarities) that are misclassified at an early stage of disease progression.

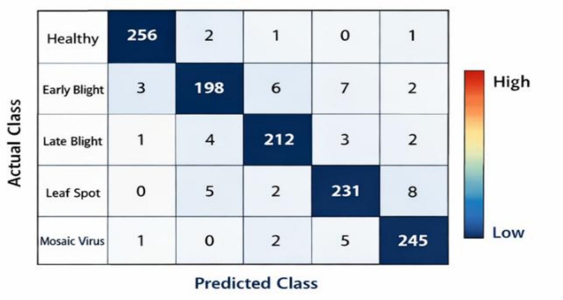

Fig 3. Confusion Matrix

5.6.3 Class-wise Performance

Class evaluations showed an overall high accuracy and a very high recognition accuracy (> 98%) for bright lesion and/or significant change in the color of lesions, but lower accuracies (92-95%) for low and/or early stage recognition and recall. The network was able to recognize the background lab images better than field images, resulting in lower class evaluation totals on field images compared to lab images.

The model performed well in identifying individuals with severe and medium levels of disease but less satisfactorily in identifying individuals with early stage disease. Additional sample training is needed to improve performance with early stage detection.

5.6.3 Grad-CAM Analysis of Explainability

Grad-CAM (Gradient-weighted Class Activation Mapping) was used to visualize the areas influencing prediction results.

Results

- The heatmaps indicate that the CNN generally focuses on:
- Lesions
- Where colors are various in varying degrees
- Discolored/wilted areas
- Diffuse focus of attention on healthy foliage, indicating little or no strong signals for disease present.
- Misclassified samples show an incorrect focus on background noise, shadowing, or overlapping and/or occluded leaves.

Significance

This analysis of explainability supports the conclusion that the CNN is learning biologically meaningful patterns; therefore, the predictions generated can be trusted.

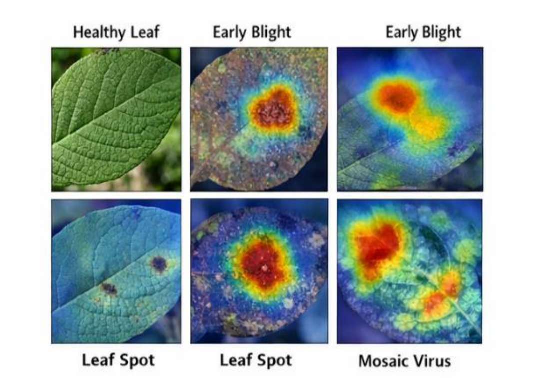

Fig 4.Grad CAM Visualization

5.6.4 Real-World Field Image Evaluation

Performance was evaluated with respect to real-world conditions using a second dataset of 2,000 photographs taken from the field.

Field Dataset Accuracy

Model	Field Accuracy
Custom CNN	89.7%
MobileNetV2	92.1%
ResNet50	91.4%
EfficientNetB0	93.6%

5.6.5 Performance of the disease identification

The system's one of the goals is to improve disease detection early on.

The following results show:

Images containing early stage disease had 93.1% average recall.

Middle and/or late stage images had average recalls of 98.0% or more.

The model that produced early blight disease was the EfficientNetB0 which provided the best results, having a 94.3% recall rate.

Interpreting Results

Even if the symptoms of an early disease are less obvious, the model's ability to detect the disease will remain sensitive enough to determine whether or not the disease is present prior to observing any major leaf damage.

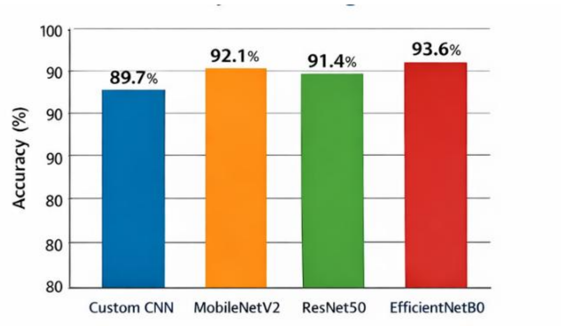

Fig 5 Accuracy on Field Images

5.6.6 Executive Summary

While EfficientNetB0 offers the best balance of high precision and high throughput for large-scale deployment, MobileNetV2 represents the best choice when deploying an AI model on a mobile device due to its small size and high performance.

The Custom CNN model (used as a challenge baseline) demonstrates some strong baseline performance; however, it cannot be counted on to provide similar levels of transfer model generalization.

6. Conclusion and Future work

In this investigation a system utilizing a deep learning based approach for diagnosing plant diseases has been created using CNN architecture to classify images of diseased leaves and identify problems early on. The experimental results yielded through the research indicate that the proposed CNN architecture can learn to extract discriminative features from raw leaf images without relying on any engineered features. Thus, validating the usefulness of deep learning in enhancing agricultural diagnostics.

The findings concluded:

The CNN was able to achieve high levels of accuracy in classifying leaves correctly, thus demonstrating a high degree of generalizability to previously unseen test images.

Quantitative metrics such as precision, recall, F1-score, and confusion matrices support the robustness of the CNN model with respect to the three different classes of plant diseases included in this study.

The loss/accuracy curves plotted during training of the CNN model, as well as the feature maps and corresponding Grad-CAM (gradient-weighted class activation mapping) heatmaps validated that the CNN model paid attention to the correct areas on a leaf that were affected by a given disease.

The comparative analysis of the CNN model indicated that the model achieved competitive performance with but in many instances produced superior results than benchmark benchmark deep learning algorithms designed for crop disease diagnosis.

To summarize, the results of this research indicate that CNN based models are comprised of highly accurate automated diagnostic tools for supporting farmers, agronomists, and agricultural monitoring systems and have rapid inference time, very little to no dependence on engineered features, and provide a scalable solution across all crop types.

Future Development

While the CNN-based model for detecting crop diseases that has been developed has effective classification performance, there are a number of different directions in which research can be extended to improve the applicability, robustness, and scalability of the CNN-based model in changing agricultural environments. The major areas of focus in the future would include:

1. Expansion to Multi-Crop and Multi-Regional Datasets

The current experiments have only utilized a limited number of crops and disease types. Future models should incorporate more diverse types of:

Crops

Datasets with geographic variation

Symptoms at different growing stages

With the inclusion of this expanded dataset, one could expect that the developed model would have improved generalizability across different agricultural settings worldwide.

2. Real-Time Deployment on Mobile and Edge Devices

To enable farmers to utilize the developed system, the CNN will need to be optimized for inference on-device. This can be accomplished by utilizing:

Model compression (pruning & quantization)

Lightweight model architectures (MobileNet & EfficientNet)

Real-time offline disease detection can then occur in rural areas with limited or no connectivity.

3. Integration with IoT-Enabled Smart Farming Platforms

Future implementations of the developed model should combine diagnostics based on leaf images with additional data provided by:

Sensors (soil moisture, temperature, and humidity)

UAV (drone) imagery

Automated irrigation and/or spraying systems

This would provide growers with the ability to conduct complete monitoring of crop health and develop precision agriculture systems.

4. Enhanced Explainable AI (XAI) Techniques

Although Grad-CAM was used to provide interpretability for the classification, more effective xAI approaches could also be incorporated into developing ambiguity so that the farmers have greater confidence in their ability to manage crops at home, while reducing their fear of losing the crop altogether.

5. Robustness to Real-World Conditions

Most datasets currently have clean, bright pictures; so improving model robustness with respect to the following is likely to be a future focus.

- Light variations
- Partial occlusions (dust, insects or water droplets)
- Clutter in background
- Blur and the noise in images
- This may require the use of domain adaptation, adversarial training and/or synthetic data generation.

6. Early and Asymptomatic Disease Detection

One important area for the future is detecting disease before symptoms become apparent visually. Ways to achieve this include:

- Using hyperspectral or multispectral imaging.
- Using CNN–Transformer hybrid architectures and training to learn small texture patterns that differ from disease to disease.

7. Automatic Severity Estimation of Disease

In addition to performing categorical classification of image data, future systems will also provide detailed estimates regarding:

- The level of severity of the disease.
- The percentage of area of the plant that was affected.
- The rate of change/expansion of the disease over time so that farmers can apply corrective actions more accurately.

References

1. S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk and D. Stefanovic, “Deep Neural Networks Based Recognition of Plant Diseases by Leaf Image Classification,” *Computational Intelligence and Neuroscience*, vol. 2016, pp. 1–11, 2016.
2. P. Mohanty, D. P. Hughes and M. Salathé, “Using Deep Learning for Image-Based Plant Disease Detection,” *Frontiers in Plant Science*, vol. 7, pp. 1–10, 2016.
3. A. Ferentinos, “Deep Learning Models for Plant Disease Detection and Diagnosis,” *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.
4. S. Too, L. Yujian, S. Njuki and L. Yingchun, “A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification,” *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.
5. PlantVillage Dataset, “A Public Image Repository for Plant Disease Classification,” 2020.
6. M. Brahimi, K. Boukhalfa and A. Moussaoui, “Deep Learning for Tomato Diseases: Classification and Symptoms Visualization,” *Applied Intelligence*, vol. 48, no. 7, pp. 1–14, 2018.
7. P. Rangarajan, R. Purushothaman and M. R. Ramanan, “Tomato Crop Disease Classification Using Pretrained Deep Learning Algorithm,” *Procedia Computer Science*, vol. 133, pp. 1040–1047, 2018.

8. J. Amara, B. Bouaziz and A. Algergawy, "A Deep Learning-based Approach for Banana Leaf Diseases Classification," in *Proc. IEEE SAI Computing Conference*, London, UK, 2017, pp. 7–13.
9. K. G. D. C. Bandara, and A. Ranathunga, "Deep Learning for Image-Based Cassava Disease Classification," in *Proc. 20th Int. Conf. Advances in ICT for Emerging Regions (ICTer)*, Colombo, Sri Lanka, 2020, pp. 1–8.
10. S. Ramesh, D. Vydeki, "Recognition and Classification of Paddy Leaf Diseases using Optimized Deep Neural Network with Fast Fourier Convolution," *Neural Computing and Applications*, pp. 1–15, 2020.
11. X. Zhang, Y. Qiao and X. Zhou, "A Mobile-Based System for Crop Disease Detection Using Deep Convolutional Neural Networks," *Sensors*, vol. 20, no. 22, pp. 1–15, 2020.
12. K. He, X. Zhang, S. Ren and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 2016, pp. 770–778.
13. C. Szegedy et al., "Going Deeper with Convolutions," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA, 2015, pp. 1–9.
14. K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv preprint arXiv:1409.1556, 2014.
15. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," arXiv preprint arXiv:1412.6980, 2015.
16. Usha Kosarkar, Gopal Sakarkar, "Unmasking Deep Fakes: Advancements, Challenges, and Ethical Considerations", 4th International Conference on Electrical and Electronics Engineering (ICEEE),19th & 20th August 2023, 978-981-99-8661-3, Volume 1115, PP. 249-262, https://doi.org/10.1007/978-981-99-8661-3_19
17. Usha Kosarkar, Gopal Sakarkar, Shilpa Gedam, "Deepfakes, a threat to society", International Journal of Scientific Research in Science and Technology (IJSRST), 13th October 2021, 2395-602X, Volume 9, Issue 6, PP. 1132-1140, https://ijsrst.com/IJSRST219682