Cureus Journal of Computer Science

Open Access Original Article

AI-Driven Agriculture: Improving Cotton Farming Through Transfer Learning-Based Disease Detection and Classification

Review began 02/03/2025
Review ended 09/22/2025
Published 09/23/2025

DOI:
https://doi.org/10.7759/s44389-024-02217-9

Varun R. Vangar [1] , Vedank Wakalkar [1] , Shiromani Waghmare [1] , Sushant Nadavade [1] , Grishma Y. Bobhate [1] , Rutuja Patil [1]

1. Computer Science and Engineering (Artificial Intelligence and Machine Learning), Vishwakarma Institute of Information Technology, Pune, IND

Corresponding authors: Varun R. Vangar, varun.252k4@gmail.com, Vedank Wakalkar, vedankw11@gmail.com, Shiromani Waghmare, shiromaniwaghmare@gmail.com, Grishma Y. Bobhate, grishma.bobhate123@gmail.com

Abstract

In the advancement of cotton farming, significant challenges related to pests and diseases have become obstacles to farming, impacting the yield and overall production of cotton. Since cotton is a cash crop, if a small percentage of the crop is affected, it can lead to a huge loss. This research paper uses various deep learning models like convolutional neural network and transfer learning models, such as the InceptionV3. Different convolutional neural network models are evaluated, and the InceptionV3, a transfer learning model, has given the best accuracy, as its complexity plays a key role in extracting features. This study used the InceptionV3 model to attain the highest validation accuracy of 96.95% in the classification of seven different diseases in cotton crops. This study is a representation of the deep learning-based capability of a disease classification system in agriculture approaches for the better accuracy of crop health and productivity. It bridges the gap between traditional farming methods and modern technology, enabling a data-driven decision-making system to ensure the sustainability of cotton farming ecosystems through healthy crops and productivity.

Categories: Deep Learning, Machine Learning (ML)
Keywords: cnn model, transfer learning, inceptionv3, data augmentation, accuracy, training accuracy, validation accuracy

Introduction

Cotton farming is an important crop in global agriculture, contributing significantly to economies and supporting millions of livelihoods. It is a versatile crop that provides the textile industry with its main source of natural fiber, generating employment and revenue in agriculture in major cotton-producing countries such as India, China, and the United States. While India has vast cultivation areas (12.07 million hectares as of 2014) [1], it lags behind in terms of production efficiency; for just 2.79 million hectares (2019) [2], China produces 4-5 times more cotton. India uses genetically modified cotton, or Bt cotton, which minimizes crop losses by fighting off pests more effectively and thereby increases productivity.

However, the conventional farming practices of most Indian farmers are becoming less effective as pests develop resistance to toxins. Diseases such as bacterial blight and cotton boll rot have become more prevalent, which are a significant challenge to productivity. These issues need innovative solutions that integrate modern technology.

Deep learning, especially convolutional neural network (CNN), has emerged as a powerful tool for image-based disease detection. By using technologies, such as CNNs, one can detect pre-symptoms of diseases or infestation through pests, for timely intervention of crop loss due to the effects of the mentioned pathogens. They improve precision with the use of mobile and IoT devices in executing pest management activities. For example, research proves that deep learning frameworks, which include InceptionV3, mask region-based CNN (RCNN), and k-means clustering, can work efficiently in diagnosing and distinguishing plant diseases effectively [3-7].

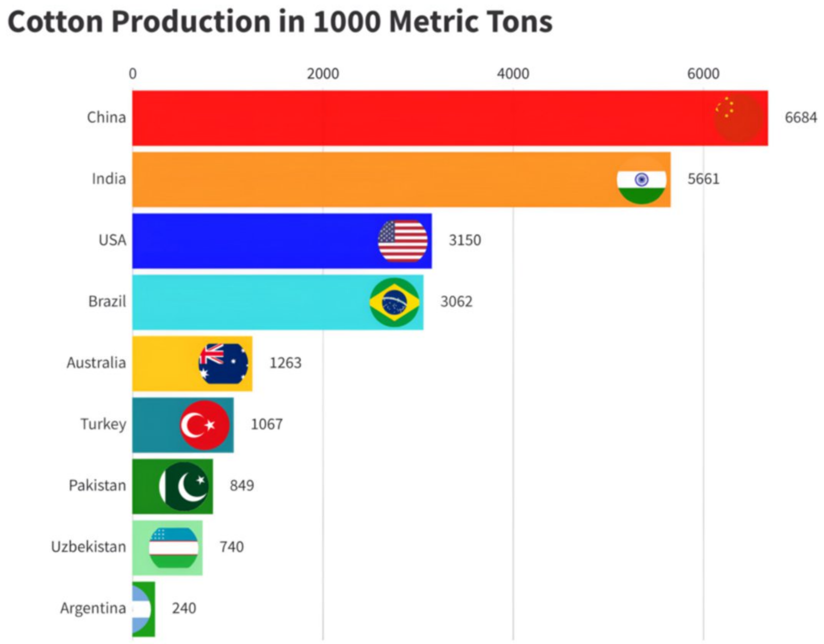

FIGURE 1: Countries With Highest Cotton Production

Source: [3]

The graph presented in Figure 1 displays the relationship between the country and its cotton production. China ranks as the country with the highest cotton production globally, followed by India. This disparity highlights the need for advanced agricultural practices in India to improve productivity.

Harakannanavar et al. [4] proposed some machine learning techniques for tomato leaf disease detection procedures based on histogram equalization, K-means clustering, and feature extraction using discrete wavelet transform and principal component analysis. The results show that their method had a high classification efficacy of 88%, 97%, and 99.6%, respectively, using support vector machine (SVM), k-nearest neighbor, and CNN. This reflects the potential use of computer vision in managing diseases in agriculture.

Recently, a good number of researches on deep learning have been forwarded for crop disease detection, such as those by Islam et al. [5], which report successful implementations of CNN, VGG-16, VGG-19, and ResNet-50 models with accuracies of 98.60%, 92.39%, 96.15%, and 98.98%, respectively [5].

Rothe and Kshirsagar [6] have proposed a pattern recognition system with an active contour model and Hu's moments for the identification of three cotton leaf diseases with an accuracy of 85% besides discussing several plant disease detection methods based on neural networks, SVM, and other machine learning algorithms applied to different crops [6]. This work shows mask RCNN with transfer learning to detect the cotton leaf diseases using 94% accuracy when trained on a diverse dataset of real-world field images, overcoming the limitations of controlled-environment models. It moves agricultural AI forward as a practical solution for multi-disease detection in a complex background, overcoming some of the most important challenges of cotton farming productivity [7].

As an interesting result of their research, Bhimte and Thool [8] have proposed an automatic cotton leaf disease diagnosis system based on the principles of image processing and an SVM classifier. Their model includes image processing and feature extraction via the gray level co-occurrence matrix, with 98.46% disease classification accuracy for diseases like Bacterial Blight and Magnesium Deficiency, showing the vast potential of an AI-oriented solution for improving cotton disease management. SVMs, random forests, and CNNs have been found to be very useful in the detection of threats against crop yields by identifying diseases and pests in plants [9].

The authors further the original work of CNNs on image classification by involving a new training criterion called M3CE-CEc, which aims to surpass cross-entropy methods and achieve better performance in the standard datasets of MNIST and CIFAR-10. Many studies have previously shown that deep learning models are effective for almost all classification tasks. More importantly, the implementations of these models surpass traditional approaches in complexity because of the higher rates of accuracy on difficult data pattern handling [10]. The literature highlights the development of deep learning models, particularly CNNs, for the accurate detection of plant diseases using leaf images, achieving impressive accuracy rates of up to 99.53% across various datasets. These models leverage advanced image processing techniques and large-scale data augmentation to enhance disease identification, demonstrating significant potential for integration into agricultural practices for early disease detection and management [11]. The paper is a comprehensive review of the detection and classification of cotton leaf diseases, focusing on the difficulties associated with the identification of diseases based on subtle symptom differences. It reviews various techniques in image processing, including segmentation and feature extraction methods, as well as machine learning classifiers, such as SVM, that have been used in previous studies to improve diagnostic accuracy [12]. The study identified *Curvularia lunata* as the causal agent of leaf spot disease in cotton, which includes its morphological characteristics and pathogenicity based on rigorous isolation and inoculation tests. Disease incidence is quite high, hence the importance of effective management to reduce yield loss in cotton crops [13].

Recently, investigators like Revathi and Hemalatha [14] have proven with advanced feature selection techniques like Enhanced Particle Swarm Optimization and Skew Divergence that machine vision systems are viable for the detection of cotton leaf diseases. Their study proves great accuracy in the classification of multiple diseases and suggests that AI-based methods could act as a tool for improving the management of cotton diseases through early and correct detection [14]. The paper presents a deep learning-based system for detecting and recognizing diseases in various plant species, achieving an impressive accuracy of 96.5% using a dataset of 35,000 images. It highlights the effectiveness of CNNs in automating plant disease diagnosis, thereby enhancing agricultural productivity and sustainability [15]. The paper revolves around the feature extraction of the crop images using EfficientNet convolution and then the input to the group-wise transformer architecture. Their proposed model achieves 99.8% accuracy on the PlantVillage dataset, 86.9% accuracy on the cassava dataset, and 99.4% accuracy on the tomato leaves dataset. The experiment suggests that this model is optimal and has best accuracy in comparison with other neural networks based on CNN [16]. In the research paper, they discussed the challenges of existing methods of the plant disease classification such as manual inspection, etc. To address these challenges, they researched the integration of CNNs and vision transformers within the ensemble learning model to increase the accuracy and robustness of plant disease detection systems [17]. In the present work of the research paper, the authors used a lightweight model ‘ConViTX’ for plant disease classification and showed improved generalizability and explainability. The diverse architecture of the ConViTX uses the breakdown of the CNNs and vision transformers to simultaneously capture local and global features. ConViTX gives 98.8% accuracy on the maize dataset and 61.42% on drone camera-captured raw images, having only 0.7 million parameters and 0.647 billion operations per second. This makes it deployable on the resource-constrained precision agriculture setups [18].

Collectively, these studies demonstrate the transformative potential of AI-driven agriculture in improving cotton farming through transfer learning-based disease detection and classification. These breakthroughs have tremendous potential to help address critical bottlenecks in disease management to improve agricultural productivity and sustainability.

Gap analysis

The literature reviewed highlights the significant progress made in applying deep learning and machine learning techniques for plant and cotton disease detection. Traditional approaches, such as image processing with feature extraction techniques (gray level co-occurrence matrix, discrete wavelet transform, principal component analysis, and Hu’s moments), combined with classifiers like SVM and k-nearest neighbors, have demonstrated promising accuracy but remain limited by their reliance on handcrafted features. These methods often fail to generalize effectively across large, diverse datasets collected from real-world conditions.

Recent studies leveraging CNNs and transfer learning frameworks (e.g., VGG-16/19, ResNet-50, Mask RCNN) show considerable improvements in classification performance, with accuracies frequently reported above 90%. However, most of these works primarily focus on binary or few-class disease classification problems. Similarly, mask RCNN-based works have shown high accuracy but are computationally heavy and not easily scalable for resource-constrained precision agriculture.

Another gap identified is that many reported studies are based on controlled-environment datasets, such as the PlantVillage dataset, which do not fully capture the complexity of real-world field conditions like background noise, lighting variation, and overlapping symptoms. This makes their direct deployment in large-scale cotton farming contexts challenging. Moreover, while novel architectures such as EfficientNet and ConViTX show strong generalizability, their application has been broader to multiple crops and not specifically optimized for the multiple disease of cotton.

To address these gaps, the present work develops a fine-tuned InceptionV3 model, optimized for the classification of seven distinct cotton leaf disease classes. By leveraging transfer learning, careful fine-tuning, and data augmentation strategies, the proposed model achieves superior F1 score and validation accuracy, outperforming conventional CNN architectures reported in prior studies. Unlike prior works restricted to fewer disease categories or controlled settings, our approach demonstrates practical applicability in real-world field conditions, thereby contributing a scalable and accurate solution for cotton disease management.

Materials And Methods

Diseases

Table 1 lists the most common diseases affecting cotton, along with their causes. In summary, this work addresses various diseases commonly found in cotton farming. A thorough understanding of each disease's etiology is essential for effective management and mitigation of risks to cotton crops. Farmers who actively seek proper knowledge can better withstand disease attacks by planting disease-resistant cotton varieties, resulting in consistently high yields year after year.

Serial No.	Diseases	Causes
1	Aphids	Aphids (Pest)
2	Armyworms	Armyworms (Pest)
3	Bacterial Blight	Caused by bacteria
4	Cotton Boll Rot	Fungal infection
5	Green Cotton Boll	Pest damage
6	Powdery Mildew	Fungal infection
7	Target Spot	Fungal infection

TABLE 1: List of Cotton Diseases and Their Causes

Aphids

Aphis gossypii, commonly known as the aphid, is a significant pest of cotton and other crops. It spreads viral diseases from plant to plant. As a sap-sucking insect, aphids inject saliva into the plant's phloem, weakening its defences and creating an environment beneficial to disease development.

Armyworms

Armyworms, especially the fall armyworm (*Spodoptera frugiperda*), are a serious threat to cotton crops. They feed on the leaves and bolls, causing significant damage and reduced yield. Their feeding also creates irregular holes in the leaves, which negatively affect plant growth and reduce photosynthesis.

Bacterial Blight

Bacterial blight in cotton is caused by the bacterium *Xanthomonas citri* pv. *malvacearum*, which negatively impacts the crop by forming angular, water-soaked spots on the leaves. These lesions eventually darken, leading to leaf shedding, reduced photosynthesis, and stunted plant growth.

Cotton Boll Rot

Cotton boll rot is a complex disease caused by various fungal pathogens, including *Fusarium moniliforme*, *Colletotrichum capsici*, and *Aspergillus flavus*. The infection initially appears as small brown or black spots that expand to cover the entire boll. The disease then infiltrates the internal tissues, leading to seed decay and tissue buildup. Affected bolls may fail to open fully and often drop prematurely. In some instances, the decay is confined to the outer layers, damaging the pericarp while leaving the inner tissues unaffected.

Green Cotton Boll

This condition affects the immature phase of the cotton plant, which causes stunted growth, failure of bolls to mature, and resulting in underdeveloped, nutrient-deficient bolls. It can be triggered by various factors such as drought, excessive moisture, or nutrient deficiencies. If left untreated, affected bolls often drop prematurely or remain on the plant without producing viable fiber.

Powdery Mildew

Powdery mildew begins as small, powdery patches on the leaves, stems, and sometimes bolls of cotton plants. These patches gradually expand, covering larger areas with dense, powdery fungal growth. The disease progresses by spreading into the inner tissues, weakening the plant. Infection can cause bolls to become underdeveloped and fall off prematurely.

Target Spot

Target spot, also known as *Corynespora* leaf spot, initially forms lesions on mature leaves along the main stem and on lower leaves. Under favorable conditions, the disease rapidly spreads upward through the plant's canopy. The lesions, characterized by necrotic spots with light to dark brown concentric rings, give the disease its name.

This study makes use of a deep learning model for the purpose of identifying different types of diseases afflicting cotton plants. The approach works on a dataset featuring images categorized under seven specific categories: aphids, armyworm, bacterial blight, cotton boll rot, green cotton boll, powdery mildew, and target spot. The dataset offers plenty of images for each disease type and examples of healthy crops stored categorically in separate folders, which facilitates easier labelling and access to data. Applying the data augmentation approach increased the model's accuracy during its training phase for the validation process. Below is a step-by-step overview of the model development process.

Data collection

The dataset is obtained from Kaggle for training the model [19]. This dataset contains 800 images of each disease and healthy data. The dataset also separates the validation data, which makes labelling easier.

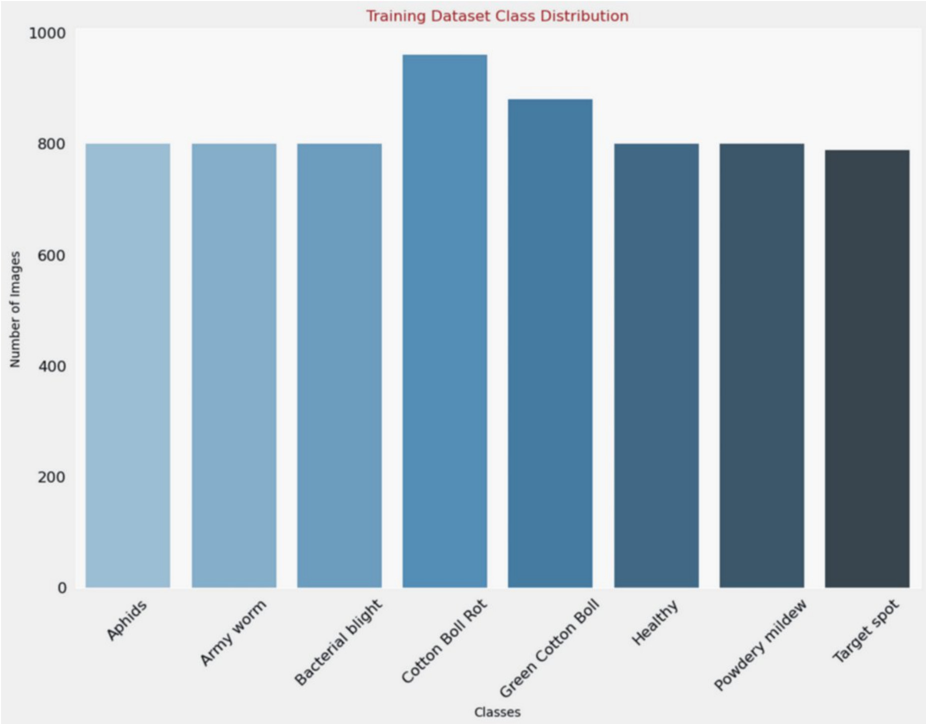

FIGURE 2: Balanced Dataset Distribution of Training Data

As shown in Figure *2*, a specific number of images were used to train the model for each type of cotton disease. The figure indicates that the training dataset is balanced, with nearly 800 images per disease category. Only cotton boll rot and green cotton boll have slightly more images than the other diseases. Using a balanced dataset helps the model generalize better across different diseases and reduces the risk of misclassifying one disease as another, since the model is trained on balanced data of all categories.

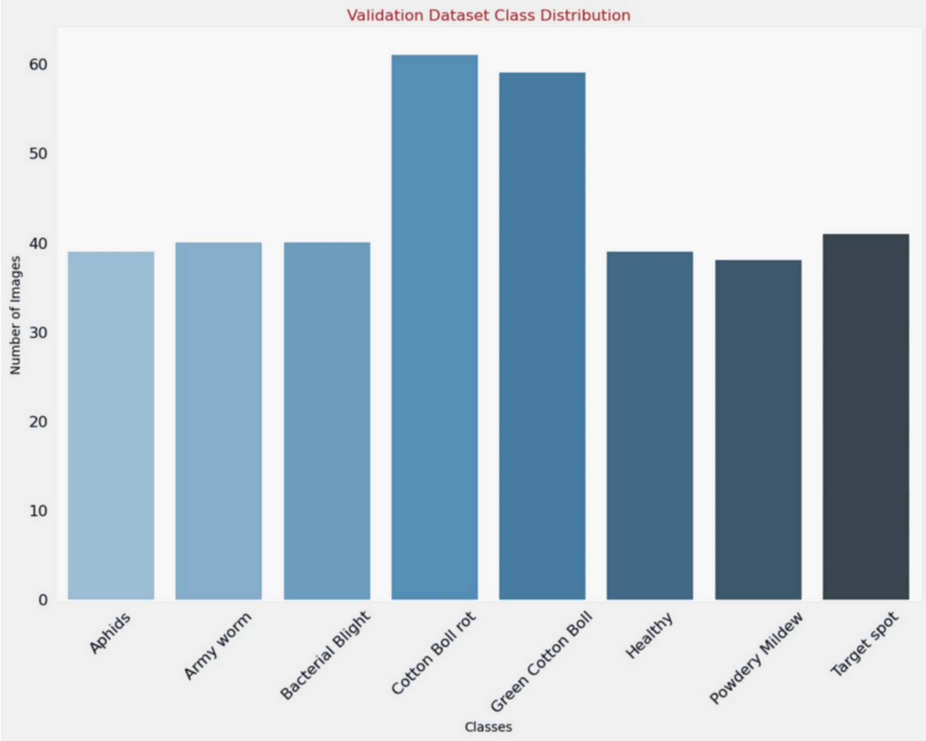

FIGURE 3: Balanced Dataset Distribution of Validation Data

Similarly, Figure *3* shows the validation dataset used to evaluate the model. Like the training data, the validation data is balanced, with slightly more images of cotton boll rot and green cotton boll.

Data preprocessing and augmentation

The dataset was divided into three subsets: training, validation, and testing. Data augmentation on the training set was performed using the ImageDataGenerator library to enhance the model's diversity and generalization capabilities [5]. The augmentations applied included random rotations up to 30°, width and height shifts up to 20%, shear transformations, zooming, horizontal flipping, and a fill mode using nearest-neighbor to handle pixels introduced by these transformations. All pixel values were normalized by rescaling them to the range 0-1 using a factor of $\frac{1}{255}$ [16].

Validation and test sets experienced only rescaling, with no augmentation artifacts possibly affecting model performance. Images have been resized into 224 × 224 to meet the required input of the InceptionV3 model.

Handling class imbalance

Class weights were computed using the compute_class_weight method from the sklearn library to address the class imbalance problem. These weights assigned higher importance to under-represented classes during the training phase, thereby preventing the model's predictions from being biased towards majority classes.

The formula used to calculate the class weight for a particular class *c* is

$$\text{weight}_c = \frac{n_{\text{samples}}}{n_{\text{classes}} \times n_c}$$

where, weight_c = class weight for class *c*, n_{samples} = total number of training samples, n_{classes} = total number of unique classes, and n_c = number of training samples in class *c*.

This formula ensures that classes with fewer samples are given more importance during training, effectively balancing the learning process across all classes.

Model architecture

For a transfer learning framework, the InceptionV3 base architecture pre-trained on ImageNet is used as a base. The model's last layers were removed and replaced with layers corresponding to our classification task, while the pre-trained features in the initial layers of the base architecture were frozen to prevent modification during early training.

The input image size was standardized to 224 × 224 × 3, as shown in Figure 4, although InceptionV3 natively accepts 299 × 299 × 3 inputs. The base model was modified to accommodate 224 × 224 × 3-sized images to align with the dataset constraints.

The first 249 layers of the InceptionV3 base were frozen to preserve generic low-level features from ImageNet, while all layers with 250 and onward were fine-tuned to adapt the model to the specific dataset.

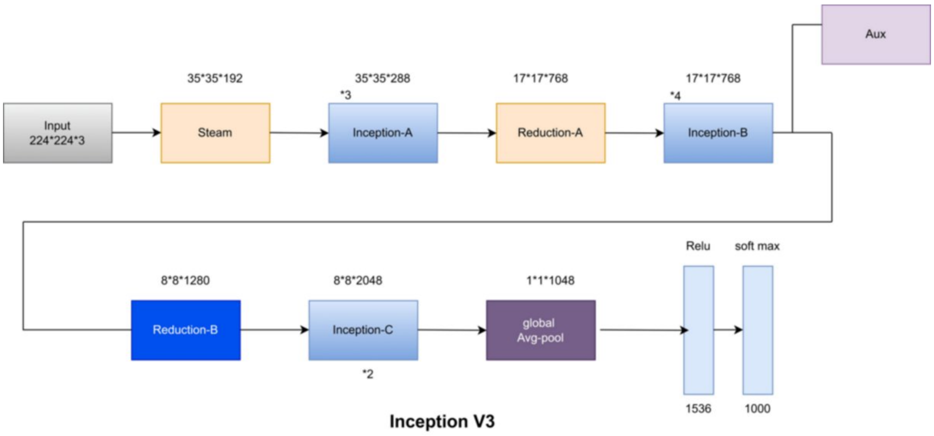

FIGURE 4: InceptionV3 Architecture

As shown in Figure 4, the architecture of InceptionV3 is primarily used due to its inception modules. These specialized convolutional blocks are designed to capture features at various spatial scales. Within one module, convolution layers of several different sizes are used, allowing the network to learn features at multiple resolutions effectively. Besides the Inception modules, InceptionV3 also contains standard convolutional and pooling layers and fully connected layers for classification.

InceptionV3 has multiple versions of inception modules, all designed to capture information at different scales. This makes it possible for the network to extract features well from images that pose problems of different sizes. Unlike traditional fully connected layers, InceptionV3 uses global average pooling as the final layer before outputting. This minimizes overfitting in the model [11].

The blocks up to the "Inception-C" module (shown in blue) represent the original base model. Following the global average pooling (GAP) layer, custom layers were added. These include two fully and properly connected (dense) layers with 512 and 256 neurons, respectively, using ReLU activations:

$$z^{(l)} = W^{(l)} a^{(l-1)} + b^{(l)}, \quad a^{(l)} = \max(0, z^{(l)})$$

GAP: Downscaling the feature maps spatially and getting the number of individual vectors of average of every feature map.

$$GAP_k = \frac{1}{H \times W} \sum_{i=1}^{H} \sum_{j=1}^{W} F_{i,j,k}$$

Batch normalization layers were also inserted after the dense layers insertion to stabilize learning:

$$\hat{x} = \frac{x - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}}, \quad y = \gamma \hat{x} + \beta$$

Dropout layers with rates of 0.5 and 0.3 were applied after each dense layer to reduce overfitting:

$$\tilde{a}_i^{(l)} = \begin{cases} a_i^{(l)}, & \text{with probability } (1-p) \\ 0, & \text{with probability } p \end{cases}$$

Finally, here the softmax layer (which is at the end of the pipeline) was used to produce output class probabilities:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^{C} e^{z_j}}$$

Model training

First, the base layers were frozen, and the model was trained. The categorical cross-entropy loss function was minimized using the Adam optimizer (learning rate = 0.001). Categorical cross-entropy loss function is defined as

$$\mathcal{L}_{\text{CCE}} = -\sum_{i=1}^{C} y_i \log(\hat{y}_i)$$

where C is the number of classes, y_i is the true label (one-hot encoded), and $\hat{y}_i$ is the predicted probability for class i.

For training, the following callbacks were used:

EarlyStopping: Monitored the validation loss and stopped training if it did not improve for five consecutive epochs, saving the best-performing model.

ReduceLROnPlateau: Reduced the learning rate by a factor of 0.5 if the validation loss plateaued for three consecutive epochs, with a minimum learning rate of 1×10^{-6}.

ModelCheckpoint: Saved the model weights corresponding to the lowest validation loss.

Fine-Tuning

After the initial training phase, the layers of the base model were unfrozen, making all layers trainable except for the first 50. Fine-tuning was performed using the same loss function and callbacks, but with a reduced learning rate of 1×10^{-5}. This enabled the model to adjust the features of its deeper layers for the dataset it was fine-tuning, which helped mitigate overfitting.

Model Evaluation

We assessed the model's performance on the test set with multiple metrics:

Accuracy: The ratio of the activity where samples were classified correctly.

Classification Report: Provides precision, recall, and F1-score for each class.

Confusion Matrix: An illustrated matrix that helps understand the utility of the model, such as false positives, false negatives, etc.

To make the confusion matrix more interpretable and assist qualitative analysis, we used an in-house heatmap visualization.

The experiments were implemented using TensorFlow. A batch size of 32 was used for training, with 20 epochs for the initial training stage and 10 epochs for fine-tuning. The best model weights, based on validation performance, were used for inference. Thus, such an approach allowed for the development of a more generalizable model that could overcome the issues of class imbalances and achieve state-of-the-art performance for multi-class image classification tasks.

Before, selecting a specific model that will have good accuracy and outperform other models, a comparison was conducted among transfer learning architectures like inceptionV3 and deep learning model CNN. The following table shows the training and validation accuracy of different models.

Algorithm	Training Accuracy	Validation Accuracy
CNN	98.22%	90.62%
InceptionV3	97.51%	96.92%

TABLE 2: Accuracy of CNN and InceptionV3 Model

CNN, Convolutional Neural Network

Table 2 shows that the InceptionV3 model was the best-performing model, outperforming others; therefore, it was selected for use for the task. InceptionV3 is a state-of-the-art CNN architecture. It implements techniques such as label smoothing, factorized 7 × 7 convolutions, and a summary of label information to earlier layers with an auxiliary classifier. Batch normalization is also applied to the layers in the auxiliary network.

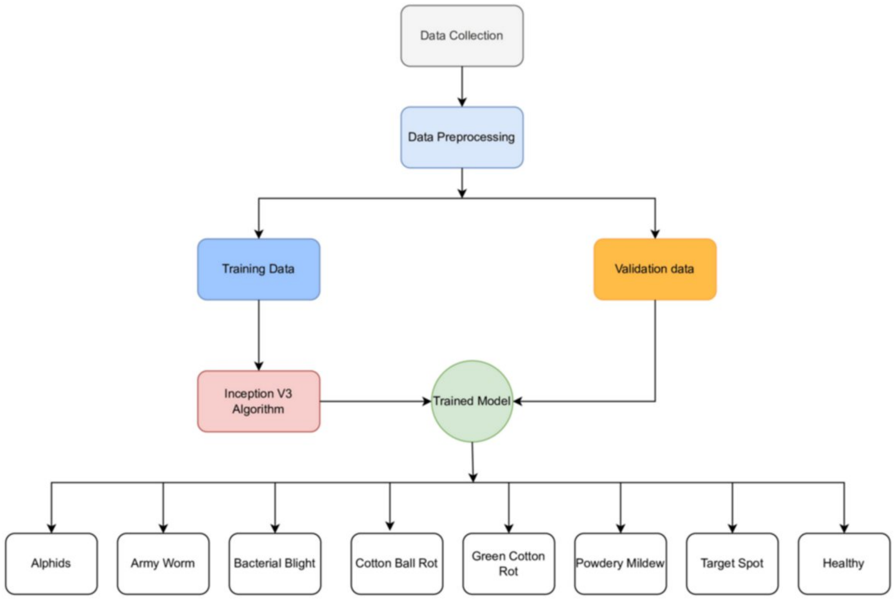

FIGURE 5: Workflow of the Model

The model workflow is illustrated in Figure 5. It begins with the data collection step, where images of cotton leaves affected by various diseases are collected. The collected data then undergoes pre-processing to improve quality, normalize formats, and prepare it for training and validation. After preprocessing, the raw dataset is divided into training and validation data. The model is then trained using the InceptionV3 on the training data. Finally, the trained model classifies input images into eight categories: Aphids, Armyworm, Bacterial Blight, Cotton Boll Rot, Green Cotton Boll, Powdery Mildew, Target Spot, and Healthy. This hierarchical workflow ensures accurate detection and classification of cotton diseases to promote timely agricultural intervention.

Transfer learning

Transfer learning in computer vision is an advanced technique, where a pre-trained model is retrained to do a specific task. This method leverages the pre-trained model's knowledge in solving new training data and enhancing the performance of a new problem, which requires less data and computational resources. Traditional machine learning approaches rely on large amounts of labelled data to generate good performance. However, data collection and model training can be very slow and computationally intensive in agriculture. The problem of this overfitting is solved via transfer learning by using a model that has already been trained on a large image dataset such as ImageNet. These models are already trained to detect a variety of features, including shapes, textures and colors. In cotton farming, the pretrained model can be fine-tuned with relatively small datasets of images of diseased cotton plants. During this fine-tuning process, the model's parameters are slightly changed for it to accurately identify diseases affecting cotton plants.

Results

The cotton industry has been significantly impacted by diseases such as Aphids, Army Worm, Bacterial Blight, Cotton Boll Rot, Green Cotton Boll, Powdery Mildew, and Target Spot. These diseases negatively affect crop yield and quality, making early detection crucial for preventing losses and enabling efficient management strategies.

Experimental setup

The experimental setup involved training two deep learning models - CNN and InceptionV3 - on a dataset specifically designed for cotton disease detection. The dataset, sourced from Kaggle [16], contained 800 images per class and was categorized into seven disease types and healthy samples. To maintain input quality, images were preprocessed in the following steps:

Resizing: Add codes to booklet images while keeping the aspect ratio.

Rescaling: Normalizing pixel values from 0 to 255 to 0 to 1.

Previous datasets are augmented with techniques that induce more variability using shearing, zooming, and horizontal flipping, achieved in this implementation to improve the generalization of the model.

The training process was performed on a system with an NVIDIA RTX 4060 GPU to speed up the computations. They were built using libraries in TensorFlow and Keras and accelerated using the GTX 1060 GPU to process large data accordingly.

Model performance and analysis

Both models performed well for the detection of diseases, as shown in Figure *6*. The blue line represents driving the training accuracy up (both models), and the orange line validates accuracy in a consistent range.

CNN Model: The training accuracy comes out to be 98.22%, whilst the validation accuracy comes out to be 90.62%. As such, the performances both on the training and validation sets are significantly high, but there is a huge gap between the training and validation performances, indicating the model is being overfitted.

The InceptionV3 Model: This transfer learning model outperformed models trained from scratch with a training accuracy of 97.51% and a validation accuracy of 96.92%. To generalize well across unseen validation data, the model architecture was designed with inception modules and auxiliary classifiers.

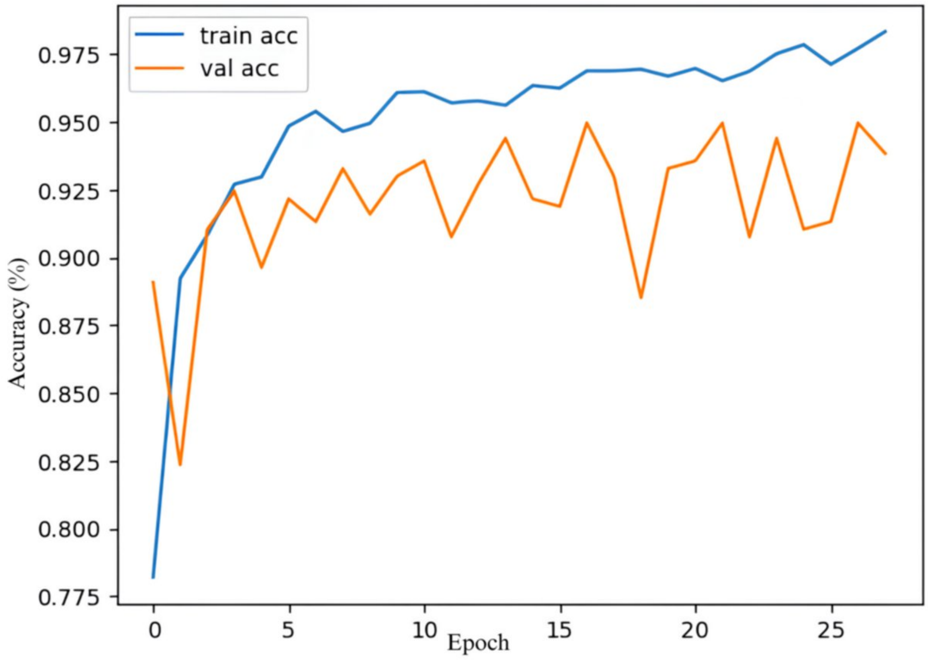

FIGURE 6: Training and Validation Accuracy

As seen in Figure 6, the accuracy curves for the model are shown for training and validation. The blue line here is the training accuracy, which increases and finally stands off at 98%, while the orange line is the validation accuracy, which ranges from around 93% to 96%. This is indicative of the success of the InceptionV3 architecture. These observations position InceptionV3 as the most favorable model for detecting cotton diseases, showcasing its strong ability to generalize and contactable performance, thus rendering it a strong candidate for real-time deployment in agricultural systems.

Class	Precision	Recall	F1-Score	Support
Aphids	0.97	0.85	0.9	39
Army worm	0.97	0.95	0.96	40
Bacterial Blight	0.97	0.95	0.96	40
Cotton Boll rot	1	0.84	0.91	61
Green Cotton Boll	0.87	1	0.93	59
Healthy leaf	0.91	1	0.95	39
Powdery Mildew	0.93	1	0.96	38
Target spot	0.9	0.93	0.92	41

TABLE 3: Performance Metrics of Cotton Disease and Pest Classification Model Across Different Classes

The classification performance of the InceptionV3 model for cotton disease detection is shown in Table 3, which proves that it is capable of accurately distinguishing between healthy leaves and different disease categories. The results presented in the table clearly reveal the model's capability of achieving the best precision, recall, and F1-scores for all classes.

Classification report

High Classification Accuracy Across All Classes

F1-scores of the model remained higher, being in the range of 0.90 to 0.96, confirming that the model is robust and can be used to detect both diseases in cotton plants and healthy leaves.

The categories Aphids, Army Worm, and Bacterial Blight have similar high precision of 97%, recall and F1-scores of about 0.96 or more.

Perfect Metrics for Specific Categories

Healthy Leaf and Green Cotton Boll achieved perfect recall (1.0), ensuring all instances of these categories were correctly identified.

Cotton Boll Rot achieved perfect precision (1.0), reflecting the model's ability to minimize false positives for this category.

Insights Into Visually Similar Categories

It indicates the potential visual similarity between Green Cotton Boll and Cotton Boll Rot with subtler overlapping features, as highlighted with higher metrics such as F1-score across both groups.

Despite this, the model still worked impressively well, with F1-scores of 0.91 (Cotton Boll Rot) and 0.93 (Green Cotton Boll).

Balanced Performance Across All Classes

The model showed stability in its performance, with an F1-score range of 0.90 to 0.96, further confirming its fidelity to be deployed in practical usage in the field.

In the Support column, which indicates the number of samples per class in our dataset, we can observe a relatively balanced dataset, which ensures that our model's performance is not biased by overrepresented or under-represented classes.

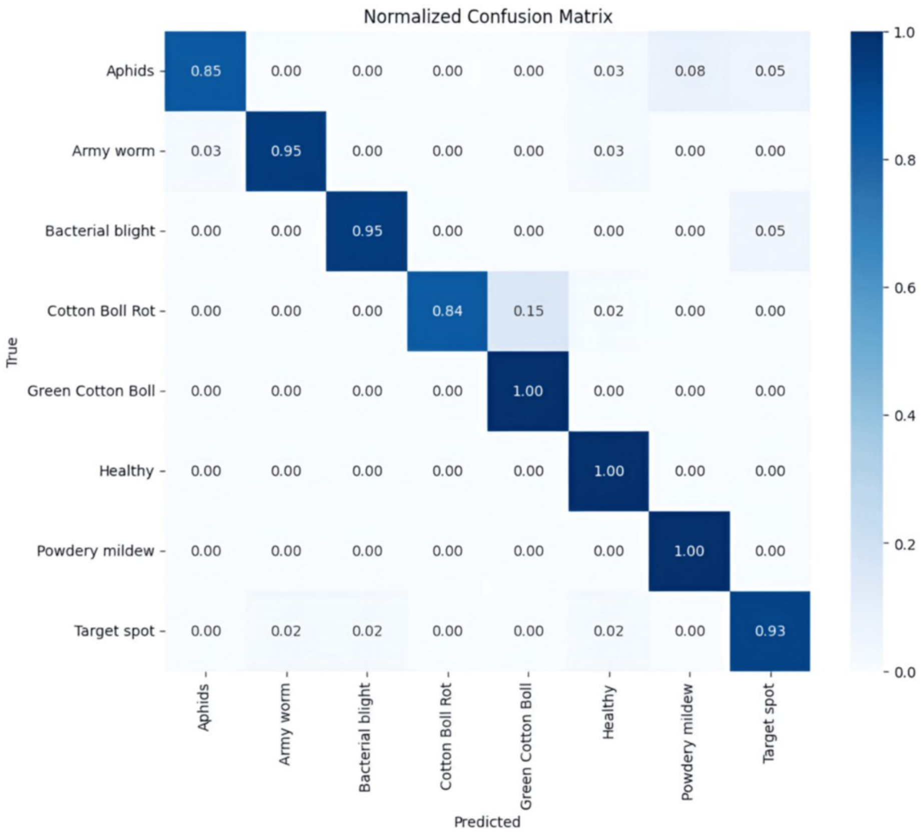

FIGURE 7: Normalized Confusion Matrix of Cotton Disease and Pest Classification Model

Figure 7 shows the confusion matrix of the normalized cotton disease and pest classification model, indicating whether it correctly predicted the different categories. The diagonal entries show true positive rates, where the classifier has predicted correctly each time. All other off-diagonal entries are the misclassifications/errors made by the model, which are very less. The higher the accuracy of classification is, the darker color appears on the diagonal of the matrix. The matrix indicates that the model is able to perform well in all classes with minimal confusion between categories, thus being reliable for practical deployment in disease and pest identification tasks.

Discussion

This study highlights how effective transfer learning with InceptionV3 can be for detecting diseases in cotton. Given the economic importance of cotton as a cash crop, it is crucial to tackle this issue. Even slight damage from diseases can lead to major financial setbacks for farmers, particularly for smallholders. The goal of this research is to meet the urgent demand for reliable, scalable, and efficient solutions to identify these diseases. The dataset utilized in this study includes a diverse set of images showcasing different diseases alongside healthy crops, allowing for the creation of a strong model.

Key objectives

Automation in Cotton Disease Detection

This research was done to develop an automated tool that can detect cotton diseases with high accuracy and low-resource usage. The solution enables farmers and agribusiness to respond quickly and effectively to reduce crop loss.

Economic Relevance of Cotton

Cotton, as a cash crop, is ubiquitous and supports the livelihoods of millions of farmers worldwide. Disease outbreaks can disrupt market supply, leading to poor yields and significant financial losses. This highlights the urgent need for a reliable system for early disease diagnosis to help minimize these losses.

Key observations

Performance of InceptionV3

InceptionV3 achieved a validation accuracy of 96.92%, which is significantly higher than that of a traditional CNN model, which achieved 90.62%. The superior performance of InceptionV3 can be attributed to:

Architectural Advantages: Inception modules were able to induce the model to learn multi-scale features, while auxiliary classifiers allowed the flow of gradients during training and ultimately optimized the resulting network.

Generalization Capability: Unlike the CNN model, which achieved a higher training accuracy of 98.22% but showed overfitting due to a gap between training and validation accuracies, InceptionV3 exhibited better generalization with minimal overfitting.

Impact of Preprocessing and Data Augmentation

Essential preprocessing and data augmentation techniques were used to improve the variability of the dataset and to facilitate generalization.

Preprocessing Steps: Resizing and rescaling normalized the images, ensuring consistency across the dataset.

Data Augmentation: Techniques like shearing, zooming, and horizontal flipping added some diversity to the dataset. These techniques increased the ability of the model to generalize and led to robust performance on unseen data.

However, despite these efforts, similar-looking diseases like bacterial blight and target spot were often misclassified, which suggests that more sophisticated feature extraction techniques or additional data modalities might be needed.

Analysis of Misclassification

The confusion matrix was normalized (Figure 7) and revealed important patterns in misclassifications:

Visual Similarity: It was difficult to tell apart diseases whose symptoms took on overlapping symptoms, including green cotton boll and cotton boll rot, due to similar patterns of discoloration.

Solution: These problems could be solved with high-resolution images or using multi-spectral data that would give deeper insights for classifications in these cases.

Challenges and drawbacks

Small Dataset Size

The small dataset size for the tested model was a limitation of generalization for different environmental circumstances, although data augmentation was used. To achieve better robustness, the dataset should first be extended with more samples from diverse datasets (different field crops and environments).

Computational Constraints

Although InceptionV3 was highly computationally efficient, the approach of training the larger dataset and deploying even advanced architectures will involve additional resources, and techniques, such as pruning and quantization, might bring deployment on edge devices closer.

Misclassification in Similar Classes

There were challenges in correct classification due to the profound resemblance between some diseases. These issues may be addressed by improving the input data through higher-resolution images or by adding multi-modal features.

Overcoming challenges

Improved Data Collection

In future work, we will focus on gathering a more representative and balanced dataset through outreach and collaboration with agricultural research institutions. Such an effort would improve representation with respect to different disease categories and environmental conditions.

Model Optimization

Optimization techniques such as pruning and quantization will be explored for the purpose of deploying on mobile devices and field monitoring systems. Such optimizations seek a balance between computational efficiency and model performance.

Real-World Deployment

The study highlights the potential for transfer learning using InceptionV3 in real-world applications, such as agriculture. The model is highly accurate and computationally efficient, suitable for deployment in resource-constrained settings. This approach could provide real-time disease detection, which would empower farmers to take appropriate action in a timely manner to reduce crop losses and improve yields. Future work will be dedicated to integrating the model into drone or smartphone-based platforms to increase accessibility and scalability.

Conclusions

The study proves that deep learning, and specifically transfer learning models, like the InceptionV3 architecture, is able to identify cotton diseases successfully, providing a useful tool for preventing crop yield loss. Through the use of pre-trained weights and the appropriate fine-tuning of the model, high validation accuracy was obtained despite sparse data coverage, indicating the effectiveness of transfer learning in agriculture. In comparison to a baseline CNN, InceptionV3 performed steadily better at all times, suggesting robustness and reliability for deployment in real-world applications. The capacity to preserve pre-trained hierarchical feature representations enabled the model to generalize more from a low amount of data, rendering it an effective solution for resource-limited cases where huge amounts of data collection might not be practical. The findings also indicate that augmenting the dataset with higher resolution imaging and using class-balancing methods can make the model even more accurate. Notably, the system proposed here is expected to find many applications in field monitoring using drones, smartphones, and other edge devices, allowing farmers to monitor crop health in real time and make data-informed decisions to avoid disease loss.

While these encouraging results are seen, there are a number of limitations to be taken into account. While the dataset was balanced, it only consisted of about 800 images per class, which restricts training diversity compared to actual farm conditions. Data augmentation did increase variability, but could not adequately reproduce variations in light, backdrop, and plant growth stage seen in field environments. This limitation specifically impacts the model's capability for visually similar disease differentiation, e.g., Cotton Boll Rot and Green Cotton Boll, and hence causes occasional misclassifications. Besides, InceptionV3 requires substantial computational demands, which can restrict its deployment on resource-limited platforms such as drones or smartphones without further optimization. Future efforts should be on increasing both quantitatively and qualitatively the dataset via collaborations with agricultural research institutions and farm data collection in real time. Model optimization methods, such as pruning, quantization, and knowledge distillation, can minimize computational requirements while maintaining accuracy, making it possible to create real-time, on-field applications. Eventually, incorporating the model into complete mobile or drone-based systems could give farmers real-time actionable insights to keep crops healthy and boost yield.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Varun R. Vangar, Vedank Wakalkar, Sushant Nadavade, Shiromani Waghmare, Grishma Y. Bobhate, Rutuja Patil

Acquisition, analysis, or interpretation of data: Varun R. Vangar, Vedank Wakalkar, Sushant Nadavade, Shiromani Waghmare

Drafting of the manuscript: Varun R. Vangar, Vedank Wakalkar, Sushant Nadavade, Shiromani Waghmare

Critical review of the manuscript for important intellectual content: Varun R. Vangar, Vedank Wakalkar, Sushant Nadavade, Shiromani Waghmare, Grishma Y. Bobhate, Rutuja Patil

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

References

1. USDA Foreign Agricultural Service: India: Cotton and products update - June 2024. (2024). Accessed: September 6, 2024: https://fas.usda.gov/data/india-cotton-and-products-update-june-2024.
2. Statista: Cotton acreage in China 2013-2023. (2024). Accessed: September 6, 2024: https://www.statista.com/statistics/241915/cotton-acreage-in-china/.
3. Statista: Leading cotton producing countries worldwide in 2022/2023. (2023). Accessed: September 6, 2024: https://www.statista.com/statistics/263055/cotton-production-worldwide-by-top-countries/.
4. Harakannanavar SS, Rudagi JM, Puranikmath VI, Siddiqua A, Pramodhini R: Plant leaf disease detection using computer vision and machine learning algorithms. Global Transitions Proceedings. 2022, 3:305-310. 10.1016/j.gltp.2022.03.016
5. Islam MM, Adil MAA, Talukder MA: DeepCrop: Deep learning-based crop disease prediction with web application. Journal of Agriculture and Food Research. 2023, 14:100764. 10.1016/j.jafr.2023.100764
6. Rothe PR, Kshirsagar RV: Cotton leaf disease identification using pattern recognition techniques. 2015 International Conference on Pervasive Computing (ICPC), Pune, India. 2015, 1-6. 10.1109/PERVASIVE.2015.7086983
7. Prashant U, Pravin S: Cotton leaf disease detection using instance segmentation. Journal of Cases on Information Technology. 2022, 24:1-10. 10.4018/jcit.296721
8. Bhimte NR, Thool VR: Diseases detection of cotton leaf spot using image processing and SVM classifier. 2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS), Madurai, India. 2018, 340-344. 10.1109/ICCONS.2018.8662906
9. Deepa AR, Chaurasia MA, Vamsi SBN, Kumar BM, Reddy VS, Anand KT: Plant diseases and pests detection using machine learning. 2023 3rd Asian Conference on Innovation in Technology (ASIANCON), Ravet IN, India. 2023, 1-4. 10.1109/ASIANCON58793.2023.10270264
10. Xin M, Wang Y: Research on image classification model based on deep convolution neural network. EURASIP Journal on Image and Video Processing. 2019, 2019:40. 10.1186/s13640-019-0417-8
11. Monigari V, Khyathi Sri G, Prathima T: Plant leaf disease prediction. International Journal for Research in Applied Science and Engineering Technology. 2021, 9:1295-1305. 10.22214/ijraset.2021.36582
12. Prajapati BS, Dabhi VK, Prajapati HB: A survey on detection and classification of cotton leaf diseases. 2016

international conference on electrical, electronics, and optimization techniques (ICEEOT), Chennai, India. 2016, 2499-2506. 10.1109/ICEEOT.2016.7755143

13. Joshi SH, Pandya JR, Chaudhary DH: Symptomatology of leaf spot disease of cotton caused by Curvularia lunata (Wakker) Boedijn. The Pharma Innovation. 2023, 12:363-365. 10.22271/tpi.2023.v12.i2e.18444
14. Revathi P, Hemalatha M: Cotton leaf spot diseases detection utilizing feature selection with skew divergence method. International Journal of Scientific Engineering and Technology. 2014, 3:22-30.
15. Sammy VM, Bobby DG, Nanette VD: Plant leaf detection and disease recognition using deep learning. 2019 IEEE Eurasia Conference on IOT, Communication and Engineering (ECICE), Yunlin, Taiwan. 2019, 579-582. 10.1109/ECICE47484.2019.8942686
16. Feng J, Ong WE, Teh WC, Zhang R: Enhanced crop disease detection with EfficientNet convolutional group-wise transformer. IEEE Access. 2024, 12:44147-44162. 10.1109/access.2024.3379303
17. Aderinwale DA, Sah M: Plant disease classification using an ensemble learning model of convolutional neural networks and vision transformers. 2025 7th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (ICHORA), Ankara, Turkiye. 2025, 1-10. 10.1109/ICHORA65333.2025.11017080
18. Thakur PS, Chaturvedi S, Seal A, Khanna P, Sheorey T, Ojha A: An ultra lightweight interpretable convolution-vision transformer fusion model for plant disease identification: ConViTX. IEEE Transactions on Computational Biology and Bioinformatics. 2025, 22:310-321. 10.1109/tcbbio.2024.3515149
19. Customized Cotton Disease Dataset. (2023). Accessed: September 6, 2024: https://www.kaggle.com/datasets/saeedazfar/customized-cotton-disease-dataset.