

**UNDERSTANDING ARTIFICIAL INTELLIGENCE'S PROGRESS IN PLANT
DISEASE DIAGNOSIS: A COMPARATIVE SURVEY ON COTTON PLANT**

Jesal Desai^{1*}, Dr. Amit Ganatra²

^{1*}Research Scholar, Department of Computer Engineering, Chandubhai S. Patel Institute of Technology, Faculty of Technology and Engineering, Charotar University of Science and Technology, Changa, India. Orcid Id- 0000-0003-4920-9246, Mail-jaddesai654@gmail.com

²Provost, Parul University (PU), Waghodia 391760, Vadodara, India, Orcid Id- 0000-0002-9993-9901

Abstract

Cross-domain transfer learning is becoming increasingly popular as an effective technique for disease detection in cotton crops with the help of artificial intelligence. Though the concept promises a lot, there are still a number of constraints and research loopholes that need to be filled. One big challenge is that many models are trained on clean, high-quality images from areas like medical scans or face detection, but real farm images are messy, natural, and often lower in quality. This mismatch makes it hard for the models to perform well in agriculture. Another challenge is the limited number of large labeled datasets. Particularly for cotton plant diseases, this hinders the ability to fine-tune or train models effectively. There are also no standardized benchmarks to assess and compare various models on an equitable basis. These gaps show why we need better models for agriculture, bigger datasets, and common testing methods, so AI tools can truly help farmers in the field. In this study, the performance of different YOLO versions was analyzed for cotton plant disease detection. Experimental results show that YOLOv8 achieved the best overall accuracy (95.14%) for cotton disease detection. The results provide valuable insights for researchers to advance further in this field and explore new opportunities.

Keywords: Cross-Domain Transfer Learning, Cotton Plant Disease Detection, Artificial Intelligence in Agriculture, Deep Learning Models, Agricultural Image Analysis

1. Introduction

The latest and the most effective variant of the YOLO (You Only Look Once) family of object detection, called YOLOv8, will be launched in January 2023 by Ultralytics. It builds upon the architecture of YOLOv5 with important performance enhancements, flexibility, and generalization of tasks. YOLOv8 is built on the adapted CSPDarknet53 network, where the traditional CspLayer is replaced with the more effective C2f module (cross-stage partial two convolutional layers), that improves the feature fusion and space learning abilities. One important use of the latest YOLOv8 is its anchor-free architecture, and with a second brain that performs objectivity, classification, and regression operations in independent branches. Through such a decoupling, tasks are more precisely and highly specialized. Even the architecture employs a Spatial Pyramid Pooling Fast (SPPF) layer that incorporates multiscale features into a fixed-size map to boost receptive field and detection accuracy. YOLOv8 is stabilized with the use of SiLU activation and batch normalization after each convolution. In the final layers of the output sections, the sigmoid function and the SoftMax are used on the objectness and the class probabilities respectively to obtain finer prediction confidence. YOLOv8 performs loss functions with Complete Intersection over Union (CIoU) and Distribution Focal Loss (DFL), to regress bounding boxes and Binary Cross-Entropy (BCE) to classify which perform better on small and overlapping objects. Overall, YOLOv8 offers the latest performance and retains high inference rates, that is why it can be implemented on real-time tasks in such spheres as agriculture, healthcare, and autonomous systems. Architecture of YOLOv8 is illustrated in figure 3.. Besides this, Azad Bakht et al. (2019) also used machine learning techniques to identify wheat rust at canopy level. These articles highlight how revolutionary AI can be in enhancing the early detection of diseases, the management of our crops and precision agriculture.

Although the field of artificial intelligence (AI) was promising as an instrument to detect plant diseases, the area of AI in cotton plant disease detection is underexplored. This is more so considering that cotton is economically a very significant global cash crop and highly susceptible to numerous pests and diseases, including Cotton Leaf Curl Virus (CLCuV),

bacterial blight and boll rot. Studies, such as that by Gagliuliano, et al. (2018), have shown that AI may be helpful in diagnosing cotton disease (Atila et al., 2021), which showed the effective diagnosis of leaf disease when sophisticated neural networks, like Squeeze-and-Excitation Networks (SE-Nets or SkillNets), are used, implying that AI may be useful in diagnosing various types of crops, including cotton. However, one of the biggest obstacles that Barbedo (2018a, 2018b) highlights that negatively affects the Deep learning and transfer learning model performance in plant pathology is the poor availability, quality, and diversity of annotated datasets. Against this background, the current project aims to provide an overall and comparative overview of the advancement and current status of plant disease diagnosis using AI, focusing on cotton. Through a systematic examination of the major methods, namely classical machine learning methods, deep learning models, transfer learning methods, and real-time deployment packages, this review attempts to highlight the research gaps, solve enduring problems, and suggest lessons for the development of reliable, scalable, and farmer-centric AI technologies specific to the cotton cultivation industry. The results not only present the vast possibilities for AI incorporation in cotton disease management but also pave the way for upcoming innovations that are capable of boosting agricultural production and sustainability.

1.1 Plant Diseases and Their Impacts on Cotton Yields

One of the economically important fiber crops in the world that is especially at stake is cotton, which suffers a series of plant Diseases that lower the quantity and fiber content of crops. Among the most widespread threats are Disease like Cotton Leaf Curl Virus, Bacterial Blight, Boll Rot, Fusarium Wilt, and Alternaria Leaf Spot that lead to retardation of the growth, rotting, and quite significant loss of crops unless they are addressed properly and in time. These conditions largely exist in the form of delicate or overlapping visual indications that are hard to detect with the naked eye particularly at the early stages and therefore, the necessity to possess computerized detection procedures. The introduction of AI-based equipment, especially equipment based on Deep learning and computer vision, has been accessible as having the capability of removing such diagnostic challenges. Brahimi et al. (2018) provided one such example and explained how learning can not only diagnose plant diseases but also enhance the capacity to explain the mechanism through which decisions are made by the AI model, which introduces confidence and interpretability to the latter.

Although the recent advances in Artificial Intelligence (AI) have been largely focused on high-value horticultural produce like tomatoes, grapes, and apples, the basic technologies possess a huge scope to be applied to cotton pathology also. Jointly, the research works of Bhandari et al., Burman et al., and Bok et al. (2020) each highlight moving away from subjective visual observations towards sensor-instrumented and computerized diagnostic systems. These innovations can improve the accuracy, consistency, and repeatability of the intensity of disease measurements, which has not been studied yet in cotton production. Moreover, Tu et al. (2019) invented a built-in artificial intelligence-based diagnostics platform tailored to the specific situation in the small-scale agricultural environment. This type of technology is especially important in helping to diagnose cotton disease in the area where medical treatment is not readily available particularly in rural and remote areas. These are valuable additions that are towards a more and more feasible application of AI in cotton production systems and demand more studies in this under-investigated field.

Moreover, the remote sensing and the deep neural networks are further streamlining the methods by which the disease pattern is mapped and predicted at the macro-level. Bazi et al. (2021) demonstrated that the Vision Transformer turned out to be applicable to image classification in remote sensing, and can be applied to large-scale agro-climatic zone cotton crop monitoring. In a similar manner, Cap et al. (2018) created an on-site and image-based leaf-detecting system that they wish to replace real-time and in-field disease detection systems. Together like this, even though in the majority of cases in non-cotton production, such technologies demonstrate the latent value of AI tools both to the preservation of the crop and to timely intervention. To address this technological gap in cotton disease management, there is an immediate need to simplify AI models, improve high-quality and cotton-specific datasets, and create solutions that are affordable and accessible and can be implemented by farmers. Such initiatives are important in facilitating an early diagnosis and achieving long-term cotton disease control in low-resource and remote agricultural regions.

2. Literature Review

Diagnostic Artificial Intelligence (AI) in the plant disease diagnosis has increasingly been a disruptive force behind the modern agriculture and sustainable farming due to its capacity to increase crop yield and support early disease detection. The recent two years of newly developed technologies through high-resolution images, a powerful computer, and Large-scale agricultural datasets have necessitated the installation of AI systems capable of detecting and classifying plant diseases with high rates of accuracy. The AI models have significantly contributed to large amounts of annotated data (i.e., tomatoes, rice, and wheat, among others) being used in detecting diseases, predicting yields, and rationally applying pesticides.

In comparison with this, AI application in cotton crop diseases detection is comparatively new. Cotton, is an important cash crop and one that is meaningful to the textile industry and millions of employees. However, it is prone to infections like cotton leaf curl virus, bacterial blight and boll rot. The diagnosis is extremely difficult due to both the convoluted nature of the symptomatology of the disease, and the local forms which it takes. Despite these difficulties, the body of AI-based cotton disease detection literature is limited, and it is an unexplored area how to boost yields, pest management, and it might be related to the economic performance of cotton producing regions.

This has been taken care of in this review by the structural classification of the available literature into the four overriding thematic categories. First, we discuss traditional machine learning (ML) algorithms: Support Vector Machines (SVMs), Decision Trees, K-Nearest Neighbors (K-NN) and Random Forests. These algorithms perform manual feature extraction, and are trained using domain knowledge. Next, the inquiry focuses on deep learning (DL) models like Convolutional Neural Networks (CNNs) since they can process more complicated image data and automation of feature extraction and decent diagnostic results. Third, the review addresses transfer learning and cross-domain models that help to overcome the lack of data by using pre-trained networks like VGG, ResNet, Inception, and MobileNet. The models are optimized to fit cotton data, and knowledge reuse aids in this. Cross-domain models also offer adaptation of other crops by taking advantage of similarities in disease characteristics. Finally, real-time deployment solutions like mobile applications, drone-based surveillance, edge-AI solutions, and IoT platforms are analyzed, which can offer farmers and agricultural workers timely and convenient Diagnosis.




Fig: Hierarchical Classification of AI Methods in Cotton Disease Detection

Each of the fields is thoroughly examined in terms of the accomplishments and the long-term concerns such as the lack of diversity in datasets, the overfitting of the DL models, the inconsistent performance in the real-world scenarios, the

insufficient multilingualism, and the lack of integration with the consulting systems. The review highlights the potential of AI in supporting cotton farming especially in low-income areas. By identifying the current limitations and knowledge gaps, this review preconditions the future research, which will encourage interdisciplinary cooperation and the development of credible, scalable, and easy-to-use AI tools to detect cotton diseases.

2.1 Traditional ML Approaches

Traditional machine learning (ML) algorithms such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Random Forests (RF) played a significant role in boosting early advances in automatic plant disease diagnosis. Such models were based on manually designed characteristics, and statistical characteristics such as color histogram, shape patterns, and texture-based matrices were manually feature-engineered out of leaf images (Mohan et al., 2016; Barbedo, 2018). The performance of these ML models was highly dependent on the discriminative capability, predictivity, and quality of the input features, and domain knowledge, as well as manual feature engineering, generally determined such features. Whereas these models provided computational efficiency and behaved very well in controlled laboratory environments, they were greatly diminished in performance in the Real-world. The various light intensities and clutter in the background, the posing of the image and the high intra-class variability were the key factors to model strength. Such constraints made the application of classical ML techniques to real world scenarios quite challenging in practice, as input data is, in most instances, non-homogenous, noisy and hard to normalize in an open field environment.

Nevertheless, these traditional machine learning methods provided a key base to developing present day AI-based Agricultural systems, particularly in the data or resource-constrained environment. The first algorithm-based classification models, including SVM and k-NN, played a crucial role in the development of baseline accuracy models in the detection of Diseases in crops like rice, wheat and tomatoes. These were then compared later as benchmarks on how well the operations of more advanced deep learning models were doing. Chakravarti et al. (2022) note that these early ML models were important to establish the potential of AI interventions in agriculture and developed a momentum to keep the research and use. The biggest drawback of these methods, however, was that they could not learn contextual dependencies or spatial hierarchies on the basis of image data directly since they were purely based on manually specified features and they did not learn directly with the raw pixel values. This was overcome to an extent with the introduction of deep learning (DL), particularly Convolutional Neural Networks (CNNs) that facilitate end-to-end Learning and automatic feature extraction of image information based on massive data sources. CNNs reinvented the concept of plant disease diagnosis by eliminating the need to design features manually and delivering superior results when operating under challenging, real world conditions.

Recent research points out the flaws of existing machine learning approaches and the increasing importance of advanced deep learning approaches to the diagnosis of plant diseases. Chandel et al. (2022) applied state-of-the-art image-based AI platforms with thermal and RGB spectral analysis of crop health monitoring tasks, which would have been highly challenging using the conventional ML approaches due to their low capabilities to handle complex, multi-modal data. Moreover, Qian et al. (2018) pointed out that the Traditional object detection algorithms, such as the region-based convolutional neural networks, were also weak.

To physical adversarial attacks, (R-CNN). This underscores the need to have more concrete and strong deep learning-based methods in agricultural diagnostics. Even in the best performing applications like wheat disease detection, Cheng et al. (2023) demonstrated that the addition of contextual information like environmental conditions to deep learning models significantly enhanced the classification performance and model flexibility, outperforming traditional methods.

Simply put, despite the fact that Traditional machine learning algorithms were essential in the early days of automated plant disease detection systems, their reliance on human-intensive feature engineering, limited robustness when faced with complex or changing environments, and failure to generalize to a wide range of conditions have led to their steadily being replaced by more flexible, scalable, and data-driven deep learning models. Nevertheless, the historical significance of traditional ML techniques and their further usefulness in low-resource settings where large annotated data sets or complex computational resources may be unavailable remain topical. Such models will still be valid in applications where lightweight, explainable, and low-cost AI deployment is needed.

2.2 Deep Learning in Plant Disease Diagnosis

Convolutional Neural Networks (CNNs) have significantly revolutionized the process of diagnosing plant diseases especially when compared to other machine learning techniques. Unlike the earlier models which rely on hand designed attributes, CNNs are capable of learning and extracting spatial features by using raw image data. They find great application in agriculture in visual classification problems. Popular CNN models include VGG16, ResNet, and Efficient Net and have achieved very good accuracy, and in most cases more than 90 percent accuracy when it comes to disease detection in crops like rice, tomatoes and wheat (Mohanty et al., 2016; Ferentinos, 2018; Das et al., 2020). These models can be especially effective at isolating finer disease symptoms, as opposed to background features, that are not very well isolated by handcrafted features. In a particular example, Das et al. (2020) showed that not only do deep learning models increase the predictive accuracy, but they also make automated disease detection systems highly scalable in real-world application.

Object detection models such as YOLO (You Only Look Once) have become particularly popular when it comes to recognizing and pinpointing particular diseased regions on a leaf in real time. The more recent models, YOLOv3 to YOLOv8, can identify and pinpoint several disease spots within one image at rapid speed and accuracy. YOLOv8, specifically, adds anchor-free detection, C2f modules (which enhance feature reuse), and independent heads for regression and classification. Such improvements have placed it among the quickest and most accurate models available today (Contributors, 2023). Its performance has been validated in recent research with cotton leaf datasets, where YOLOv8 produced great real-time performance with high accuracy and precision (Li et al., 2024; Zhang et al., 2022). This kind of speed and precision can be priceless in the agricultural sector, where early detection of diseases is beneficial in a big way in terms of whether or not to be able to keep the crop crop and reduce the use of pesticides.

The most significant drawback is that they rely on large training data sets that are well-labeled. CNNs and other deep learning models require large amounts of annotated data to work. With small or unbalanced data set training, they may experience overfitting or poor generalization (Cho et al., 2015). This presents a major problem in identifying cotton diseases where the annotated datasets are few and the images can be very diverse due to light, leaf orientation and disease development. Moreover, the computational requirements of these models are enormous, typically state-of-the-art GPUs and high-performance back-end infrastructure, so small-holder farmers or institutions with low-resource settings cannot afford the technology.

In order to address these limitations, Future studies may consider methods such as Data augmentation, synthetic data generation, and few-shot learning, so that models can be trained to be more accurate on small datasets. The accuracy and interpretability of AI systems can be enhanced with the help of combining hyperspectral imaging and multi-modal data fusion, as proposed by Duan et al. (2019). As lightweight CNNs and transformer-based vision architectures continue to improve, there is an ever-growing possibility of making AI-based plant health monitoring more affordable and practical to the mainstream in agriculture.

2.3 Cotton-specific Studies

The past few years have been characterized by development of artificial intelligence (AI) especially in its capability to enhance precision and speed of diagnosis of diseases and pests in cotton crops. The exemplary study that illustrates this point well is the study by Li et al. (2024), who created a state-of-the-art deep learning model named CFNet-VoV-GCSP-LSKNet-YOLOv8s. This hybrid model uses attention mechanisms, global context encoders, and light-weight convolutional layers to correctly detect and classify the numerous cotton diseases and pests, even in the field setting with noise, shadows, and overlapping symptoms that typically hinder traditional models. It is a tremendous step to precision farming and provides robust, high-performance solution to problematic the diagnostic environments.

Li et al. (2020) suggested the Generalized Focal Loss, a generalization of traditional loss functions used in dense object detection, in a prior but still current study. This method, though not originally intended to work with cotton, has been successfully applied to agricultural detection systems, including cotton pest detection. Training models to learn using small or clustered disease regions can increase the quality of the bounding box and the confidence that the model has in its predictions.

The other exciting discovery is by Mo and Wei (2025) who used StyleGAN2-ADA to augment data to solve the issue of limited annotated data. They also added a Decoupled Focused Self-Attention mechanism so that their light model can

concentrate on the most symptomatic regions of a cotton leaf. What is even more remarkable is that their system achieves excellent accuracy with relatively low processing needs, which makes it suitable to be used on mobile devices and edge computing platforms that can be directly implemented in the field by farmers or agronomists. Cotton diagnostics has even received a ripple effect in studies in other fields. As an indicator, in the research of Lotter, Sorensen, and Cox (2017) of the mammogram classification, their concepts of multi-scale CNNs and curriculum learning have influenced more tangible training approaches of agricultural models, especially those necessitating fine-grained information coverage or multi-level categorization framework.

All these advances were provided by the previous studies such as Mohan et al. (2016) and Mohanty, Hughes, and Salathé (2016). These papers proved that one could train convolutional neural networks (CNNs) with massive amounts of diseased plant leaves and obtain impressive results across a wide range of crops, including cotton. Their solution introduced image-based deep learning to the radar of viable solutions to plant pathology and preconditioned the creation of cotton-specific datasets and benchmarking protocols.

Table: Comparison of Machine Learning and Deep Learning Models for Plant Disease Detection

Model	Detection Type	Dataset Used	Accuracy / mAP	Strengths	Limitations
SVM (Support Vector Machine)	Classification	Custom dataset (cotton leaf images)	~85–90% accuracy	Simple to implement, effective for small datasets	Poor generalization in complex backgrounds; needs manual feature extraction
Random Forest (RF)	Classification	Image dataset of cotton leaf diseases	~88% accuracy	Robust to overfitting, interpretable	Limited scalability; less accurate than DL for large image data
CNN (Basic)	Image Classification	PlantVillage (adapted for cotton in some studies)	90–95% accuracy	Learns hierarchical features; better than ML in visual tasks	Needs a large labeled dataset; may overfit small data
ResNet-50	Image Classification (DL)	PlantVillage + Custom Cotton Dataset	~96% accuracy	Deep residual learning allows accurate feature extraction	Computationally expensive; slower training
VGG-16 / VGG-19	Image Classification	Custom Cotton Dataset	92–94% accuracy	Good at capturing spatial features	High parameter count; not suitable for edge devices
InceptionV3	Transfer Learning	Pre-trained + fine-tuned on cotton data	~95–96% accuracy	Efficient; fewer parameters; better for mobile deployment	Limited explainability; requires preprocessing
YOLOv5	Object Detection	Custom annotated cotton dataset	mAP ~87–92%	Fast, real-time detection; good for field deployment	Struggles with small object detection or poor lighting
Faster R-CNN	Object Detection	Annotated leaf dataset (general & cotton diseases)	mAP ~85–90%	High detection accuracy; region proposal networks enhance precision	Computationally heavy; not ideal for real-time
MobileNetV2	Lightweight Classification	PlantVillage (cotton subset)	~90% accuracy	Suitable for smartphones and IoT devices; efficient	Less accurate than deeper CNNs
Vision Transformer (ViT)	Classification	Synthetic & real cotton datasets	~97% accuracy (in some studies)	High performance on complex datasets; attention-based interpretability	Requires a large dataset and training time
EfficientNet	Classification	Cotton leaf images (custom)	~95–96% accuracy	Excellent trade-off between accuracy and efficiency	Training from scratch is resource-intensive

Together, these contributions mark a shift from rule-based systems and manually engineered features toward advanced, end-to-end AI pipelines that are capable of Real-time illness detection in the field. As AI models become better at identifying subtle symptoms, handling imbalanced data, and operating on low-power devices, the future of scalable, accessible cotton crop monitoring looks increasingly promising.

2.4 Cross-domain Transfer (Medical, Facial Recognition)

Considering cross-domain transfer learning as one of the efficient tools of creating the field of AI-based plant disease diagnostics, the latter is specifically applicable in the scenarios where the information regarding the agriculture is inadequate or lacking in diversity. Plant pathology is achieving remarkable success in other methods initially developed elsewhere, including biometrics and medical imaging.

Other tasks like skin lesion classification (Pham et al., 2018) applied to detect plant diseases are performed in large quantities using deep convolutional neural networks (CNNs). The models are applicable in poor conditions too, e.g. different lighting and multifaceted background effects, which are seen in medical and agricultural image processing. In much the same vein, new biometric security features (even in a decentralized farm) can be used by intelligent and privacy-conscious plant monitoring systems (including cancelable iris templates with random projections and phase encoding) (Rajasekar, Premalatha, and Sathya, 2021). Rajasekar et al. (2021) demonstrated this aspect of crossovers by demonstrating the training of deep learning models, which are trained on large image databases like ImageNet on the particular task of cotton plant disease classification. Their study has determined that the method in generalization and accuracy of the models classification is greatly improved and can help in overcoming the weaknesses that are usually common in overfitting and lack of data that are common in agricultural data.

In another important example, real-time object detection models like YOLO (You Only Look Once) (Redmon et al., 2016), originally developed for applications like facial recognition and object tracking, have been successfully repurposed for the detection of crop diseases. Such systems provide the necessary speed and scalability to be used in an actual field environment. The second step of this process as outlined by Munawar et al. (2024) is to refine such models technologically, yet to make them applicable in practice to farmers through the use of convenient interfaces, multilingual services, and the provision of real-time agricultural advice services.

On the macro level, this convergence of ideas between various disciplines demonstrates that not only is the technical transfer learning between domains feasible, but there is also a potential of interdisciplinary innovation due to this convergence. Scientists can create more specific and cost-effective systems of diagnosis by adjusting such high-accuracy models to the unique visual attributes of cotton-plant diseases such as leaf texture, color-alteration, and patterns of symptoms. It will be vital to consider ethical and usability factors in the future as well to ensure that these AI tools do not harm but become useful and scalable in low-resource agricultural societies.

2.5 Limitations and Research Gaps

Although cross-domain transfer learning can be very promising in enhancing cotton plant disease detection, it has a number of critical limitations and research gaps that must be overcome. A major issue is that there is a huge disparity between the source domains, such as medical imaging or facial recognition, and the target domain of agricultural diagnostics. An example of this is that medical or facial recognition images are typically of high quality, captured under controlled settings, and have accurate annotations. Conversely, cotton leaf images are commonly taken in natural settings with unpredictable lighting, different backgrounds, and lower resolution, which makes them more difficult to train to work with without fine-tuning. This domain mismatch has an impact on the generalization capabilities of the model when used in cotton diseases.

Moreover, the majority of current transfer learning models, including the ones that are trained on skin lesion detection or iris and face recognition, are trained to recognize regular anatomical features. Nevertheless, the diseases of plants frequently have irregular forms, different colors, and different textures, which are not optimized in these models. The other important gap is the absence of domain-specific adaptation methodologies and the limited access to massive annotated datasets of cotton disease images. Transfer learning is not as effective in the absence of sufficient high-quality, labeled data.

In addition, even though models such as YOLO (You Only Look Once) have shown promising real-time object detection performance, their performance remains largely reliant on the availability of domain-specific training data. The existing studies are also characterized by the lack of standardized criteria that can be used to assess the effectiveness of cross-domain models in agricultural settings. This causes uneven outcomes and complicates the comparison of the effectiveness of various methods. The way forward is to work on improving domain adaptation methods, on large and varied datasets unique to cotton diseases, and on shared evaluation metrics. Such attempts will contribute to the enhancement of the reliability and scalability of AI systems in real-life agricultural settings.

3. Model Architecture: YOLOv8 for Detection of Cotton Plant Disease

The University of Washington's Ali Farhadi and Joseph Redmon developed YOLO (You Only Look Once) in 2016. It altered the paradigm of object detection, in which the detection strategy is a single stage of regressed problem rather than multi-stage techniques, which are founded on sliding window or region proposal networks. This was an innovativeness that enabled Real-time Accurate and efficient performance. The first model to perform both classification and localization in a single pass of the neural network was YOLOv1. However, it had problems in small and variable sized object detection. In 2016, YOLOv2 was released, which eliminated these limitations by applying batch normalization, anchor boxes and dimension clustering, which enabled significantly higher performance on objects of different sizes. Released in 2018, YOLOv3 enhanced the detection with a more detailed and stronger backbone, alternate anchor boxes and multi-scale detection and spatial pyramid pooling to represent features in a better manner. YOLOv4 in 2020 was a chronological extension of the history of the architecture YOLOv3, introducing additional convolutional layers, and additional training and regularization methods and increased robustness and generalization. The other change that was significant was the change to YOLOv5 that ceased being implemented on the Darknet, but in PyTorch, and this enabled the model to be more accessible and adjustable. It also offered automatic hyperparameter optimization, inbuilt experiment monitoring and simplified export to simplified training and deployment processes. With architectural improvements in YOLOv6, the model became trainable in less time and detected more objects with a more beneficial accuracy computation efficiency tradeoff. New tasks such as the estimation of postures were added to the model based on the COCO key point dataset in YOLOv7.

The latest version of the series, YOLOv8, comes with considerable speed and accuracy improvements. It incorporates modern architectural changes and optimization techniques, which make it suitable for a broad scope of real-life applications, such as Agriculture automation and cotton plant disease detection.

Across its versions, the YOLO framework has demonstrated continuous improvements in several core areas: backbone architecture, loss function design, anchor box strategies, and input resolution handling. Each version has progressively enhanced object detection performance while reducing inference time, enabling deployment in both research and real-time field applications. Modern object detectors, including the YOLO family, are typically. The backbone, neck, and head are the three basic components of the body. As shown in Figure 1, the backbone (typically a convolutional neural network) captures multiscale information from the input image. This streamlined architecture enabled real-time object detection. The backbone (typically a CNN) extracts hierarchical features; the neck refines these features across scales using modules such as the Feature Pyramid Network (FPN); and the head generates final predictions, followed by non-maximum suppression to eliminate redundant detections.




Figure 2: Modular architecture of object detectors comprising backbone, neck, and head for feature extraction, refinement, and prediction

3.1 Architecture of YOLOv8

Ultralytics is set to introduce the newest and the succeeding best version of the YOLO (You Only Look Once) family of object detectors in January 2023, known as YOLOv8. It expands on the architecture of YOLOv5 and offers valuable performance improvements, flexibility and tasks generalization. YOLOv8 is based on the modified CSPDarknet53 network, in which the standard CspLayer is substituted by the more functional C2f module (cross-stage partial two convolutional layers), which enhances the capability of fusion of features and space learning. Another application of the more recent YOLOv8 is its anchor-free architecture, and having a second brain with objectivity, classification, and regression tasks in separate operations. By means of such a decoupling tasks become highly specialized and more precise. Not to mention the architecture itself uses a Spatial Pyramid Pooling Fast (SPPF) layer which adds multiscale property to a constant size map in order to enhance receptive field and detection accuracy. Each convolution is followed by the stabilization of YOLOv8 with the help of SiLU activation and the application of batch normalization. The sigmoid function and the SoftMax are applied at the last levels of the output parts respectively on the objectness and the probabilities of the classes to derive more accurate prediction confidence. To regress the bounding boxes and Binary Cross-Entropy (BCE) to classify which yield better results on small and overlapping objects, YOLOv8 uses both Complete Intersection over Union (CIoU) and Distribution Focal Loss (DFL) loss functions. On the whole, YOLOv8 has the best up-to-date performance and can retain high inference rates, that is why it could be applied to real-time tasks in such areas as agriculture, healthcare, and autonomous systems. YOLOv8 architecture is shown in figure 3.




Figure 3 YOLOv8 architecture with a modified CSPDarknet53 backbone, featuring C2f modules, SPPF layer for multiscale pooling, and a detached head for objectivity, classification, and regression problems.

3.2 Experimental Setup

To develop and train our models, we used the Google Colab, the cloud-based Jupyter notebook environment, which gives access to an NVIDIA Tesla T4 with 15,110 MiB memory. Its connection with Google Drive allows the ease of data management and provides remote access. The experiments were conducted using the Ultralytics YOLOv8 PyTorch implementation (v8.0.100+), which offers efficient training pipelines and pretrained models. A custom cotton plant disease dataset was curated from a variety of open-source platforms, including Kaggle, MS COCO, Roboflow, and Google Images. The final dataset consists of 2,047 labeled images, annotated in YOLO format (.txt) and compatible with .xml (Pascal VOC) for broader framework adaptability (Darknet, TensorFlow, and PyTorch). The dataset encompasses various cotton plant leaf classes, including healthy, bacterial blight, leaf curl virus, mealybug, and leaf spot. Annotation was carried out using bounding boxes with class indices. Several data augmentation techniques were used to increase model generalization, including horizontal flipping, brightness modulation, mosaic augmentation, rotation, and scaling. The data was split into an 80:20 ratio for training and validation. The model was trained for 50 epochs with a batch size of 16, and the Adam W optimizer was used at a learning rate of 1e-3. The training was conducted on a 640×640 image resolution using CUDA-enabled GPU acceleration.

Table 1. Experimental Parameter Settings

Parameters	Values
Learning Rate	0.001
Dropout	0.02
Optimizer	AdamW
Shearing	-0.2 to +0.2
Horizontal Flipping	True
Rotating	-20° to +20°
Zooming	0.8 to 1.5
Validation Split	0.2

4. Results and Discussion

4.1 Performance Evaluation

A set of Standard performance measurements has been used in order to measure the performance of the suggested cotton disease-detecting system to the fullest extent possible. They include some of precision, recall, F1-score, accuracy, and mean average precision (mAP). These are measures obtained based on the confusion matrix and they include True Positive (TP), False Positive (FP), True Negative (TN), and False Negative (FN) values which are used to measure the classification capability of the detection model. Accuracy: This is computed by the ratio of cases that were predicted correctly to the overall number of predictions made as a percentage:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN} \quad (1)$$

Precision measures the number of projected positive samples that were really positive.

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

Recall (also called sensitivity) measures how many of the actual positive cases were correctly predicted:

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

The F1-score, the harmonic mean of precision and recall, balances the two metrics and is given by:

$$F1 - Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (4)$$

In object detection tasks, Intersection over Union (IoU) is used to evaluate the overlap between the predicted bounding boxes and the ground truth, with values ranging from 0 (no overlap) to 1 (perfect match). The Average Precision (AP) is determined as the area under the accuracy-recall curve for each class, and the mean average precision (mAP) is the mean of AP values over all courses.

$$mAP = \frac{1}{C} \sum_{i=1}^C AP_i \quad (5)$$

This collection of metrics gives a solid basis to the analysis of the model reliability and its applicability to the deployment in precision agriculture projects, including early detection of cotton leaf diseases.

4.2 Experimental Results

We used a custom dataset with multiple leaf conditions to train our deep learning model: Bacterial Blight, Curl Virus, Herbicide Growth Damage, Leaf Hopper Jassids, Leaf Redding, Leaf Variegation, Healthy Leaf, and Background to assess the efficacy of our model at detecting cotton plant diseases. It was trained with 50 epochs. Figure 3 illustrates the change in training and validation performance with the number of epochs (50) with the help of four major plots: training and validation loss, top 1 and top 5 accuracy. The training loss curve demonstrates a rapid decrease and reaches a level of zero and has effective learning. The loss of Validation occurs quickly. it converges with some fluctuations but still low in general, thus meaning that the model can be used on new data. The highest accuracy of the top 1 rapidly rises to above the 95 percent mark after approximately 20 epochs and flattens around 96 percent, which is an indication of good classification. The top-five accuracy is both high and constant, reaching up to 100 percent in the first part of the training, which again demonstrates the reliability of the model in pushing correct predictions to the top of its predictions.




Figure 3a: Accuracy Curve




Figure 3b: Loss Curve




Figure 3: Training and validation performance of the model across 50 epochs.

Accuracy curve showing the improvement and stabilization of classification accuracy.

Validation loss curve depicting loss reduction and convergence behavior.

These trends validate that the model was learning the distinguishing characteristics of every one of the classes and remained strong on the validation set. This ability to generalize to new information is indicated by the use of the model on new data in a consistent manner across epochs.

The confusion matrix analysis also confirms the model resilience. The normalized confusion matrix, in Figure 4, shows that most of the classes were correctly predicted. Leaf Redding, Herbicide Growth Damage and Bacterial Blight had a 100 percent accuracy in classification. Leaf Variegation and Leaf Hopper Jassids also performed with a good accuracy of 91 and 86 respectively, although the Healthy Leaf class achieved slightly lower performance with 84% accuracy, and was mainly misclassified as Curl Virus, perhaps because of visual similarity between leaves with and without the virus. Likewise, the cases of Curl Virus were occasionally confounded with Healthy Leaf, which means that there could be some overlap between their visual characteristics or a minor imbalance in the data.




Figure 4: Confusion matrix for the cotton disease classification model, showing per-class prediction accuracy.

Overall, the model has a top-1 testing accuracy of 95.14% and a training precision of 94.46%. These findings validate the model's great capability in recognizing and categorizing diverse cotton. Leaf diseases. The few observed misclassifications suggest that further improvement may be achieved through additional data augmentation or refinement of visually similar class boundaries. Nevertheless, the current performance is promising and supports the application of the model for real-time disease identification in precision agriculture.

To validate, we assess the efficiency of our suggested approach by performing a comparison study using recent YOLO-based approaches used in cotton disease detection tasks. Table 2 presents the major performance parameters, including accuracy, precision, and model characteristics from previous studies.

Table 2: Metrics including accuracy, precision, and model characteristics

Model Type	Accuracy/mAp	Precision	Features
YOLOv8	95.14	94.46	Trained on 8 Class Cotton data set, Stable Accuracy 50 epochs
YOLOv3	94.6	94.4	Used for Pest/Disease recognition in cotton leaves
YOLOv5	91.1	89.4	Enhance with small target detection module
YOLOv8 (CFNet-VOV-GCSP)	93.7	89.9	Lightweight model using an advanced backbone structure

These findings prove the excellence of our method in the Accurate detection and classification of cotton leaf diseases. The consistent convergence, top-1 and top-5 accuracy, and the consistent performance of the model on the validation performance are factors why it is a more consistent and generalizable solution to real-time precision agriculture applications.

5. Conclusion

In this study, we developed a deep learning model of YOLOv8 that is used to detect and classify diseases in cotton plants automatically. The model trained on a custom dataset of eight leaf condition classes showed good generalization with a top-1 testing accuracy of 95.14 and a precision of 94.46. Performance metrics and confusion matrix analysis confirmed its robustness, with near-perfect classification for several classes and minimal confusion between visually similar categories. YOLOv3 also performed strongly, and YOLOv5 and lightweight YOLOv8 variants showed trade-offs in precision and efficiency. Comparative results show that our model outperforms prior YOLO-based models, validating its superiority in accuracy, learning stability, and real-time applicability. These findings emphasize the model's potential for integration into precision agricultural techniques provide fast and precise disease management.

References

- [1] Ashqar, B., & Abu-Naser, S. (2019). Image-based tomato leaves disease detection using deep learning. *International Journal of Engineering Research*, 2(12), 10–16.
- [2] Atanassova, S., Nikolov, P., Valchev, N., Masheva, S., & Yorgov, D. (2019). Early detection of powdery mildew (*Podosphaera xanthii*) on cucumber leaves based on visible and near-infrared spectroscopy. *AIP Conference Proceedings*, 2075, 160014.
- [3] Atila, U., Uçar, M., Akyol, K., & Uçar, E. (2021). Plant leaf disease classification using an EfficientNet deep learning model. *Ecological Informatics*, 61, 101182.
- [4] Atole, R. R., & Park, D. (2018). A multiclass deep convolutional neural network classifier for the detection of common rice plant anomalies. *International Journal of Advanced Computer Science and Applications*.

- [5] Azadbakht, M., Ashourloo, D., Aghighi, H., Radiom, S., & Alimohammadi, A. (2019). Wheat leaf rust detection at canopy scale under different LAI levels using machine learning techniques. *Computers and Electronics in Agriculture*. <https://doi.org/10.1016/j.compag.2018.11.016>
- [6] Barbedo, J. G. A. (2018a). Factors influencing the use of deep learning for plant disease recognition. *Biosystems Engineering*, 172, 84–91.
- [7] Barbedo, J. G. A. (2018b). Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification. *Computers and Electronics in Agriculture*, 153, 46–53.
- [8] Bazi, Y., Bashmal, L., Al Rahhal, M. M., Dayil, R. A., & Ajlan, N. A. (2021). Vision transformers for remote sensing image classification. *Remote Sensing*, 13(3), 1–20. <https://doi.org/10.3390/rs13030516>
- [9] Bhandari, M., Shahi, T. B., Neupane, A., & Walsh, K. B. (2023). BotanicX-AI: Identification of tomato leaf diseases using an explanation-driven deep-learning model. *Journal of Imaging*, 9(2). <https://doi.org/10.3390/jimaging9020053>
- [10] Bi, S., Zhang, Y., Dong, M., & Min, H. (2019). An embedded inference framework for convolutional neural network applications. *IEEE Access*. <https://doi.org/10.1109/ACCESS.2019.2956080>
- [11] Bierman, A., LaPlumm, T., Cadle-Davidson, L., Gadoury, D., Martinez, D., & Sapkota, S., et al. (2019). A high-throughput phenotyping system using machine vision to quantify the severity of grapevine powdery mildew. *Plant Phenomics*, 2019, 1–13. <https://doi.org/10.34133/2019/9209727>
- [12] Bock, C. H., Barbedo, J. G., Del Ponte, E. M., Bohnenkamp, D., & Mahlein, A. K. (2020). From visual estimates to fully automated sensor-based measurements of plant disease severity: Status and challenges for improving accuracy. *Phytopathology Research*, 2, 1–30.
- [13] Brahimi, M., Arsenovic, M., Laraba, S., Sladojevic, S., Boukhalfa, K., & Moussaoui, A. (2018). Deep learning for plant diseases: Detection and saliency map visualization. In J. Zhou & F. Chen (Eds.), *Human and Machine Learning* (pp. 93–117). Springer. https://doi.org/10.1007/978-3-319-90403-0_6
- [14] Cap, H. Q., Suwa, K., Fujita, E., Kagiwada, S., Uga, H., & Iyatomi, H. (2018). A deep learning approach for on-site plant leaf detection. In *Proceedings of the IEEE 14th International Colloquium on Signal Processing & Its Applications (CSPA)* (pp. 118–122). IEEE.
- [15] Chakraborty, S. K., Chandel, N. S., Jat, D., Tiwari, M. K., Rajwade, Y. A., & Subeesh, A. (2022). Deep learning approaches and interventions for futuristic engineering in agriculture. *Neural Computing and Applications*, 34(23), 20539–20573.
- [16] Chandel, N. S., Chakraborty, S. K., Chandel, A. K., Dubey, K., Subeesh, A., Jat, D., & Rajwade, Y. A. (2024). State-of-the-art AI-enabled mobile device for real-time water stress detection of field crops. *Engineering Applications of Artificial Intelligence*, 131, 107863.
- [17] Chandel, N. S., Rajwade, Y. A., Dubey, K., Chandel, A. K., Subeesh, A., & Tiwari, M. K. (2022). Water stress identification of winter wheat crop with state-of-the-art AI techniques and high-resolution thermal-RGB imagery. *Plants*, 11(23), 3344.
- [18] Chen, S.-T., Cornelius, C., Martin, J., & Chau, D. H. (2018). Robust physical adversarial attack on Faster R-CNN object detector. In the *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.
- [19] Cheng, S., Cheng, H., Yang, R., Zhou, J., Li, Z., Shi, B., Lee, M., & Ma, Q. (2023). A high-performance wheat disease detection based on position information. *Plants*, 12(5). <https://doi.org/10.3390/plants12051191>
- [20] Cho, J., Lee, K., Shin, E., Choy, G., & Do, S. (2015). How much data is needed to train a medical image deep learning system to achieve the necessary high accuracy? *arXiv preprint arXiv:1511.06348*.
- [21] Contributors, M. (2023). *YOLOv8 by MMYOLO*. <https://github.com/open-mmlab/mmyolo/tree/main/configs/yolov8>
- [22] Das, A., Mallick, C., & Dutta, S. (2020). Deep learning-based automated feature engineering for rice leaf disease

- [23] Davidson, M. W., & Haim, D. A. (2020). The promise and peril of using big data for social science research. In S. R. Cotten (Ed.), *Future directions in social science research* (pp. 91–112). Springer. https://doi.org/10.1007/978-3-030-40266-6_5
- [24] Dong, Y., Zhao, Z., Wang, Z., & Liu, B. (2022). Real-time cucumber disease recognition using deep learning. *Computers and Electronics in Agriculture*, 193, 106689.
- [25] Duan, P., Yang, G., Sun, J., Xu, B., Wang, G., & Yang, H. (2019). A review of hyperspectral imaging for plant disease detection in agriculture. *Remote Sensing*, 11(10), 1180. <https://doi.org/10.3390/rs11101180>
- [26] Ferentinos, K. P. (2018). Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, 145, 311–318. <https://doi.org/10.1016/j.compag.2018.01.009>
- [27] Fuentes, A., Im, D. H., Yoon, S., & Park, D. S. (2017). Spectral analysis of CNN for tomato disease identification. In *International Conference on Artificial Intelligence and Soft Computing (ICAISC 2017)* (pp. 40–51). Springer.
- [28] Greenspan, H., van Ginneken, B., & Summers, R. M. (Eds.). (2017). *Deep learning in medical image analysis and multimodal learning for clinical decision support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017*. Springer.
- [29] Hakani, R., & Rawat, A. (2024). Edge computing-driven real-time drone detection using YOLOv9 and NVIDIA Jetson Nano. *Drones*, 8(11), 680. <https://doi.org/10.3390/drones8110680>
- [30] Han, L., Haleem, M. S., & Taylor, M. (2017). A novel computer vision-based approach to automatic detection and severity assessment of crop diseases. *International Journal of Applied Engineering Research*, 12(18), 7169–7175.
- [31] Hatami, N., Gavet, Y., & Debayle, J. (2017). Classification of time-series images using deep convolutional neural networks. In *Proceedings of the 10th International Conference on Machine Vision (ICMV)* (Vol. 10696). Vienna, Austria.
- [32] Jocher, G., Chaurasia, A., & Qiu, J. (2023). *YOLO by Ultralytics*. <https://github.com/ultralytics/ultralytics>
- [33] Karki, M., Cho, J., Lee, E., Hahm, M. H., & Yoon, S. Y. (2020). CT window trainable neural network for improving intracranial hemorrhage detection by combining multiple settings. *Artificial Intelligence in Medicine*, 106, 101874.
- [34] Kawahara, J., & Hamarneh, G. (2017). Multi-resolution-tract CNN with hybrid pretrained and skin-lesion trained layers. In the *Machine Learning in Medical Imaging (MLMI) Workshop*, part of MICCAI.
- [35] Li, R., He, Y., Li, Y., Qin, W., Abbas, A., Ji, R., Li, S., Wu, Y., Sun, X., & Yang, J. (2024). Identification of cotton pest and disease based on CFNet-VoV-GCSP-LSKNet-YOLOv8s: A new era of precision agriculture. *Frontiers in Plant Science*, 15, 1348402. <https://doi.org/10.3389/fpls.2024.1348402>
- [36] Li, X., Wang, W., Wu, L., Chen, S., Hu, X., Li, J., Tang, J., & Yang, J. (2020). Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in Neural Information Processing Systems*, 33, 21002–21012.
- [37] Lotter, W., Sorensen, G., & Cox, D. (2017). A multi-scale CNN and curriculum learning strategy for mammogram classification. *arXiv preprint arXiv:1707.06978*.
- [38] Mo, H., & Wei, L. (2025). Lightweight detection of cotton leaf diseases using StyleGAN2-ADA and decoupled focused self-attention. *Journal of King Saud University - Computer and Information Sciences*, 37(41). <https://doi.org/10.1007/s44443-025-00054-x>
- [39] Mohan, K. J., Balasubramanian, M., & Palanivel, S. (2016). Detection and recognition of diseases from paddy plant leaf image. *International Journal of Computer Applications*, 144(12), 1–5.
- [40] Mohanty, S., Hughes, D., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7, Article 1419. <https://doi.org/10.3389/fpls.2016.01419>

- [41] Munawar, S. M., Rajendiran, D., & Sabjan, K. B. (2024). *Plant diseases diagnosis with artificial intelligence (AI)*. In P. S. Babu, M. Balasubramanian, & M. Arul Murugan (Eds.), *Microbial data intelligence and computational techniques for sustainable computing* (pp. 187–193). Springer. https://doi.org/10.1007/978-981-99-9621-6_12
- [42] Pham, T.-C., Luong, C.-M., Visani, M., & Hoang, V.-D. (2018). Deep CNN and data augmentation for skin lesion classification. Springer International Publishing.
- [43] Rajasekar, V., Premalatha, J., & Sathya, K. (2021). Cancelable iris template for secure authentication based on random projection and double random phase encoding. *Peer-to-Peer Networking and Applications*, 14(2), 747–762.
- [44] Rajasekar, V., Venu, K., Jena, S. R., & Varthini, R. J. (2021). Detection of cotton plant diseases using deep transfer learning. *Journal of Mobile Multimedia*, 18(2), [page numbers if available]. <https://doi.org/10.13052/jmm1550-4646.1828>
- [45] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788.
- [46] Tee, C., Al-Shabi, M., Cheah, W. P., & Goh, M. (2017). Facial expression recognition using a hybrid CNN–SIFT aggregator. In *Multi-disciplinary Trends in Artificial Intelligence*. Springer.
- [47] Terven, J., Córdova-Esparza, D.-M., & Romero-González, J.-A. (2023). A comprehensive review of YOLO architectures in computer vision: From YOLOv1 to YOLOv8 and YOLO-NAS. *Machine Learning and Knowledge Extraction*, 5(4), 1680–1716.
- [48] Terven, J., Córdova-Esparza, D.-M., & Romero-González, J.-A. (2023). A comprehensive review of YOLO architectures in computer vision: From YOLOv1 to YOLOv8 and YOLO-NAS. *Machine Learning and Knowledge Extraction*, 5(4), 1680–1716. <https://doi.org/10.3390/make5040083>
- [49] Venugopal, D., Jayasankar, T., Sikkandar, M. Y., Waly, M. I., Pustokhina, I. V., Pustokhin, D. A., & Shankar, K. (2021). A novel deep neural network for intracranial haemorrhage detection and classification. *Computers, Materials & Continua*, 67(2), 1509–1523. <https://doi.org/10.32604/cmc.2021.015480>
- [50] Zhang, Y., Ma, B., Hu, Y., Li, C., & Li, Y. (2022). Accurate cotton diseases and pests detection in complex background based on an improved YOLOX model. *Computers and Electronics in Agriculture*, 203, 107484. <https://doi.org/10.1016/j.compag.2022.107484>
- [51] Zheng, Z., Wang, P., Liu, W., Li, J., Ye, R., & Ren, D. (2020). Distance-IoU loss: Faster and better learning for bounding box regression. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34, 12993–13000.