International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

Cotton Disease Detection Using Deep Transfer Learning with RESNET152V2 and Integrated LLM-Based Diagnostic Assistant

Jayesh Deore, Prof. A. S. Chaudhari, Umesh Bhamare, Yogesh Patil
Departmentof Computer Engineering,R C Patel Institute of Technology, Shirpur, India

***Abstract* - Cotton contributes significantly to agricultural income, yet diseases continue to affect productivity and farmer profit. Manual inspection of crops is still common practice, however it demands time, attention and becomes difficult when dealing with large agricultural fields. With this motivation, a computer-based solution for cotton plant disease identification was developed using deep learning. The dataset used for experimentation was sourced from Kaggle and consisted of 1957 training samples, 324 validation samples and 18 images for testing.**

Three transfer learning architectures were tested in the study namely ResNet50, InceptionV3 and ResNet152V2. Among these, ResNet152V2 produced the most promising results, achieving an accuracy of 98.10 percent with a loss of 0.2066, and the validation accuracy reached 100 percent within 20 epochs. The trained model is capable of classifying four categories that include diseased cotton leaf, diseased cotton plant, fresh cotton leaf and fresh cotton plant.

A frontend interface was built using React so that users such as farmers can upload images captured from a mobile camera or drone. The model processes the input and returns the prediction on the spot. To improve usability, an LLM based module is integrated to provide simple explanations and preventive suggestions for the detected disease. The overall study shows that the system can assist in early diagnosis of cotton diseases and help reduce crop loss by enabling timely intervention.

1. INTRODUCTION

Cotton has long been at the heart of the textile sector. In many developing countries, it is more than just another plant grown in the fields — it is the main source of income for a lot of farming families. Specialists used to be consulted to examine the cotton directly, which is good enough for small-scale cultivation, but if the ground is vast enough, taking a walk to examine every single plant is tiring and inefficient. Oftentimes, diseased crops are identified when it is already too late.[1] [2]

But in the last few years, agriculture has begun to embrace technology in a real, tangible way. Cameras, mobile phones, and even drones enable pictures to be taken without having to have someone in every nook and cranny on the farm. Using machine learning, those pictures enable analysis of whether or not the plant is healthy or has a disease. It is this observation that has guided this current project where I developed a system to detect cotton diseases using transfer learning and made this system so simple to use that even a farmer could do so. All they have to do is upload an image, and in a short time, they have their result.[2]

The aim here is to encourage smarter practices in agriculture. Early intervention means pesticides can be employed only when needed, the loss of crops can be managed, and the farmer has a better chance of getting a good yield. A device which carries out a speedy diagnosis can really help with efficient crop management.[3]

2. BACKGROUND AND RELATED WORK

Plant disease image datasets for disease identification have received considerable attention in precise agriculture studies over the last decade. Among various computationally intelligent techniques, deep convolutional neural networks (CNNs) have proven themselves superior to other machine learning techniques for image classification. Traditional machine learning techniques require extensive use of human-crafted feature extraction in image analysis. In contrast to classical techniques, CNNs can learn hierarchical features such as distortions in leaf colors, irregularities in leaf textures, boundaries of lesions, and abnormalities in leaf structures, which are treated in diseased leaves. Such techniques are highly suitable for agricultural image classification tasks, where disease symptoms can be varied in size, location, and type.[2]

Nevertheless, the training of deep models from a scratch requires a very large set of labeled images, along with sufficient computing power, which may not be easily accessible in research environments related to agriculture. This issue has been tackled by the advent of transfer learning, where the problem can be handled by using a set of pre-trained models named ResNet, Inception, VGG, and DenseNet, initially trained on a very large dataset named ImageNet. These models possess the ability of general visual learning,

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

making it easy for a fast retraining of the model on a small set of images for crop disease classification tasks for various crops like tomato, rice, potato, maize, along with a few more crops, as mentioned in a few research papers.

Although there is significant literature on plant disease classification, there is less related work focused on cotton. There are fewer publicly accessible datasets for cotton plant diseases, and most of the current contributions tend to be lab-focused, concentrating on lab tests rather than field deployment. Many contributions tend to mainly address leaf disease, but there is scope for addressing whole plant disease, which is equally important under real-world settings. There is a considerable gap related to making contributions more transferable to practical solutions accessible by farmers via smartphones, drones, or field inspection systems.

The major research gaps which can be identified are:

Comparatively less study has been done on different models of transfer learning on the datasets of cotton diseases.

As there is no end-to-end deployable solution available in many of these studies, as they tend to finish their assessments of models and lack implementation of prediction systems for real-time deployment.

Very limited work is done on explanation-based feedback, keeping farmers ignorant about why a given plant is marked as diseased and what should be done further.

Few examples of full plant images, potentially better indicating plant disease distribution than separate leaf samples. The proposed study will take on these challenges head-on. The proposed study will compare different state-of-the-art models and then implement the best model on an interactive web interface. The model will be developed in such a way that it can operate using normal camera pictures or drone pictures. Therefore, it will provide a farmer-friendly diagnostic model. The main purpose of proposing this model is to connect research and implementation.

3. METHODOLOGY

3.1 Dataset Description

The cotton disease dataset employed in this research was obtained from the Kaggle dataset repository, which has a cumulative sum of 2299 marked images classified into four different classes, namely Diseased Cotton Leaf, Diseased Cotton Plant, Fresh Cotton Leaf, and Fresh Cotton Plant. The diseased leaf images portray visible symptoms brought about by the disease in form of discoloration, spots, and irregular surfaces, while the diseased plants portray the overall structural modification that takes place in the cotton plants that have caught the disease, while the fresh images portray images that appear normal with no signs or symptoms of the disease whatsoever. This promotes a clear demarcation between the diseased images and those that appear to be intact or fresh. From the overall number of images provided in the dataset, 1957 images are involved in training, while image classification involves 324 images with the final 18 images conducting tests. The dataset has a natural variability in the images captured using lighting conditions, complexity levels of the background objects, location of the image with respect to the leaf, and angle directed towards the camera capture. These levels provide natural visual patterns offering highly generalized features.[8]

The pretraining processing included many activities, which helped in making similar images and enhancing model generality. Images were all resized to allow compatibility with either ResNet or Inception architectures. Image pixel intensity for all images was normalized to lie between 0 and 1. Other methods of data augmentation included image rotations, zooming, and flipping. Image data augmentation presented many variations to images synthetically. This could be seen as part of the approaches used to combat overfitting and enhance the generality of models in distinguishing real-time images of crops from all fields of high variability. This variability is experienced because of high intensity shadows and sunlight in cotton fields. Another field variability is experienced from soils and leaf overlay.[5]

In an effort to promote a fair comparison, there are three of the most popular transfer learning models identified: ResNet50, InceptionV3, and ResNet152V2. ResNet50 has been utilized as a baseline medium-depth model, whereas InceptionV3 featured parallel convolutional layers to support equal processing of images of different scales. Also, based on having a significantly deeper structure than the above-mentioned models, ResNet152V2 offered even more enhanced representational capacity for extracting complicated textures and disease patterns at the cotton leaf surface as well as at cotton plant specimens. For models like ResNet50, InceptionV3, and ResNet152V2, their weight components belonging to feature extraction tasks have been retained. But new classification layers are now reduced to a dense connection model ended with a four-unit softmax layer designed for dataset-defined categories. Both early stopping criteria and a learning rate mechanism are implemented to alleviate potential overfitting. There is a 20-epoch training cycle using the Adam optimizer and cross-entropy loss. 3.4 Finally, after completing training, the accuracy of these models for training and validation, along with how well they can validate, were evaluated. Although ResNet50 performed well in terms of accuracy, training capabilities of InceptionV3 were better with a moderate loss in training speed, but ResNet152V2 performed much better than these two models. This model achieved a training accuracy of 98.10% with a validation accuracy of 100% for identifying differences in symptoms, which were subtle, in images of somewhat similar classes of cotton. This model has been selected on the basis of these attributes for the implementation of this application for determining cotton disease.[6]

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

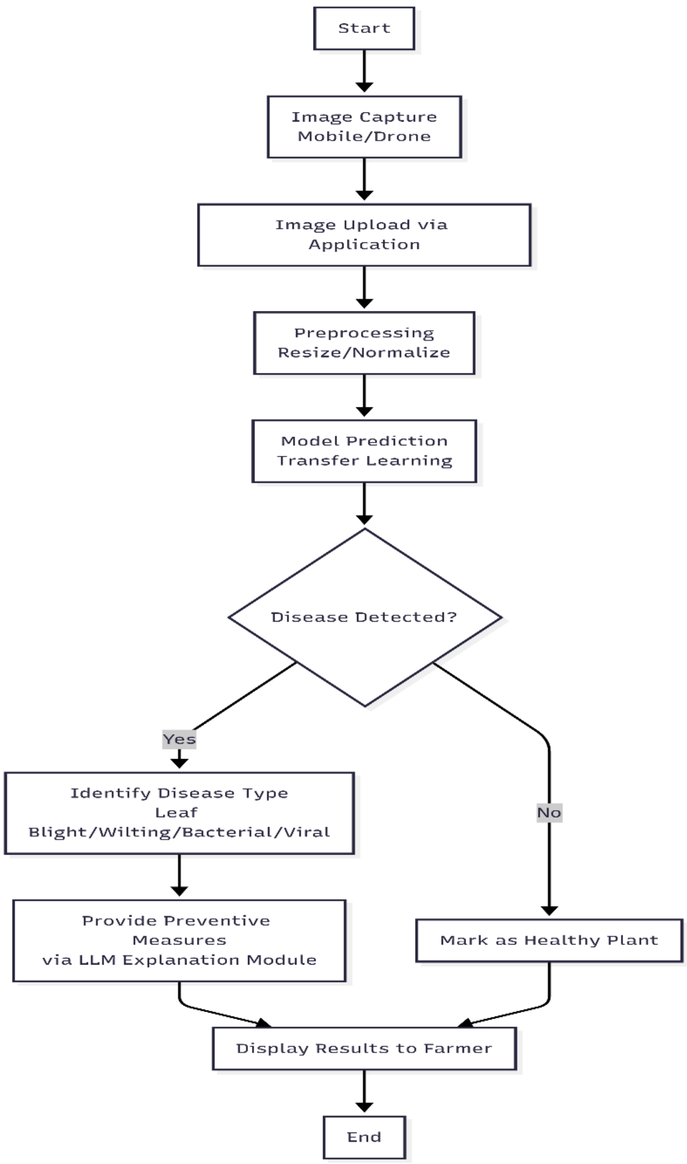

4. IMPLEMENTATION

The proposed cotton disease classification model was not just intended for development but was developed into an entirely functional model that can really help farmers and agricultural workers. The model was planned to work with three key components: backend model development and deployment, frontend development for ease of accessibility, and an interpretation model for deriving key insights. Each of these modules helps form an end-to-end model that can identify diseases for cotton plants from photographs taken using mobile cameras or drone technology. The key idea behind this model was to develop something that can fill the gap between research and practical usability for farmers. This implied that anyone with no knowledge of technology could now avail themselves of advanced cotton disease classifications.

4.1 Backend Model Deployment

After model training and selection of the best model output with ResNet152V2, the model was converted into a deployable form and implemented with a Python-based model inference server. The model inference server is capable of performing multiple operations such as receiving input images, pre-processing the images into the required format for model output, and making predictions based on the model with results delivered as part of a formatted JSON output. The model is implemented with the necessary optimizations for rapid execution of predictions for quick turnaround time between the input and the output, making the system relevant for agriculture-based decisions. For the purpose of enhancing model robustness and practical usage for agricultural model implementation and deployment, the model is designed with the possible requirements for scaling and implementing extensions such as the addition of multiple models and the implementation and execution capability for large volumes at the farm level and cloud implementation.

4.2 Frontend

The final software layer of interaction with the users is developed with React.js. The choice of React.js is informed by its performance and scalability. The software is designed to be minimal and efficient. The system will enable farmers to easily interact with the system. The users will upload pictures from their galleries or take pictures with either a laptop computer or a smartphone. The system will then display the image and its corresponding classification from the model. The system has been designed with responsiveness on smartphones to enable farmers to easily access and interact with the system while working. The design of software has been made simple to enable non-technical people to access and interact with the system. The system communicates with its backside API in real time.

4.3 Explanation & Recommendation Layer

However, disease classification may not be enough for farmers. Over and above that, knowledge of the reason and the remedy is equally valuable. To bridge that gap, an explanation and recommendation interface is also embedded. Based on the disease type determined, appropriate information and pictorial symptoms that may be detected upon observation of the affected plant are also generated. For example, where a leaf may appear with brown circular marks showing fungal disease, an appropriate diagnosis may read, “Fungal disease. Remove the leaf before progression and use recommended organic fungicide.” This provides added meaning through explanation and makes the tool less of a classification tool and more of a farm advisory tool. This provides added meaning through explanation and makes the tool less of a classification tool and more of a farm advisory tool. 4.4 Practical Application Scope In addition to diagnosing a single plant, it appears that the system has promising contexts in which it can be expanded for use in agricultural environments. Coupled with a drone technology platform, it would be applicable for scanning an entire field with rapid identification of diseased regions. In relation to specific growth levels in crops, this technology would be valuable for farmers who would benefit in monitoring processes where initial symptoms of disease can be identified before it progresses. Even pre-harvest analysis for

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 12 , December - 2025

quality can be done to check the health of a field during the final phase of plant growth. This technology can be used in an agricultural advice center, agricultural extension services, and an institution for research on crop monitoring. Further extensions can be done for it to become a system with capabilities in multiple disease detection, soil condition feedback, irrigation assistance, and pesticide guides.

5. CONCLUSION

The use of transfer learning to design a system to detect the presence of the cotton disease, based on images provided, has proven to be a promising solution to aid farmers in efficiently detecting the infections at a preliminary stage. With the ability to instantly detect the images provided by the user, the system has proven to be an efficient alternative to current methods used to detect the cotton disease, which require time and expert analysis to identify the same. With some improvements, the system has the ability to dramatically aid in smart farming technology to efficiently manage crop data.

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 770–778, 2016.

[2] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 2818–2826, 2016.

[3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems (NIPS)*, pp. 1097–1105, 2012.

[4] E. C. Too, Y. Li, S. Njuki, and L. Yingchun, "A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.

[5] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using Deep Learning for Image-Based Plant Disease Detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016.

[6] K. P. Ferentinos, "Deep Learning Models for Plant Disease Detection and Diagnosis," *Comput. Electron. Agric.*, vol. 145, pp. 311–318, 2018.

[7] M. Brahimi, K. Boukhalfa, and A. Moussaoui, "Deep Learning for Tomato Diseases: Classification and Symptoms Visualization," *Applied Artificial Intelligence*, vol. 31, no. 4, pp. 299–315, 2017.

[8] V. S. Dhaka et al., "A Systematic Review on Machine Learning and Deep Learning Techniques for Plant Disease Classification," *Agronomy*, vol. 11, no. 6, p. 1116, 2021.

[9] Cotton Disease Dataset, Kaggle. Available: https://www.kaggle.com/janmejaybhoi/cotton-disease-dataset, 2020.

[10] S. E. A. Raza, S. A. Prince, J. P. Clarkson, and N. M. Rajpoot, "Automatic Detection of Diseased Tomato Plants Using Deep Convolutional Neural Networks," *Machine Vision and Applications*, vol. 27, pp. 1–12, 2015.

[11] V. Singh and A. K. Misra, "Detection of Plant Leaf Diseases Using Image Segmentation and Soft Computing Techniques," *Information Processing in Agriculture*, vol. 4, no. 1, pp. 41–49, 2017.