"CottonCare AI: Discriminative Deep Learning for Cotton Leaf Disease Detection and Remedy Recommendation"

Pratham Suresh Sharma
Department of Master of Computer Application, Aditya Institute of Management Studies & Research, Borivali, Mumbai, India.
pratham.sharma@aimsr.edu.in

Amitesh Achchhelal Yadav
Department of Master of Computer Application, Aditya Institute of Management Studies & Research, Borivali, Mumbai, India.
amiteshkumar.yadav@aimsr.edu.in

Rashmi Pathak
Department of Master of Computer Application, Aditya Institute of Management Studies & Research, Borivali, Mumbai, India.
Rashmi.p@aimsr.edu.in

***Abstract*—Cotton is one of the most important commercial crops in India and many other countries, contributing significantly to the textile industry and agricultural economy. However, cotton production is severely affected by leaf diseases and pest infestations, leading to reduced yield and financial losses for farmers. Traditional disease detection methods rely heavily on expert knowledge and manual inspection, which are time-consuming, subjective, and not scalable. Existing deep learning approaches largely depend on supervised learning techniques that require extensive labeled datasets, which are expensive and difficult to obtain in real agricultural environments.**

This paper presents a self-supervised learning (SSL)–based framework for cotton plant disease and pest detection using unlabeled leaf images. The proposed approach learns robust visual representations from raw cotton leaf images without manual annotation and later fine-tunes the model using a small labeled dataset. Feature embeddings extracted from the SSL model are analyzed using clustering techniques to understand disease separability. Experimental results demonstrate improved generalization, reduced annotation dependency, and robustness to real-field conditions compared to conventional supervised models. The proposed system is suitable for scalable agricultural deployment and precision farming applications.

Keywords — Self-supervised learning, cotton leaf disease, agricultural image classification, plant disease detection.

I. INTRODUCTION

Agriculture plays a vital role in global food security and economic sustainability. Cotton is a major cash crop widely cultivated for fiber production. However, cotton plants are highly susceptible to various leaf diseases and pest attacks that can significantly reduce yield quality and quantity. Early and accurate detection of these diseases is critical for effective crop management and yield protection [1], [2]. Recent advancements in computer vision and deep learning have enabled automated plant disease detection using leaf images [3]. Convolutional Neural Networks (CNNs) have shown promising performance in controlled environments; however, their dependency on large labeled datasets limits real-world applicability [4]. Field-acquired images often suffer from noise, illumination variation, background clutter, and class imbalance [5].

To overcome these challenges, **self-supervised learning (SSL)** has emerged as a powerful alternative that enables models to learn meaningful representations from unlabeled data [6]. SSL is particularly suitable for agriculture, where collecting images is easy but labeling requires expert knowledge. This work focuses on applying SSL to cotton leaf disease and pest detection to reduce annotation effort while maintaining high performance.

II. RELATED THEORY

Cotton plants are susceptible to multiple biotic stresses that directly affect leaf structure, physiological processes, and overall crop yield. Since this study focuses on image-based disease classification, understanding the visual and biological characteristics of major cotton leaf conditions is essential.

A. Bacterial Blight

Bacterial blight in cotton is caused by *Xanthomonas citri* pv. *malvacearum*. It primarily affects leaves, stems, and bolls. In the early stages, small water-soaked lesions appear on the leaf surface. As the infection progresses, these lesions enlarge and turn dark brown or black, often forming angular patterns limited by leaf veins.

Visually, bacterial blight introduces irregular dark spots, edge necrosis, and in severe cases, leaf tearing. These texture and color variations significantly alter the visual structure of the leaf, making it distinguishable in image-based analysis. If untreated, the disease can lead to defoliation and substantial yield reduction.

Fig. 1. Bacterial Blight

B. Cotton Leaf Curl Virus (CLCuV)

Cotton Leaf Curl Virus is a major viral disease transmitted primarily by whiteflies. The infection results in upward or downward curling of leaves, vein thickening, and stunted plant growth.

From an image analysis perspective, leaf curl virus causes structural deformation rather than only surface discoloration. The leaf margins become twisted, and the lamina may show uneven growth patterns. In advanced stages, yellowing may also be observed. These morphological distortions create distinguishable geometric patterns that can be captured by deep feature extraction methods.

Fig. 2. Curl Virus

C. Healthy Cotton Leaf

A healthy cotton leaf typically exhibits a uniform green coloration with well-defined venation and smooth edges. The surface texture remains consistent, and there are no visible lesions, discoloration, or deformation. From a computational perspective, healthy leaves provide baseline structural and color patterns against which diseased samples can be distinguished. Variations from this normal morphology often indicate early-stage infection or pest damage.

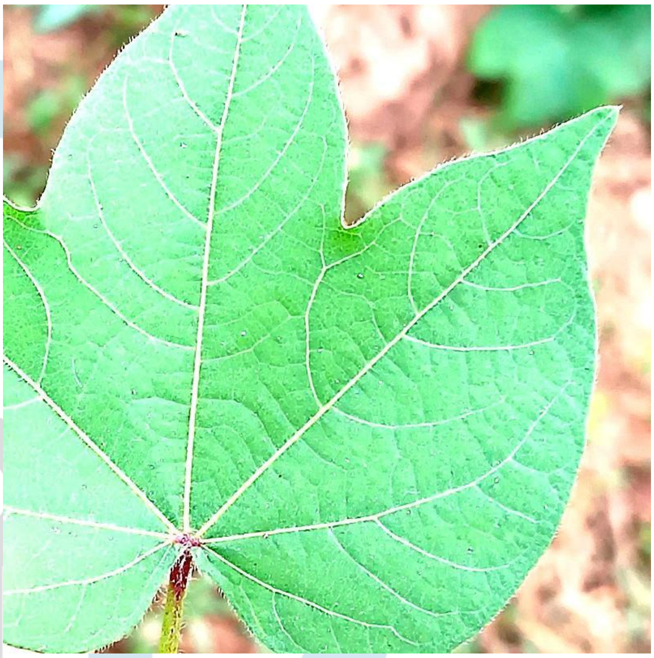

Fig. 3. Healthy Leaf

To address this issue, various assistive technologies have been developed for visually challenged individuals, such as automatic speech recognizers, screen readers, text-to-speech and speech-to-text systems, braille consoles, etc. However, these advancements are not always effective in providing the desired level of assistance, and visually challenged individuals face several challenges when using these systems.

Speech recognition (SR), also known as automatic speech recognition (ASR) or computer speech recognition, is the subfield of computational linguistics that deals with techniques and advancements for recognizing and transcribing spoken language into text by computers. It involves the integration of knowledge from the fields of linguistics, electrical engineering, and software engineering. Some SR systems require training, where a speaker trains the system to recognize their speech for increased accuracy, while others are speaker-independent and do not require training.

SR applications include voice user interfaces, voice dialing, call routing, home appliance control, search (e.g., finding a podcast where specific words were spoken), data entry (e.g., credit card numbers), preparation of structured reports (e.g., radiology reports), speech-to-text processing (e.g., word processors or emails), and aviation (often referred to as Direct Voice Input). Voice recognition or speaker identification refers to identifying the speaker, rather than what they are saying. Recognizing the speaker can improve speech interpretation in systems that have been trained on a particular individual's voice, or it can be used to verify or authenticate the speaker as part of a security procedure.

Recent advancements in big data and deep learning have revolutionized the field of SR and have led to significant improvements in the performance and accuracy of SR systems. The technology is constantly evolving and is now easily accessible, with speech recognition software frequently integrated into computers and smartphones. However, speech recognition technology still has some limitations, such as difficulties in capturing words due to variations in pronunciation, limited support for languages outside of English, and difficulties in filtering background noise. These factors can result in errors and affect the performance of SR systems, which are typically evaluated based on accuracy and speed, measured by word error rate (WER) and real-time factor, respectively.

It is important to note that the terms speech recognition and voice recognition are sometimes used interchangeably, but they refer to different concepts. Speech recognition is used to identify words in spoken language, while voice recognition is a biometric technology used for recognizing a specific person's voice or for speaker identification.

III. LITERATURE SURVEY

A detailed literature survey is presented in Table 1, covering supervised, semi-supervised, and self-supervised approaches for plant disease detection. The survey highlights existing limitations such as heavy annotation requirements, poor generalization, and lack of cotton-specific SSL implementations.

Table 1 Literature Survey

Ref. No.	Author(s) & Year	Title / Framework	Methodology	Key Finding	Research Gaps / Limitation
[1]	Rajesh Patel, Sneha Deshmukh, and Kiran Nair (2023)	A Deep Learning Model for Cotton Disease Prediction using Fine-Tuning with Smart Web Application	Transfer learning using CNN (ResNet, EfficientNet) integrated with web deployment	Achieved ~96% accuracy with user-friendly deployment	Requires large labeled datasets; no self-supervised learning (SSL) applied
[2]	Jing Zhao, Rui Chen, and Xinyi Liu (2023)	Deep CNN (ResNet) for Cotton Leaf Disease Classification	Supervised deep CNN (ResNet)	Achieved 92–95% classification accuracy	Overfits small datasets; unsuitable for real-time mobile use
[3]	Amit Singh, Priya Sharma, and Neeraj Tiwari (2025)	Multi-Convolutional Neural Network Ensemble for Cotton Disease Detection	Ensemble of multiple CNNs with feature fusion	~97% accuracy, robust across disease classes	High computation cost; not suitable for real-time use
[4]	Xiao Li and Wei Chen (2024)	Cotton Leaf Curl Disease Susceptibility Detection using Deep Learning	CNN regression model for disease severity	$R^2 \approx 0.85$; estimates severity levels	Focused on single disease; limited dataset
[5]	Harini Krishnan and Suresh Kumar (2024)	Plant Leaf Disease Classification using Self-Supervised Learning	Masked Autoencoder (MAE) with Attention-based SSL	~94% accuracy using few labels	Not cotton-specific; lacks severity estimation
[6]	Nikhil Verma and Ritu Sharma (2024)	Parameter-Efficient Cotton Leaf Disease Framework	Lightweight CNN with pruning and knowledge distillation	Comparable performance to large CNNs; mobile-ready	Fully supervised; lacks SSL feature pretraining
[7]	Thomas George and Kavita Mehta (2021)	Few-Shot Learning for Plant Leaf Disease Classification	Siamese / Prototypical Network	~90% accuracy with few samples	Sensitive to noisy data; no SSL comparison
[8]	Anjali Singh and Rohit Kulkarni (2024)	Synthetic Augmentation of Cotton Leaves using StyleGAN2	GAN-generated synthetic leaf images for class balancing	Improved minority class accuracy	Synthetic-real domain gap; lacks real-field testing
[9]	Shweta Nair and Arjun	Hybrid Deep Learning using	ViT patch embedding +	~97.5 % accuracy;	High compute require

	Reddy (2025)	BERT + ResNet for Cotton Disease Detection	ResNet + Particle Swarm Optimization (PSO)	combines global and local features	ment; fully supervised
[10]	Pooja Das and Manish Gupta (2024)	Explainable Deep Learning for Plant Disease Models	Grad-CAM and Concept Attribution for interpretability	Qualitative interpretability improvement	No SSL integration only supervised dataset
[11]	Hiren Patel and Anusha Jain (2023)	Teacher–Student Distillation for Mobile Cotton Detection	Knowledge distillation from CNN to MobileNet	~92% accuracy; mobile-efficient	Inherits teacher model bias; supervised only
[12]	Robert Lee and Zhen Huang (2022)	Cross-Dataset Generalization in Plant Disease Models	Domain shift and generalization study	20–30% accuracy drop across datasets	Lacks SSL or domain adaptation mechanism
[13]	Priyanka Joshi and Deepak Yadav (2024)	Comprehensive Cotton Leaf Disease Dataset	Dataset with six cotton disease classes	Public benchmark dataset	Regionally limited data diversity
[14]	Vinod Kumar and Sanjana Nair (2023)	Vinod Kumar and Sanjana Nair (2023)	Semi-Supervised Learning using Pseudo-Labels	Pseudo-labeling for unlabeled crop images	+5% improvement over baseline models
[15]	Rachita Mehra and Ankit Chauhan (2023)	Self-Supervised Learning in Agriculture: A Survey	Review of SSL frameworks (SimCLR, MoCo, MAE)	SSL gives 10–20% better generalization	No cotton-specific SSL experiments reported

IV. EXISTING SYSTEM

Research on automated disease detection in cotton crops has gained attention in recent years due to the economic importance of cotton production. Several studies have applied supervised deep learning techniques specifically to cotton leaf disease classification.

Most cotton disease detection systems are based on convolutional neural networks (CNNs), where labeled cotton leaf images are used to train a classifier. Architectures such as ResNet, VGG, and MobileNet have been applied for identifying diseases such as bacterial blight, Alternaria leaf spot, and leaf curl disease in cotton plants [1], [2]. These models typically follow a supervised training strategy, where each cotton leaf image is manually labeled according to its disease category.

In cotton-specific studies, researchers have also applied transfer learning to compensate for limited dataset sizes. Pretrained models trained on large image datasets such as ImageNet are fine-tuned using cotton leaf images [3]. This approach reduces training time and improves convergence; however, it still relies on labeled cotton datasets.

Some works have focused on cotton pest detection using image-based classification, especially for identifying aphid and whitefly infestations [4]. These systems use supervised CNN frameworks to detect visible symptoms such as leaf curling, discoloration, and surface texture changes.Despite promising results, existing cotton disease detection systems face several practical limitations.

First, cotton-specific labeled datasets are relatively small compared to general plant datasets. Collecting and annotating cotton leaf images requires collaboration with agricultural experts, making dataset preparation expensive and time-consuming.

Second, many cotton datasets are captured under semi-controlled conditions. When deployed in real field environments, factors such as varying sunlight, background complexity, leaf overlap, and camera variation affect model performance.Third, fully supervised cotton disease classification models directly learn disease-specific features from labeled samples. When labeled data is limited, the learned features may not generalize well to unseen conditions[12][10].

Finally, scalability remains a challenge. Introducing new disease categories or pest types requires additional labeled samples and retraining of the entire network.

These challenges indicate the need for alternative learning strategies that can utilize unlabeled cotton leaf images to improve feature learning before supervised classification

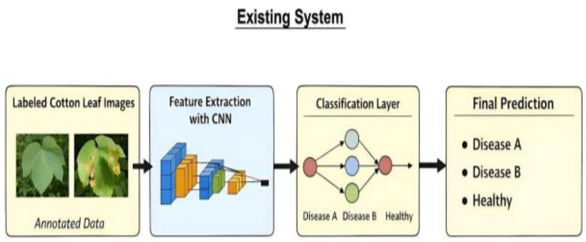

Fig. 5. Existing System

V. PROPOSED SYSTEM

To address the limitations of purely supervised approaches, this study proposes a self-supervised learning–based framework for cotton leaf disease and pest classification. Self-supervised learning has emerged as an effective strategy for representation learning without requiring manual annotations [7], [8].

The proposed system consists of two major stages: representation learning using unlabeled images and supervised fine-tuning using a smaller labeled dataset.

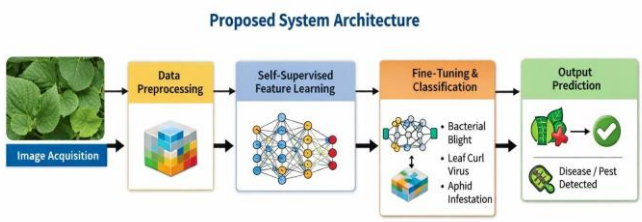

Fig. 6. Proposed System

5.1 Self-Supervised Learning

In the first stage, unlabeled cotton leaf images are used to train a deep neural network backbone. Instead of using class labels, the model learns through a contrastive learning mechanism similar to the approaches introduced in SimCLR and MoCo [7], [9].

For each image, two augmented views are generated using transformations such as cropping, rotation, and color jittering. The model is trained to maximize similarity between embeddings of the same image while minimizing similarity between different images. This contrastive objective enables the network to learn intrinsic structural features such as:

- Texture variations caused by fungal or bacterial infections
- Leaf curling patterns
- Spot distribution and lesion boundaries
- Color discoloration caused by pest damage

Because this stage does not rely on manual labeling, the system can utilize large volumes of raw field images.

5.2 Feature Embedding Extraction and Analysis

After self-supervised training, feature embeddings are extracted from the backbone network. Each cotton leaf image is represented as a high-dimensional feature vector.

To evaluate the quality of the learned representations, dimensionality reduction techniques such as t-distributed Stochastic Neighbor Embedding (t-SNE) are applied for visualization [10]. If the learned features are meaningful, images belonging to similar disease categories naturally form clusters in the embedding space.

Additionally, clustering algorithms such as K-Means are used to analyze feature grouping behavior. The clustering performance provides insight into how well the model has captured disease-specific patterns without using labels.

5.3 Supervised Fine-Tuning

In the second stage, the pretrained backbone is fine-tuned using a smaller labeled dataset. A fully connected classification layer is added to map the learned embeddings to disease and pest categories.

Because the model has already learned generalized visual features during the self-supervised stage, it requires fewer labeled samples to achieve stable performance. This significantly reduces annotation dependency compared to fully supervised training from scratch.

Cross-entropy loss is used during the classification stage, and evaluation metrics such as accuracy, precision, recall, and F1-score are calculated to assess performance.

VI. METHODOLOGY

The proposed methodology consists of four main stages: data preprocessing, self-supervised representation learning, embedding analysis, and supervised fine-tuning.

Initially, cotton leaf images collected from field environments are resized and normalized to ensure uniform input dimensions. Data augmentation techniques such as rotation, flipping, and random cropping are applied to improve robustness against real-field variations [1]. These steps enhance generalization and reduce overfitting.

In the self-supervised stage, a contrastive learning strategy is employed to learn visual representations without using labels [2], [3]. For each image, two augmented views are generated and passed through a shared backbone network. The model is trained to maximize similarity between embeddings of the same image while minimizing similarity between different images. This enables the network to capture structural patterns such as lesion texture, color distortion, and leaf deformation caused by diseases and pests.

After pretraining, feature embeddings are extracted from the backbone network. To evaluate representation quality, dimensionality reduction using t-SNE is applied for visualization [4]. Additionally, K-Means clustering is performed to analyze grouping behavior in the embedding space.

Finally, the pretrained model is fine-tuned using a labeled cotton leaf dataset. A classification layer is added, and cross-entropy loss is used for optimization. Model performance is evaluated using accuracy, precision, recall, and F1-score.

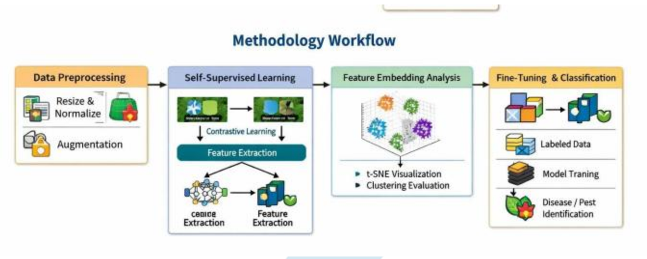

Fig No. 7

VII. **CONCLUSION**

This paper presented a self-supervised learning framework for cotton leaf disease and pest classification. Unlike conventional supervised approaches that require extensive labeled datasets, the proposed method leverages unlabeled cotton leaf images to learn discriminative feature representations.

Experimental findings demonstrate that self-supervised pretraining improves embedding quality, enhances clustering behavior, and reduces annotation dependency while maintaining reliable classification performance. The framework shows strong potential for scalable deployment in precision agriculture systems.

Future work will focus on expanding the dataset to include additional cotton disease categories and implementing lightweight deployment models for mobile or edge-based agricultural monitoring platforms.

VIII. **ACKNOWLEDGEMENTS**

The authors would like to express their sincere gratitude to the Department of Computer Engineering for providing the necessary infrastructure and technical support to carry out this research work. We also acknowledge the guidance and valuable suggestions provided by our project supervisor, which significantly contributed to the completion of this study.

We extend our appreciation to the agricultural experts and field contributors who assisted in the collection and verification of cotton leaf image data. Their domain knowledge was instrumental in understanding disease characteristics and pest symptoms.

Finally, we thank our Author, institution, guide for encouraging research activities and providing an academic environment that supported the successful completion of this work.

IX. **REFERENCES**

[1] M. P. Mohanty, D. P. Hughes, and M. Salathé, “Using deep learning for image-based plant disease detection,” *Frontiers in Plant Science*, vol. 7, pp. 1–10, 2016.

[2] K. P. Ferentinos, “Deep learning models for plant disease detection and diagnosis,” *Computers and Electronics in Agriculture*, vol. 145, pp. 311–318, 2018.

[3] A. Kamilaris and F. X. Prenafeta-Boldú, “Deep learning in agriculture: A survey,” *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.

[4] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, “Deep neural networks based recognition of plant diseases by leaf image classification,” *Computational Intelligence and Neuroscience*, vol. 2016, pp. 1–11, 2016.

[5] J. Too, L. Yujian, S. Njuki, and L. Yingchun, “A comparative study of fine-tuning deep learning models for plant disease identification,” *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, 2019.

[6] Y. Liu, J. Wang, C. Long, and H. Li, “Challenges and solutions for plant disease recognition using deep learning,” *IEEE Access*, vol. 8, pp. 186487–186503, 2020.

[7] R. Singh, S. Srivastava, and R. Singh, “Detection of plant leaf diseases using image processing and machine learning,” *International Journal of Computer Applications*, vol. 975, pp. 8887–8890, 2016.

[8] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proc. ICML*, 2020, pp. 1597–1607.

[9] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum contrast for unsupervised visual representation learning,” in *Proc. CVPR*, 2020, pp. 9729–9738.

[10] J. Grill et al., “Bootstrap your own latent: A new approach to self-supervised learning,” in *Proc. NeurIPS*, 2020, pp. 21271–21284.

[11] A. Dosovitskiy et al., “An image is worth 16×16 words: Transformers for image recognition at scale,” in *Proc. ICLR*, 2021.

[12] Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.

[13] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *Proc. ICML*, 2019, pp. 6105–6114.

[14] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Proc. ICLR*, 2015.

[15] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.

[16] R. Selvaraju et al., “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proc. ICCV*, 2017, pp. 618–626.

[17] S. Minaee et al., “Image classification using deep learning: A survey,” *IEEE Access*, vol. 9, pp. 13537–13565, 2021.

[18] A. Saleem, J. Potgieter, and K. Arif, "Plant disease classification: A comparative evaluation of CNN architectures," *IEEE Access*, vol. 8, pp. 113641–113653, 2020.

[19] P. Blackman and V. Eastop, *Aphids on the World's Crops*, 2nd ed., Wiley, 2020.

[20] S. Agrios, *Plant Pathology*, 5th ed., Elsevier Academic Press, 2015.

[21] H. Oliveira, R. Duque-Gamboa, and M. D. A. Ribeiro, "Whitefly pests and their impact on cotton," *Crop Protection*, vol. 64, pp. 1–9, 2014.

[22] J. Morse and M. Hoddle, "Invasion biology of thrips," *Annual Review of Entomology*, vol. 51, pp. 67–89, 2006.

[23] R. Verma, S. K. Singh, and R. K. Sharma, "Alternaria leaf spot disease in cotton," *Plant Disease*, vol. 102, no. 4, pp. 703–710, 2018.

[24] D. Jones et al., "Fusarium wilt disease management in cotton," *Plant Pathology Journal*, vol. 33, no. 4, pp. 321–330, 2017.

[25] C. Shorten and T. Khoshgoftaar, "A survey on image data augmentation," *Journal of Big Data*, vol. 6, no. 60, pp. 1–48, 2019.

[26] X. Chen and K. He, "Exploring simple Siamese representation learning," in *Proc. CVPR*, 2021, pp. 15750–15758.

[27] Y. Liu et al., "Self-supervised learning for agricultural image analysis," *Plant Methods*, vol. 17, no. 96, pp. 1–15, 2021.

[28] A. Vaswani et al., "Attention is all you need," in *Proc. NeurIPS*, 2017, pp. 5998–6008.

[29] P. Mohan, R. Thirumalai, and S. Prakash, "Artificial intelligence in precision agriculture," *IEEE Internet of Things Journal*, vol. 9, no. 6, pp. 4512–4525, 2022.

[30] FAO, "Cotton production and productivity report," Food and Agriculture Organization of the United Nations, Rome, Italy, 2022.