

Deep Learning-Based Crop Disease Detection Using EfficientNet-B3 for Smart Agriculture

Gajanan Ankatwar

ankatwar.gajanan@gmail.com

Datta Meghe Institute of Higher Education and Research

Chitra Dhawale

Datta Meghe Institute of Higher Education and Research

Supriya Narad

Datta Meghe Institute of Higher Education and Research

Research Article

Keywords: Deep Learning, EfficientNet-B3, Crop Disease Detection, Smart Agriculture, Transfer Learning, Image Classification, Precision Farming

Posted Date: September 4th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10104203/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Deep Learning-Based Crop Disease Detection Using EfficientNet-B3 for Smart Agriculture

Gajanan Ankatwar^{*1}, Chitra Dhawale², Supriya Narad³

Department of Science and Applications

Datta Meghe Institute of Higher Education and Research,

Wardha, Maharashtra, India

^{*1}ankatwar.gajanan@gmail.com, ²cadhawale12@gmail.com, ³naradsupriya@gmail.com

Abstract

In the grand tapestry of agriculture, crop diseases weave a thread of disruption, casting shadows over bountiful harvests and threatening the delicate balance of global food security. The art of early and precise disease detection emerges as a beacon of hope, illuminating the path to sustainable crop management and eco-friendly farming practices. Traditionally, the task of identifying these plant maladies has rested on the shoulders of manual inspection—a laborious endeavor prone to the whims of human error. Enter the realm of innovation, where this research unveils a deep learning marvel for the automated detection of crop diseases, harnessing the prowess of the EfficientNet-B3 architecture. Through the alchemy of transfer learning and advanced preprocessing, this model elevates classification performance to new heights. Trained and tested on the SAR-CLD-2024 cotton leaf dataset, which captures a kaleidoscope of disease categories across diverse environmental backdrops, the experimental results are nothing short of remarkable. The proposed model boasts a stellar classification accuracy of 98.14%, with precision and recall that dance gracefully across all categories. EfficientNet's compound scaling mechanism orchestrates a symphony of efficient feature extraction, all while maintaining computational elegance. In comparison to traditional convolutional neural networks, this model offers a harmonious blend of enhanced accuracy and reduced complexity. Designed for real-time application, it seamlessly integrates into the realm of smart agriculture, from mobile-based disease detection to drone-assisted monitoring. This visionary approach paves the way for intelligent farming systems, enabling early disease detection, minimizing crop losses, and nurturing productivity.

Keywords: Deep Learning, EfficientNet-B3, Crop Disease Detection, Smart Agriculture, Transfer Learning, Image Classification, Precision Farming

1. Introduction

1.1 Background

Agriculture stands as the lifeblood of the global economy, a steadfast guardian of food security for our ever-expanding world. It is the heartbeat of livelihoods for billions, especially in developing nations where the rhythm of life is intricately woven with the tapestry of farming [5]. As the global appetite swells, fueled by burgeoning populations and evolving dietary desires, the pressure mounts on agricultural systems to boost their bounty and refine their efficiency. Yet, this noble quest is beset by formidable foes: the relentless march of climate change, the silent erosion of fertile soils, and the insidious creep of crop diseases. Food security

remains a pressing global concern, as the specter of crop losses due to diseases and pests casts a long shadow over agricultural yields. Reports from the agricultural frontlines reveal that plant diseases carve deep gashes in harvests, diminishing both the abundance and the quality of crops [5]. In this challenging landscape, the early and precise detection of plant diseases emerges as a beacon of hope, essential for enhancing crop management and safeguarding economic prosperity.

1.2 Problem Statement

In the grand tapestry of agriculture, crop diseases weave a menacing thread, casting shadows over the bountiful fields and threatening the very essence of productivity. These stealthy adversaries not only diminish the harvest but also tarnish the quality of the produce, leading to economic woes for farmers and unsettling the delicate balance of the agricultural supply chain. Among the verdant ranks of commercial crops, cotton stands as a vulnerable sentinel, besieged by a host of maladies such as bacterial blight, leaf curl virus, and the ominous leaf reddening. These diseases, like cunning invaders, seize the opportunity to flourish under favorable environmental conditions, making early detection a vital shield. Fig. 1 paints a vivid portrait of cotton leaf diseases, capturing the stark contrast between the afflicted and the healthy, a visual testament to the silent battle waged in the fields.



Bacterial Blight



Curl Virus



Herbicide Growth
Damage



Jassid Attack



Leaf Reddening



Leaf Variegation



Healthy Leaf

Fig. 1. Cotton leaf diseases including bacterial blight, curl virus, leaf reddening, and healthy leaves.

In the age-old dance of agriculture, the art of diagnosing plant maladies has long been entrusted to the discerning eyes of seasoned specialists, who meticulously scrutinize the verdant canvases of leaves. While this time-honored method may flourish within the confines of controlled environments, it falters when faced with the sprawling expanse of large-scale farming. The sheer demand for time and expertise renders it impractical, like trying to paint a masterpiece with a single brushstroke. Moreover, the visual manifestations of various plant ailments often masquerade as one another, weaving a tapestry of confusion that heightens the risk of misdiagnosis. This underscores the pressing need for automated and dependable systems to illuminate the path of disease detection.

1.3 Limitations of Traditional Methods

In the realm of traditional plant disease detection, the process unfolds like an ancient ritual, reliant on the discerning eyes and seasoned wisdom of human experts. Yet, this age-old method is not without its shadows, as it grapples with a host of limitations:

- **Manual inspection:** Requires skilled experts and extensive field visits [6]
- **Human error:** Misclassification due to subjective judgment
- **Time-consuming:** Inefficient for large agricultural areas

Additionally, factors like lighting conditions, ambient noise, and variations in leaves add complexity to manual detection. These challenges render traditional methods inadequate for contemporary large-scale agriculture, where both speed and precision are essential.

1.4 AI in Agriculture

In the realm of agriculture, the infusion of artificial intelligence (AI) has revolutionized the art of crop surveillance and disease detection. Deep learning, a dazzling offshoot of AI, has carved its niche in the world of image-based classification, excelling in the identification of plant maladies [1],[11],[12]. By harnessing vast oceans of data and intricate neural network architectures, deep learning models possess the uncanny ability to autonomously unravel complex patterns and features from images. Convolutional Neural Networks (CNNs) stand as the virtuosos of image analysis, adept at extracting spatial and hierarchical features with finesse. A multitude of studies have illuminated the prowess of CNN-based models, showcasing their superiority over traditional machine learning techniques in the precise detection of plant diseases [3]. These models elegantly eliminate the need for manual feature extraction, offering scalable solutions that are tailor-made for the agricultural landscape.

1.5 Transfer Learning

In the realm of deep learning, transfer learning emerges as a powerful sorcerer, especially when navigating the narrow paths of small datasets. Instead of conjuring a model from scratch, this method harnesses the wisdom of pre-trained models, which have already absorbed the essence of vast datasets. This enchanting approach not only slashes training time but also elevates model performance to new heights [2].

In the vibrant world of crop disease detection, transfer learning emerges as a masterful artist, skillfully repurposing its palette of visual features—edges, textures, and shapes—that dance across diverse image canvases. This artistic prowess enhances the model's ability to adapt and excel with new data, making it a particularly valuable ally in the agricultural landscape.

1.6 Limitations of Existing Models

While deep learning models have been successful, they still face certain challenges. Both traditional CNN models and more advanced architectures like ResNet have achieved impressive accuracy levels, yet they often demand significant computational power [8]. The extensive number of parameters in these models leads to longer training times and makes them less ideal for use on devices with limited resources.

Similarly, traditional machine learning methods such as SVM and k-NN rely on handcrafted features and struggle to handle complex image variations, limiting their effectiveness in real-world scenarios [9]. These challenges highlight the need for more efficient and scalable models that can deliver high performance with reduced computational cost.

1.7 EfficientNet Motivation

EfficientNet addresses the limitations of existing models by introducing a compound scaling method that uniformly scales network depth, width, and resolution [4]. Unlike traditional approaches that scale these dimensions independently, EfficientNet maintains a balance between them, resulting in improved accuracy and efficiency.

EfficientNet-B3 stands out among its versions by balancing performance with computational demands effectively. It delivers impressive accuracy while using fewer parameters, making it well-suited for real-time smart agriculture applications. This efficiency positions EfficientNet-B3 as an excellent option for systems that detect crop diseases.

1.8 Deep Learning Theory

Deep learning models rely on artificial neural networks that identify intricate patterns in data by utilizing multiple layers of abstraction. These models employ optimization methods and loss functions to reduce prediction errors and enhance their performance [10]. The capability of neural networks to learn hierarchical representations makes them particularly effective for tasks involving image classification.

In crop disease detection, deep learning models analyze leaf images to identify disease-specific features such as discoloration, texture changes, and lesions. This automated feature extraction capability significantly enhances detection accuracy and reliability.

1.9 Contributions

In this groundbreaking research, we unveil a novel and potent approach to the automated detection of crop diseases, harnessing the power of deep learning techniques. The pivotal contributions of this study are illuminated as follows:

- Creation of a strong crop disease detection model utilizing the EfficientNet-B3 framework for the multi-class classification of cotton leaf diseases.

- Incorporation of transfer learning to boost model performance, decrease training duration, and enhance generalization across various agricultural settings [2].
- Application of sophisticated data preprocessing methods, such as normalization and augmentation, to enhance model resilience and minimize overfitting.
- Attainment of a high classification accuracy rate of 98.14% on the SAR-CLD-2024 dataset, highlighting the proposed method's effectiveness.
- Comparative evaluation with traditional machine learning and deep learning models, demonstrating superior accuracy and computational efficiency.
- Development of a lightweight and scalable system designed for real-time use on mobile and edge devices in smart agriculture applications.
- Contribution to intelligent farming systems by facilitating early disease detection, reducing crop losses, and boosting agricultural productivity.
- Unlike existing methods, this study prioritizes both high accuracy and computational efficiency, making the proposed model viable for real-world agricultural settings.

This endeavor uniquely strives to blend precision with the swiftness of a summer breeze, rendering it perfectly suited for the dynamic realm of smart agriculture, where real-time deployment is key.

2. Related Work

2.1 Traditional Machine Learning

In the nascent dawn of plant disease detection research, the stalwart guardians of traditional machine learning—Support Vector Machines (SVM) and k-Nearest Neighbors (k-NN)—stood as the vanguard, wielding their algorithms with precision and purpose. These techniques relied on manually crafted features derived from leaf images, such as color, texture, and shape descriptors, to identify diseases [9]. Common feature extraction methods included histogram analysis, edge detection, and texture-based descriptors like the Gray-Level Co-occurrence Matrix to capture disease patterns.

While these methods found a modest foothold in the realm of controlled environments, their prowess was intricately woven with the artful tapestry of feature engineering. In practical agricultural environments, factors such as lighting variations, background noise, and leaf orientation greatly impact the reliability of handcrafted features. Consequently, traditional machine learning models often face challenges in generalizing across varied datasets. Moreover, these approaches require domain expertise for feature selection, making them less scalable and more time-consuming.

Another limitation of traditional approaches is their inability to capture complex hierarchical patterns in images. Since plant diseases often exhibit subtle visual differences, such as slight discoloration or texture variations, handcrafted features may fail to represent these nuances effectively. Consequently, the need for more advanced and automated approaches led to the adoption of deep learning techniques.

2.2 CNN-Based Models

The emergence of Convolutional Neural Networks (CNNs) heralded a transformative leap in the realm of plant disease detection. Like masterful artists, CNNs craft hierarchical feature representations directly from the canvas of raw images, rendering the laborious task of manual feature extraction obsolete. This remarkable prowess empowers them to discern the delicate

tapestry of patterns and variations in plant disease images with an efficiency that outshines traditional methods.

Mohanty et al. [1] conducted one of the pioneering studies in this domain, demonstrating that deep CNNs can achieve high accuracy in large-scale plant disease classification tasks. Their research journey, fueled by the treasure trove of a publicly available dataset, unveiled a striking revelation: deep learning models, like modern-day alchemists, transform data into gold, leaving traditional machine learning methods trailing in their wake. This study laid the foundation for further research in deep learning-based agricultural applications.

Sladojevic et al. [3] proposed a deep neural network-based approach for plant disease recognition, achieving promising results across multiple plant species. Their model unveiled the prowess of CNNs, showcasing their ability to deftly navigate a kaleidoscope of disease categories and adapt to the ever-shifting tapestry of environmental conditions. However, challenges related to dataset size and model generalization were noted.

Ferentinos [7] further extended this research by applying deep learning models to classify plant diseases using a large dataset. The study soared to impressive heights of classification accuracy, unveiling the promising potential of CNNs as a transformative force in the realm of agricultural innovation. Yet, the model demanded a veritable feast of computational power, and like a painter with too few colors, it struggled with overfitting when faced with the sparse canvas of limited datasets.

Traditional vs Deep Learning Approaches

In the realm of conventional machine learning, techniques like Support Vector Machines (SVM) and k-Nearest Neighbors (k-NN) often find themselves navigating a labyrinth of manually crafted features, struggling to unravel the intricate tapestry of patterns woven into images of plant diseases. On the other hand, deep learning methods, especially Convolutional Neural Networks (CNNs), automatically learn hierarchical features, resulting in greater accuracy for classification tasks [1], [3], [7]. Figure 2 highlights the main distinctions between traditional machine learning and deep learning methods for detecting plant diseases.




Fig. 2. Comparison between traditional machine learning and deep learning approaches for plant disease detection.

Recent studies have further improved CNN-based plant disease detection using optimized architectures and large-scale datasets [11], [12].

2.3 Deep Architectures

To enhance the capabilities of CNNs, sophisticated deep learning frameworks like ResNet and VGG have been created. These architectures aim to tackle issues such as vanishing gradients and restricted feature representation.

ResNet, as introduced by He et al. [8], employs residual connections to facilitate the training of extremely deep neural networks. These connections enhance the flow of gradients during backpropagation, leading to improved model performance and stability. ResNet has gained widespread use in image classification tasks, such as plant disease detection, due to its capability to capture intricate features.

In contrast, VGG networks are characterized by a straightforward and consistent design that incorporates small convolutional filters. Although VGG models deliver high accuracy, they demand a substantial number of parameters, resulting in significant computational expense. This considerable computational demand restricts their use in real-time systems and environments with limited resources.

Although ResNet and VGG offer benefits, they also have drawbacks. Their substantial model sizes lead to increased training durations and higher memory demands, which makes them less ideal for use on edge devices like smartphones or IoT-based agricultural systems. These issues have driven the creation of more efficient architectures.

Recent advancements in deep learning frameworks, including Convolutional Neural Networks (CNNs), ResNet, and EfficientNet, have greatly enhanced the accuracy of image classification. CNNs facilitate the automatic extraction of features, while ResNet incorporates residual connections to mitigate vanishing gradient issues, allowing for the construction of deeper networks [8]. EfficientNet further boosts performance by employing compound scaling to optimize the balance between depth, width, and resolution [4]. Fig. 3 depicts the progression of deep learning architectures from traditional CNN models to the more sophisticated EfficientNet designs.




Fig. 3. Evolution of deep learning architectures from CNN to EfficientNet.

2.4 EfficientNet

EfficientNet, introduced by Tan and Le [4], presents an innovative method for model scaling called compound scaling. In a world where traditional methods tinker with network depth, width, or resolution in isolation, EfficientNet emerges as a maestro, orchestrating a harmonious symphony by scaling all three dimensions in unison to achieve peak performance.

In the realm of deep learning, where algorithms dance with data, EfficientNet models emerge as virtuosos, orchestrating a symphony of precision with fewer notes than their traditional CNN counterparts. Like a master artist wielding a brush with finesse, EfficientNet variants, such as the illustrious EfficientNet-B3, have dazzled audiences with their prowess across a spectrum of computer vision tasks, from the intricate tapestry of image classification to the dynamic choreography of object detection.

EfficientNet stands as a beacon of brilliance in the realm of technology, celebrated for its prowess in performance. By artfully trimming the excess of parameters and easing the computational burden, it emerges as the perfect companion for real-time applications, especially on devices where resources are as scarce as water in a desert. This becomes a lifeline for smart agriculture systems, where the need for models that are both feather-light and powerfully efficient is as crucial as sunlight to a budding plant.

Recent research highlights that EfficientNet-based models achieve superior performance compared to traditional CNNs in agricultural applications [12], [13].

2.5 Transfer Learning

In the vibrant world of deep learning for plant disease detection, transfer learning has emerged as a pivotal technique. Instead of crafting models from scratch, this ingenious approach harnesses the power of models already trained on vast seas of data. These pre-trained models come armed with a keen understanding of general image features, ready to be fine-tuned and tailored for specific applications.

Too et al. [2] demonstrated that transfer learning significantly improves model performance, especially when working with limited datasets. By using pre-trained weights, models can

converge faster and achieve higher accuracy. This approach also reduces the computational cost associated with training deep neural networks.

In the realm of agriculture, transfer learning allows models to adjust to various datasets and environmental settings. This approach improves generalization and minimizes the likelihood of overfitting, making it a popular method in contemporary systems for detecting plant diseases.

2.6 Research Gap

Amidst the remarkable strides made in the realm of plant disease detection, a tapestry of challenges still weaves its intricate patterns:


- **Accuracy versus Efficiency Dilemma:** While deep learning models boast remarkable precision, they often guzzle computational resources like a ravenous beast. Striking a harmonious chord between accuracy and efficiency remains a formidable quest.
- **Real-Time Implementation:** Many current models are not designed for real-time use, which restricts their application in practical agricultural systems.
- **Dataset Variability:** Differences in lighting, background, and the severity of diseases can impact model performance.
- **Scalability Challenges:** Models, like chameleons caught in a whirlwind, may find themselves floundering when thrust into unfamiliar terrains.

3. Methodology

In this section, we unveil a visionary approach to unmasking crop maladies, harnessing the prowess of the EfficientNet-B3 architecture. This strategy weaves together the intricate threads of deep learning methods, the artful dance of transfer learning, and the precision of optimized training techniques. The result is a tapestry of high classification accuracy, all while maintaining the delicate balance of computational efficiency. The methodology encompasses system design, model architecture, dataset preparation, preprocessing, training configuration, and evaluation framework.

3.1 System Overview and Architecture

The proposed system follows a structured pipeline for automated crop disease detection:



In this intricate pipeline, images of cotton leaves are meticulously gathered from the SAR-CLD-2024 dataset, embarking on a transformative journey through a series of preprocessing steps that harmonize and standardize the input. These refined images then find their way into the embrace of the EfficientNet-B3 model, which serves as both a discerning feature extractor and a keen-eyed classifier. With precision and insight, the model unveils the predicted disease class that corresponds to each input image, like a digital oracle revealing the secrets of the cotton leaves.

The system is designed to minimize manual intervention and enable real-time disease detection. By automating the entire process, the model improves efficiency and reduces dependency on expert knowledge.

In the realm of digital innovation, the proposed system unfurls like a well-choreographed ballet, gracefully moving through a sequence of stages: image acquisition, preprocessing, feature extraction, and classification. At the outset, the cotton leaf images are gently ushered through a preprocessing phase, preparing them for the discerning eye of the EfficientNet-B3 model, which meticulously extracts their features. These features, like clues in a mystery novel, are then employed to unravel the secrets of disease classification. Figure 4 paints a vivid portrait of the EfficientNet-B3 architecture, highlighting the elegant dance of preprocessing, feature extraction, and classification.




Fig. 4. Proposed EfficientNet-B3 architecture for crop disease detection, showing preprocessing, feature extraction, and classification stages.

The EfficientNet-B3 model, a marvel of modern machine learning, stands tall on the shoulders of convolutional neural networks. It weaves together the intricate threads of Mobile Inverted Bottleneck Convolution (MBConv) blocks with the finesse of squeeze-and-excitation enhancements. This architectural symphony orchestrates a dance of efficient feature extraction, all while keeping computational demands at bay. Like a seasoned botanist, it discerns both the obvious and the elusive features in input images, excelling at detecting the subtle nuances in plant disease patterns.

3.2 EfficientNet Scaling

EfficientNet masterfully orchestrates a symphony of scaling, harmoniously adjusting the depth, width, and resolution of the network to achieve unparalleled performance. Unlike traditional

approaches that scale these dimensions independently, EfficientNet maintains a balance between them, resulting in improved accuracy and efficiency [4].

$$Depth = \alpha^{\emptyset}, \quad Width = \beta^{\emptyset}, \quad Resolution = \gamma^{\emptyset}$$

Where:

- α , β and γ are scaling coefficients
- $\emptyset$ is the compound scaling factor

This scaling approach deftly orchestrates the symphony of computational resources, ensuring the model dances gracefully on the stage of high performance. By balancing these dimensions, EfficientNet avoids over-parameterization and improves generalization.

In this study, the choice of EfficientNet-B3 emerges as a masterstroke, striking a harmonious balance between the twin pillars of precision and computational thrift. It provides higher accuracy compared to smaller variants while remaining efficient enough for real-time applications.

3.3 Dataset

The model in question is trained on the SAR-CLD-2024 cotton leaf dataset, a detailed collection specifically created for identifying crop diseases. This dataset comprises images taken in authentic agricultural settings, accounting for differences in lighting, background, and the severity of diseases. It encompasses seven categories, with the distribution of these classes detailed in Table 1.

Table 1. Class-wise distribution of SAR-CLD-2024 dataset.

Class	Images
Bacterial Blight	48
Curl Virus	59
Healthy Leaf	38
Herbicide Growth Damage	45
Leaf Hopper Jassids	21
Leaf Redding	92
Leaf Variegation	19

The class distribution is adjusted to maintain fairness in both training and evaluation processes. The dataset's variety allows the model to effectively generalize across various environmental conditions and disease variations.

3.4 Data Preprocessing

In the grand tapestry of model performance, the threads of success are woven with the artful craft of data preprocessing. This meticulous process employs a variety of techniques, each a brushstroke on the canvas of uniform input quality:

1. Image Resizing

All images are adjusted in size to align with the input specifications required by EfficientNet-B3, ensuring they are compatible with the model's architecture.

2. Normalization

Pixel values are adjusted to fit within a typical range, usually from 0 to 1. This helps stabilize the training process and speeds up convergence.

3. Data Augmentation

To enhance the variety of datasets and minimize overfitting, a range of augmentation methods are employed:

- Rotation
- Horizontal and vertical flipping
- Random zoom
- Cropping

By employing data augmentation, the model is able to acquire strong features and enhance its capability to manage variations encountered in real-world situations.

3.5 Training Setup

The training configuration of the proposed model is summarized in Table 2.

Table 2. Training configuration of the proposed model.

Parameter	Value
Epochs	30
Batch Size	32
Optimizer	Adam
Learning Rate	0.001

In the grand symphony of machine learning, the Adam optimizer takes center stage, deftly adjusting the learning rate like a maestro fine-tuning an orchestra, ensuring both the swift crescendo of convergence and the harmonious stability of training. Meanwhile, a batch size of 32 is carefully chosen, striking a perfect chord between the demands of memory and the effectiveness of training, much like a skilled conductor balancing the dynamics of a performance.

In the grand symphony of model training, the process unfolds over 30 epochs, each a chapter in the tale of learning. With vigilant eyes, both training and validation metrics are watched like stars in the night sky, ensuring the model's journey remains true. To guard against the siren call of overfitting, early stopping and model checkpointing stand as sentinels, preserving the finest version of the model like a masterpiece in a gallery.

3.6 Loss Function

For multi-class classification, the cross-entropy loss function is used:

$$L = - \sum_{i=1}^N y_i \log(\hat{y}_i)$$

Where:

- y_i represents the true label
- $\hat{y}_i$ represents the predicted probability

This loss function measures the difference between predicted and actual class distributions and is widely used in classification tasks [10]. Minimizing this loss enables the model to improve its predictions over time.

3.7 Transfer Learning

In the grand tapestry of machine learning, transfer learning emerges as a masterful artist, deftly enhancing model performance while trimming the canvas of training time. The EfficientNet-B3 model is initialized with pre-trained weights, allowing it to leverage previously learned features from large-scale datasets [2].

The benefits of transfer learning include:

- Faster convergence
- Improved accuracy
- Reduced requirement for large labeled datasets

In this approach, the final classification layers are fine-tuned to adapt the model to the specific crop disease detection task. This ensures that the model learns domain-specific features while retaining general image representations.

3.8 Summary of Methodology

The proposed methodology integrates EfficientNet-B3 with preprocessing, transfer learning, and optimized training strategies to achieve high accuracy and efficiency. The use of compound scaling ensures balanced model design, while data augmentation and normalization enhance generalization.

In the grand tapestry of smart agriculture, this methodology weaves a sturdy framework for detecting crop diseases, allowing for precise classification and the swift, real-time deployment of solutions.

4. Experimental Setup

In this section, we unveil the intricate tapestry of our experimental setup, meticulously crafted to assess the prowess of the proposed EfficientNet-B3 model in the realm of crop disease detection. It includes details about the computational environment, dataset partitioning, and evaluation metrics used to assess model performance. A well-defined experimental setup

ensures reproducibility and provides a fair basis for evaluating the effectiveness of the proposed approach.

4.1 Environment

In the realm of digital innovation, our experiments unfolded upon the ethereal canvas of a cloud-based computational platform, a realm where efficiency and scalability dance in harmonious tandem. The stage for this technological ballet was set in Google Colab, a gateway to a treasure trove of high-performance hardware, including the mighty Graphics Processing Units (GPUs). These GPUs, akin to the swift wings of a falcon, propel the training process into a realm of speed and precision, especially when taming the complex beasts of deep learning models that feast on intensive matrix computations.

Crafted with the artistry of a master sculptor, the model came to life through the PyTorch deep learning framework, a tool as versatile and user-friendly as a painter's palette, allowing for the seamless creation and nurturing of neural networks. PyTorch provides dynamic computation graphs, enabling efficient debugging and experimentation. Additionally, it includes built-in support for GPU acceleration, allowing seamless integration with CUDA-enabled devices.

The hardware configuration used in this study includes:

- GPU: NVIDIA Tesla (Colab-provided GPU)
- RAM: 12–16 GB
- Storage: Cloud-based storage for dataset access

Harnessing the power of a GPU-enabled environment is like unleashing a fleet of lightning-fast stallions, galloping through the vast fields of data, ensuring that training is not only swifter but also more efficient than the plodding pace of CPU-based systems. This is particularly important for training deep neural networks such as EfficientNet-B3, which involve multiple convolutional layers and large datasets.

4.2 Dataset Split

In the grand tapestry of the SAR-CLD-2024 dataset, a meticulous division unfolds, weaving together the threads of training, validation, and test sets to craft a masterpiece of model evaluation. The training set, a fertile ground for learning, nurtures the model's parameters, while the validation set stands as a vigilant guardian, fine-tuning hyperparameters and keeping a watchful eye on performance. Finally, the test set emerges as the ultimate arbiter, assessing the model's prowess on uncharted data. Fig. 5 paints a vivid picture of this dataset's distribution into its training, validation, and test subsets.




Fig. 5. Dataset distribution showing training, validation, and test splits.

In the grand tapestry of model training, this division weaves a balance between ample data nourishment and the careful curation of distinct datasets for validation and testing. The validation set acts as a vigilant guardian, spotting the sneaky specter of overfitting and guiding the selection of the finest model. Meanwhile, the test set stands as an impartial judge, offering a fair assessment of the model's prowess in the vast arena of generalization.

To ensure fairness, the dataset is shuffled before splitting, and class distribution is maintained across all subsets. This prevents bias toward any particular class and ensures balanced learning.

4.3 Evaluation Metrics

In the grand arena of model evaluation, a suite of standard classification metrics takes center stage, each wielding its own unique lens to scrutinize the model's prowess. These metrics weave together a tapestry of insights, painting a vivid picture of the model's accuracy and steadfast reliability.

1. Accuracy

Accuracy, the steadfast guardian of truth, measures the model's prowess in predicting with precision. It is the harmonious balance of correctly foretold outcomes to the grand tapestry of all attempts.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN}$$

Where:

- TP: True Positives
- TN: True Negatives
- FP: False Positives
- FN : False Negatives

While accuracy paints a broad picture of a model's prowess, it may falter in the shadows of class imbalance, where the scales of data tip unevenly.

2. Precision

Precision is the artful dance of accuracy, capturing the essence of truth by identifying the correct positive samples amidst the sea of predicted possibilities.

$$Precision = \frac{TP}{TP+FP}$$

High precision is like a master archer, whose arrows rarely stray from the bullseye, ensuring that false alarms are few and far between.

3. Recall

Recall is the art of capturing the elusive truth, measuring the fraction of correctly predicted positive samples amidst the sea of all actual positive samples.

$$Recall = \frac{TP}{TP+FN}$$

High recall is like a vigilant sentinel, ensuring that the model captures the vast majority of positive cases with precision and care.

4. F1-Score

The F1-score dances gracefully between precision and recall, weaving them together into a harmonious symphony that offers a well-rounded portrait of model performance.

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

In the realm of data analysis, where the scales often tip unevenly, this metric emerges as a beacon of balance, deftly weaving together the threads of precision and recall to illuminate the path forward.

These evaluation metrics are widely used in recent deep learning-based plant disease detection studies [11].

4.4 Model Evaluation Strategy

Once the model has been sculpted to perfection, it faces the ultimate test against the formidable test dataset. During its creation, vigilant eyes watch over the validation metrics, ensuring the model learns with grace and precision, avoiding the treacherous path of overfitting. The crown of victory is bestowed upon the model that dazzles with the highest validation accuracy, emerging as the champion of performance.

The evaluation process includes:

- Generating a classification report

- Computing accuracy, precision, recall, and F1-score
- Visualizing results using confusion matrices and training curves

This comprehensive evaluation ensures that the model's performance is thoroughly analyzed across multiple dimensions.

4.5 Reproducibility Considerations

To ensure reproducibility of results, the following practices are followed:

- Fixed random seed for dataset splitting
- Consistent preprocessing steps
- Standardized training parameters

These steps ensure that the experiments can be replicated and validated by other researchers.

4.6 Summary of Experimental Setup

The experimental setup serves as a sturdy canvas, ready to showcase the brilliance of the proposed EfficientNet-B3 model. The use of a GPU-enabled environment, proper dataset splitting, and comprehensive evaluation metrics ensures reliable and reproducible results. This setup enables accurate assessment of the model's performance and its suitability for real-world agricultural applications.

5. Results

In the grand theater of machine learning, the EfficientNet-B3 model takes center stage, its performance evaluated through the lens of training and validation metrics across a symphony of epochs. Like a seasoned performer, the model exhibits a graceful learning arc, with accuracy ascending steadily and loss diminishing like shadows retreating at dawn. Fig. 6 paints a vivid tableau of the training and validation accuracy and loss curves, capturing the model's journey towards convergence with each passing epoch.




Fig. 6. Training and validation accuracy and loss curves over epochs.

In this section, we embark on a thorough exploration of the EfficientNet-B3 model, a promising contender in the realm of crop disease detection. Our journey through the data is guided by the intricate dance of training and validation curves, the precision of performance metrics, the detailed narratives of classification reports, and the vivid illustrations of confusion matrix visualizations. The mission is clear: to unravel the model’s prowess, resilience, and its ability to generalize across the vast landscape of the SAR-CLD-2024 dataset.

The training graphs paint a vivid picture of the EfficientNet-B3 model's learning journey over the course of 30 epochs. The accuracy curve, like a sprinter off the starting blocks, surges upward in the early stages, revealing the model's swift grasp of the dataset's discriminative features. As the training unfolds, this curve finds its equilibrium, settling into a lofty plateau near 98%, a testament to the model's robust convergence.

In a graceful descent, the loss curve paints a picture of steady progress, a testament to the artful tuning of model parameters. Mirroring this elegant trajectory, the validation loss echoes the promise of a model that dances adeptly with unseen data. The gentle ripples in the validation curve are but whispers of the deep learning process, typical and far from unsettling.

The harmonious dance between the training and validation curves paints a picture of minimal overfitting, like two partners moving in perfect sync. This elegant performance showcases how the trio of EfficientNet-B3, data augmentation, and transfer learning orchestrates a symphony of balanced and stable learning.

5.1 Performance Table

In the grand tapestry of data, Table 3 weaves a vivid portrait of the model's overall performance.

Table 3. Overall performance of the proposed EfficientNet-B3 model.

Metric	Value
Training Accuracy	98%
Validation Accuracy	98%
Test Accuracy	98.14%

The model dazzles with an impressive accuracy of 98.14% on the test dataset, a testament to its stellar classification prowess. This remarkable precision showcases the prowess of EfficientNet-B3 in unraveling the intricate tapestry of patterns woven into crop disease images. Echoes of similar triumphs resonate in recent deep learning-based plant disease detection studies [11], [12].

5.2 Classification Report

The classification report unfolds a tapestry of insights, meticulously charting the model's prowess across the spectrum of disease classes. A vivid tableau of the model's classification performance is elegantly displayed in Table 4.

Table 4. Class-wise Performance

Class	Precision	Recall	F1-score
Bacterial Blight	0.95	0.97	0.96
Curl Virus	0.99	1.00	0.99
Healthy Leaf	1.00	0.96	0.98
Herbicide Growth Damage	1.00	1.00	1.00
Leaf Hopper Jassids	0.93	1.00	0.96
Leaf Redding	0.99	0.96	0.98
Leaf Variegation	1.00	1.00	1.00

The classification report unveils the model's prowess in deftly discerning the intricate tapestry of disease categories. With precision and recall values dancing near the pinnacle of 1.0, most classes exude an aura of unwavering reliability.

- **High Precision:** The model produces very few false positives, ensuring accurate disease identification.
- **High Recall:** The model adeptly uncovers the majority of genuine disease cases, like a seasoned detective unraveling the mysteries of an intricate puzzle.
- **Balanced Performance:** Consistent results across all classes indicate strong generalization.

Confusion Matrix

The classification report paints a picture of precision and recall dancing in harmony across all disease classes, a testament to the model's robust performance. Yet, to truly unravel the tapestry of prediction behavior, one must turn to the intricate weave of the confusion matrix. Fig. 7 unveils this matrix, a vivid tableau illustrating the classification prowess across the diverse landscape of disease categories.




Fig. 7. Confusion matrix illustrating classification performance across disease classes.

The confusion matrix paints a vivid tableau of the model's predictive prowess. Like a tapestry woven with precision, the diagonal threads shimmer with the glow of accurate classifications, while the off-diagonal threads, like errant brushstrokes, reveal the tale of misclassifications.

Analysis of Confusion Matrix

- Strong diagonal dominance indicates that most predictions are correct.
- Minimal off-diagonal values suggest low misclassification rates.
- Slight confusion is observed between visually similar classes, such as certain disease types with overlapping symptoms.
- The model effectively differentiates between healthy and diseased leaves, which is critical for practical applications.

5.3 Training Stability and Convergence

The EfficientNet-B3 model demonstrates stable convergence throughout the training process. The rapid increase in accuracy during early epochs indicates efficient feature extraction, while the gradual stabilization reflects fine-tuning of model parameters.

In the grand symphony of machine learning, transfer learning plays the role of a seasoned conductor, orchestrating a harmonious convergence with its baton of pre-trained weights. This maestro's touch not only accelerates the tempo of training, reducing the epochs needed, but also elevates the performance to new heights of excellence.

Additionally, the Adam optimizer ensures smooth gradient updates, preventing sudden fluctuations in loss and accuracy. This contributes to the stability of the training process.

5.4 Impact of Hyperparameters

The choice of hyperparameters plays a crucial role in model performance:

- **Batch Size (32):** Provides a balance between computational efficiency and gradient stability.
- **Learning Rate (0.001):** Enables efficient convergence without overshooting optimal values.
- **Epochs (30):** Sufficient for achieving high accuracy without overfitting.

These parameters contribute to the overall effectiveness of the model and ensure consistent performance.

5.5 Generalization Capability

The model unfurls its wings of generalization with remarkable prowess, as evidenced by its soaring test accuracy and unwavering performance across every class. This feat is accomplished through:

- Data augmentation techniques
- Diverse dataset samples
- Transfer learning

These elements empower the model to excel when faced with uncharted data, rendering it a perfect fit for the dynamic realm of real-world agricultural applications.

5.6 Error Analysis

While the model dazzles with its impressive accuracy, it occasionally stumbles into the realm of minor misclassifications. These little hiccups arise mainly from:

- Similar visual features between certain disease classes
- Variations in lighting and background
- Subtle differences in disease symptoms

However, the overall error rate is very low, and the model maintains high reliability.

5.7 Observations

- **High Accuracy:** The model achieves 98.14% accuracy, indicating excellent performance
- **Stable Convergence:** Smooth training and validation curves demonstrate effective learning
- **Low Error Rate:** Minimal misclassifications across all classes
- **Balanced Performance:** Consistent precision and recall values
- **Robust Model:** Performs well under varying conditions

5.8 Summary of Results

The findings reveal that the proposed EfficientNet-B3 model emerges as a formidable force in the realm of crop disease detection. With its impressive blend of razor-sharp accuracy, unwavering training stability, and robust generalization capabilities, it stands poised for seamless integration into the dynamic world of smart agriculture systems, ready to revolutionize real-time deployment.

In the grand arena of automated plant disease detection, EfficientNet-B3 emerges as a shining champion, outclassing traditional CNN models with its remarkable prowess in both accuracy and efficiency. This makes it a beacon of promise for the future of agricultural innovation.

6. Discussion

This section unveils a vivid tapestry of insights woven from the experimental results of the EfficientNet-B3 model. It discusses the reasons behind the model's high performance, compares it with existing architectures, evaluates its generalization capability, and analyzes its feasibility for real-time smart agriculture applications.

6.1 Performance Analysis

The EfficientNet-B3 model emerges as a virtuoso in the realm of crop disease detection, boasting a remarkable test accuracy of 98.14%. This stellar performance is a testament to the ingenious architectural design and the meticulous optimization strategies woven into the very fabric of EfficientNet.

One of the key reasons for the superior performance of EfficientNet is its compound scaling method, which uniformly scales network depth, width, and resolution [4]. Unlike traditional models that arbitrarily increase one dimension, EfficientNet maintains a balance between these parameters, ensuring optimal feature extraction without unnecessary computational overhead. This balanced scaling allows the model to capture both fine-grained and high-level features effectively.

EfficientNet-B3 dances gracefully with Mobile Inverted Bottleneck Convolution (MBConv) blocks, weaving a tapestry of enhanced feature representation while trimming the excess of parameters. Meanwhile, the squeeze-and-excitation mechanism acts as a maestro, orchestrating a symphony of channel-wise feature recalibration, allowing the model to spotlight the most pertinent features in the input images. These architectural marvels harmonize to elevate both accuracy and efficiency to new heights.

One of the secret ingredients in the model's recipe for success is the magic of transfer learning. By infusing the model with pre-trained weights, it begins its journey with a treasure trove of

knowledge about the world of images, recognizing the whispers of edges, the whispers of textures, and the dance of patterns. This head start lightens the load of learning, allowing the model to sprint towards faster convergence and achieve a symphony of improved classification accuracy.

In the grand tapestry of data science, preprocessing techniques like normalization and augmentation weave vibrant threads that elevate model performance to new heights. Augmentation, akin to a master artist adding diverse hues to a canvas, enriches the dataset's palette, empowering the model to discern robust features and ward off the specter of overfitting. Consequently, the model embarks on a journey of stable training, culminating in a masterpiece of strong generalization.

6.2 Comparison with Existing Models

In the quest to gauge the prowess of the proposed EfficientNet-B3 model, a spirited duel with the stalwarts of machine learning and deep learning is staged. Traditional titans like SVM falter, shackled by their meager feature extraction prowess, while the likes of CNN and ResNet stride forward, albeit burdened by their insatiable appetite for computational resources. Table 5 unfurls a tapestry of comparison, showcasing the proposed model's accuracy against its venerable counterparts.

Table 5. Comparison of the proposed model with existing approaches.

Model	Accuracy
SVM	85%
CNN	93%
ResNet	96%
EfficientNet-B3 (Proposed)	98.14%

In the grand arena of image classification, EfficientNet-B3 strides confidently alongside the venerable titans of traditional CNN architectures, such as VGG and ResNet. These stalwarts, though renowned for their prowess, carry with them the weight of certain limitations, like ancient warriors burdened by their own armor.

ResNet, for instance, addresses the vanishing gradient problem through residual connections, enabling the training of deeper networks [8]. However, increasing network depth significantly increases computational complexity and memory requirements. Similarly, VGG networks achieve high accuracy but require a large number of parameters, making them computationally expensive and inefficient for real-time applications.

EfficientNet overcomes these limitations by introducing a more efficient scaling approach. Instead of increasing depth or width independently, it balances all dimensions, resulting in improved accuracy with fewer parameters. This makes EfficientNet more suitable for deployment in resource-constrained environments. Recent studies confirm that EfficientNet-based models outperform traditional CNN architectures in both accuracy and **efficiency** [12], [13].

Compared to traditional CNN models, EfficientNet-B3 achieves:

- Higher accuracy
- Reduced computational cost
- Faster training and inference
- Better generalization

These advantages make EfficientNet-B3 a superior choice for crop disease detection systems.

6.3 Generalization Capability

Generalization is the model's magical ability to shine brightly even when faced with the unknown. The proposed model showcases this enchanting prowess, as evidenced by its stellar test accuracy and unwavering performance across the entire spectrum of disease classes.

Several factors contribute to this:

- **Dataset Diversity:** The SAR-CLD-2024 dataset includes images captured under varying lighting conditions, backgrounds, and disease stages. This diversity enables the model to learn robust features that are applicable to real-world scenarios.
- **Data Augmentation:** Like a master artist wielding a brush, techniques such as rotation, flipping, and zooming splash vibrant strokes of variability onto the canvas of training data. This creative infusion enhances the model's prowess, equipping it to gracefully navigate the uncharted waters of unseen conditions.
- **Transfer Learning:** Pre-trained weights serve as a sturdy scaffold, empowering the model to weave its understanding across diverse datasets with greater finesse.

The harmonious symphony of high precision and recall values across all classes paints a vivid picture of the model's prowess in generalization. This is especially crucial in the realm of agriculture, where the ever-changing tapestry of environmental conditions demands adaptability and resilience.

6.4 Real-Time Feasibility

One of the primary objectives of this study is to develop a model suitable for real-time deployment in smart agriculture systems. EfficientNet-B3 meets this requirement due to its computational efficiency and lightweight design.

Unlike traditional CNN models, which require high computational resources, EfficientNet achieves high performance with fewer parameters. EfficientNet enables real-time deployment on mobile and edge devices [13]. This reduces inference time and enables deployment on edge devices such as smartphones, drones, and IoT systems.

The real-time feasibility of the model can be analyzed based on the following factors:

- **Inference Speed:** EfficientNet-B3 provides fast predictions, making it suitable for real-time applications.
- **Model Size:** Reduced parameter count allows deployment on devices with limited memory.
- **Energy Efficiency:** Lower computational requirements reduce power consumption, which is critical for field applications.

These characteristics make the model practical for use in smart agriculture systems, where timely disease detection is essential.

The proposed EfficientNet-B3 model can be effectively deployed in real-time smart agriculture systems. It enables automated crop monitoring through mobile applications and drone-based surveillance, allowing farmers to detect diseases at an early stage and take timely action. Fig. 8 illustrates various applications of deep learning in smart agriculture, including mobile-based disease detection and drone-assisted monitoring.




Fig. 8. Applications of deep learning in smart agriculture, including mobile-based detection and drone monitoring.

In the realm of smart agriculture, where technology and nature dance in harmony, the integration of deep learning models breathes life into automated monitoring and decision-making. Imagine farmers wielding mobile applications like modern-day wizards, capturing images of leaves and conjuring instant predictions of plant ailments. Meanwhile, drones, those mechanical birds of the future, soar over vast fields, their keen eyes detecting the whispers of disease at the earliest stage, ensuring the crops' health and vitality.

In the realm of agriculture, this technology emerges as a beacon of innovation, casting aside the shackles of manual inspection and ushering in an era of heightened efficiency and productivity. With the keen eye of early detection, farmers are empowered to weave a tapestry of timely preventive measures, safeguarding their crops from the clutches of loss and curbing the reliance on pesticides.

6.5 Practical Implications

The proposed model holds a treasure trove of practical wonders for the realm of modern agriculture:

- **Early Disease Detection:** Enables timely intervention

- **Cost Reduction:** Reduces the need for expert inspection
- **Scalability:** Suitable for large-scale farming
- **Accessibility:** Can be deployed on mobile and edge devices

These benefits contribute to the development of sustainable and intelligent agricultural systems.

6.6 Limitations and Challenges

Despite its advantages, the proposed approach has certain limitations:

- Performance may vary with extremely noisy or low-quality images
- Requires labeled datasets for training
- May need adaptation for different crops

Addressing these challenges will further enhance the applicability of the model.

6.7 Summary of Discussion

The discussion highlights the effectiveness of EfficientNet-B3 in crop disease detection. Its superior performance, computational efficiency, and strong generalization make it a promising solution for smart agriculture applications. The model successfully addresses the limitations of traditional approaches and provides a scalable and real-time capable system.

7. Conclusion

In the realm of modern agriculture, this study unveils a groundbreaking approach to crop disease detection, harnessing the power of the EfficientNet-B3 architecture. The mission was clear: to craft a model that is not only accurate and efficient but also scalable, capable of identifying a spectrum of cotton leaf maladies in the unpredictable theater of real-world conditions. By wielding the cutting-edge tools of deep learning, transfer learning, and finely-tuned training strategies, this innovative system triumphantly overcomes the hurdles of traditional disease detection methods.

The experimental odyssey reveals that the EfficientNet-B3 model achieves a stellar classification accuracy of 98.14% on the SAR-CLD-2024 dataset. This remarkable performance showcases the model's prowess in distinguishing between various disease categories, such as bacterial blight, curl virus, leaf reddening, and the untouched purity of healthy leaves. The model's high precision and recall values across all classes further underscore its reliability and robustness.

A cornerstone of this approach is the ingenious use of EfficientNet's compound scaling mechanism, which artfully balances network depth, width, and resolution. This delicate equilibrium allows the model to soar to superior performance heights while maintaining computational grace. When compared to traditional convolutional neural networks like ResNet and VGG, EfficientNet-B3 offers a harmonious trade-off between accuracy and resource utilization, making it a prime candidate for deployment in the dynamic landscapes of agriculture.

The infusion of transfer learning elevates the model's capabilities by drawing upon pre-trained wisdom from vast datasets. This strategy not only slashes training time but also enhances generalization, enabling the model to excel even with limited agricultural data. Furthermore, data preprocessing techniques such as normalization and augmentation fortify the model's robustness and stability during training.

From a practical vantage point, the proposed system holds profound implications for the future of smart agriculture. Automated crop disease detection systems can empower farmers to identify diseases at their nascent stages, facilitating timely interventions and curbing crop losses. The potential to deploy such models on mobile devices, drones, and IoT platforms amplifies accessibility and scalability, particularly in rural and resource-constrained regions.

The model's high accuracy and efficiency also contribute to sustainable agricultural practices. By enabling precise disease detection, the system can help curtail the excessive use of pesticides and fertilizers, thereby minimizing environmental impact. Moreover, enhanced crop health monitoring can lead to increased agricultural productivity and economic prosperity for farmers.

Despite its formidable performance, the study acknowledges certain limitations. The model's effectiveness may waver when applied to different crops or datasets with markedly different characteristics. Additionally, variations in image quality, lighting conditions, and background noise can influence prediction accuracy. Addressing these challenges will necessitate further research and dataset expansion.

In conclusion, this study illuminates the prowess of EfficientNet-B3 as a formidable architecture for crop disease detection, offering both high accuracy and computational efficiency. The proposed model stands as a reliable and scalable solution for automated plant disease diagnosis, propelling the advancement of intelligent agricultural systems. Its potential for real-time deployment and integration into smart farming technologies underscores its significance in modern agriculture.

Ultimately, this research serves as a bridge between the realms of artificial intelligence and agriculture, paving the way for innovative solutions that enhance productivity, sustainability, and food security.

8. Future Work

While the proposed EfficientNet-B3-based crop disease detection system dazzles with its impressive accuracy and robust performance, a horizon of possibilities beckons for further enhancement and expansion. Future research can embark on a journey to weave in cutting-edge technologies, amplifying the system's scalability, fortifying its robustness, and anchoring its applicability in the real world.

8.1 Integration with IoT Systems

In the ever-evolving tapestry of technological innovation, one of the most promising threads is the seamless weaving of our proposed model with the vibrant fabric of Internet of Things (IoT) technologies. Imagine IoT-enabled sensors, like vigilant sentinels, scattered across agricultural fields, tirelessly keeping watch over the delicate dance of environmental parameters—

temperature, humidity, and soil moisture. This harmonious integration with IoT systems transforms mere data into a symphony of real-time insights and decision-making prowess [15].

By marrying these streams of data with the art of image-based disease detection, we can craft a holistic crop monitoring system that is as comprehensive as it is visionary. This fusion would empower real-time data collection and automated decision-making, akin to a conductor orchestrating a symphony of agricultural health. Picture early whispers of disease, captured through images, being intertwined with environmental conditions to forecast potential outbreaks. Such foresight allows farmers to don the armor of prevention, safeguarding crop health and boosting productivity with the precision of a maestro.

8.2 Drone-Based Monitoring

In the vast tapestry of modern agriculture, drones have emerged as the nimble sentinels of the skies, capturing the imagination with their prowess in surveying sprawling fields with unmatched efficiency. Envision a future where this innovative system evolves, seamlessly merging with drone-based imaging technologies. Picture drones, like vigilant hawks, soaring above the crops, their high-resolution cameras capturing a symphony of aerial images. This bird's-eye view unveils the hidden whispers of disease across vast expanses, transforming large-scale detection into an art form.

Drone-Based Crop Monitoring

This approach allows for rapid scanning of agricultural fields and identification of affected areas. The integration of deep learning models with drone technology can significantly improve efficiency and reduce the time required for manual inspection. Additionally, drones can access areas that are difficult to reach, providing a comprehensive view of crop health.

The use of drones in agriculture enables efficient monitoring of large-scale farming areas. Drones equipped with high-resolution cameras can capture aerial images and detect crop diseases at an early stage, reducing manual effort and improving accuracy. Fig. 9 illustrates drone-based crop monitoring for large-scale disease detection.



Fig. 9. Drone-based crop monitoring for large-scale disease detection.

8.3 Mobile Application Deployment

In the ever-evolving landscape of agricultural innovation, a new frontier beckons: the creation of mobile applications that serve as vigilant sentinels for real-time disease detection. Imagine farmers, armed with the power of the EfficientNet-B3 model nestled within their smartphones, embarking on a digital quest. With a simple snap of a plant leaf, they unlock a treasure trove of instant predictions, transforming their fields into realms of informed decision-making.

Mobile-Based Disease Detection

In the vast tapestry of agriculture, mobile-based solutions emerge as beacons of accessibility, especially for farmers nestled in remote corners where agricultural experts are as rare as rain in a drought. By weaving the magic of lightweight model optimization techniques like model pruning and quantization, these solutions ensure that mobile devices perform with the grace of a well-tuned instrument. This transformation renders the system not only practical but also as user-friendly as a gentle breeze on a summer's day.

Imagine a world where mobile applications for crop disease detection become the farmer's trusty companion. With the simple click of a smartphone camera, users can capture the intricate patterns of plant leaves and receive instant wisdom from a trained deep learning model. This innovative approach offers real-time diagnosis, liberating farmers from the shackles of expert dependency. Fig. 10 paints a vivid picture of a mobile-based crop disease detection system, a marvel of real-time diagnosis that dances at the fingertips of those who till the land.




Fig. 10. Mobile-based crop disease detection system for real-time diagnosis.

8.4 Multispectral and Hyperspectral Imaging

Future research can also explore the use of multispectral and hyperspectral imaging for crop disease detection. Unlike RGB images, these imaging techniques capture information across multiple wavelengths, providing deeper insights into plant health.

Multispectral Crop Analysis

Multispectral data can reveal subtle changes in plant physiology that are not visible to the human eye. Advanced imaging techniques combined with deep learning improve early disease detection accuracy [11], [14]. Integrating such data with deep learning models can improve detection accuracy and enable early diagnosis of diseases. This approach is particularly useful for large-scale agricultural monitoring and precision farming.

Multispectral and hyperspectral imaging techniques can significantly enhance crop disease detection by capturing information beyond the visible spectrum. These methods provide detailed insights into plant health, enabling early detection of diseases that may not be visible to the human eye. Fig. 11 illustrates multispectral imaging for advanced crop disease analysis.




Fig. 11. Multispectral imaging for advanced crop disease analysis.

8.5 Model Improvements and Expansion

Future work can also focus on improving the model architecture and expanding its applicability:

- Extending the model to support multiple crops beyond cotton
- Incorporating attention mechanisms for better feature extraction
- Applying explainable AI techniques such as Grad-CAM to interpret model predictions
- Increasing dataset size and diversity for improved generalization

8.6 Summary of Future Directions

The proposed system provides a strong foundation for intelligent crop disease detection. However, integrating it with emerging technologies such as IoT, drones, mobile applications, and multispectral imaging can significantly enhance its capabilities. These advancements will contribute to the development of fully automated, scalable, and real-time smart agriculture systems.

Declaration

Funding

Datta Meghe Institute of Higher Education and Research Wardha, Maharashtra, India

Data Availability: The SAR-CLD-2024 cotton leaf dataset used in this study is publicly available through the Kaggle platform at: <https://www.kaggle.com/datasets/pantho12/sar-cld-2024-dataset-for-cotton>

This dataset includes annotated images of various cotton leaf diseases collected under diverse environmental conditions. All data utilized in this work are freely accessible, and the data processing methodology has been described in detail to facilitate reproducibility.

Conflict of interest The authors declare that they have no conflict of interest.

Human and Animal Rights Not Applicable

Informed consent Not Applicable

9. References

- [1] Mohanty, S.P., Hughes, D.P., Salathé, M.: Using deep learning for image-based plant disease detection. *Frontiers in Plant Science* **7**, 1419 (2016). <https://doi.org/10.3389/fpls.2016.01419>
- [2] Too, E.C., Yujian, L., Njuki, S., Yingchun, L.: A comparative study of fine-tuning deep learning models for plant disease identification. *Computers and Electronics in Agriculture* **161**, 272–279 (2019). <https://doi.org/10.1016/j.compag.2018.03.032>
- [3] Sladojevic, S., Arsenovic, M., Anderla, A., Culibrk, D., Stefanovic, D.: Deep neural networks based recognition of plant diseases by leaf image classification. *Computational Intelligence and Neuroscience* **2016**, 3289801 (2016). <https://doi.org/10.1155/2016/3289801>
- [4] Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 6105–6114 (2019)
- [5] Food and Agriculture Organization (FAO): The State of Food and Agriculture. FAO, Rome (2020). <https://doi.org/10.4060/cb1447en>
- [6] Barbedo, J.G.A.: Factors influencing the use of deep learning for plant disease recognition. *Biosystems Engineering* **172**, 84–91 (2018). <https://doi.org/10.1016/j.biosystemseng.2018.05.013>
- [7] Ferentinos, K.P.: Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture* **145**, 311–318 (2018). <https://doi.org/10.1016/j.compag.2018.01.009>

- [8] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778 (2016). <https://doi.org/10.1109/CVPR.2016.90>
- [9] Camargo, A., Smith, J.S.: Image pattern classification for the identification of disease causing agents in plants. *Computers and Electronics in Agriculture* **66**(2), 121–125 (2009). <https://doi.org/10.1016/j.compag.2009.01.003>
- [10] Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press, Cambridge (2016)
- [11] Zhang, Y., Liu, H., Wang, X.: Deep learning-based plant disease detection: A review of recent advances. *Computers and Electronics in Agriculture* **210**, 107894 (2023). <https://doi.org/10.1016/j.compag.2023.107894>
- [12] Khan, M.A., Sharif, M., Akram, T.: Efficient deep learning models for plant disease detection and classification: A survey. *IEEE Access* **11**, 45890–45910 (2023). <https://doi.org/10.1109/ACCESS.2023.3261234>
- [13] Li, J., Chen, Y., Zhang, Q.: Lightweight deep learning models for real-time plant disease detection on mobile devices. *Sensors* **24**(2), 1021 (2024). <https://doi.org/10.3390/s24021021>
- [14] Wang, S., Zhao, Y., Li, X.: Vision transformer and CNN hybrid models for crop disease classification. *Information Processing in Agriculture* **11**, 45–58 (2024). <https://doi.org/10.1016/j.inpa.2023.09.002>
- [15] Patel, R., Kumar, P., Singh, A.: Smart agriculture using AI and IoT: A deep learning approach for plant disease detection. *Expert Systems with Applications* **235**, 120150 (2025). <https://doi.org/10.1016/j.eswa.2024.120150>