



Multi-Crop Disease Detection and Management Recommendation System Using Deep Learning

K. Hemalatha¹, Prof. K. Venkata Rao², P. Swathi³

M. Tech Student, Department of Computer Science and Systems Engineering, Andhra University College of Engineering (A), Visakhapatnam, Andhra Pradesh, India¹

Professor, Department of Computer Science and Systems Engineering, Andhra University College of Engineering (A), Visakhapatnam, Andhra Pradesh, India²

Research Scholar, Department of Computer Science and Systems Engineering, Andhra University College of Engineering (A), Visakhapatnam, Andhra Pradesh, India³

Abstract: Crop diseases are a major challenge in agriculture, as late or incorrect identification can cause crop damage and reduce productivity. Existing crop disease detection systems mainly focus on identifying diseases from plant leaf images, but provide limited support for disease management. This creates a gap between disease identification and practical management. The proposed Multi-Crop Disease Detection and Management Recommendation System Using Deep Learning addresses this gap by combining disease detection with management recommendations. The system uses deep learning and transfer learning models to classify crop diseases from leaf images. It provides chemical and organic management approaches based on the identified crop and disease. Thus, the system connects disease identification with suitable management practices, making it more useful for practical crop disease management.

Keywords: Crop Disease Detection, Deep Learning, Management Recommendation System, Multi-Crop Classification, Transfer Learning.

I. INTRODUCTION

Plant diseases remain one of the most persistent threats to agricultural output and food security across the world. When a disease goes unnoticed or is identified wrongly, the consequences range from lower yield and poorer produce quality to the complete loss of a crop during severe outbreaks. Because the number of known crop diseases keeps growing, farmers and agronomists need automated tools that are reliable and can scale.

Precision agriculture uses many tools to track crop health and limit production losses, and recognizing a disease quickly from a photograph of a leaf is one of the most valuable among them. Convolutional Neural Networks (CNNs) and transfer learning have become popular for this task because they can learn the visual patterns that separate one disease from another. An image-based approach lets farmers act early, before an infection spreads through the field.

Many deep learning models have been proposed for this purpose. Most start from CNN backbones pre-trained on ImageNet, such as EfficientNet and MobileNet, and fine-tune them on leaf images of a particular crop. Accuracy alone, however, does not tell a farmer what to do next. A practical system must also be easy to use and must turn its prediction into advice, which is why web applications built with frameworks such as Streamlit are valuable: they place both the diagnosis and the management guidance in one simple interface.

The need for disease information is not limited to one crop or region. Smallholder farmers, extension workers and large agribusinesses all depend on timely diagnosis, from rice paddies to cotton fields, to protect yields and keep food supplies steady.

This paper presents a multi-crop disease detection and management recommendation system built on transfer learning for five crops. The trained models are deployed in a web application in which the predicted disease is combined with the crop's growth stage and the severity of infection to produce structured chemical and organic management recommendations.



The main contributions of this work are as follows:

- A single application that classifies leaf diseases of five crops (Rice, Wheat, Chilli, Cotton and Maize) across twenty disease and healthy classes.
- A comparison of three transfer learning architectures (EfficientNetB0, EfficientNetB2 and MobileNetV2) using the same pipeline for every crop.
- A rule-based recommendation engine that gives chemical and organic management guidance based on the predicted disease, the crop's growth stage and the severity of infection.
- A Streamlit web application that brings image upload, disease prediction and management recommendation into one workflow.

II. LITERATURE SURVEY

Image-based detection of crop diseases has become a key part of precision agriculture as crop varieties and disease strains diversify [1]. This section reviews existing work on detecting diseases from leaf images across different crops.

Transfer Learning and DenseNet-Based Detection: Mithul et al. [1] used transfer learning to detect paddy crop diseases, fine-tuning pre-trained CNN backbones on rice leaf images with strong results. Reha et al. [2] identified plant leaf diseases with DenseNet, whose dense connections encourage feature reuse and help gradients flow. Chandu et al. [3] compared machine learning and deep learning approaches to crop disease detection and found that deep transfer-learning models consistently outperform classical classifiers.

Image-Based Deep Learning Studies: Amritavarshini et al. [4] detected plant leaf diseases from images with deep learning and reached high accuracy on curated datasets. Rozaqi and Sunyoto [5] used CNNs to identify diseases in potato leaves, and Andrew et al. [6] surveyed deep learning techniques for leaf disease detection across several agricultural applications, noting that most systems are trained and evaluated on a single crop.

Hybrid and Comparative Architectures: Chen et al. [7] combined a CNN with a Transformer for tomato leaf disease detection and reported high accuracy with strong generalization. Sanida et al. [8] proposed an efficient hybrid CNN that merges EfficientNet and MobileNet variants for tomato disease classification. Together, these studies suggest that deeper, compound-scaled or densely connected networks tend to be more accurate than lightweight ones such as MobileNetV2, at the price of slower inference and larger models.

Large-Scale and EfficientNet-Based Studies: Mohanty et al. [9] trained a deep CNN on 54,306 PlantVillage images covering 14 crop species and 26 diseases and reported 99.35% accuracy on a held-out test set. Their accuracy fell to 31.4% on images taken under different conditions, which shows how strongly such models depend on the conditions of the training data. Atila et al. [10] applied EfficientNet with transfer learning to the same dataset, and the B5 and B4 variants gave the best results, with accuracies of 99.91% on the original dataset and 99.97% on the augmented dataset.

Lightweight Models and Management Suggestions: Prattasha and Reza [11] proposed a depthwise separable CNN for 12 rice diseases that reached 95.3% test accuracy with a substantially smaller model size. Sritharan et al. [12] developed a mobile application that detects plant diseases from leaf images and suggests pesticides and management practices for the detected disease, reporting 98% accuracy. This work shows the value of pairing detection with treatment advice. In the proposed system, the recommendation additionally depends on the crop's growth stage and the severity of infection. Taken together, the reviewed studies achieve high accuracy on curated datasets, most of them for a single crop, and few of them connect the detected disease to management guidance. This motivates the research gap described next.

III. RESEARCH GAP

The reviewed literature reveals the following gaps:

- **Limited multi-crop integration:** most systems are trained and deployed for a single crop, forcing farmers who grow several crops to use separate tools.
- **Detection without actionable recommendation:** most systems output only a disease label (and sometimes a confidence score) without chemical or organic management guidance.
- **Lack of growth-stage and severity awareness:** very few systems use the crop's growth stage or the observed severity of infection, both critical to selecting a real-world treatment.
- **Insufficient comparative evaluation in a unified pipeline:** few projects compare the same architectures (EfficientNetB0, EfficientNetB2, MobileNetV2) consistently across multiple crops within one integrated system.



The proposed system directly addresses each of these gaps by supporting five crops and twenty disease/healthy classes within a single application, comparatively evaluating three transfer learning architectures per crop, and integrating a growth-stage- and severity-aware chemical/organic recommendation engine into the detection workflow.

IV. METHODOLOGY

The overall workflow of the proposed system is shown in Fig. 1 and described in the following subsections.




Fig. 1: System flowchart depicting the complete workflow.

A. Data Acquisition and Preprocessing

Crop leaf images for Rice, Wheat, Chilli, Cotton and Maize were sourced from publicly available Kaggle datasets [13]. Each dataset contains labelled images for the crop's disease classes and a healthy class. Preprocessing ensured quality and consistency:

- *Image Resizing*: images were resized to the resolution required by each architecture (224×224 for EfficientNetB0/MobileNetV2, 260×260 for EfficientNetB2).
- *Normalization*: pixel values were rescaled using each architecture's dedicated preprocessing function for stable training.
- *Data Augmentation*: random rotation, flipping, zoom and brightness modulation were applied to the training set to increase diversity and reduce overfitting.

B. Exploratory Data Analysis

Class distribution analysis identified class imbalance, which was corrected through augmentation. Visual inspection of sample images from each class confirmed labelling consistency and identified visually similar disease pairs likely to cause classification confusion.



C. Feature Extraction (Transfer Learning)

Pre-trained EfficientNetB0 and EfficientNetB2 [14] and MobileNetV2 [15] backbones, initialized with ImageNet weights, were used to extract high-level spatial features from the preprocessed leaf images. A Global Average Pooling layer, a Dropout layer and a Dense (softmax) output layer were attached to each backbone to form the crop-specific classification head (Fig. 2).




Fig. 2: Transfer Learning Architecture.

D. Model Training and Evaluation

Each architecture was trained independently per crop using a two-phase strategy: an initial feature-extraction phase with the backbone frozen, followed by fine-tuning in which the top backbone layers were unfrozen and retrained at a reduced learning rate. Models were evaluated using accuracy on a held-out validation split, and the confusion matrix was used to examine per-class performance.

E. Model Selection and Comparison

The performance of EfficientNetB0, EfficientNetB2 and MobileNetV2 was compared on the basis of accuracy, and the best-performing model per crop was selected for deployment in the final system.

F. Recommendation Engine Design

The predicted disease, together with the user-specified crop growth stage and disease severity, is passed to a rule-based recommendation engine that returns structured chemical and organic/natural management guidance, connecting detection directly with actionable treatment steps (Table I).

TABLE I INPUTS AND OUTPUT OF THE RECOMMENDATION ENGINE

Parameter	Source	Role in recommendation
Crop	User selection	Selects the crop-specific model
Disease	Model prediction	Identifies the disease class
Growth stage	User input	Determines whether chemical advice is suitable
Severity	User input	Influences the recommended treatment
Output	Rule base	Chemical and organic guidance

G. Ethical Considerations

To ensure responsible use of the recommendation engine, chemical treatment suggestions are withheld where severity and growth-stage combinations warrant caution (e.g., late-stage flowering crops), and organic guidance is favoured in such cases to avoid potentially harmful chemical application.



H. Novelty of the Project

- 1) *Multi-Crop, Multi-Architecture Integration*: Unlike traditional single-crop approaches, our methodology integrates five crops (Rice, Wheat, Chilli, Cotton and Maize) within one system and comparatively evaluates three transfer learning architectures per crop, enabling selection of the most suitable architecture per deployment context.
- 2) *Growth-Stage- and Severity-Aware Recommendation*: The recommendation engine is sensitive to both the crop's growth stage and the observed disease severity, rather than issuing a single static recommendation per disease, and can withhold or adjust chemical guidance under sensitive conditions.
- 3) *Detection-to-Action Pipeline*: Beyond classification, the system connects detection outputs directly to actionable chemical and organic management steps, closing the loop that most detection-only systems leave open.

I. Dataset Analysis and Description

The combined dataset comprises leaf images for five crops spanning twenty disease and healthy classes, each labelled with crop type and disease category (Table II):

TABLE II DATASET DESCRIPTION: CROPS AND CLASSES

Crop	Classes	Healthy Class	Disease Classes
Rice	5	Yes	Bacterial Blight, Blast, Brown Spot, Tungro
Wheat	3	Yes	Septoria, Stripe Rust
Chilli	4	Yes	Anthrachnose, Cercospora Leaf Spot, Leaf Curl
Cotton	4	Yes	Bacterial Blight, Curl Virus, Fusarium Wilt
Maize	4	Yes	Blight, Common Rust, Gray Leaf Spot
Total	20	5	15 Disease Classes

This dataset covers cereal, fibre and vegetable crops with diverse lesion shapes, colour patterns and severity levels, enabling comparative evaluation of classifiers that generalize across visually distinct disease presentations.

J. Algorithms and Their Architectures

Three transfer learning architectures were employed. All three are convolutional neural networks pre-trained on ImageNet, and each was used as a feature-extraction backbone followed by the classification head shown in Fig. 2. Table III summarizes their main characteristics, which follow the original architecture definitions [14], [15]. The classification head of every model ends with a softmax output layer.

TABLE III COMPARISON OF THE BACKBONE ARCHITECTURES

Feature	EfficientNetB0	EfficientNetB2	MobileNetV2
Input size	224×224	260×260	224×224
Building block	MBConv + SE	MBConv + SE	Inverted residual
Backbone activation	Swish	Swish	ReLU6
Output activation	Softmax	Softmax	Softmax
Blocks	16	23	17
Final features	1280	1408	1280
Parameters (approx.)	5.3 M	9.2 M	3.4 M

SE: squeeze-and-excitation; M: million. Softmax is the output activation of the classification head used in this work. Parameter counts are for the ImageNet-trained networks.

1) *EfficientNetB0*: a compound-scaled CNN with a strong balance of accuracy and computational efficiency, used as the primary high-accuracy baseline. Its architecture starts with a 3×3 convolution stem, followed by 16 mobile inverted bottleneck (MBConv) blocks arranged in seven stages, and ends with a 1×1 convolution that produces 1280 features. Each MBConv block expands the channels with a 1×1 convolution, applies a depthwise convolution, recalibrates the channels with a squeeze-and-excitation module and projects them back with a 1×1 convolution, using the Swish activation (Fig. 3a) [14].

2) *EfficientNetB2*: a moderately scaled-up EfficientNet variant with a larger input resolution and a modest increase in computation. It uses the same MBConv design as EfficientNetB0, but the network is deeper (23 blocks), wider (1408



output features) and takes a 260×260 input, because the compound scaling rule increases depth, width and resolution together. The higher resolution and capacity allow finer lesion texture to be captured [14].

3) *MobileNetV2*: a lightweight architecture using depth-wise separable convolutions, selected as the efficiency-focused baseline for resource-constrained deployment. Its main building block is the inverted residual with a linear bottleneck, made of a 1×1 expansion convolution, a 3×3 depthwise convolution and a 1×1 linear projection, with ReLU6 activation and a shortcut connection when the input and output shapes match (Fig. 3b). The network has a 3×3 convolution stem, 17 such blocks and a final 1×1 convolution that gives 1280 features. Depthwise separable convolutions greatly reduce parameters and computation, which makes the model efficient but limits its representational capacity [15].




Fig. 3: Building blocks of the backbones: (a) MBConv block of EfficientNet; (b) inverted residual block of MobileNetV2.

Each architecture offers distinct advantages, allowing a comprehensive comparison of accuracy versus computational efficiency. The complete detection and recommendation procedure that uses these models is given in Algorithm 1.

Algorithm 1: Disease Detection and Management Recommendation

Step 1: Take the crop type and the leaf image as input.

Step 2: Load the trained model of the selected crop.

Step 3: Resize the image to the model input size and normalize the pixel values.

Step 4: Pass the image through the backbone, Global Average Pooling, Dropout and the softmax layer to obtain the class probabilities.

Step 5: Take the class with the highest probability as the predicted disease.

Step 6: Read the growth stage and the infection severity entered by the user.

Step 7: Query the rule base with the crop, disease, growth stage and severity.

Step 8: Display the predicted disease with the chemical and organic management recommendation.

K. System Architecture

The application follows a client-server architecture implemented with the Streamlit framework, and consists of the following components (Fig. 4):

1) *Client side*: a five-step interface (crop selection, leaf image upload, disease prediction, growth-stage/severity entry and recommendation display) through which farmers or agronomists interact with the system.

2) *Server side*: handles image upload, invokes the appropriate per-crop trained model for inference and queries the recommendation rule base.

3) *Model storage*: trained models for each crop-architecture combination are serialized (.keras/SavedModel) and loaded on demand with caching.

4) *Recommendation rule base*: chemical and organic rules keyed by crop, disease, growth stage and severity are queried at inference time.

5) *Data flow*: the uploaded image and crop selection are preprocessed and passed to the model; the predicted disease, growth stage and severity are passed to the recommendation engine, and the result is returned to the client.

6) *Scalability*: the modular per-crop design allows new crops or updated models to be added without redesigning the full system.

7) *Result display*: the predicted disease and the corresponding chemical and organic management recommendation are shown together on the results page of the client interface.




Fig. 4: System Architecture.

V. RESULTS

We evaluated EfficientNetB0, EfficientNetB2 and MobileNetV2 for crop disease classification using the Rice dataset as a representative example. Accuracy, the proportion of correctly predicted instances, is used as the primary evaluation metric.

A. Performance Comparison

EfficientNetB0 achieved the highest accuracy of 97%, followed by EfficientNetB2 at 96% (Table IV; Fig. 5). MobileNetV2 recorded the lowest accuracy of about 72%, with noticeable variation across training runs, and remained below both EfficientNet models.

TABLE IV ACCURACY COMPARISON OF MODELS (RICE DATASET)

Model	Input size	Accuracy (%)
EfficientNetB0	224×224	97
EfficientNetB2	260×260	96
MobileNetV2	224×224	72*

*Approximate; varied across runs.




Fig. 5: Bar graph of accuracy across models.



B. Per-Class Analysis

To examine per-class performance, Fig. 6 presents the confusion matrix for the EfficientNetB0 model on the Chilli dataset, covering Anthracnose, Cercospora Leaf Spot, Fresh Leaf and Leaf Curl Disease. The model correctly classified 46 of 51 Anthracnose samples, with the remaining 5 misclassified as Cercospora Leaf Spot. It classified 90 of 91 Cercospora Leaf Spot and all 76 Leaf Curl Disease samples correctly, and 62 of 63 Fresh Leaf (healthy) samples, with one confused for Cercospora Leaf Spot. Confusion was thus limited to Anthracnose and Cercospora Leaf Spot, which can present visually similar early-stage lesions.




Fig. 6: Confusion matrix of EfficientNetB0 on the Chilli dataset.

C. Model Output

The Streamlit application integrates the pipeline in a user-friendly interface. After a leaf image is uploaded, a confirmation message is shown, followed by the disease prediction. The user then selects the crop's growth stage and severity, and the application displays the corresponding chemical and organic management recommendation (Figs. 7–9). In the example shown, the uploaded leaf is predicted as Cercospora Leaf Spot, after which the growth stage is entered and both chemical and organic management recommendations are displayed.




Fig. 7: Input Page – Crop Selection and Leaf Image Upload.



Prediction Complete! 🎉

Prediction Result

Cercospora Leaf Spot ⇄

Fig. 8: Disease Prediction.

Step 4: Crop Growth Information

Select Growth Stage
▼

Seed / Nursery

Disease Severity
▼

Mild

Step 5: Management Recommendation

Recommendations generated successfully! 🎉

Chemical Management

• Use an appropriate fungicide containing Mancozeb, Chlorothalonil, or another fungicide suitable for chilli Cercospora leaf spot.

Organic / Natural Management

• Remove severely infected leaves and fallen diseased debris, reduce prolonged leaf wetness, and use Trichoderma or Bacillus biological management.

Crop: Chilli | Disease: Cercospora Leaf Spot | Growth Stage: Seed / Nursery | Severity: Mild

Fig. 9: Output Page – Predicted Disease and Management Recommendation.

D. Discussion

EfficientNetB0 and EfficientNetB2 performed almost equally, with EfficientNetB0 higher by one percentage point. As EfficientNetB0 is the smaller of the two, it is the more economical choice for deployment. MobileNetV2 is a lightweight network with fewer parameters than either EfficientNet model; it reached about 72% and its accuracy varied between runs, so it is less suitable where high accuracy is required. Because the recommendation engine acts on the predicted disease, the reliability of the management advice depends directly on the accuracy of the classifier.

VI. CONCLUSION

This work developed a multi-crop disease detection and management recommendation system using transfer learning, and evaluated EfficientNetB0, EfficientNetB2 and MobileNetV2 across five crops. EfficientNetB0 emerged as the top-performing model with 97% accuracy, closely followed by EfficientNetB2 at 96%. MobileNetV2 was less accurate (about 72%) and less consistent across runs, but its smaller model size and faster inference make it suited to resource-constrained, real-time deployment where some loss of accuracy is acceptable.

Beyond model comparison, the growth-stage- and severity-aware recommendation engine provides practical chemical and organic guidance directly within the detection pipeline, closing the gap between diagnosis and real-world crop disease management left by detection-only systems.

The study has some limitations. The models were trained on publicly available Kaggle images, which may not reflect the variety of real field conditions (models trained on controlled images can perform much worse on images taken under other conditions [9]), growth stage and severity are entered manually by the user, and the recommendations are generated from a rule base rather than learned from data.



VII. FUTURE SCOPE

The system can be extended to additional crops and disease types to serve a wider range of farming communities. Offline, mobile-based inference would allow use in areas with limited connectivity, and automated severity estimation directly from the leaf image would reduce dependence on manual input. The recommendation engine can also be extended to suggest dosage based on the field area and the affected yield, so that the quantity of chemical or organic treatment is matched to the extent of infection. Incorporating real-time weather and soil data, multilingual support for farmers, and a continuous feedback and retraining mechanism can further improve accuracy, transparency and reliability.

REFERENCES

- [1] D. Mithul K. V., N. Mohankumar, A. A. K. Ashok, and E. P. Prakash, "Deep learning-based paddy crop disease detection using transfer learning approaches," in *Proc. 5th Int. Conf. Communication, Computing and Electronics Systems (ICCCES-2026)*, 2026.
- [2] P. Reha, B. Dhanam, P. Marichamy, and P. Lingeswari, "Deep learning based approach for plant leaf disease identification using DenseNet," in *Proc. Int. Conf. Computer Networks and Inventive Communication Technologies (ICCNCT-2026)*, 2026.
- [3] H. P. Chandu et al., "Machine and deep learning approaches for crop disease detection: An in-depth analysis," in *2024 IEEE 9th Int. Conf. Convergence in Technology (I2CT)*, Pune, India, 2024.
- [4] S. Amritavarshini, K. Sai Ganesh, M. Arvind, K. Ravi Teja, and D. Vijayan, "Image-based plant leaf disease detection using deep learning," in *2024 Int. Conf. Signal Processing, Computation, Electronics, Power and Telecommunication (IconSCEPT)*, Karaikal, India, 2024.
- [5] A. J. Rozaqi and A. Sunyoto, "Identification of disease in potato leaves using convolutional neural network (CNN) algorithm," in *3rd Int. Conf. Information and Communications Technology (ICOIACT)*, 2020, pp. 72–76.
- [6] J. Andrew, J. Eunice, D. E. Popescu, M. K. Chowdary, and J. Hemanth, "Deep learning-based leaf disease detection in crops using images for agricultural applications," *Agronomy*, vol. 12, no. 10, Art. no. 2395, 2022.
- [7] X. Chen, G. Zhang, T. Lava, and X. Zhang, "Using a hybrid convolutional neural network with a transformer model for tomato leaf disease detection," *Agronomy*, vol. 14, no. 4, art. 673, 2024.
- [8] M. V. Sanida, T. Sanida, A. Sideris, and M. Dasygenis, "An efficient hybrid CNN classification model for tomato crop disease," *Technologies*, vol. 11, no. 1, art. 10, 2023.
- [9] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Front. Plant Sci.*, vol. 7, Art. no. 1419, 2016.
- [10] Ü. Atila, M. Uçar, K. Akyol, and E. Uçar, "Plant leaf disease classification using EfficientNet deep learning model," *Ecol. Informatics*, vol. 61, Art. no. 101182, 2021.
- [11] M. S. I. Prottasha and S. M. S. Reza, "A classification model based on depthwise separable convolutional neural network to identify rice plant diseases," *Int. J. Electr. Comput. Eng.*, vol. 12, no. 4, pp. 3642–3654, 2022.
- [12] L. Sritharan, M. Anjanan, and A. Gamage, "Plant diseases detection using image processing and suggest pesticides and managements," in *Proc. 7th Int. Conf. Convergence in Technology (I2CT)*, 2022, pp. 1–8.
- [13] A. S. Kini, P. K. V. Reddy, et al, "Techniques of Deep Learning and Image Processing in Plant Leaf Disease Detection," *International Journal of Electrical and Computer Engineering*, vol. 13(3), 3029-3040, 2023.
- [14] J. Chen, D. Zhang, and Y. A. Nanekaran, "Using deep transfer learning for image-based plant identification," *Computers and Electronics in Agriculture*, vol. 173, Art. no. 105393, 2020.
- [15] Kaggle, "Crop leaf disease datasets – Rice, Wheat, Chilli, Cotton, Maize," [Online]. Available: <https://www.kaggle.com/datasets>.