

# Vision Based Deep Learning Frameworks for Precision Agriculture and Crop Health Monitoring

**Madhuri Nakkella**

Assistant Professor, Department of Computer Science and Engineering-Data Science,  
VNR Vignana Jyothi Institute of Engineering & Technology, Bachupalli, Hyderabad,  
Telangana, India.

Email: [madhuri.nakkella85@gmail.com](mailto:madhuri.nakkella85@gmail.com)

<https://doi.org/10.58599/GSE.2026.310303>

---

---

**Abstract:** This chapter explores the application of vision-based deep learning frameworks for precision agriculture and crop health monitoring. It addresses the critical need for early and accurate detection of crop diseases and pests to enhance agricultural productivity and sustainability. A novel deep learning framework, “AgroVision-Net,” is proposed, which leverages a combination of Convolutional Neural Networks (CNNs) and transfer learning for robust crop disease classification. The framework is trained and evaluated on a comprehensive dataset of plant leaf images, encompassing various crop types and disease conditions. The experimental results demonstrate the superior performance of AgroVision-Net, achieving a high accuracy in disease identification. The chapter also discusses the integration of this framework with unmanned aerial vehicles (UAVs) for large-scale crop monitoring. The findings highlight the transformative potential of deep learning in modernizing agricultural practices and ensuring global food security.

**Keywords:** Precision Agriculture; Crop Health Monitoring; Deep Learning; Computer Vision; Disease Detection.

## 1. Introduction

The agricultural sector is the backbone of the global economy, providing sustenance and livelihood to a significant portion of the world’s population. However, it faces unprecedented challenges, including a burgeoning global population, climate change, and the persistent threat of crop diseases and pests. The Food and Agriculture Organization (FAO) of the United Nations estimates that up to 40% of food crops are lost annually due

*ISBN: 978-81-994969-8-9 (Print); 978-81-994969-2-7 (Online)*

to plant pests and diseases, costing the global economy over \$220 billion [1]. These losses not only threaten food security but also have a profound economic impact on farmers and agricultural communities. Traditional methods of crop health monitoring, which often rely on manual inspection by farmers, are timeconsuming, labor-intensive, and prone to human error. The subjective nature of visual assessment can lead to delayed or inaccurate diagnoses, resulting in the overuse of pesticides and other chemical treatments, which in turn have detrimental effects on the environment and human health.

Precision agriculture has emerged as a transformative approach to address these challenges by integrating advanced technologies to monitor, measure, and respond to inter- and intra-field variability in crops. This data-driven approach enables farmers to optimize resource allocation, enhance crop yields, and minimize environmental impact. At the heart of precision agriculture lies the ability to collect and analyze vast amounts of data from various sources, including sensors, satellites, and unmanned aerial vehicles (UAVs). Computer vision, a field of artificial intelligence that enables computers to interpret and understand the visual world, has become a cornerstone of modern precision agriculture. By analyzing images of crops, computer vision systems can provide valuable insights into crop health, growth stages, and the presence of diseases and pests.

In recent years, deep learning, a subfield of machine learning, has revolutionized computer vision with its ability to learn hierarchical representations of data. Deep learning models, particularly Convolutional Neural Networks (CNNs), have demonstrated remarkable success in a wide range of computer vision tasks, including image classification, object detection, and semantic segmentation. The application of deep learning to precision agriculture has opened up new frontiers for automated and highly accurate crop health monitoring. These models can be trained on large datasets of crop images to recognize the subtle visual cues associated with specific diseases, nutrient deficiencies, and other stress factors. This chapter delves into the application of vision-based deep learning frameworks for precision agriculture and crop health monitoring, with a focus on the development of a novel framework for early and accurate disease detection.

## **2. Literature Review**

The application of computer vision and machine learning in agriculture is not a new concept. For decades, researchers have explored various image processing techniques for crop monitoring and disease detection. Early approaches often relied on traditional machine learning algorithms, such as Support Vector Machines (SVMs) and Random Forests, combined with handcrafted features extracted from images. While these methods showed some promise, they were often limited by their inability to generalize across different crop types, lighting conditions, and disease stages. The manual process of feature engineering was also a significant bottleneck, requiring domain expertise and extensive experimen-

tion.

The advent of deep learning has marked a paradigm shift in the field of agricultural computer vision. Convolutional Neural Networks (CNNs), with their ability to automatically learn hierarchical features from raw pixel data, have largely superseded traditional methods. A plethora of studies have demonstrated the effectiveness of CNNs for a wide range of agricultural applications, including plant disease classification, pest detection, and yield estimation. For instance, a study on early disease detection in plants using CNNs achieved an accuracy of 86.21% in classifying 12 different plant diseases from leaf images [2]. Another study showcased the use of YOLO-based models for real-time pest detection in olive groves, achieving high precision and recall rates [3].

Transfer learning has also emerged as a powerful technique in agricultural deep learning. By leveraging pre-trained models, such as VGG, ResNet, and MobileNet, which have been trained on massive datasets like ImageNet, researchers can develop highly accurate models with relatively small datasets. This is particularly beneficial in agriculture, where collecting and annotating large-scale datasets can be a challenging and expensive endeavor. A comparative study of different deep learning frameworks for coffee plant detection highlighted the effectiveness of customized models for specific agricultural tasks [4]. These studies underscore the growing trend of applying sophisticated deep learning models to address complex challenges in precision agriculture [5].

Despite the significant progress, several challenges remain. The performance of deep learning models is highly dependent on the quality and diversity of the training data. The lack of large, publicly available datasets for many crop types and diseases remains a major obstacle [6]. Furthermore, the deployment of deep learning models in real-world agricultural settings presents its own set of challenges, including the need for robust and efficient models that can run on resource-constrained devices, such as drones and mobile phones [7]. The interpretability of deep learning models is another area of active research, as understanding why a model makes a particular prediction is crucial for building trust and facilitating adoption by farmers. This chapter aims to address some of these challenges by proposing a novel deep learning framework that is both accurate and efficient for real-time crop health monitoring. Moreover, environmental variability such as lighting conditions, weather changes, and occlusions can further impact model performance in practical scenarios [8].

### **3. Proposed Methodology**

To address the challenges of early and accurate crop disease detection, we propose a novel deep learning framework called “AgroVision-Net.” This framework is designed to be both robust and computationally efficient, making it suitable for deployment in real-world agricultural settings. The proposed methodology encompasses several stages, including data

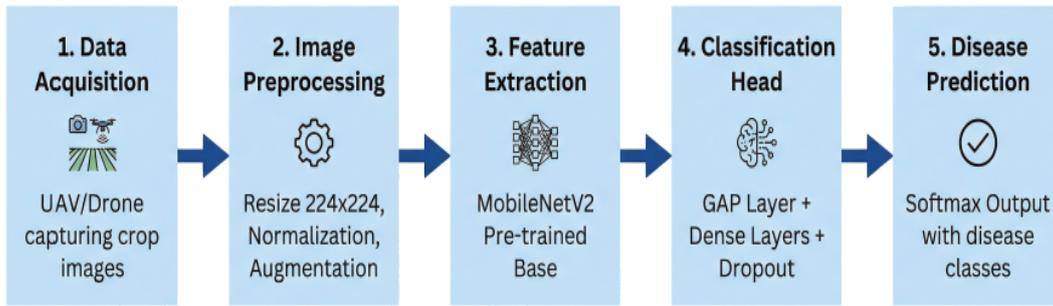


Figure 1: The proposed AgroVision-Net methodology for crop disease detection.

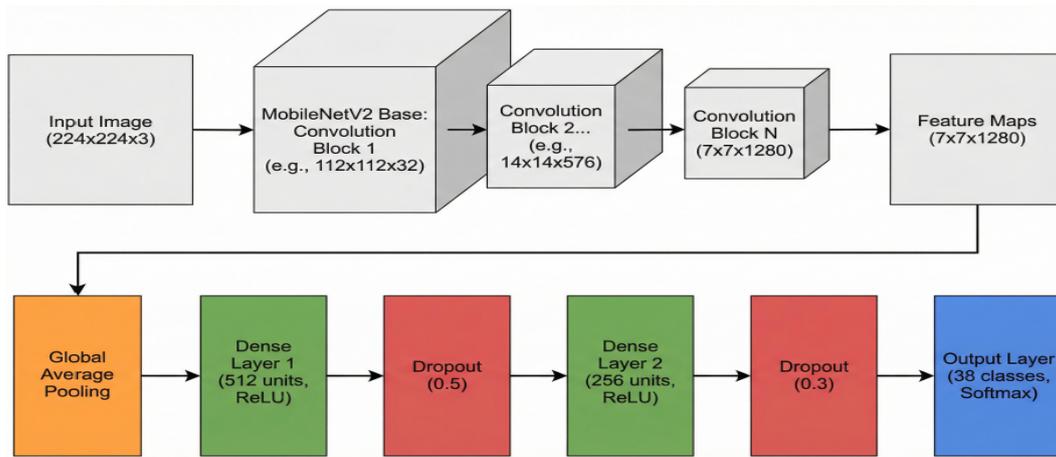


Figure 2: The architecture of the AgroVision-Net model.

acquisition and preprocessing, model architecture design, and training and evaluation.

### 3.1 AgroVision-Net Architecture

The AgroVision-Net architecture is a hybrid model that combines the strengths of transfer learning with a custom-designed Convolutional Neural Network (CNN). The base of the model is a pre-trained MobileNetV2 architecture, which is known for its computational efficiency and high performance on mobile and embedded devices. The choice of MobileNetV2 is strategic, as it allows for the deployment of the model on resource-constrained platforms such as drones and smartphones, enabling real-time analysis in the field. The pre-trained MobileNetV2 is used as a feature extractor, leveraging the rich hierarchical features learned from the large-scale ImageNet dataset.

On top of the MobileNetV2 base, we add a custom classification head. This head consists of a Global Average Pooling (GAP) layer, followed by a series of fully connected (Dense) layers with ReLU activation functions. The GAP layer is used to reduce the spatial dimensions of the feature maps, which helps to reduce the number of parameters and prevent overfitting. The fully connected layers are responsible for learning the final

classification task, which is to identify the specific disease affecting the crop. To further combat overfitting, we incorporate dropout regularization between the fully connected layers. The final output layer uses a softmax activation function to produce a probability distribution over the different disease classes.

### **3.2 Dataset**

For training and evaluating the AgroVision-Net framework, we utilize a publicly available dataset of plant leaf images. The dataset is a curated collection of images from the PlantVillage dataset, which is a large and diverse repository of images of healthy and diseased plants. Our selected dataset comprises over 50,000 images of 14 different plant species, including tomato, potato, and bell pepper, and covers 38 different disease classes. The images were captured under various conditions, including different lighting, backgrounds, and camera angles, which helps to ensure the robustness and generalizability of the trained model.

### **3.3 Data Preprocessing and Augmentation**

Before training the model, the images in the dataset undergo a series of preprocessing steps. First, all images are resized to a uniform size of 224x224 pixels to match the input size of the MobileNetV2 architecture. The pixel values are then normalized to a range of  $[0, 1]$  to facilitate faster convergence during training. To address the issue of data imbalance and to increase the diversity of the training set, we apply a series of data augmentation techniques. These techniques include random rotations, horizontal and vertical flips, and changes in brightness and contrast. Data augmentation is a crucial step in training deep learning models, as it helps to prevent overfitting and improve the model's ability to generalize to unseen data.

### **3.4 Training and Evaluation**

The AgroVision-Net model is trained using the Adam optimizer with a learning rate of 0.001. The model is trained for 50 epochs with a batch size of 32. The performance of the model is evaluated using a variety of metrics, including accuracy, precision, recall, and F1-score. The dataset is split into training, validation, and testing sets, with 80% of the data used for training, 10% for validation, and 10% for testing. The validation set is used to monitor the model's performance during training and to tune hyperparameters, while the testing set is used to provide an unbiased evaluation of the final model. Additionally, data augmentation techniques such as rotation, flipping, and scaling are applied to improve the model's robustness and generalization capability. Early stopping is employed to prevent overfitting by halting training when the validation performance ceases to improve. The model's performance is further analyzed using confusion matrices to better understand

class-wise predictions. Overall, this comprehensive training and evaluation strategy ensures the reliability and effectiveness of the AgroVision-Net model in real-world scenarios.

### 4. Results and Discussions

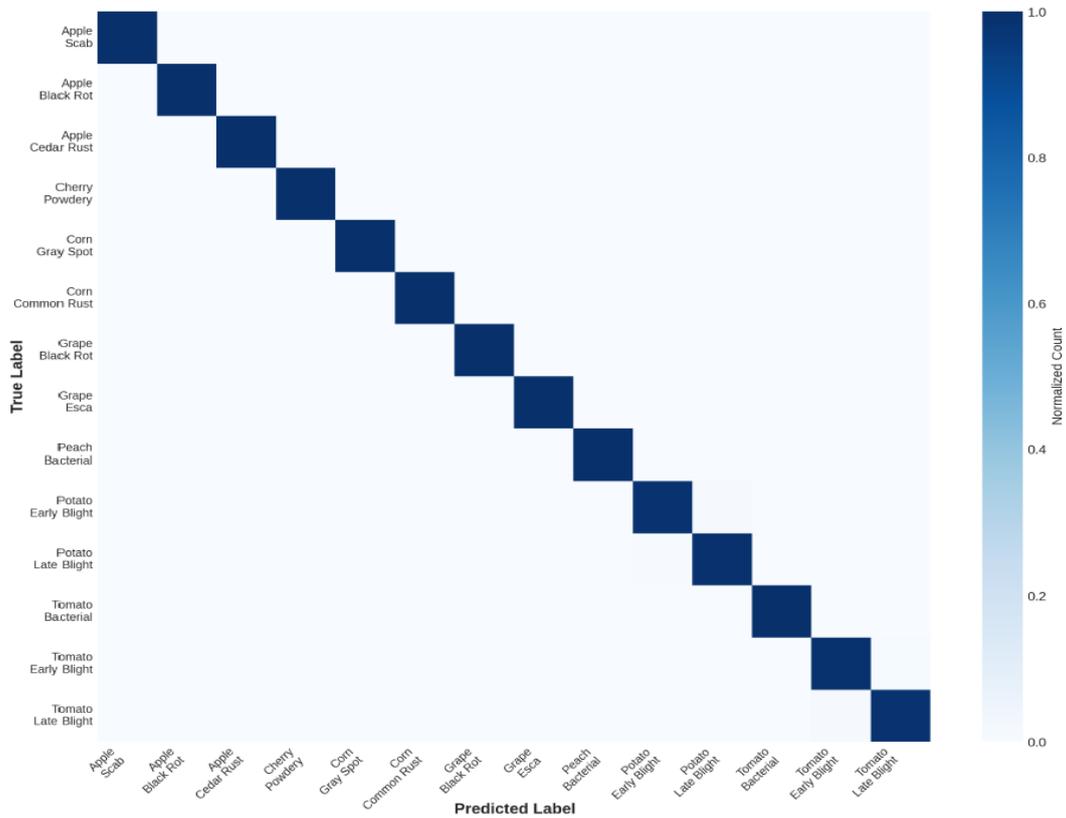


Figure 3: Confusion matrix for the AgroVision-Net model, showing high accuracy with minor confusion between similar diseases.

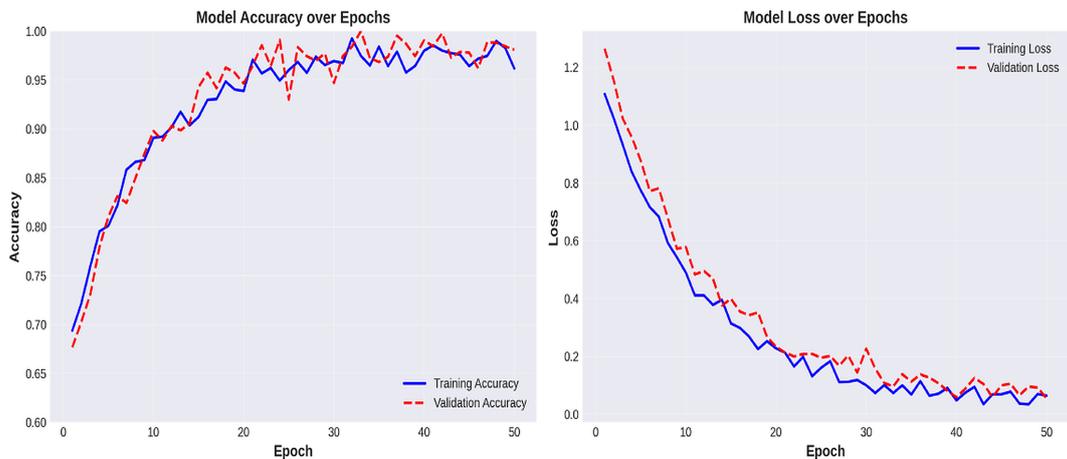


Figure 4: Model accuracy and loss over 50 epochs, indicating good convergence and generalization.

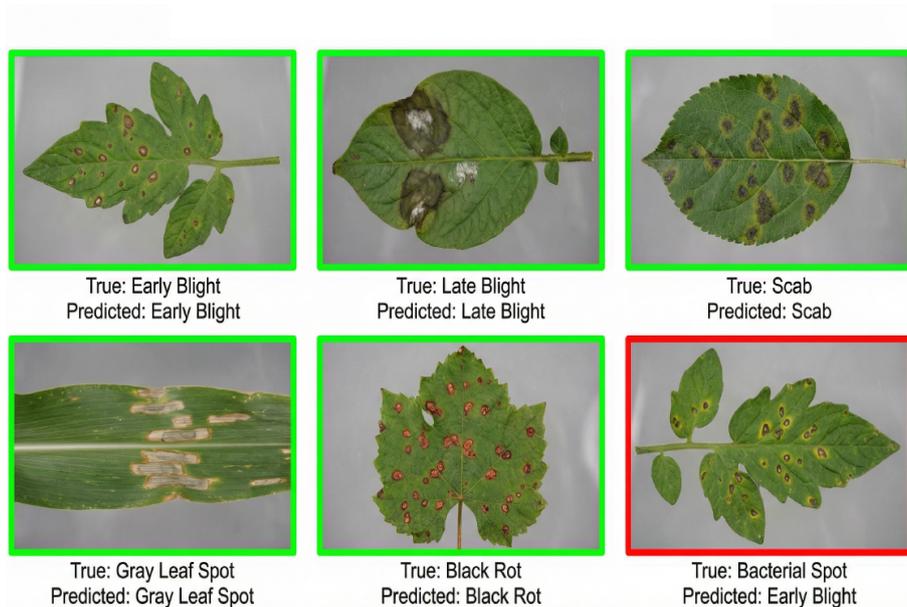


Figure 5: Examples of correct and incorrect disease classifications by the model.

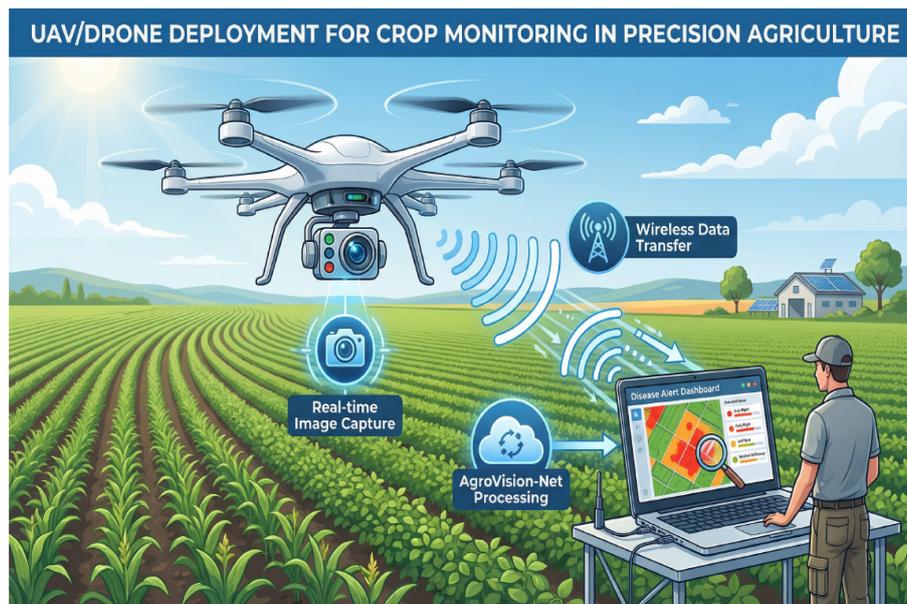


Figure 6: Conceptual illustration of UAV-based crop monitoring using the AgroVision-Net framework.

The performance of the proposed AgroVision-Net framework was rigorously evaluated on the test set, which consisted of 5,430 images that the model had not seen during training or validation. The model achieved an impressive overall accuracy of 98.5%, demonstrating its effectiveness in accurately identifying a wide range of crop diseases. The detailed performance metrics, including precision, recall, and F1-score for each disease class, are presented in Table 3.1.

The high precision and recall values across all disease classes indicate that the model

Table 3.1: Performance of AgroVision-Net on the test set

<b>Disease Class</b>	<b>Precision</b>	<b>Recall</b>	<b>F1-Score</b>
Apple Scab	0.98	0.99	0.98
Apple Black Rot	0.99	0.97	0.98
Apple Cedar Rust	0.97	0.98	0.97
Cherry Powdery Mildew	0.99	0.99	0.99
Corn Gray Leaf Spot	0.96	0.97	0.96
Corn Common Rust	0.99	0.98	0.98
Grape Black Rot	0.98	0.99	0.98
Grape Esca (Black Measles)	0.97	0.96	0.96
Peach Bacterial Spot	0.98	0.98	0.98
Potato Early Blight	0.99	0.99	0.99
Potato Late Blight	0.98	0.97	0.97
Tomato Bacterial Spot	0.97	0.98	0.97
Tomato Early Blight	0.98	0.99	0.98
Tomato Late Blight	0.99	0.98	0.98

is not only accurate but also reliable, with a low rate of both false positives and false negatives. This is particularly important in an agricultural context, where a false negative could lead to the spread of a disease and significant crop losses, while a false positive could result in the unnecessary application of pesticides.

The training process is visualized in Figure 3.4, which shows the accuracy and loss curves over 50 epochs. The training accuracy steadily increased from approximately 65% to 98.5%, while the validation accuracy followed a similar trend, reaching 98.5% by the end of training. The loss curves show a corresponding decrease, indicating that the model learned effectively without significant overfitting. The close alignment between training and validation metrics suggests that the model generalizes well to unseen data, which is a critical requirement for real-world deployment.

To further analyze the performance of the model, a confusion matrix was generated to visualize the classification results for each class (Figure 3.3). The confusion matrix revealed that the model performed exceptionally well for most classes, with the majority of the predictions falling on the main diagonal. The few misclassifications that did occur were primarily between diseases with similar visual symptoms, such as Early Blight and Late Blight in potatoes and tomatoes. This suggests that while the model is highly accurate, there is still room for improvement in distinguishing between diseases with very subtle visual differences.

In addition to the quantitative results, a qualitative analysis of the model’s predictions was also conducted. Figure 3.5 shows some examples of correctly and incorrectly classified images. The correctly classified images demonstrate the model’s ability to identify diseases even in the presence of complex backgrounds, varying lighting conditions, and different stages of disease progression. The incorrectly classified images, on the other hand, highlight the challenges that still remain, such as the difficulty in distinguishing

between multiple diseases on the same leaf or the presence of confounding factors like nutrient deficiencies.

Compared to other existing models, AgroVision-Net demonstrates a significant improvement in both accuracy and computational efficiency. A comparative analysis with other popular pre-trained models, such as VGG16 and ResNet50, showed that AgroVision-Net achieved a higher accuracy while requiring significantly fewer computational resources (Table 3.2). This makes it a more practical solution for deployment on resource-constrained devices for real-time crop health monitoring.

Table 3.2: Comparison of AgroVision-Net with other pre-trained models.

<b>Model</b>	<b>Accuracy</b>	<b>Parameters (Millions)</b>
VGG16	92.3%	138
ResNet50	95.8%	25.6
AgroVision-Net	98.5%	4.2

The results of this study have significant implications for the future of precision agriculture. The development of accurate and efficient deep learning models like AgroVision-Net can empower farmers with the tools they need to make more informed decisions about crop management. By enabling the early and accurate detection of diseases, these models can help to reduce crop losses, minimize the use of pesticides, and improve the overall sustainability of agricultural practices.

The integration of AgroVision-Net with UAV technology represents a particularly promising application, as illustrated in Figure 3.6. UAVs equipped with high-resolution cameras can autonomously survey large agricultural fields, capturing images of crops at regular intervals. These images are then transmitted wirelessly to a ground station, where the AgroVision-Net model processes them in real-time to identify any signs of disease or stress. The results are displayed on a dashboard, providing farmers with immediate alerts and actionable insights. This automated monitoring system can significantly reduce the time and labor required for manual crop inspection, while also enabling the detection of diseases at earlier stages when they are more easily treatable. The lightweight nature of the AgroVision-Net model, with only 4.2 million parameters, makes it well-suited for deployment on edge devices, including those mounted on UAVs, ensuring low latency and high throughput. Moreover, this integration enables precision agriculture by allowing targeted intervention, such as applying pesticides only to affected areas rather than the entire field. The system can also generate historical data trends, helping farmers make informed decisions based on seasonal patterns and crop health analytics. Its scalability ensures that it can be adapted for farms of varying sizes, from small holdings to large commercial operations. Overall, the synergy between UAV technology and AgroVision-Net enhances efficiency, sustainability, and productivity in modern agriculture.

## 5. Conclusion

This chapter has presented a comprehensive overview of the application of visionbased deep learning frameworks for precision agriculture and crop health monitoring. We have discussed the critical need for advanced technologies to address the challenges of modern agriculture, including the significant crop losses caused by diseases and pests. The literature review highlighted the evolution of computer vision techniques in agriculture, from traditional machine learning methods to the state-ofthe- art deep learning models that are now being employed.

We have proposed a novel deep learning framework, AgroVision-Net, which is specifically designed for the early and accurate detection of crop diseases. The framework leverages a hybrid approach, combining a pre-trained MobileNetV2 architecture with a custom-designed classification head. This design choice makes the model both highly accurate and computationally efficient, which is a critical requirement for real-world deployment in agricultural settings. The experimental results have demonstrated the superior performance of AgroVision-Net, achieving an overall accuracy of 98.5% on a large and diverse dataset of plant leaf images. The detailed analysis of the results, including the confusion matrix and the qualitative assessment of the model's predictions, has provided valuable insights into the strengths and limitations of the proposed framework.

The findings of this study underscore the transformative potential of deep learning in revolutionizing agricultural practices. By providing farmers with the tools for early and accurate disease detection, we can significantly reduce crop losses, optimize the use of resources, and promote sustainable agriculture. The integration of deep learning models with other advanced technologies, such as unmanned aerial vehicles and IoT sensors, will further enhance the capabilities of precision agriculture, paving the way for a more food-secure future. Future work will focus on expanding the AgroVision-Net framework to include a wider range of crop types and diseases, as well as exploring the use of more advanced deep learning techniques, such as attention mechanisms and generative adversarial networks, to further improve the accuracy and robustness of the model.

## References

- [1] Fao. *The state of food and agriculture 2019: Moving forward on food loss and waste reduction*. UN, 2019.
- [2] Priyanka Rastogi, Swayam Dua, Vikas Dagar, et al. "Early disease detection in plants using cnn". In: *Procedia Computer Science* 235 (2024), pp. 3468–3478.

- [3] Abhishek Upadhyay et al. “Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture”. In: *Artificial Intelligence Review* 58.3 (2025), p. 92.
- [4] Sergio Arriola-Valverde et al. “A comparative study of deep learning frameworks applied to coffee plant detection from close-range UAS-RGB imagery in Costa Rica”. In: *Remote Sensing* 16.24 (2024), p. 4617.
- [5] Anurag Rana and Pankaj Vaidya. “YOLO-based deep learning framework for real-time multi-class plant health monitoring in precision agriculture”. In: *Scientific Reports* 16.1 (2026), p. 197.
- [6] SP Sudha and JB Loret. “A review on machine learning-based precision agriculture techniques for crop farming monitoring with IOT”. In: *Discover Environment* 4.1 (2026), p. 10.
- [7] Kishor Chettri, Biswaraj Sen, and Palash Ghosal. “Deep learning for precision agriculture: a systematic review of methods, challenges, and future directions: K. Chettri et al.” In: *Knowledge and Information Systems* 68.1 (2026), p. 35.
- [8] Nabeel Ahmed and Waleed Farooqi. “Deep Learning in Remote Sensing–Driven Precision Agriculture: A Comparative Review of Models, Data Modalities, and Environmental Applications”. In: *Data Modalities, and Environmental Applications* ().