Plant Direct

WILEY

American Society of Plant Biologists
SEB SOCIETY FOR EXPERIMENTAL BIOLOGY

RESEARCH ARTICLE OPEN ACCESS

InsightNet: A Deep Learning Framework for Enhanced Plant Disease Detection and Explainable Insights

Mubasshar U. I. Tamim | Sultanul A. Hamim | Sumaiya Malik | M. F. Mridha | Sharfuddin Mahmood

Department of Computer Science and Engineering, American International University–Bangladesh, Dhaka, Bangladesh

Correspondence: M. F. Mridha (firoz.mridha@aiub.edu)

Received: 20 November 2024 | **Revised:** 12 April 2025 | **Accepted:** 17 April 2025

Funding: The authors thank AIUB authority for their financial support.

Keywords: deep learning | InsightNet | leaf disease classification | XAI

ABSTRACT

Sustainable agriculture holds the key in meeting food production requirements for a rapidly growing population without exacerbating environmental degradation. Plant leaf diseases pose a critical threat to crop yield and quality. Existing inspection methods are labor-intensive and prone to human errors, while lacking support for large-scale agriculture. This research aims to enhance plant health by developing advanced deep learning models for the detection and classification of plant diseases across a variety of species. A deep learning model based on the paradigm of the MobileNet architecture is proposed, which employs a dedicated design through deeper convolutional layers, dropout regularization, and fully connected layers. This results in significant improvements in disease classification in tomato, bean, and chili plants, with accuracy rates of 97.90%, 98.12%, and 97.95%, respectively. Moreover, Grad-CAM is used to shed light on the decision-making process of the proposed model. The work contributes to the advancement of precision farming and sustainable agricultural practices, supporting timely and accurate plant disease diagnosis.

1 | Introduction

Sustainable agriculture is key in addressing the challenges posed by the continuous growth of the global population and the requirement to increase food production by 70% by 2050. Given the urgent need within this limited timeframe, the adoption of Agriculture 4.0 technologies, e.g., smart sensors, robotics, and precision agriculture, is key in helping to meet the target of increased agricultural output without a substantial impact in the environmental footprint. However, despite the need for increased food production, the agricultural share of the global GDP dropped to 3%, potentially indicating both the lack of sufficient innovation within the sector and the adoption of current smart technologies. There is a mismatch between the supply of Agriculture 4.0 technologies and the requirements of the end-users, which manifests itself in terms of limited scalability, increased costs, and lack of accessibility, which are very important for smallholder farmers in developing countries.

Plant leaf diseases, caused by pathogens, environmental stress, and outdated agricultural practices, are a major threat to crop yield and quality. Existing techniques for leaf disease detection are based on human inspection, which makes them slow, labor-intensive, and prone to human errors. Moreover, given that large-scale agriculture is needed to meet growing demand, the use of manual methods is deemed ineffective and incompatible for integration within digital agricultural farms. Automated plant inspection can be enabled via the IoT-based smart sensors (Tirkey et al. 2023; Ramanjot et al. 2023), e.g., harnessing the power of computer vision and deep learning (DL) algorithms.

Abbreviations: ABC, a black cat; DEF, does not ever fret; GHI, goes home immediately.
All authors contributed equally.

Moreover, these intelligent machines can offer accurate, rapid, and scalable deployment, supporting real-time crop monitoring for early disease detection. Such a proactive strategy can provide further benefits in reducing crop losses, optimizing land use and use of pesticides, leading to more sustainable farming practices (Krishnan et al. 2022).

Detection and classification of plant leaf diseases involve the use of technologies, which analyze digital images to accurately identify relevant information; i.e., features, and subsequently, classify foliage diseases affecting overall plant health (Saleem et al. 2019). Early identification and intervention are key to maintaining healthy crops and increasing yields. With the help of artificial intelligence (AI), agronomists and farmers can easily detect and manage disease outbreaks in vast agricultural areas (Ferentinos 2018).

Detection and classification of plant diseases pose significant challenges in agriculture technology, due to the complex and diverse nature of pathogens, environmental conditions, and the crop species involved. Rapid and reliable detection techniques are required to avoid severe losses in production quality and quantity (Saleem et al. 2019). The increased adoption of deep learning techniques has ushered in a new era of deep learning-driven automation in plant disease detection, enabling proactive, technology-based interventions in controlling infections. The precision of these interventions has also seen a boost, thanks to the use of deep learning models in agriculture, paving the way for precision farming that harnesses technology to combat plant diseases and enhance agricultural health and sustainability (Ferentinos 2018; Mohanty et al. 2016). Different plant species, such as beans, chili, and tomatoes, exhibit unique symptoms and expressions of diseases, making it challenging to develop diagnostic models that can be used across the board. Previous methods exhibit adequate performance; however, more scalability is required in practical applications, supported by higher accuracy. Researchers utilized DL architectures such as ResNet50, ResNet101, Xception, InceptionV3, MobileNetV2, NASNetLarge, and EfficientNet series to address these issues (Saleem et al. 2019; Mohanty et al. 2016; Too et al. 2019). These models are fundamental in developing plant disease detection using pre-trained networks on large datasets and employing transfer learning to fine-tune models for this specific task.

Although base models are efficient, they may not perform well when used alone to detect plant diseases in different species. This can limit the model's generalization because different species may have symptomatic disease differences. Thus, specialized solutions are required for maximum accuracy and performance. This highlights the importance of using novel approaches that can cross the conventional boundaries of existing models (Too et al. 2019; Saleem et al. 2020; Hamim and Jony 2024). A low-cost real-time inspection system, with modest computational resource requirements, for identifying leaf diseases would be highly beneficial for farmers, enabling them to quickly treat or isolate affected plants. Following an analysis of DL architecture, it was determined that the MobileNet model is both efficient and resource-light, supporting its use in real-time scenarios and deployment on mobile or edge devices. Thus, in this research, we adopt this base architecture and design enhancements to increase its accuracy, while preserving its useful attributes. The proposed InsightNet model uses deep convolutional neural networks and innovative training techniques, including extensive data augmentation and fine-tuning protocols, to effectively learn from diverse plant disease images.

The aim of this research is to introduce enhanced DL models integrated with XAI called InsightNet for the solution of classification problems in plant leaf diseases and also the detection with the help of XAI, specifically, disease detection, i.e., whether a disease is present or not in various plant species, e.g., tomato, bean, and chili. Following model development and training, its performance is compared with state-of-the-art methods in terms of accuracy, efficiency, and adaptability in cross-species disease detection and classification. By enhancing the architecture of MobileNet in plant disease detection and classification, we aim to establish new benchmarks in the field, providing a robust solution that supports precision farming and ensures agricultural sustainability. The contributions of this research are as follows:

- Utilization of a comprehensive cross-species dataset: We employ a diverse image dataset, encompassing multiple plant species, including tomatoes, beans, and chili plants. The dataset is used to train the InsightNet model to classify a variety of plant diseases, thereby addressing the challenge of cross-species disease detection.
- Enhanced deep learning architecture: The system architecture of MobileNet is tailored to the problem of plant disease classification. The specific enhancements include depth increases in later convolutional layers, reaching 1024 channels from Conv5 onward, and the introduction of two fully connected layers (fc1 and fc2 with 1024 neurons each), along with 0.5 dropout regularization to prevent overfitting. During testing, an accuracy of around 98% is obtained for the three plant species of tomato, bean, and chili, which represents a substantial improvement compared to state-of-the-art.
- Comparison with existing models: The performance of the proposed InsightNet is evaluated against traditional DL models such as ResNet50, Xception, etc. The comparative analysis highlights improvements in accuracy and efficiency for the proposed model.
- Explainable AI (XAI) techniques: Grad-CAM is utilized to support the interpretability of the proposed solution. This provides insights into the model's decision-making, making it possible to understand which features of the plant images are most indicative of specific diseases.

The rest of the paper is organized as follows: Section 2 provides an overview of the relevant literature, emphasizing the advancements and challenges in plant disease detection. Section 3 details the materials and methods for developing the Advanced MobileNet architecture, including the dataset preparation, model training, and validation processes. Section 4 presents the outcomes and discussions regarding the model's effectiveness and aspects of explainable AI applied in this study. Finally, Section 5 provides an overview of the research, summarizing the findings and suggesting future research directions in agricultural technology and plant disease management.

2 | Literature Review

The literature on plant disease detection has significantly evolved with the advent of deep learning technologies, particularly CNNs. These advanced models have demonstrated remarkable efficacy in accurately identifying healthy and diseased plant leaves, enabling timely interventions to mitigate crop losses. This review synthesizes recent studies focusing on various deep learning architectures, including CNNs, hybrid models, and transfer learning approaches, highlighting their effectiveness in enhancing diagnostic accuracy and operational efficiency. By examining these advancements, the review aims to provide insights into the current state of research in plant disease detection and its implications for improving agricultural practices.

2.1 | Convolutional Neural Networks (CNNs)

Recent research demonstrates that CNN architectures are essential in advancing plant disease detection. The use of CNNs has rapidly grown in agricultural technology, particularly in detecting plant diseases. Advanced architectures, such as MobileNetV3 and InceptionV3, have proven to be remarkably versatile and computationally efficient, offering significant improvements in creating models that excel in accuracy and deployment feasibility across various platforms, including mobile and edge computing devices. Several studies (Kumar et al. 2023; Singh et al. 2023; Shah et al. 2023; Mahesh et al. 2023; Shaheed et al. 2023) have highlighted the significance of these developments, emphasizing the potential of these models to seamlessly integrate into resource-constrained environments without sacrificing their disease detection capabilities. MobileNetV3 and its variants stand out for their ability to handle nuanced features in leaf images through the use of dilated convolution and the efficient channel attention (ECA) module.

Researchers have made significant strides in developing plant disease detection systems using deep learning, aiming to effectively differentiate between healthy and diseased plant leaves. Building on this foundation, another investigation designed a specialized leaf blight recognition model that not only reached an accuracy of 96.3% (Praveen et al. 2022). The ResNet18 model demonstrated its capability by identifying leaf types and their associated diseases with an accuracy of 96% (Abd Algani et al. 2023). These achievements illustrate the growing potential of deep learning technologies in enhancing agricultural practices and disease management.

2.2 | Hybrid Models

In the realm of agriculture, the timely prediction of crop yields is crucial for farmers, particularly in developing countries where traditional methods are often slow and prone to error (Shastry and Sanjay 2021). To revolutionize this process, researchers have introduced a hybrid prediction strategy combining weighted principal component analysis (w-PCA) and a sophisticated artificial neural network (ANN) enhanced by modified-particle swarm optimization (m-PSO). This innovative approach, dubbed hybrid-ANN (H-ANN), has demonstrated remarkable improvements in accuracy, outperforming conventional models significantly. Meanwhile, the threat of plant diseases looms large, prompting the need for rapid identification solutions (Kannadasan et al. 2021). A novel hybrid model utilizing k-means clustering and a convolutional neural network (CNN) has emerged, achieving a commendable 92.6% accuracy in classifying leaf diseases. Specifically for olive trees, combining CNNs with machine learning classifiers resulted in deep hybrid models that reached an impressive accuracy of 96.14% (El Akhal et al. 2023). These advancements herald a new era in agricultural productivity and disease management.

2.3 | Transfer Learning Approaches

Transfer learning has emerged as a pivotal strategy for enhancing model performance in plant disease detection, particularly in the classification of various leaf diseases. One innovative study harnessed feature concatenation alongside transfer learning techniques using pre-trained MobileNetV2 and NASNetMobile kernels, resulting in an impressive accuracy of 95% for tomato leaf disease classification (Al-gaashani et al. 2022). Building on this approach, EfficientLeafNetB4, a lightweight model based on MobileNet, showcased its capabilities by achieving a robust AUC score of 0.912 in classifying diseases in chili leaves (Pratap and Kumar 2023). Furthermore, another investigation employed transfer learning with the lightweight NASNetMobile architecture, achieving accuracy rates of 93.20% for RGB photos and 83.60% for grayscale images in the context of tomato leaf diseases (Ghodekar and Kumar 2023). These advancements underline the significant role of transfer learning in improving the accuracy and efficiency of disease detection systems, paving the way for more effective agricultural practices.

2.4 | Ensemble Models

Research comparing different machine learning systems, such as support vector machines and random forests, with deep learning architectures like Inception v3, VGG16, and VGG19, demonstrated that deep learning methods outperformed conventional machine learning techniques in disease classification (Sujatha et al. 2021). This highlights the potential of ensemble models to improve diagnostic accuracy in agricultural applications.

The literature on plant disease detection illustrates a significant advancement in the application of deep learning technologies, particularly CNNs, which have become essential tools for accurately classifying healthy and diseased plant leaves. Various studies highlight the successful implementation of CNN architectures, achieving high accuracy rates and demonstrating their capability to handle large and complex datasets. Transfer learning has emerged as a crucial strategy, allowing researchers to leverage pre-trained models for specific agricultural applications, which not only improves diagnostic performance but also reduces the need for extensive labeled data. Hybrid models combining different architectures have shown promising results in enhancing classification capabilities, while ensemble approaches further optimize performance by integrating multiple models. Collectively, these advancements underscore the potential of deep learning in revolutionizing plant disease management, paving the way for more efficient and sustainable agricultural practices. The ongoing research in this field points towards the necessity of continued innovation and integration of these technologies to address the challenges posed by plant diseases in a rapidly changing agricultural landscape.

3 | Model and Architecture

This section offers an in-depth exploration of the system's underlying model and architectural framework. It meticulously examines its components, design principles, and structural organization, providing a holistic view of how diverse elements synergize to shape the complete system. Figure 1 shows a complete overview of this research method.

3.1 | Data Collection

The data for this study were collected from three separate datasets, each focusing on a specific type of crop disease. The datasets were meticulously assembled to encompass various disease manifestations, providing a comprehensive basis for training our deep learning models.

3.2 | Dataset Descriptions

The diseases affecting bean, chili, and tomato plants exhibit various symptoms and characteristics. Bean plants can suffer from Angular Leaf Spot, which manifests as small, angular water-soaked lesions on leaves turning brown with a yellow halo, and Bean Rust, identified by rust-colored spores on leaf undersides leading to defoliation. Chili plants are susceptible to leaf curl, a viral infection causing distorted, curled leaves with yellowing, leaf spot resulting in small, dark lesions merging to premature defoliation, and whitefly infestations, causing yellowing, wilting, and sooty mold growth. Tomatoes can contract bacterial spot, forming dark, water-soaked lesions with a yellow halo on leaves, stems, and fruit, and septoria leaf spot, forming small, dark spots that enlarge, merge, and reduce yield.

3.2.1 | Dataset

The study utilizes three distinct datasets to evaluate plant disease detection. The first dataset (D1) consists of images capturing chili plant leaves under various disease conditions, as detailed in Table 1. The second dataset (D2) is a relatively large collection comprising 4000 images representing all common diseases affecting tomato leaves, as shown in Table 2. Finally, the third dataset (D3) depicts a variety of diseases affecting the leaves of bean plants, with specific details provided in Table 3. Together, these datasets provide a comprehensive foundation for assessing the performance of plant disease detection models across different crop species.

These datasets were used to ensure diversity in the conditions presented, which is essential for training robust and generalizable deep-learning models for plant disease detection and classification. Images were gathered using standard procedures for quality and resolution to maintain consistency and provide a reliable basis for the model development and validation stages. Sample images of Datasets D1, D2, and D3 are shown in Figure 2.

3.3 | Data Preprocessing

The beans, chili, and tomato datasets underwent meticulous preprocessing to prepare them for practical deep-learning analysis. Each dataset was transformed to ensure optimal model input quality, starting with the resizing of all images to 224 × 224 pixels. This resizing step is vital for maintaining consistency across the datasets and facilitating the learning process of deep learning models. The target dimension was determined

TABLE 1 | Dataset D1 image count for chili.

Class	Total images
Healthy	115
Leaf curl	40
Leaf spot	82
Whitefly	42
Total	279

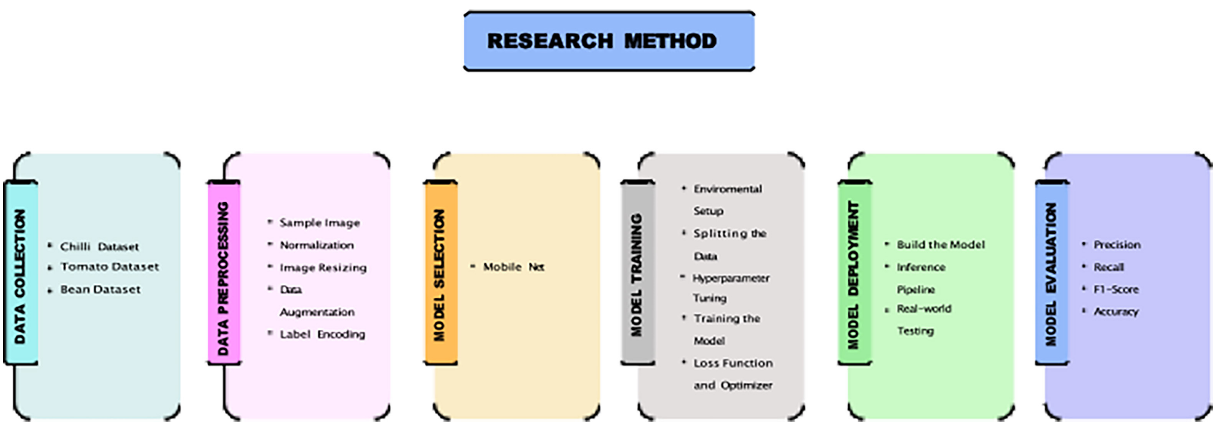

FIGURE 1 | Research methodology.

based on the input requirements of the base deep learning models, balancing the preservation of relevant features with computational efficiency. To artificially expand the datasets, each image underwent a series of augmentations, including random rotations within a 40° range, brightness adjustments by a factor of 0.2, and contrast modifications by a factor of 0.2. Two augmentations were generated and stored in each image's designated augmented dataset folder, ensuring that the aspect ratio was maintained without losing critical information.

The data augmentation techniques used were chosen specifically to optimize the diversity of the training set and prevent overfitting, a common issue in deep learning models, especially when training on small datasets. Augmentation is critical in synthetically expanding the dataset by generating new samples from the original images with various transformations.

1. Rotation and Flipping: Random rotations were applied to all images in a 40° range to add rotational invariance. This avoids overfitting to a single orientation of the plant features. Additionally, random horizontal and vertical flips were employed to mimic different viewpoints, providing additional robustness to the model.
2. Brightness and Contrast Adjustments: Adjustments of brightness (by a factor of 0.2) and contrast (by a factor of 0.2) were performed to simulate changes in lighting conditions and environmental conditions. These conversions allow the model to generalize better under varying real-world conditions, as field-setting images usually contain exposure and illumination changes.
3. Aspect Ratio Preservation: Aspect ratio of the images remained unchanged during augmentation, which is necessary to maintain the integrity of the plant features. By resizing the image dimensions without distorting their proportions, the model is provided with more realistic and varied inputs and is capable of learning the identification characteristics of the plant diseases better.
4. Randomization and Data Splitting: Once the augmentations were applied, the data were randomized to eliminate any sequence bias that could exist due to the original ordering of the images. The data were then divided into training, validation, and test sets, 70% for training and 30% for validation. Randomization is necessary to avoid the model learning any unwanted patterns due to the order in which the images were being presented.
5. One-Hot Encoding: The categorical labels were represented in one-hot encoded form, which is a binary matrix representation that is suitable for the categorical cross-entropy loss function used during training. This encoding scheme enables the model to predict multiple classes and is optimized for classification tasks

TABLE 2 | Dataset D2 image count for tomato.

Class	Total images
Bacterial spot	1000
Healthy	1000
Septoria leaf spot	1000
Septoria leaf spot	1000
Total	4000

TABLE 3 | Dataset D3 image count for bean.

Class	Total images
Angular leaf spot	345
Bean rust	348
Healthy	342
Total	1035

To further enhance the model's training process, the dataset was randomized to prevent any order bias that might impact learning. After shuffling, the dataset was split into training, validation, and test sets, allocating 70% for training and 30% for validation. Class labels for the datasets were encoded as integers and then transformed into one-hot encoded format, which is necessary for the categorical cross-entropy loss function used during model training. This encoding translates the categorical labels into a binary matrix representation suitable for the

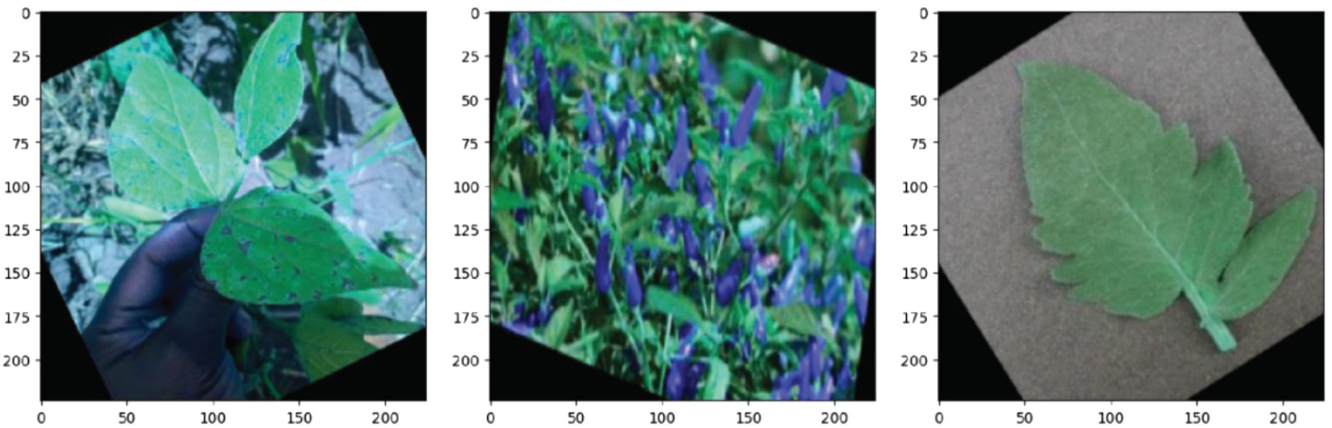

FIGURE 2 | Sample images of leaves from datasets D1, D2, and D3, demonstrating various stages of processing and analysis.

prediction task. Finally, the dimensions and distributions of the training, validation, and test sets were confirmed to ensure the integrity of the split, guaranteeing that each subset accurately represents the overall dataset and is crucial for an unbiased evaluation of the deep learning model's performance.

3.4 | Environment Setup

Our research was grounded in the unique and robust capabilities of Python, which we harnessed to craft the detection algorithms central to this study. By harnessing the synergistic power of the TensorFlow and Keras libraries, we were able to train a diverse range of pre-trained deep-learning models. We chose Google Colab as the development environment due to its seamless integration with advanced computational resources, including Tesla GPUs, and its easy access to machine learning libraries.

The experimental configuration was meticulously designed to meet the diverse requirements of the selected models. Specifically, we paired the MobileNet architecture with the NVIDIA A100 GPU, a combination that provided the computational acuity necessary for intensive deep-learning tasks. Similarly, we entrusted other models, such as DenseNet121 and XceptionV3, to the NVIDIA T4 GPU, known for its efficiency and adaptability.

The models were fed with input data standardized to specific dimensions, depending on the model's intrinsic architecture. The standardization process ensured that each model received optimally sized inputs, a crucial factor in determining the performance of CNN-based models. The trade-off between computational resource utilization and model performance was addressed by calibrating batch sizes, which varied between 32 and 64. The Adam optimizer was the driving force behind the training processes, with its learning rate finessed to ensuring a deliberate progression toward the global minimum. To cultivate robustness in the models, a training regimen of 15 epochs was set for MobileNet, while other architectures were limited to 10 epochs. The activation landscape within our networks was sculpted using the ReLU function, which provided the necessary computational efficiency and gradient propagation.

Data augmentation was applied to models trained with the A100 GPU, to enhance the ability to generalize across varied agricultural scenarios. Conversely, this process was eschewed for models trained with the T4 GPU to delineate the impact of data augmentation on model efficacy. A methodical approach was taken to avoid overfitting, which relied on a carefully crafted early stopping protocol that relied on the "val_loss" metric. This protocol halted the training process if there was no progress for three epochs. The "restor_best_weights" function also worked with this protocol to ensure the model's performance remained consistent throughout all iterations.

To ensure transparency and clarity in the machine learning model, gradient-weighted class activation mapping (GradCAM) was used to visualize the reasoning behind its predictive decisions. Moreover, Table 4 provides a comprehensive overview of the parameters used in computational experiments.

3.5 | Model Architecture

In this study, we present the proposed InsightNet model, which is based on the base model of the MobileNet architecture for the identification of leaf diseases, aiming to balance computational efficiency and high classification accuracy. The model retains the core structure of the MobileNet, with appropriate modifications, e.g., the addition of fully connected layers to enhance its discriminative capabilities.

TABLE 4 | Experimental setup summary.

Model	Input size (Px)	Batch size	Learning rate	Epochs	Activation function
InsightNet model for D1	224	32	0.0001	15	ReLU
InsightNet model for D2	224	64	0.0001	15	ReLU
InsightNet model for D3	224	32	0.0001	15	ReLU
MobileNetV2	224	32	0.01	10	ReLU
EfficientNetB2	260	32	0.01	10	ReLU
DenseNet121	224	32	0.01	10	ReLU
XceptionV3	299	32	0.01	10	ReLU
ResNet	224	32	0.01	10	ReLU
DenseNet169	224	32	0.01	10	ReLU
DenseNet201	224	32	0.01	10	ReLU
VGG16	224	32	0.01	10	ReLU
VGG19	224	32	0.01	10	ReLU
MobileNet	224	32	0.01	10	ReLU

3.5.1 | Architecture Overview

The input to the model is an image of size $224 \times 224 \times 3$ representing the height, width, and color channels of the input image. The initial layer, Conv1, applies a 3×3 convolution with a stride of 2, resulting in an output of $112 \times 112 \times 32$. This layer is followed by a series of Depthwise separable convolutional blocks, a hallmark of MobileNet, which efficiently reduce the number of parameters and computational cost while maintaining performance. Each block consists of a Depthwise convolution followed by a pointwise convolution, interspersed with ReLU activation functions. These blocks progressively reduce the spatial dimensions of the feature maps and increase the depth, culminating in an output of $7 \times 7 \times 1024$ after the final convolutional block.

Following the convolutional base, a global average pooling layer condenses the spatial dimensions, yielding a $1 \times 1 \times 1024$ output, which serves as the input to the fully connected layers. The first fully connected layer (fc1) contains 1024 neurons, activated by a ReLU function, and is followed by a dropout layer with a rate of 0.5 to mitigate overfitting. This pattern is repeated with a second fully connected layer (fc2) and dropout layer, before finally reaching the output layer (fc3). The output layer consists of 3 neurons, corresponding to the number of leaf disease classes, and employs a softmax activation function to produce class probabilities. The detailed structure of the proposed InsightNet model, including its constituent layers, output shapes, parameter counts, and layer connections, is summarized in Table 5. This provides a clear understanding of the architectural complexity and configuration of InsightNet comparison of their design choices. This comparison underscores the specific modifications in InsightNet aimed at enhancing task-specific performance while maintaining computational efficiency.

3.5.2 | Working Mechanism

The proposed InsightNet model has several working processes, all starting with the preprocessing of images to the desired $224 \times 224 \times 3$ dimensions. During a forward pass, an image fed as input passes through convolutional layers where feature extraction takes place. In each separable convolutional block, increasingly abstract and complex features are extracted, reducing the spatial dimensions while expanding the depth of feature maps. After the last convolutional block, the global average pooling layer pools the feature maps into one single vector of size $1 \times 1 \times 1024$, thus summarizing the presence of the features detected in the previous layers. This goes through a fully connected layer that does high-level reasoning and classification. Regularization is provided by dropout layers placed in between fully connected layers, which stop the model from fitting too strongly to the training data by randomly deactivating a fraction of neurons during every training iteration. The last output layer applies the softmax activation function to create probabilities for each class, specifying the likelihood of the input image being of each disease category. In this case, the class with the maximum value of the probability function would be selected for the predicted label.

Therefore, this InsightNet model of leaf disease identification integrates the strengths of the lightweight architecture by MobileNet with an improvement in discriminative power added by extra fully connected layers. This hybrid approach will ensure computational efficiency with high accuracy in classification—this is going to be an effective tool for real-world applications in agriculture. Figure 3 shows the model architecture of the InsightNet model to elucidate the distinctions between the MobileNet base model and the InsightNet architecture; Table 6 presents a detailed model architecture of the proposed InsightNet model.

Fine-tuning is a necessary step to improve the model's performance, especially if one is leveraging pre-trained models as the base. The base model, MobileNet, provides lightweight and efficient architecture that is a very good feature extractor for plant disease classification. The following outlines the specific fine-tuning strategies employed in the InsightNet model to improve its discriminative ability:

1. Preprocessing and Feature Extraction: All input images were resized to the size of $224 \times 224 \times 3$ pixels to match the input requirements of MobileNet. The image passes through multiple layers of convolutional blocks that progressively extract increasingly higher level features. At each convolutional phase, filters identify various patterns such as edges, textures, and shapes, which are necessary for recognizing different types of plant diseases. Depthwise separable convolutions were employed in order to reduce the number of parameters and calculations, making the model feasible to be deployed in resource-constrained environments.
2. Global Average Pooling (GAP): After the convolutional layers, the global average pooling layer reduces the feature maps to a single vector. This is accomplished by computing the average of all points in the spatial feature maps, summarizing the most important features in a condensed manner. The resulting output vector of size $1 \times 1 \times 1024$ is passed through fully connected layers for classification.

TABLE 5 | Model architecture summary.

Layer (type)	Output shape	Param #	Connected to
Mobilenet_0.50_224 (functional)	(None, 7, 7, 512)	829,536	Input layer
Flatten (flatten)	(None, 25,088)	0	Mobilenet_0.50_224
Dense (dense)	(None, 256)	6,422,784	Flatten
Dropout (dropout)	(None, 256)	0	Dense
Dense_1 (dense)	(None, 4)	1028	Dropout

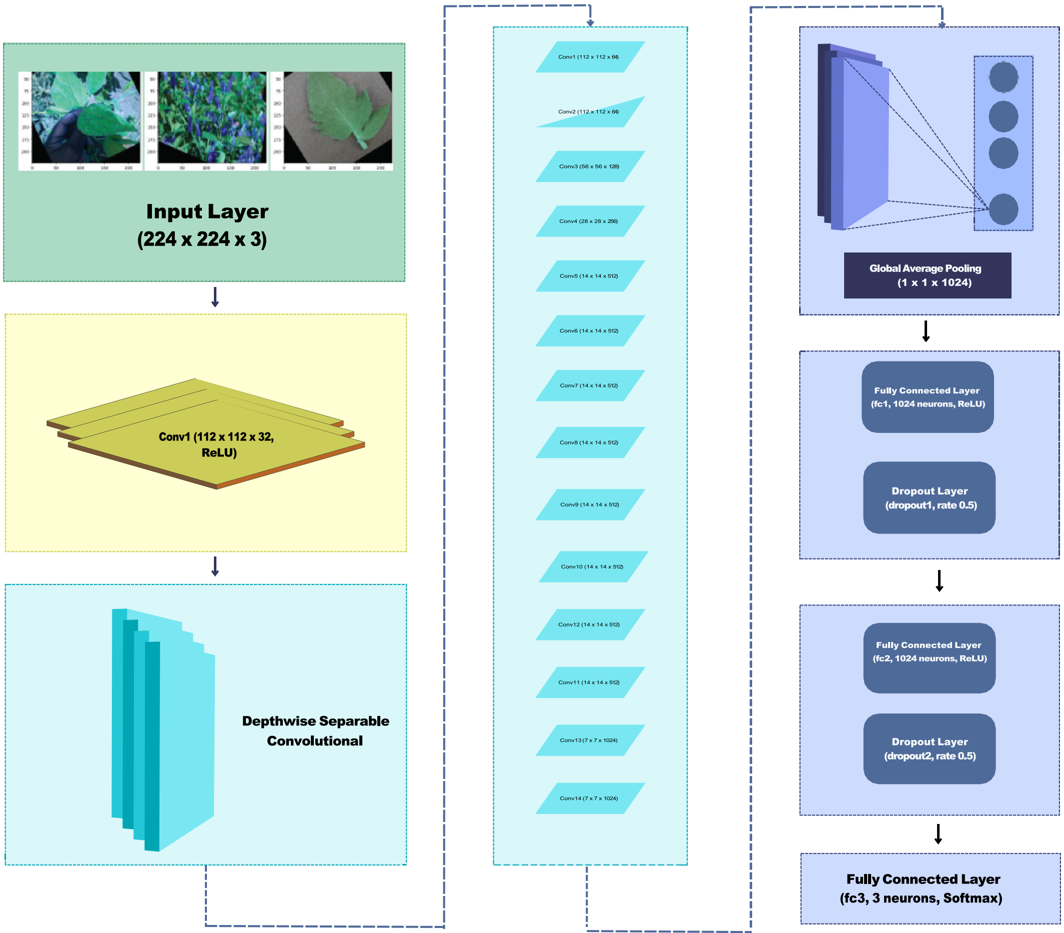

FIGURE 3 | Model architecture of the proposed InsightNet model.

TABLE 6 | Comparison between MobileNet base model and InsightNet model.

Feature	MobileNet base model	InsightNet model
Input layer	224 × 224 × 3	224 × 224 × 3
Initial convolution layer	Conv1 (3 × 3, 32, Stride 2, ReLU6)	Conv1 (3 × 3, 32, ReLU)
Convolutional layers	Depthwise separable convolutions with ReLU6 across all layers	Depthwise separable convolutions with ReLU (Conv2 to Conv14)
Global pooling layer	Global average pooling (1 × 1 × 1024)	Global average pooling (1 × 1 × 1024)
Fully connected layers	Typically none or minimal (task-dependent)	Two dense layers (fc1, fc2) with 1024 neurons each, ReLU activation
Dropout layers	Not explicitly used	Dropout layers (rate 0.5) after fc1 and fc2
Output layer	Task-specific (softmax with required classes)	Fully connected (fc3) with 3 neurons, softmax activation
Specialization	General-purpose model optimized for mobile and edge devices	Specific to plant disease detection tasks

3. Fully Connected Layers: In the InsightNet model, we introduced two fully connected layers after the global pooling layer, with the objective of performing high-level reasoning. These layers further process the features extracted by the convolutional blocks and enhance the model's capacity to make accurate predictions. The number of neurons in each of the fully connected layers (256 for the first and 4 for the second) was experimentally selected to balance model complexity and computational efficiency.

4. Dropout Regularization: Dropout layers were added between the fully connected layers to avoid overfitting by randomly disabling a portion of neurons at each iteration during training. This enables the model to generalize better on new, unseen data by not allowing the model to depend too much on any one neuron, which otherwise would result in memorization instead of learning helpful features.
5. Softmax Activation: The final output layer uses the softmax activation function to produce class probabilities. It converts the model's output into a probability distribution over the disease classes, and the class with the highest probability is selected as the predicted class. This is required for multiclass classification problems like plant disease identification, where each image can be one among many classes.
6. Task-Specific Improvements: For additional optimization of the model for the classification of plant diseases, we performed certain task-specific changes to the architecture of MobileNet. This involved tweaking the convolutional layers and adding dropout regularization after every fully connected layer. These task-specific improvements enable the model to perform even better on the particular task of leaf disease recognition without compromising on computational efficiency, which is paramount for real-world usage.
7. Comparison with MobileNet: Unlike the baseline MobileNet model that uses few fully connected layers, the InsightNet model introduces two additional dense layers, consisting of 1024 neurons each, which play a crucial role in introducing the necessary complexity to discern various plant diseases. The point-wise comparison between the MobileNet and InsightNet models (as given in Table 6) brings forth such structural differences and demonstrates the task-specific improvements included in InsightNet for achieving better accuracy.

4 | Result

This study developed and evaluated neural network models on three datasets: bean, chili, and tomato to improve plant disease identification. Ten base models and three uniquely modified MobileNet models comprised the models. Notably, the updated MobileNet models outperformed the base models regarding accuracy and data loss reduction across all three datasets.

TABLE 7 | The performance evaluation for Dataset D1.

Model	Loss	Accuracy
MobileNetV2	0.23	89.7%
EfficientNetB2	0.18	90.1%
DenseNet121	0.25	88.4%
XceptionV3	0.21	89.9%
ResNet	0.27	87.2%
DenseNet169	0.20	89.3%
DenseNet201	0.19	90.5%
VGG16	0.30	85.7%
VGG19	0.32	84.9%
MobileNet	0.28	92.5%
InsightNet model	0.0466	98.07%

Tables 7, 8, and 9 represent the performance evaluation for Datasets D1 (chili), D2 (tomato), and D3 (bean). Table 7 shows comparisons between different models on the chili dataset, clearly establishing that the proposed InsightNet model has an accuracy of 98.07% and a very low loss of 0.0466. This greatly improves on models like MobileNetV2 and DenseNet201, thus proving that the modified MobileNet model is much more efficient and reliable for plant disease detection in chili plants. In the Tomato dataset, the modified MobileNet's accuracy stood at 97.9% with a loss of 0.0750, once again outshining the base MobileNet's 93.3% accuracy and 0.30 loss. Similarly in the Bean dataset, the modified MobileNet model achieved a remarkable accuracy of 98.07% with a loss of 0.0466, significantly outperforming the base MobileNet model's accuracy of 92.5% and loss of 0.28. Similarly, the Chili dataset saw the modified MobileNet achieving an impressive accuracy

TABLE 8 | The performance evaluation for Dataset D2.

Model	Loss	Accuracy
MobileNetV2	0.27	89.20%
EfficientNetB2	0.22	90.60%
DenseNet121	0.24	88.90%
XceptionV3	0.19	89.50%
ResNet	0.31	86.40%
DenseNet169	0.23	89.00%
DenseNet201	0.21	90.20%
VGG16	0.34	85.30%
VGG19	0.33	84.60%
MobileNet	0.29	91.10%
InsightNet Model	0.1363	97.95%

TABLE 9 | The performance evaluation for Dataset D3.

Model	Loss	Accuracy
MobileNetV2	0.25	88.50%
EfficientNetB2	0.17	90.90%
DenseNet121	0.26	87.70%
XceptionV3	0.22	89.30%
ResNet	0.29	86.80%
DenseNet169	0.24	88.10%
DenseNet201	0.20	90.30%
VGG16	0.35	85.00%
VGG19	0.36	84.20%
MobileNet	0.30	93.30%
InsightNet model	0.075	97.90%

of 97.95% against a loss of 0.1363, a substantial improvement over the base MobileNet's 91.1% accuracy and 0.29 loss. The consistent outperformance of the modified MobileNet models is attributed to the fine-tuning of their architecture. Modifications were carefully crafted to enhance feature extraction and generalization capabilities, allowing for a deeper and more nuanced understanding of the plant disease indicators in the imagery data. These improvements lead to a superior ability to distinguish between healthy and diseased plant conditions.

The performance outcomes of the base and modified MobileNet models, shown in Table 10, highlight the effectiveness of the proposed InsightNet. Here, the modified model shows a significant improvement for the Chili dataset, soaring to 98.12% across precision, recall, F1-score, and accuracy metrics from the base model's livery of 92.50%. The Bean dataset images a marked advancement as well, with the InsightNet model reaching a 4.60% increase in all metrics, moving the envelope from 93.30% to 97.90%. The Tomato dataset's analysis substantiates the modified model's raised performance, where a 6.28% increase is observed, leading to an impressive 97.95% in comparison to the base model's 91.67%. Table 10 not only highlights the quantifiable advancements of the proposed InsightNet model but also attests to its augmented predictive proficiency, positioning it as a formidable tool in the domain of agricultural disease surveillance. Figure 4 shows a comparative analysis between the base and proposed InsightNet models on the 3 datasets.

TABLE 10 | Performance evaluation of the base and proposed models for different crops.

Model	Precision	Recall	F1-score	Accuracy
Base model (chili)	92.50%	92.50%	92.50%	92.50%
Proposed model (chili)	97.96%	97.45%	97.69%	98.07%
Base model (tomato)	94.67%	94.67%	94.67%	91.67%
Proposed model (tomato)	97.50%	97.50%	97.50%	97.95%
Base model (bean)	94.97%	94.97%	94.97%	93.30%
Proposed model (bean)	96.45%	96.46%	96.40%	97.90%

Figures 5–10 show the learning curves that describe the architectural optimization and enhanced predictive performance of the MobileNet models in the assessment of the models over the three datasets: Bean, Chili, and Tomato. The modified MobileNet shows an expeditious elevation to peak training accuracy, with a notable generality to validation data, as seen in the Bean dataset (Figures 5 and 6). This model's training and validation accuracies combine closely, eschewing the overfitting marked in the base MobileNet, where a noticeable difference between training and validation curves shows a less generalizable approach. Similarly, in the Chili dataset (Figures 7 and 8), the modified MobileNet's early and secured confluence of accuracy metrics highlights its robustness, a stark comparison to the base model's wider accuracy gap. Loss trends across these datasets accentuate the modified model's stable validation loss, representing a more tailored fit and consistent learning. A similar pattern appears in the Tomato dataset (Figures 9 and 10), with the modified MobileNet maintaining lower validation loss and displaying less disparity between training and validation accuracies.

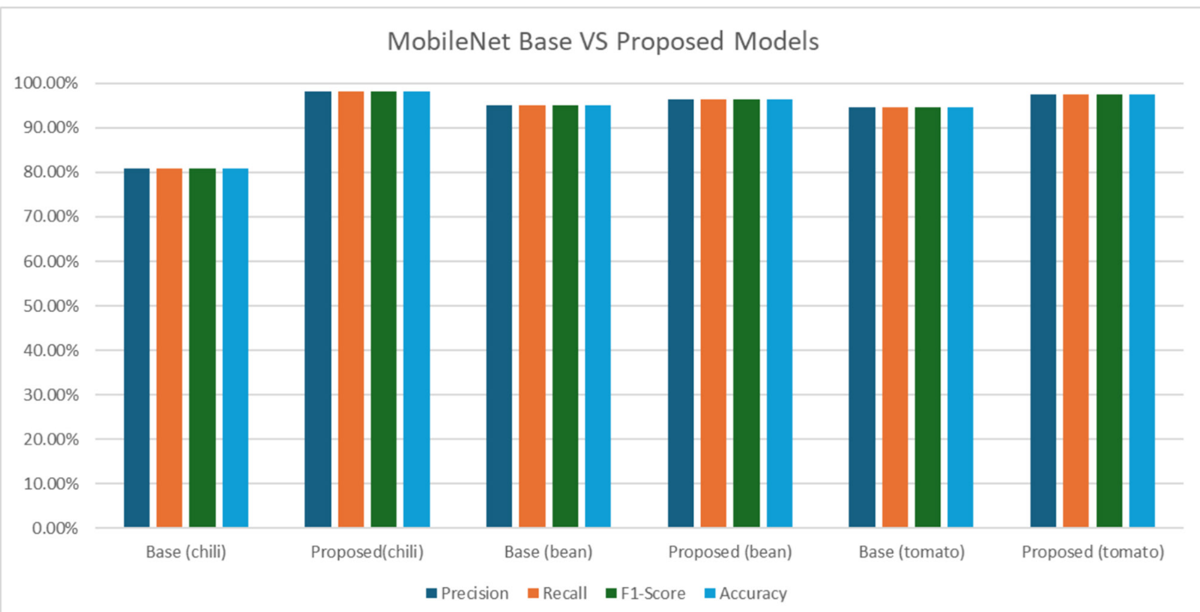

FIGURE 4 | A comparative analysis of different datasets between the base and proposed models.

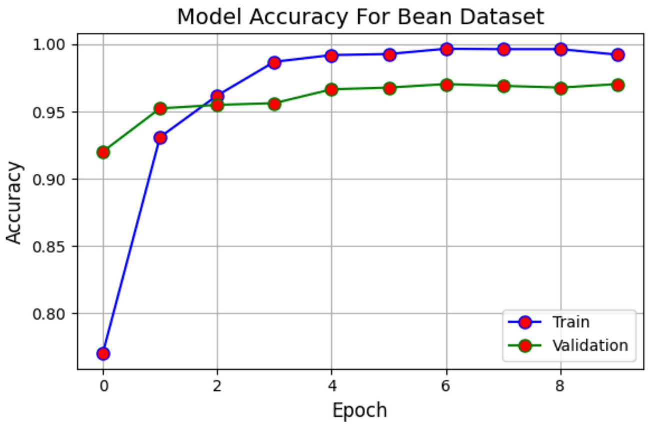

FIGURE 5 | Training and validation accuracy for the bean proposed InsightNet model across different epochs.

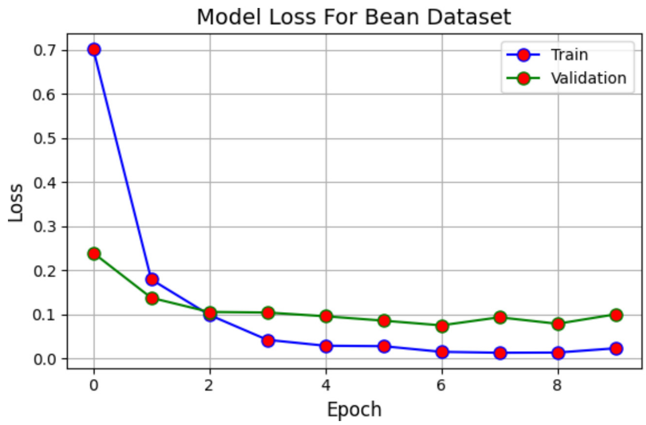

FIGURE 6 | Training and validation loss for the bean proposed InsightNet model across different epochs.

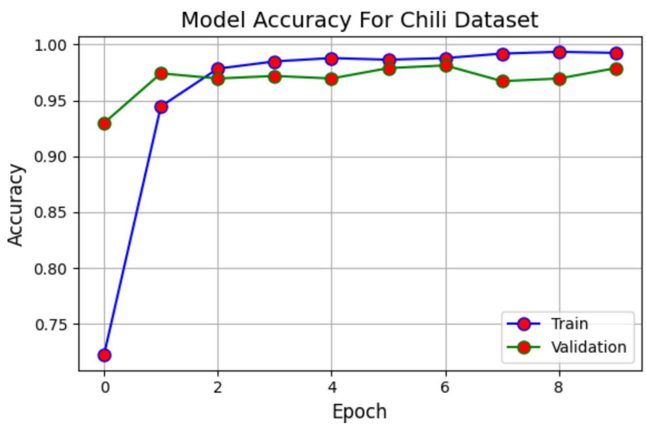

FIGURE 7 | Training accuracy for the chili proposed model across different epochs.

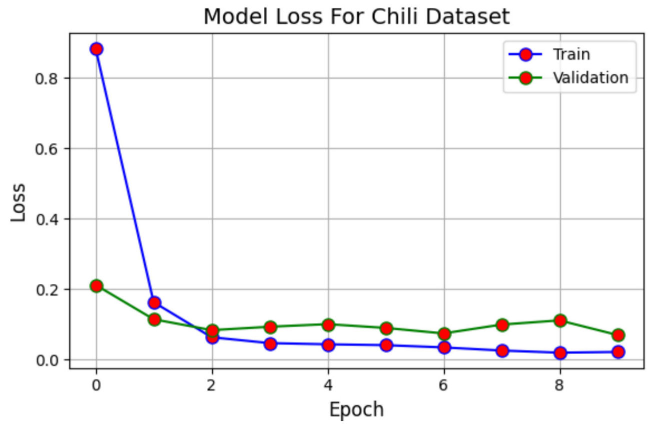

FIGURE 8 | Training loss for the chili proposed model across different epochs.

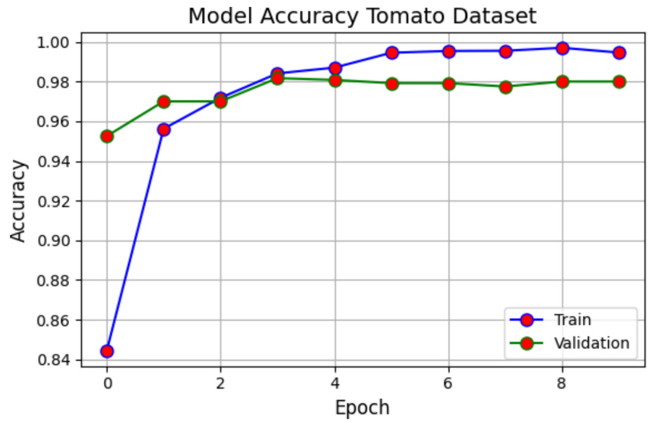

FIGURE 9 | Training accuracy for the tomato proposed model across different epochs.

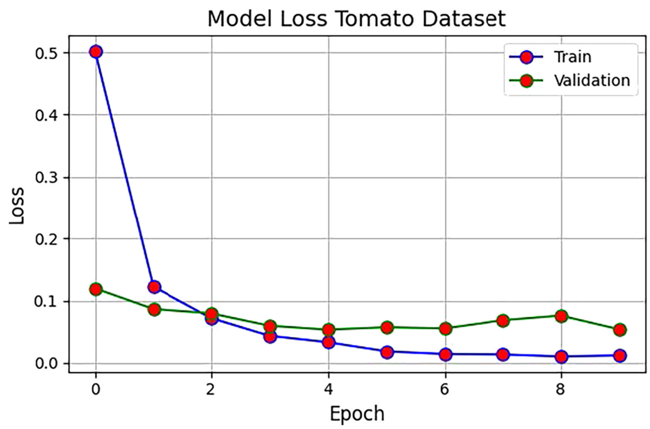

FIGURE 10 | Training loss for the tomato proposed model across different epochs.

Figures 5 and 6 represent a comparative analysis between the performance metrics of the base and proposed model across three datasets—Chili, Bean, and Tomato. Precision, recall, F1-score, and accuracy metrics quantify model performance. For the Chili dataset, the proposed model shows significant improvements over the base model, with precision, recall, F1-score, and accuracy all soaring from 92.50% to a remarkable 98.12%. In the Bean dataset, the proposed model also shows enhanced results, increasing the metrics from 93.30% to 97.90%, indicating a significant increase in the model's predictive precision and reliability. The Tomato dataset results are particularly considerable, with the proposed model achieving 97.95% across all metrics, a notable rise from the base model's 91.67%. This consistent enhancement in key performance indicators across varied datasets substantiates the proposed model's advanced capability for accurate classification, underscoring its potential utility in precision agriculture for real-time plant disease detection.

The model demonstrates an expeditious rise to peak training accuracy, with a notable generalization to validation data, as seen in the Bean Dataset (Figures 5 and 6). This model's training and validation accuracies converge closely, eschewing the overfitting observed in the base MobileNet, where a pronounced discrepancy between training and validation curves suggests a less generalized approach. Similarly, the proposed model's early and tight convergence of accuracy metrics in the Chili dataset (Figures 5 and 6 and Figures 7 and 8) emphasizes its robustness, in sharp contrast to the standard model's broader accuracy gap. Loss patterns across these datasets highlight the updated model's steady validation loss, indicating a more customized fit and reliable learning. A similar result can be seen in the Tomato dataset (Figures 9 and 10) where the proposed model shows less variation in training and validation accuracy and maintains reduced validation loss. Consolidated, these observations support the push towards sustainable crop management and precision agriculture by demonstrating the flexibility of the proposed model and validating its deployment for real-time agricultural disease detection, supported by higher precision, recall, and F1 scores.

The 3 dataset's confusion matrices are illustrated in Figures 11–13; Figure 11 shows the confusion matrix from the chili plant disease experiments. The true positive rates for the "healthy," "leaf curl," and "leaf spot" classes show the model's good diagnostic performance, as seen by the observable diagonal values (115 for "healthy," 122 for "leaf curl," and 122 for "leaf spot"). It indicates that the model can differentiate between healthy leaves and those impacted by disease. A tiny proportion of healthy leaves were mistakenly categorized as having "leaf spot" and "whitefly" infections, indicating a minor categorization error that may have resulted from the inherent similarity in visual symptoms. As depicted in Figure 12, the confusion matrix for bean plant diseases encompasses "Angular leaf spot," "Bean rust," and a "Healthy" class. The model exhibits remarkable accuracy in identifying a majority of healthy and diseased cases—246 for "Angular leaf spot," 252 for "Bean rust," and 250 for "Healthy." Although some slight declassifications indicate symptom overlap between diseases, the model's high true positive rates attest to its adeptness in learning and generalizing from the training set. Figure 13 shows the confusion matrix for tomato plant illnesses. It reveals the model's ability to accurately distinguish between the "bacterial spot," "septoria leaf spot," "yellow leaf curl virus," and "healthy" classifications. The high true positive counts (290 for "septoria leaf spot," 269 for "bacterial spot," 311 for "yellow leaf curl virus," and 300 for "healthy") indicate a strong classification capability. Although there is some ambiguity between "bacterial spot" and "yellow leaf curl virus," as well as between "septoria leaf spot" and "bacterial spot," the model appears to be highly perceptive to the unique characteristics of each illness.

Therefore, it can be said that the study's findings strongly validate the effectiveness of the proposed model. The modifications have produced notable gains in the main performance indicators in the tomato, bean, and chili datasets. Compared to the baseline models, the data demonstrate that the architectural changes made to the MobileNet models have resulted in better performance in disease classification tasks. These results show notable advancements in using deep learning models in agricultural health rather than just little steps. The research provides empirical evidence that validates the feasibility of these updated models for accurate and timely plant disease detection, hence solidifying their role as essential instruments for the progress of agricultural technologies.

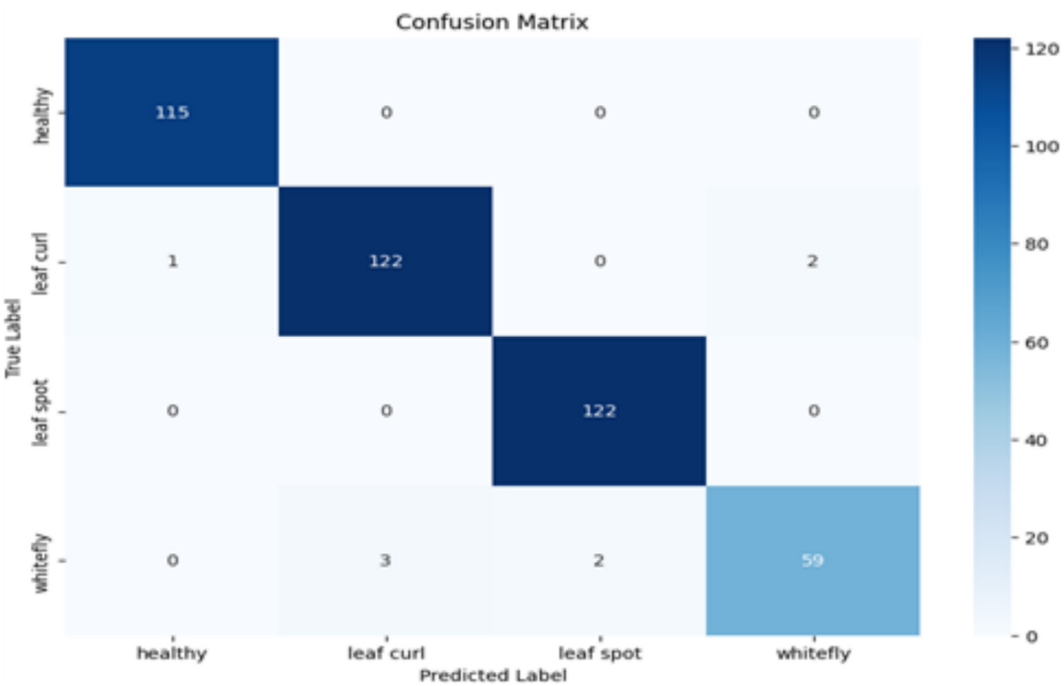

FIGURE 11 | Confusion matrix for chili.

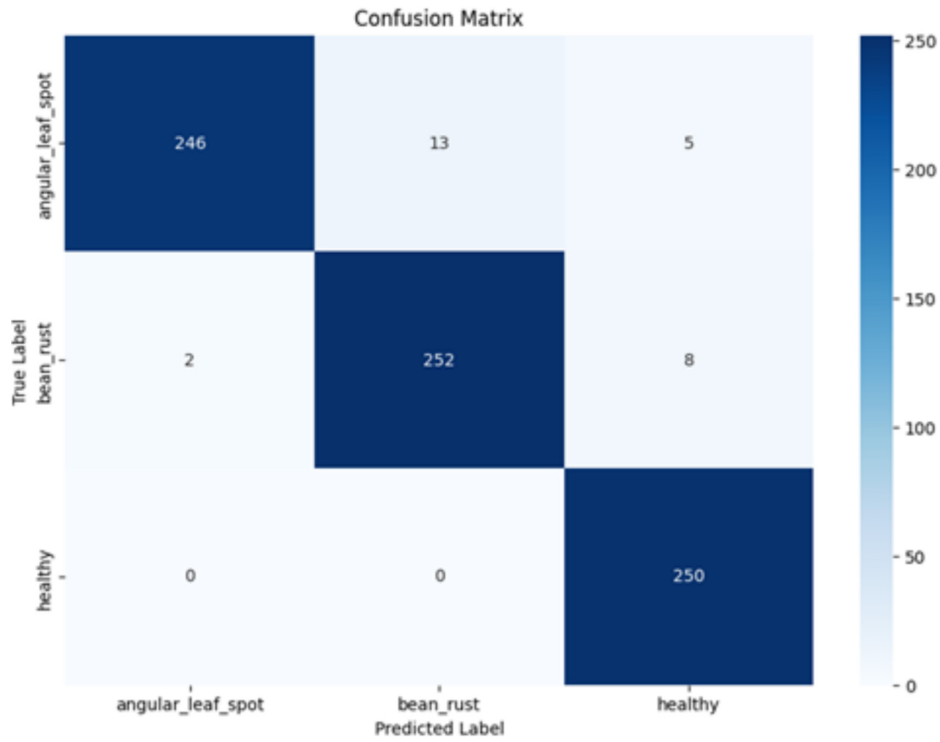

FIGURE 12 | Confusion matrix for bean.

4.1 | Grad-CAM Analysis

Grad-CAM was proposed by Selvaraju et al. (Selvaraju et al. 2017). It is a method for improving the transparency of decisions made by a wide class of CNN-based models. Grad-CAM enables us to visually evaluate the direction the model is looking in, ensuring that it is genuinely concentrating on and activating around the appropriate patterns in the image. This section will present the result of the analysis of plant leaf conditions, using original images and their corresponding Grad-CAM heatmaps for three types of leaves: chili, bean, and tomato. As evidenced in both the routines, the model focuses on disease-relevant regions; however, the correctness test is done specifically to identify the disease symptoms. The results obtained for all three are detailed below.

4.1.1 | Chili

Images in Figure 14 explain the analysis of chili leaf conditions from original images and their respective Grad-CAM heatmaps. The original image corresponds to the chili leaves captured under the symptoms of the disease. These original images are further elaborated through its heatmap generated by the Grad-CAM method. In Figure 14, the first part is a plant leaf with noticeable spotting, and the second part is different leaves with different spotty and discoloration rates. This could be a disease infection. It creates the basis for the check of the model focus on the relevancy of used features during a disease detection. The corresponding heatmaps next to each original image indicate those regions that the deep learning model considered more important during its decision for disease detection. The heatmaps of the diseased leaves indicated hot areas of intensive activation around spots and discolored regions, signifying that the model can focus on symptomatic zones. As an example, from the second and third images, leaves have clear spots, and these hot spots are also indicated in the corresponding heatmaps. This means that the model was able to find the areas expressing the symptoms of the disease, hence the ability of fine localization of disease-relevant features. It can, therefore, be assumed that the model detects similar patterns of disease because the results on the heatmaps are consistent throughout different leaves. The keenness of the model forms an excellent match with the biological markers of disease, which enhances the robustness of the model in the correct diagnosis of chili leaf diseases.

Therefore, Grad-CAM assists in understanding the model decisions. It is brought to an assurance that the model attends only to those areas of the chili leaves, which are disease-relevant. It supports visual interpretability, because it directly deduces confidence in the model and hence leads to enhancing reliability and further drives evidence in real-time scenarios, more importantly, seeing agricultural disease detection. Such an approach ensures that the model performs well not only in terms of accuracy but also gives transparent and explainable responses, hence presenting how well the experiments can be extended for other plant species and other types of disease to enhance and validate agricultural health deep learning diagnosis.

4.1.2 | Tomato

The images in Figure 15 are of the conditions found in tomato leaves; these are the images of the Image and the associated Grad-CAM heatmaps. Low symptoms in the normal leaves indicate a disease, which is visualized later using the generated

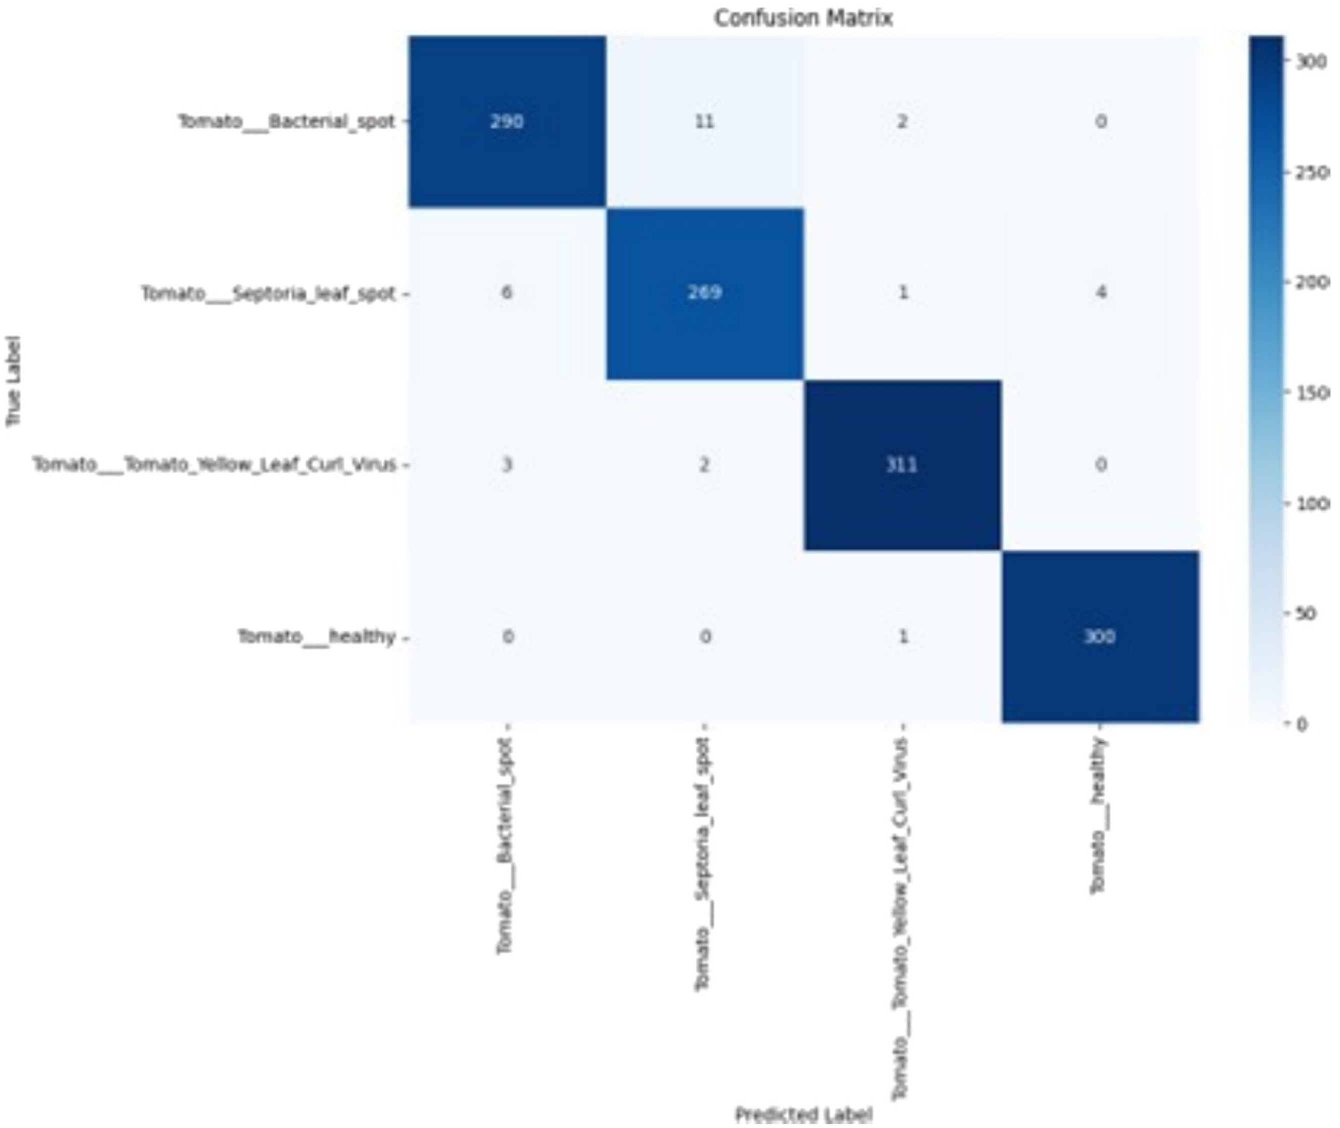

FIGURE 13 | Confusion matrix for tomato.

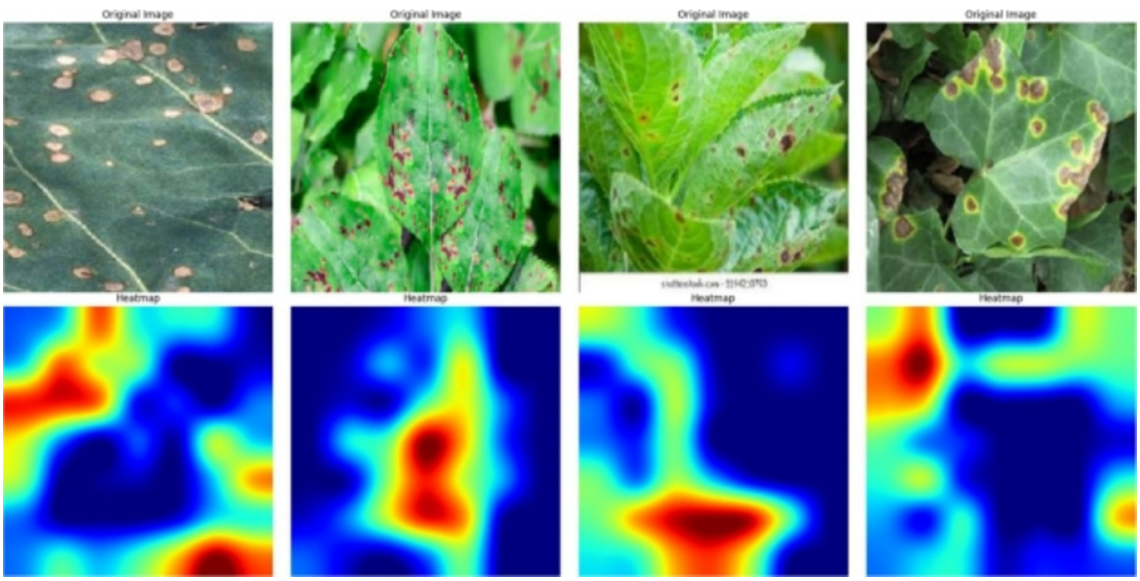

FIGURE 14 | Analysis of chili leaf conditions using original images and corresponding heatmaps. Each pair demonstrates specific areas of interest highlighted by heatmaps for different leaves, indicating potential areas of stress or disease.

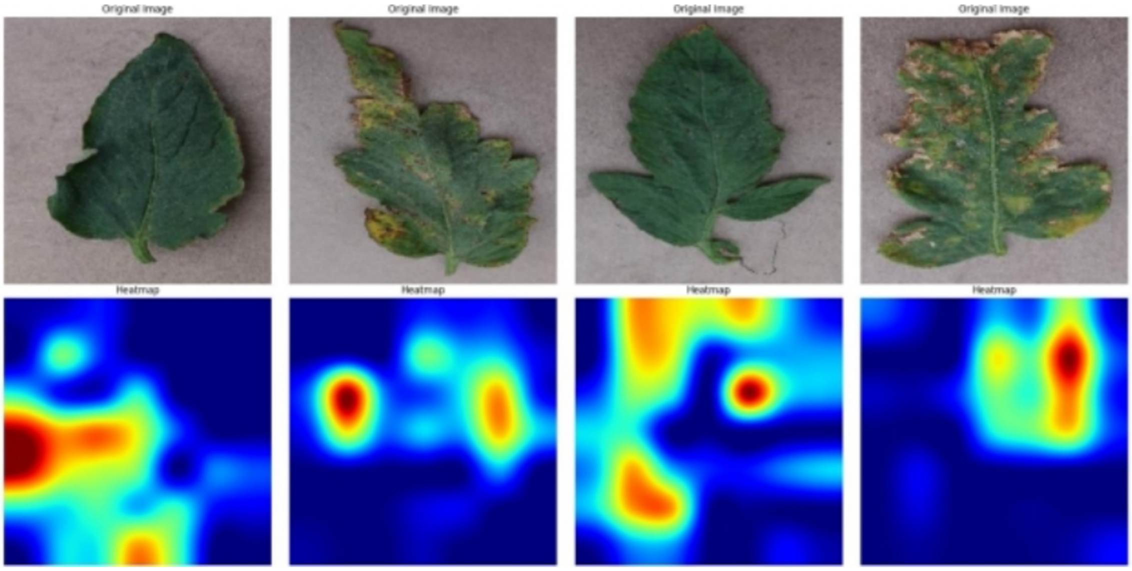

FIGURE 15 | Analysis of tomato leaf conditions using original images and corresponding heatmaps. Each pair demonstrates specific areas of interest highlighted by heatmaps for different leaves, indicating potential areas of stress or disease.

heatmaps following the theory of Grad-CAM. The first one is of a rather healthy leaf, where only minor, non-evident features of the disease are present, but the following ones have much clearer signs of spotted discoloration and observable structural deformation of the leaf tissue due to the infectious disease. These original pictures and heatmaps will provide a basis for further comparison of focus and relevance of the model toward the detection of the disease. The most important regions that the deep learning model has considered are highlighted for each of the resulting heatmaps of the original pictures. In the first image, for a healthy leaf, the heatmap shows moderate activation, reflecting which areas had importance to the model in the prediction while also confirming the presence of no very strong disease markers. In contrast, the heatmaps in the case of the diseased leaves reveal areas of intense activations all around the spots and discolored regions, showing clear focus by the model on symptomatic areas. For instance, the second and the third images show leaves with spots clearly, with the corresponding heatmaps also depicting excellent activation in the spots, implying that the model thus correctly pinpointed areas with the disease manifestations, hence assuring that its potential locational accuracy of disease-related features was quite high.

The presence of continuity across the heatmap predictions of different leaves implies that the network can pick up similarities in disease patterns. The model is robust in making fine diagnoses of tomato leaf diseases through the alignment of its focus with their biological markers; Grad-CAM analysis helps to address the decision-making of the deep learning model correctly. Such visual interpretability further improves the reliability of the model and enables its use in real-life applications, particularly in the case of disease detection in crops. If this model is accurate and has a good performance, such approaches can be easily employed universally in other plant species and types of disease, which is in a bid to improve and validate deep learning models within the agricultural health diagnosis application.

4.1.3 | Bean

The images in Figure 16 contain a display of bean leaf analysis. Here are the original images and their Grad-CAM heatmaps. Each of the original images represents the state of the leaves from bean under different conditions of disease symptom expression. Symptoms displayed in the original images suggest the possibility of the presence of diseases. The first one is a leaf that shows no signs of a disease, with an alternately viewed healthy leaf. With a more peculiar view, evident spots and discolorations mark the possibility of a disease infection being the scenario. The heatmaps corresponding to them show regions most significant that the deep learning model treated in finding diseases. In the case of the healthy leaf, the heatmap thus shows very little activation, or in other words, the model identified the leaf properly whereby it was devoid of symptoms of the disease. In the case of leaves with disease, however, activations are seen around the spots and discolored areas of the leaf, distinctly indicating the symptomatic region, as seen in the heatmaps.

The fact that the hotted areas in the heatmaps remain consistent across varied leaves implied the model can determine prototypical patterns of the disease. The focus of the model around the biological measures of disease—a characteristic of the model's capacity to diagnose bean leaf disease in a consistent and true manner—suggested model robustness. The analysis of Grad-CAM showed importance in proving the fact that the model does focus on the disease-relevant regions of the bean leaves. This visualization of explanation augments the trustworthiness of the model and shows proof of its application in the real world of disease detection in agriculture. Best applied to other species of plants and types of diseases, this approach assures that the model is not accurate or transparent in its findings, thus getting further validation and improvement of deep learning models in agricultural health diagnostics.

5 | Discussion

InsightNet is a major leap forward from all other previous deep learning models of plant disease detection, including its base variant, MobileNet. The result here shows a manifold increase in both accuracy and computational efficiency, improving the problems presented by traditional models. While these architectures, such as ResNet50, Xception, InceptionV3, and DenseNet121, have achieved great accuracy rates, they require the most computational resources; hence, that limits their reach for different plant species with diverse symptoms. For instance, ResNet18 obtained an accuracy of 96% in detecting diseases of tomato leaves and a deep CNN model obtained 97.02% validation accuracy on a dataset of 18,392 augmented images. These highly resource-intensive models may be unaffordable for small-scale farmers or those working in resource-poor settings. A direct comparison of the performance factors between the base MobileNet architecture and the InsightNet model is shown in Table 11. This table highlights key differences in terms of model complexity, learning rates, dropout usage, and specialization, underscoring the advantages of InsightNet for specific classification tasks.

InsightNet has been developed to support high performance for which efficient deployment is necessary. This resulted in superior performance on several datasets over and above the baseline MobileNet model. The results of the Bean dataset are 97.90% for InsightNet and 93.30% for MobileNet. Similarly, for the Chili dataset, the accuracy of the proposed model and base model is 98.12% and 92.50%, respectively. InsightNet outperformed with an accuracy of 97.95% on the Tomato dataset, while MobileNet was a little behind with 91.67%. These results hint that the model generalizes well across plant species and is bound to be useful in agricultural applications for a wide variety of crops.

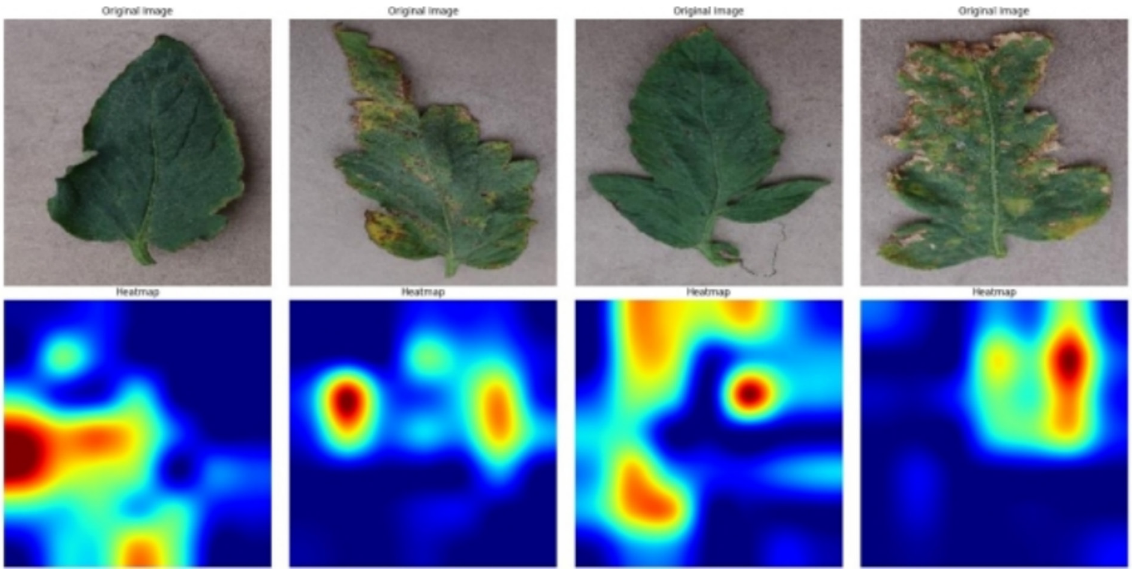

FIGURE 16 | Analysis of bean leaf conditions using original images and corresponding heatmaps. Each pair demonstrates specific areas of interest highlighted by heatmaps for different leaves, indicating potential areas of stress or disease.

TABLE 11 | Comparison of performance factors between MobileNet base model and InsightNet model.

Factor	MobileNet base model	InsightNet model
Model complexity	Single dense layer (256 neurons, fewer parameters)	Two dense layers (1024 neurons each, more parameters)
Learning rate	Higher learning rate (0.01) for faster convergence but less precision	Lower learning rate (0.0001) for finer convergence
Output classes	General-purpose design; adaptable for various tasks	Tailored for a specific 3-class plant disease classification task
Dropout regularization	Single dropout layer, less regularization	Dropout after both dense layers, reducing overfitting
Alpha value	Alpha = 1.0 (wider model, more feature channels)	Alpha = 0.5 (lighter model, fewer feature channels)
Best use case	Suitable for general tasks and efficiency-critical scenarios	Ideal for specialized, high-performance classification tasks

Further, the performance metrics of InsightNet prove its superiority over traditional models. The Chili dataset exhibited very high values in performance metrics like precision, recall, and F1-score, around 98.12%, which signifies huge uplifts over the base model value of 92.50%. This kind of similar trend also continues for the Bean and Tomato datasets, where InsightNet tends to outperform the base models in all key metrics under consideration. These enhancements further validate the real-world agricultural applications of the model and provide a reliable solution in detecting and classifying plant diseases.

Tomatoes, beans, and chili plants bring out how each different plant species adapts, insisting InsightNet is of importance in handling diverse aspects related to symptoms of the disease among different crops. The fact that accuracy increased from approximately 92% to almost 98% is critical to the improvement in the management of agriculture. Increasing such precision in the diagnosis of plant diseases within agriculture helps farmers with effective means of early disease identification, which helps protect crop production and decreases dependence on chemical interventions. Advancements like these are encouraging better agriculture practices and sustainable farming methods. The success of InsightNet bodes for a more substantial contribution to the agricultural industry's resilience and food security.

While the model has been trained and validated on well-curated datasets, its performance in real agricultural environments remains to be assessed. Real-world conditions—such as uncontrolled lighting, occlusions, and natural symptom variability—pose additional challenges do not present in controlled datasets. To address this, future work will involve testing and refining the model in collaboration with farmers and agricultural researchers in actual field settings. This step will help assess model robustness, identify deployment challenges, and adapt InsightNet for use on mobile or edge devices.

Future work will aim at the practical deployment of the InsightNet model with the help of Agriculture 4.0 technology, including Internet of Things (IoT) devices, edge computing devices like Raspberry Pi and NVIDIA Jetson Nano, and advanced sensory devices. InsightNet deployment on these platforms will enable real-time, offline plant disease classification in the field, allowing farmers to capture leaf images through a camera sensor or mobile application, receive instant classification, and access treatment recommendations. These functionalities will provide ongoing tracking of crop health, real-time alerts, and smooth integration with agricultural management systems for timely interventions and minimization of crop losses. For this purpose, model simplifications and user-friendly interfaces will be created to enable farmers and agricultural technicians, independent of their technical expertise, to utilize and profit from these tools. Besides, the enhancements will also be aimed at enhancing the capacity of the model to distinguish between visually indistinguishable diseases, i.e., “leaf spot” and “whitefly” on chili crops—misclassifications reported in the confusion matrix. Future enhancements will involve the addition of attention mechanisms, high-resolution or multi-spectral imaging, together with temporal analysis of symptom development to enhance correct disease classification under real agricultural settings. Co-pilot trials between agricultural producers and industry partners will help develop and validate the system in multiple ecological settings. InsightNet marks a significant step forward in agricultural health management by improving diagnostic accuracy and enabling scalable, efficient disease detection. Its success underscores the transformative potential of deep learning in agriculture, contributing to sustainable farming practices, food security, and the empowerment of stakeholders across the agricultural ecosystem.

6 | Conclusions

This paper therefore proposes InsightNet, a deep learning-based approach that efficiently integrates a modified MobileNet architecture with XAI in identifying and classifying plant diseases from different species. Given the problems of a high computational cost and low adaptability of traditional diagnosis models, InsightNet presents yet another leap in agricultural technology. The model architecture will allow complex disease manifestations in a way that shall lighten the burden of diagnosis, hence aiding effective functions in agriculture. InsightNet leverages the benchmark leaf image dataset for training and testing to ensure performance in practical scenarios.

Success with InsightNet identified the magic mix of knowledge standing at the confluence of plant pathology and the best computational methodologies. It will therefore not only realize very good accuracy in disease detection but also provide a strong foundation for the model to be scalable at various agricultural environments. Considering the promising performance, the current study was underpinned by some limitations which include the quality images and comprehensive training datasets representative of the full spectrum of disease symptoms that future research needs to overcome. Further studies are required to be done on the performance of InsightNet regarding the climatic and environmental conditions, which usually characterize most of the agricultural settings in reality. Overall, this study will introduce InsightNet, a platform with the potential for very important disruptive changes in agriculture by way of improved disease detection and management. This rapid and continuous improvement in AI-driven diagnostic tools will further contribute toward creating sustainable agriculture and improved food security across the world, opening up even newer frontiers in this deep learning application of AI.

Author Contributions

Mubasshar U. I. Tamim conceptualized the study, developed the architecture for the proposed model, and led to the technical implementation of the framework. Sultanul A. Hamim contributed to the model's implementation, specifically focusing on data preprocessing, feature engineering, and optimizing the architecture for heart failure prediction. He also performed experiments and data analyses to validate the performance of the proposed model. In addition, he was responsible for designing the model evaluation methodology and interpreting the results. Sumaiya Malik assisted with a critical review of the architecture and evaluation metrics, ensuring the robustness of the experiments. She was also responsible for drafting the manuscript, revising it for intellectual content, and preparing figures and tables for presentation. M. F. Mridha contributed to the development of ideas, supervised the research, and administered the project. Sharfuddin Mahmood validated the findings, reviewed, and edited the manuscript.

Acknowledgments

This work has been carried out in AIUB. The authors thank AIUB authority for their financial support. We would also like to express our sincere gratitude to the Advanced Machine Intelligence Research Lab (AMIR Lab) for their invaluable support and guidance, which significantly contributed to the successful completion of this study.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Peer Review

The peer review history for this article is available in the Supporting Information for this article.

References

Abd Algani, Y. M., O. J. M. Caro, L. M. R. Bravo, C. Kaur, M. S. Al Ansari, and B. K. Bala. 2023. "Measurement: Sensors." *Measurement* 25: 100643.

Al-gaashani, M. S., F. Shang, M. S. Muthanna, M. Khayyat, and A. A. Abd El-Latif. 2022. "Tomato Leaf Disease Classification by Exploiting Transfer Learning and Feature Concatenation." *IET Image Processing* 16, no. 3: 913–925.

El Akhal, H., A. B. Yahya, N. Moussa, and A. E. B. El Alaoui. 2023. "A Novel Approach for Image-Based Olive Leaf Diseases Classification Using a Deep Hybrid Model." *Ecological Informatics* 77: 102276.

Ferentinos, K. P. 2018. "Deep Learning Models for Plant Disease Detection and Diagnosis." *Computers and Electronics in Agriculture* 145: 311–318.

Ghodekar, A. and A. Kumar. 2023. "LightLeafNet: Lightweight and Efficient NASNetMobile Architecture for Tomato Leaf Disease Classification." Available SSRN 4570347.

Hamim, S. A., and A. I. Jony. 2024. "Enhancing Brain Tumor MRI Segmentation Accuracy and Efficiency With Optimized U-Net Architecture." *Malaysian Journal of Science and Advanced Technology* 4, no. 3: 197–202.

Kannadasan, R., M. H. Alsharif, A. Jahid, et al. 2021. "Categorizing Diseases From Leaf Images Using a Hybrid Learning Model." *Symmetry* 13, no. 11: 2073.

Krishnan, V. G., J. Deepa, P. V. Rao, V. Divya, and S. Kaviarasan. 2022. "An Automated Segmentation and Classification Model for Banana Leaf Disease Detection." *Journal of Applied Biology and Biotechnology* 10, no. 1: 213–220.

Kumar, Y., R. Singh, M. R. Moudgil, and Kamini. 2023. "A Systematic Review of Different Categories of Plant Disease Detection Using Deep Learning-Based Approaches." *Archives of Computational Methods in Engineering* 30, no. 8: 4757–4779.

Mahesh, T., R. Sivakami, I. Manimozhi, et al. 2023. "Early Predictive Model for Detection of Plant Leaf Diseases Using MobileNetV2 Architecture." *International Journal of Intelligent Systems and Applications in Engineering* 11, no. 2: 46–54.

Mohanty, S. P., D. P. Hughes, and M. Salathé. 2016. "Using Deep Learning for Image-Based Plant Disease Detection." *Frontiers in Plant Science* 7: 215232.

Pratap, V. K., and N. S. Kumar. 2023. "High-Precision Multiclass Classification of Chili Leaf Disease Through Customized EffecientNetB4 From Chili Leaf Images." *Smart Agricultural Technology* 5: 100295.

Praveen, P., M. Nischitha, C. Supriya, M. Yogitha, and A. Suryanandh. 2022. "To Detect Plant Disease Identification on Leaf Using Machine Learning Algorithms." In *Intelligent System Design: Proceedings of INDIA 2022*, 239–249. Springer.

Ramanjot, M. U., A. Wadhawan, J. Singla, N. Jhanjhi, R. M. Ghoniem, et al. 2023. "Plant Disease Detection and Classification: A Systematic Literature Review." *Sensors* 23, no. 10: 4769.

Saleem, M. H., S. Khanchi, J. Potgieter, and K. M. Arif. 2020. "Image-Based Plant Disease Identification by Deep Learning Meta-Architectures." *Plants* 9, no. 11: 1451.

Saleem, M. H., J. Potgieter, and K. M. Arif. 2019. "Plant Disease Detection and Classification by Deep Learning." *Plants* 8, no. 11: 468.

Selvaraju, R. R., M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. 2017. "Grad-Cam: Visual Explanations From Deep Networks via Gradient-Based Localization." 618–626.

Shah, S. R., S. Qadri, H. Bibi, S. M. W. Shah, M. I. Sharif, and F. Marinello. 2023. "Comparing Inception V3, VGG 16, VGG 19, CNN, and ResNet 50: A Case Study on Early Detection of a Rice Disease." *Agronomy* 13, no. 6: 1633.

Shaheed, K., I. Qureshi, F. Abbas, et al. 2023. "EfficientRMT-Net—An Efficient ResNet-50 and Vision Transformers Approach for Classifying Potato Plant Leaf Diseases." *Sensors* 23, no. 23: 9516.

Shastry, K. A., and H. Sanjay. 2021. "Hybrid Prediction Strategy to Predict Agricultural Information." *Applied Soft Computing* 98: 106811.

Singh, V., A. Chug, and A. P. Singh. 2023. "Classification of Beans Leaf Diseases Using Fine Tuned CNN Model." *Procedia Computer Science* 218: 348–356.

Sujatha, R., J. M. Chatterjee, N. Jhanjhi, and S. N. Brohi. 2021. "Performance of Deep Learning vs Machine Learning in Plant Leaf Disease Detection." *Microprocessors and Microsystems* 80: 103615.

Tirkey, D., K. K. Singh, and S. Tripathi. 2023. "Performance Analysis of AI-Based Solutions for Crop Disease Identification, Detection, and Classification." *Smart Agricultural Technology* 5: 100238.

Too, E. C., L. Yujian, S. Njuki, and L. Yingchun. 2019. "A Comparative Study of Fine-Tuning Deep Learning Models for Plant Disease Identification." *Computers and Electronics in Agriculture* 161: 272–279.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.