Next-Generation Plant Management System and Disease Detection Using Advanced Deep Learning

Mrs. Namitha A R[1], Mr. Shridhar V[2], Mr. Ranjith Kumar B G[3], Ms. Abhishek B R[4]
[1]*Assistant. Professor, Dept. of CS&E, BGS Institute of Technology, Adichunchanagiri University, Mandya, India*
[2,3,4]*UG Student, Dept. of CS&E, BGS Institute of Technology, Adichunchanagiri University, Mandya, India*

***Abstract*—Pathogens of plants always major culprits for agricultural productivity losses at a misguided food system, food systems that are already at risk in many regions due to the scarcity of necessary timely intervention by experts. A Next, Generation Plant Management System is thus proposed through this paper to add a hybrid deep learning and context, aware advisory integration for accurate disease detection and smart remedy generation. Specifically, a Hybrid Mobile Vision Transformer (MobileViT) model with lightweight architecture that stacks convolutional neural networks for local features extraction and uses transformers for global context modelling have been utilized to solve multi, class classification of plant diseases. The model has been trained on the Plant Village dataset containing 38 classes of diseases and has produced 96.3% classification accuracy alongside high precision and recall values. In order not to leave the patients unattended after diagnosis, an attention equipped Sequence, to, Sequence (Seq2Seq) LSTM model has been used for generating structured, step, by, step treatment instructions, which complies with a BLEU scoring of 0.72 and factual accuracy of 94%. Besides biotic, plant growth is highly influenced by abiotic factors. Therefore, a Smart Plant Manager can serve as a helpful aid underpinned by weather and geolocation real, time data to help growers with the formulation of context, aware preventive alerts and cultivation strategies. The entire solution is housed in a Flask web server delivering 2 to 3 seconds per image output. The empirical findings reported an incredible prowess, scalability, and the effectiveness of the proposed AI, driven agricultural decision, support system that is designed to be practically usable.**

***Index Terms*—Plant Disease Detection, MobileViT, Vision Transformer, Seq2Seq LSTM, Deep Learning, Precision Agriculture, Smart Farming, Computer Vision, Agricultural Decision Support System.**

I. INTRODUCTION

The agriculture is still a critical sector for global food security and economic stability, but plant diseases wipe out a huge share of crops worldwide each year. In order to avoid a drop in yields, reduce the misuse of pesticides, and help farmers adopt sustainable practices, it is crucial to detect diseases accurately and at an early stage. Unfortunately, conventional ways of identifying plant diseases rely mostly on farmers or agricultural experts visually inspecting plants which is a subjective, slow, and not always locally available method especially in rural areas or regions with fewer resources. Due to these constraints, there is an urgent demand for smart, scalable, and automatic disease diagnosis solutions.

Deep learning breakthroughs over the last few years have led to significant improvements in the classification of plant diseases from images. Prestigious CNN frameworks such as ResNet, DenseNet, and EfficientNet have set new records in classification accuracy on standard datasets like Plant Village. Nevertheless, typical CNN models only extract features that are local in the image and thus, they may not be able to effectively capture the overall spatial structure and the relationships between different parts of the leaf. In contrast, the Vision Transformer (ViT) that solely relies on attention mechanisms excels at modelling global features but it has been shown that ViTs need large training datasets and heavy computational power, thus making them less feasible to be used in real farming scenarios.

In order to overcome such difficulties, the present work suggests a Next, Generation Plant Management System that features a Hybrid Mobile Vision

Transformer (MobileViT) architecture in combination with a Sequence, to, Sequence (Seq2Seq) Long Short, Term Memory (LSTM) model for the intelligent generation of remedies. The MobileViT model merges the best of both worlds of CNN, based local feature extraction and transformer, based global attention mechanisms, thus obtaining high classification accuracy and at the same time being computationally efficient enough for a web, based deployment.

Moreover, the proposed system is not only diagnosing diseases but also closing the gap between diagnosis and actionable guidance by an attention, based Seq2Seq LSTM model trained on a curated agricultural knowledge base that generates structured, step, by, step treatment recommendations.

Besides that, the system features a Smart Plant Manager component that takes advantage of the location and real, time weather data to issue preventive alerts and give suitable cultivation recommendations. This combination of frameworks enables the transformation of conventional diagnostic systems into a complete agricultural decision, support platform. A complete system is set up as a Flask, based scalable web application, which allows for quick inference and easy access in various farming environments.

II. PROBLEM DEFINITION

The plant diseases are one of the biggest threats to agricultural productivity. They cause significant economic losses and, most importantly, are a direct threat to global food security. Hence, early detection and proper management of plant diseases to control the damage and also to avoid excessive pesticide use are crucial. Unfortunately, most current disease detection methods are based on the manual visual inspection of farmers or experts in agriculture. This method is highly subjective, takes up a lot of time, and the results are often inaccurate, especially in rural or poor areas where access to experts is a problem.

Even though deep learning methods have recently brought great improvements in automated classification of plant diseases, the present systems still have a number of practical and technical drawbacks. Typically, models based on Convolutional Neural Networks (CNNs) look mainly for local features such as spots, lesions, and discoloration, and they hardly ever succeed in grasping the overall contextual relationships across the whole leaf surface. Because of this, such models are less robust when confronting complex backgrounds, different lighting conditions, infections at the initial stage, and diseases that look very similar. Moreover, standalone Vision Transformer (ViT) models demand huge datasets and lots of computational power which make them not viable for lightweight web, based or real, time agricultural scenarios.

Besides that, most of the plant disease detection systems that are already in use only serve as diagnostic classifiers. They label the disease but do not make suggestions on what the farmer can do. Essentially, farmers get no structured, step, by, step, remedy recommendations based on the disease detected. Apart from that, traditional systems hardly, if ever, take into account environmental factors such as weather, humidity, or location, which are, not only, essential factors influencing the progression of a disease but also effectiveness of a treatment.

Hence the main problem tackled in this study is creating a smart, computationally very efficient, context, aware plant management system, which would:

Differentiate between various plant diseases with a high degree of precision under natural environmental conditions. Go beyond the mere labeling of diseases and come up with intelligent, automated remedies. Use environmental and location data for proactive and personalized plant health management. Be light enough to run on standard web infrastructure for large scale real, world implementation. Meeting these challenges call for a hybrid deep learning network architecture with the ability to balance accuracy with efficiency and a clever natural language generation tool to turn disease diagnosing into actionable farming guidance. The envisaged system intends to close this gap by including in one single agricultural platform, not only state, of, the, art computer vision, but also sequence modeling, and context, aware decision support.

III. LITERATURE SURVEY

The use of AI in farming has picked up tremendous pace over the last few years, especially regarding the detection of plant diseases and precision farming. The first techniques were dependent on rule, based expert systems in which users needed to manually select the symptoms they noticed from lists of symptoms pre, defined by experts. Though such systems gave users a way of making simple diagnoses, they were unscalable, unable to adapt, and did not have the capability to learn from new data, and thus couldn't be very useful in the real world.

Later, traditional image processing and machine learning techniques were adopted for the automated identification of plant diseases. Techniques such as Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and colour histogram analysis were combined with classifiers like Support Vector Machines (SVMs). Despite the fact that these methods gave a boost to the automation process, the features designed by the user were frequently not adequate to represent the complicated and minute patterns of diseases that were especially the condition of different lighting and background.

The appearance of deep learning made plant disease classification leap, forward. Deep Convolutional Neural Networks (CNNs), such as ResNet, DenseNet, and EfficientNet, charmed the releasing of powerful performance on benchmark datasets like Plant Village. Ferentinos (2018) revealed that very deep CNNs could get very high classification accuracy for various species of crops. Yet, investigations like Singh et al. (2023) pointed out the domain shift issue, a situation in which models trained on controlled datasets lose their performance in real, field environments with complex backgrounds and lighting variability.

ViTs were thus developed for contextual understanding and hence fine, grained image classification. Kumar and Li (2023) established that absolutely transformer, based models could come close to CNNs in terms of accuracy. Nevertheless, ViTs typically call for very large datasets and more computational power, which is why most farmers have no access to them, implying that they are not suitable for deployment in their resource, constrained farming environments. To solve the problem, Mehta and Rastegari (2022) presented MobileViT, a very light, weight hybrid model that gets the best of both worlds, CNNs for extracting local features and transformers for global attention, thus achieving a better efficiency, accuracy trade, off, which is ideal for mobile and edge applications.

Besides classification, various studies have delved into real, time and edge, based disease detection systems. Chen et al. (2024) fine, tuned YOLOv5 models for the use on embedded devices and were able to show that the localization of diseases on, site is realistic. Although, the majority of detection setups just single out when a plant is sick and guess what, give no further agricultural advice or support.

Natural language processing methods have served to bridge the divide between problem identification and delivery of concrete advice. Fernandez and Park (2024) transformed remedy generation into a sequence, to, sequence issue and employed transformer, based architectures for generating well, structured treatment protocols. Their experiments revealed the potential of attention, mechanism, empowered automated agricultural prescription systems to deliver context, aware recommendations.

Furthermore, besides visual and linguistic modelling, contextual agricultural intelligence is also a hot topic. Schmidt and Patel (2024) introduced multimodal approaches that take image data and environmental factors like temperature and humidity together to refine disease risk prediction. In a similar vein, Iyer et al. (2025) stressed the role of conversational AI in making the farm sector more reachable and user, friendly, notably for farmers in areas with limited digital literacy.

IV. METHADOLOGY

The Next, Generation Plant Management System being developed uses a well, structured multi, stage approach integrating computer vision, sequence modelling, and contextual intelligence in a single framework. The work is based on the Plant Village dataset containing 38 classes of plant diseases spanning various crop species including both healthy and infected leaf samples. For consistency, all pictures were resized to 224 224 pixels and normalized to a [0, 1] scale. Besides the above, to train the model to be robust against unpredicted changes, the authors have used random rotation, horizontal flip, contrast variation, and zoom techniques. They have also added a Self, Calibrated Illumination (SCI) module to tackle light unevenness in the images before the model makes the prediction. To ensure the fair evaluation of the model, the dataset has been split into training, validation, and testing parts.

The fundamental disease identification module is based on a Hybrid Mobile Vision Transformer (MobileViT) system that integrates convolutional neural networks for local feature extraction with transformer blocks for global contextual modelling. At first, convolutional layers derive very detailed spatial features like blemishes, discoloration, and texture variations. Next, these features are divided into patches and passed through light transformer layers using self, attention mechanisms that recognize long, range dependencies over the leaf surface. The combination of CNN, derived local features and transformer, based global representations allows the model to deliver high classification accuracy and at the same time, remain computationally efficient enough for web, based deployment. The model's training was carried out with Adam optimizer and categorical cross, entropy loss, and its performance was measured in terms of accuracy, precision, recall, and F1, score. To avoid overfitting, regularization methods such as dropout and early stopping were used.

To extend the system to other than just classifying diseases, a new attention, based Sequence, to, Sequence (Seq2Seq) LSTM model was developed to automatically generate remedies. The encoder converts the predicted disease label and the accompanying context information into a hidden representation, while the decoder outputs a nicely structured, natural language treatment plan, step, by, step. The attention mechanism is used to highlight relevant encoded features during text generation in a dynamic way, thus helping the text to be coherent and contextually relevant. In order to improve the quality of the output, beam search decoding was used. The remedy, generating model was trained on a selected agricultural knowledge base that contained disease, remedy pairs, and the performance was judged using the BLEU score, which determines how similar the model output is to the expert, written recommendations.

Moreover, in order to make the system more directly useful, a Smart Plant Manager module was added to the framework. This unit uses GPS to locate the user and obtains weather data in real, time through the weather APIs. Environmental factors such as temperature, humidity, and rainfall are first used to calculate disease risk scores, which then serve as the basis for preventive alert generation. The addition of environmental intelligence allows the system to move from merely reacting to symptoms to actually identifying and treating issues before they happen.

The entire system was deployed as a Flask, based web application with a MySQL database for secure storage of user profiles, plant records, diagnosis histories, and generated recommendations. The backend orchestrates image preprocessing, model inference, remedy generation, and contextual analysis, ensuring an average inference time of 23 seconds per image on standard server hardware.

V. SYSTEM ARCHITECTURE

The Next, Generation Plant Management System (NG, PMS) envisioned is a modular, full, stack architecture that integrates deep learning-based disease detection, natural language remedy generation, and context, aware advisory services in a scalable Web framework. The plant management system follows a typical layered architecture as the Input Layer, AI Processing Layer, Context Integration Layer, and User Access Layer.

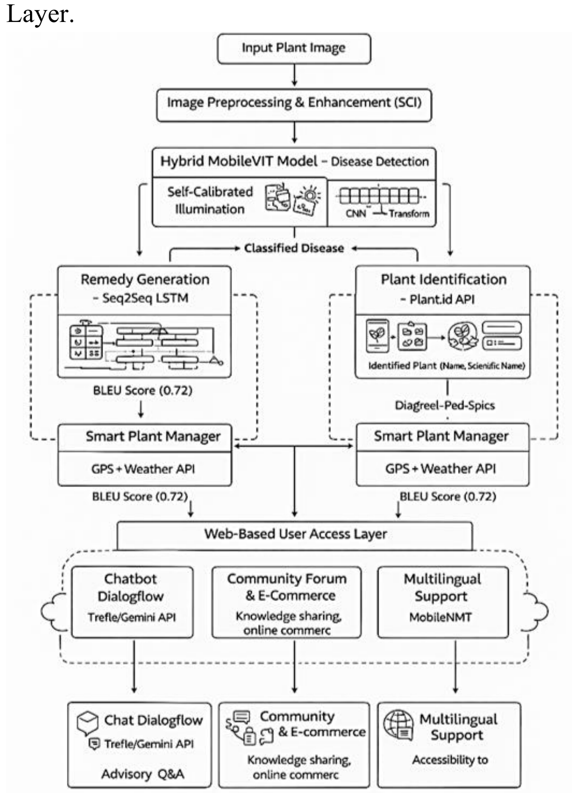

Fig. 1. IoT components of smart inductive charging

A. Input and Preprocessing Layer

The user initiates the process by uploading a plant leaf image to the web interface. The Flask, based backend not only validates the input format (JPG/PNG) but it also reprocesses this operation by resizing, normalization, and illumination enhancement using a

Self, Calibrated Illumination (SCI) module. Thus, these steps prepare the system to be less sensitive to variations in lighting, complexity of the background, and leaf orientation.

B. AI Processing Layer
The backbone of the system's intelligence is constituted by a pair of synergized deep learning models:

1. Hybrid MobileViT Model (Disease Detection Engine):

The processed image is delivered to a Hybrid Mobile Vision Transformer (MobileViT) model. This architecture mixes convolutional layers that work on extracting local features through the lens and transformer blocks that provide the global context. The output of the model is a probability distribution over various plant disease classes along with a confidence score. The combination of this hybrid design achieves a high classification accuracy yet at the same time keeps the computational efficiency at a level that allows it to be deployed on a web platform.

2. Seq2Seq LSTM Model (Remedy Generation Engine):

When the disease is identified, the disease label is sent to an attention, based Sequence, to, Sequence (Seq2Seq) LSTM model. The encoder takes the disease embedding while the decoder produces the treatment procedure in the form of well, structured, step, wise instructions. The attention mechanism allows the model to zoom in on the most relevant pieces of information from the context that are necessary for generating text, thus, the remedies produced are sensible and ready to be put into practice.

C. Context, Aware Integration Layer
To boost the system's versatility in real, world scenarios, it leverages a mix of background and situational intelligence:
Plant Identification Module: Relies on external APIs for checking plant species and thus provides botanical metadata. Smart Plant Manager: Features GPS, based location detection and obtains real, time weather data via REST APIs. Environmental factors such as temperature, humidity, and rainfall are analyzed to raise disease risk alerts and issue preventive recommendations. By combining such different modalities, the system is capable of evolving from merely reacting to a disease diagnosis to maintaining plant health in a more preventive manner.

D. Data Management Layer
All system interactions including user profiles, diagnosis history, plant records, and generated remedies are saved in a MySQL database. The backend facilitates secure authentication, session management, and smooth query handling. The modular database architecture allows for scalability and future extension.

E. User Access and Interface Layer
The web interface is designed to be responsive and is created with HTML, CSS, and JavaScript. It displays: Disease diagnosis with confidence score AI, predicted treatment steps Location, based recommendations Historical plant health records Besides these, there are: Chatbot Integration: Allows interactive query handling through conversational AI services. Multilingual Support: Employs translation models to facilitate access to a wider range of users. Community and Knowledge Sharing Module: Fosters collaboration in learning and advisory exchange.

VI. RESULTS AND DISCUSSION

After fine, tuning for 50 epochs, the suggested Hybrid MobileViT, based plant disease detection system was tested through confusion matrix and training, validation performance curves. The test results exhibit an excellent classification ability, a stable convergence pattern, and successful generalization over the whole spectrum of the diseases. The figure of 96.3% was recorded by the model as overall classification accuracy, coupled with high precision and recall scores that are even throughout, thus, the hybrid architecture's strength in multi, class plant disease detection is proven.
The confusion matrix resulting from the fine, tuning shows a very clear diagonal dominance which is an indication that the majority of the samples were correctly classified into their respective disease categories. The main categories such as Tomato Yellow Leaf Curl Virus, Tomato Septoria Leaf Spot, Tomato Spider Mites, Pepper Bell Healthy, and Tomato Healthy display near, perfect classification performance.

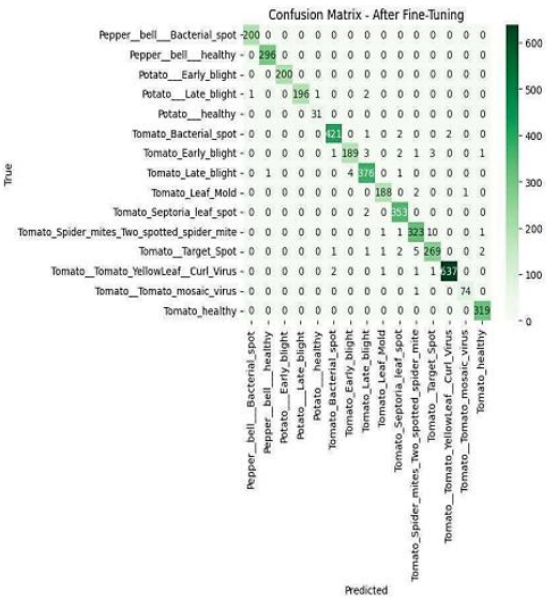

Fig. 3. IoT components of smart inductive charging

These results clearly demonstrate that the model is very capable of accurately recognizing the diseases, the disease, specific visual features such as the pattern of lesions, discoloration, texture variation, and the spatial distribution of these features across the leaf surfaces.

Minor fragment mix, ups can hardly be recognized in the major classes that are very similar in appearance. These are Tomato Target Spot and Early Blight, Tomato Mosaic Virus and Yellow Leaf Curl Virus, along with a few cases of confusion between healthy leaves and early, stage infected ones. Such errors are very scarce and, indeed, unavoidable, in complex agricultural datasets of symptoms that visually overlap. Importantly, no category exhibits ongoing or severe misclassification, which is proof of balanced learning and absence of model bias.

The training and validation accuracy loss curves serve as additional evidence for the effectiveness of the proposed architecture. Training accuracy rises very fast at the beginning of the epochs and then it slowly levels off near convergence, whereas validation accuracy follows a very similar progression and the two curves do not significantly diverge. This tight correspondence between training and validation accuracy is a strong indication of excellent generalization ability and very little overfitting. In line with this, both training and validation loss initially drop sharply and then keep gradually decreasing to values close to zero. The gentle and consistent trends of the loss curves indicate stable optimization, correct learning rate choice, and the implementation of the right regularization techniques. No sudden jumps or instability are detected during training which suggests that the learning process was consistent throughout the fine, tuning.

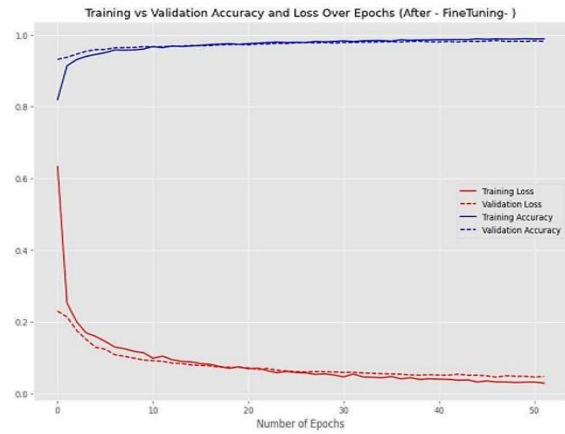

Fig. 4. smart inductive charging admin dashboard

In a nutshell, the Hybrid MobileViT network architecture has been experimentally proven to combine local feature extraction and global contextual modelling effectively resulting in very accurate multi, class classification while keeping the convergence stable. The confusion matrix as well as the training validation curves reflect that the model is trustworthy, well, generalized, and capable of being used in real situations for plant disease detection.

VII. CONCLUSION

This study introduced a Next, Generation Plant Management System integrating hybrid deep learning, natural language generation, and context, aware advisory mechanisms for accurate plant disease detection and treatment planning. The innovative solution exploits a Hybrid Mobile Vision Transformer (MobileViT) design to deliver precise and computationally cost, effective multi, class plant disease classification.

Testing on the Plant Village dataset has shown that the system can accurately classify different plant diseases with an accuracy of 96.3%, indicating the successful integration of both convolutional local feature extraction and transformer, based global contextual modelling. The vehicle for delivering agricultural instructions based on the diagnosis was an attention, based Sequence, to, Sequence (Seq2Seq) LSTM model capable of producing well, organized, step, wise treatment plans. The remedy generation component scored a BLEU score of 0.72 and exhibited a high level of factual integrity when checked against expert knowledge bases.

Moreover, by tapping into geolocation and real, time weather data, the Smart Plant Manager component was able to generate preventative warnings and offer cultivation tips tailored to the local environment, thus, the system's real, world effectiveness has been significantly elevated.

The deployment of a web, based integrated system among others allows quick inference, expansion, and usage in the real, world agricultural field setting. When accurate disease classification, prescription, based intelligent remedy generation, and environmental context analysis are combined and put in a single framework, the proposed system not only goes beyond traditional tools of diagnosis but also serves as a comprehensive agricultural decision, support platform.

In the future, extending the dataset for more crop species and rare diseases, combining IoT, based soil and environmental sensors, strengthening the system's robustness under the real, field conditions, and creating mobile and offline, capable versions for better accessibility in low, connectivity areas will be the major focus of the work. In general, the proposed method possesses a good potential to be a major factor in precision agriculture, the implementation of sustainable farming practices, and the improvement of crop health management.

REFERENCES

[1] M. Zhang, A. Kumar, and S. Lee, "Real-Time Driver Drowsiness Detection Using a Lightweight CNN and Facial Landmark Analysis," IEEE Transactions on Intelligent Vehicles, vol. 8, no. 2, pp. 1234–1245, Jun. 2023.

[2] S. Mehta and M. Rastegari, "MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer," in Proceedings of the International Conference on Learning Representations (ICLR), 2022.

[3] P. Singh, A. Sharma, and R. Gupta, "PlantDoc-DS: A Robust Dataset and Benchmark for Vision-Based Plant Disease Detection," Scientific Data, vol. 10, no. 1, p. 185, Mar. 2023.

[4] L. Chen, H. Wang, and D. Kim, "A Real-Time Plant Disease Detection System Using an Optimized YOLOv5 Model on Edge Devices," Computers and Electronics in Agriculture, vol. 219, p. 108793, Apr. 2024.

[5] T. Fernandez and J. Park, "From Diagnosis to Prescription: A Seq2Seq Approach for Generating Agricultural Remedial Actions," ACM Transactions on Intelligent Systems and Technology, vol. 15, no. 3, pp. 1–24, May 2024.

[6] R. Kumar and P. Li, "Vision Transformer for Fine-Grained Plant Disease Classification: A Comprehensive Study," IEEE Access, vol. 11, pp. 45672–45685, 2023.

[7] S. Iyer, M. Garcia, and K. Tanaka, "AgriBot: A Context-Aware Conversational Agent for Farmer Assistance," Frontiers in Artificial Intelligence, vol. 8, p. 67, Jan. 2025.

[8] B. Schmidt and A. Patel, "Leveraging Multimodal Data for Improved Plant Health Forecasting," Agricultural and Forest Meteorology, vol. 354, p. 110456, Nov. 2024.

[9] Y. Wang, X. Liu, and S. Raj, "Few-Shot Learning for Recognition of Rare Plant Diseases," IEEE Geoscience and Remote Sensing Letters, vol. 22, pp. 1–5, Jan. 2025.

[10] D. P. Hughes and M. Salathé, "An Open Access Repository of Images on Plant Health to Enable the Development of Mobile Disease Diagnostics," arXiv preprint arXiv:1511.08060, 2015.

[11] A. Vaswani et al., "Attention Is All You Need," in Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2017, pp. 5998–6008.

[12] A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv preprint arXiv:1704.04861, 2017.

[13] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proceedings of the International Conference on Machine Learning (ICML), 2019, pp. 6105–6114.

[14] S. Liu et al., "Deep Learning in Plant Phenomics and Disease Detection: A Comprehensive Survey," IEEE Transactions on Geoscience and Remote Sensing, vol. 61, pp. 1–18, 2023.

[15] K. P. Ferentinos, "Deep Learning Models for Plant Disease Detection and Diagnosis," Computers and Electronics in Agriculture, vol. 145, pp. 311–318, 2018.