

# Deep Learning-Based Corn Disease Classification Using VGG16 Transfer Learning

Alaa Kh. I. AlDammagh, Samy S. Abu-Naser

Department of Information Technology,

Faculty of Engineering & Information Technology,

Al-Azhar University - Gaza, Palestine

**Abstract:** Corn (maize) is a critical global crop, yet its production is significantly threatened by various leaf diseases, leading to substantial economic losses and food security concerns. Traditional manual disease detection methods are time-consuming, subjective, and impractical for large-scale agricultural applications. This study presents a deep learning-based approach for automated classification of corn leaf diseases using transfer learning with the VGG16 convolutional neural network architecture. The dataset comprises 5,976 images across four categories: Blight (1,303 images), Common Rust (1,047 images), Gray Leaf Spot (1,643 images), and Healthy leaves (1,983 images), standardized to 256×256pixel resolution. The methodology employs a two-phase training strategy: initial training with a frozen VGG16 base followed by fine-tuning of the last eight layers using a reduced learning rate. Real-time data augmentation techniques including random rotations, width and height shifts, and horizontal and vertical flipping were applied to enhance model generalization and mitigate overfitting. The model achieved an overall test accuracy of 95.52%, correctly classifying 1,419 out of 1,494 test images. Per-class evaluation demonstrated exceptional performance for Healthy leaves (100% precision, 99.80% recall) and Common Rust (96.59% precision, 97.70% recall), with F1-scores of 0.9990 and 0.9714, respectively. The Blight and Gray Leaf Spot classes achieved F1-scores of 0.9053 and 0.9318, with minimal confusion attributed to visual similarities between lesion patterns. The minimal gap between training and validation loss curves confirmed robust generalization without significant overfitting. These findings demonstrate that transfer learning with VGG16 provides a reliable, computationally efficient solution for automated corn disease diagnosis, offering practical potential for deployment in resource-constrained agricultural settings through mobile or smartphone-based applications to support farmers in early disease intervention and crop management.

**Keywords:** Corn Disease, Deep Learning, Convolutional Neural Networks (CNNs), DenseNet201, ImageNet, VGG19, Image Classification

## 1. Introduction:

Corn, commonly known as maize, is a crop that can be used in a variety of ways. It grows in desert, tropical, and subtropical climates. It is the world's second most widely farmed crop. With a global average yield of more than 5 tons/ha, more than 170 nations are producing approximately 1147.7 MM tons of maize on a total area of 193.7 MM. Global consumption of maize primarily involves feeding (61%), foodstuffs (17%), and the manufacturing industry (22%). With 83% of worldwide production going towards animal feed, starch, and biofuels, it has become a major commercial crop [1]. It is inherently vulnerable to numerous diseases that can harm the plant's leaves, trunk, and fruit at any stage of development [2]. This can result in significant financial loss because it directly affects the corn harvest's yield. Globally, food shortages, hunger, or even famine may result from reduced agricultural production of staple crops like maize. The most serious of these diseases are those that affect corn leaves when they are growing. Thus, three prevalent leaf diseases are examined in this work. Cercosporin leaf spot, common rust, and northern leaf blight [3].

This research introduces an innovative convolutional neural network framework that leverages deep learning to automate the detection of three major corn leaf disease Blight, Gray Leaf Spot, and Common Rus alongside a healthy control class. By transitioning from labor-intensive manual inspections to an automated system, the framework employs transfer learning via a pre-trained VGG16 model[4-7], originally trained on ImageNet. The architecture is specifically modified by replacing its top layers with a GlobalMaxPooling2D layer and a dense softmax classifier [8-10]. To ensure high reliability and combat overfitting, the model is trained for 20 epochs using the Adam optimizer and extensive data augmentation such as rotation, shifting, and flipping—applied to a dataset partitioned into 80% training and 20% testing, with a 20% validation split [11-14]. Performance is monitored through a checkpoint callback that saves the optimal weights based on validation loss, ultimately demonstrating a robust capability for agricultural diagnostics.

The contributions of this paper are as follows:

- Using corn leaf images as input, develop a deep transfer learning model based on VGG16 for the classification of corn diseases. Three diseases were considered in this study: Blight, Gray Leaf Spot, and Common Rust. In addition, a fourth healthy class was included.
- Implement transfer learning with the VGG16 convolutional neural network architecture, fine-tuning the model for the four-class classification task.
- Evaluate the performance of the model using accuracy metrics on a held-out test set, and report training and validation curves. The results show that deploying such models on mobile or smartphone devices can aid farmers in quickly and correctly identifying diseases.




Figure 1. Sample images of the Corn Leaf

## 2. Related work

The role of deep learning in agriculture, especially in detecting corn leaf diseases, has been growing rapidly over the years, the authors in [15] proposed a DenseNet121 optimization for the categorization of maize diseases. They gathered 12,332 with 250x250 resolution photos from various locations. They trained the improved model against four well-known CNN models: VGG, EfficientNet, XceptionNet, and NASNet. The accuracy of the suggested optimal model in classifying photos of maize leaves was 98.06%. Karmacharya et al. [5] evaluated DenseNet201 and Xception for corn leaf disease classification using a Mendeley dataset of 4,721 images across four classes (Common Rust, Northern Leaf Blight, Gray Leaf Spot, and Healthy). Their methodology incorporated data augmentation and transfer learning with ImageNet-pretrained weights. DenseNet201 achieved superior performance with 98.69% test accuracy and precision, recall, and F1-scores above 98%, while Xception attained 96.61% accuracy with better computational efficiency. Both models demonstrated excellent discriminatory capability with near-perfect AUC scores. The study highlights the effectiveness of modern CNN architectures for automated disease detection and provides valuable benchmark comparisons for transfer learning approaches in agricultural applications.

In [16], Common rust infections in the early, middle, and late stages of maize leaf development were identified using the VGG16 model. They used fuzzy decision criteria to extract features and Otsu's thresholds for image segmentation. On the validation set, the model's accuracy was 95.63%, while on the testing set, it was 89%. The VGG19 model's initial pre-trained layers were kept by [17], who substituted a convolutional bloc which comprises of a convolution layer, a batch normalization layer, and Swish activation for the end layers. Two Inception blocks come after this convolutional block.

In [18], they conducted a comprehensive study on corn leaf disease classification using deep transfer learning, focusing on three diseases (Cercospora leaf spot, common rust, and northern leaf blight) plus healthy leaves across 3,852 images. They evaluated ten CNN architectures through transfer learning with rigorous methodology involving three data split ratios (50/50, 70/30, 90/10) and ten repeated runs per experiment. DarkNet-53 achieved the highest accuracy at 98.6% with 90% training data, while ResNet-18 offered the best speed-accuracy compromise (96.9-97.8% accuracy, 0.724 ms inference time). The study identified primary misclassification between northern leaf blight and Cercospora leaf spot due to visual similarities, providing valuable benchmarks for transfer learning approaches in agricultural applications.

In [19], they developed an end-to-end deep learning model for corn leaf disease classification using feature fusion from multiple pre-trained CNNs (EfficientNetB0 and DenseNet121) to balance accuracy and computational efficiency. Using 15,408 Plant Village images across four classes (Northern Leaf Blight, Common Rust, Gray Leaf Spot, and Healthy), they applied data augmentation and an 80/20 train-test split with 20% of training samples for validation. Their dual-branch architecture fused extracted features via concatenation, achieving 98.56% accuracy with 16.20 million parameters outperforming larger models like ResNet152 (98.37%) and InceptionV3 (96.26%). The model achieved perfect scores for Common Rust and Healthy classes, with minor misclassifications between visually similar Gray Leaf Spot and Northern Leaf Blight. This work demonstrates that feature fusion from multiple architectures produces robust models exceeding single-network performance, offering valuable insights for optimizing feature extraction and architectural design in resource-constrained applications.

In [20], they demonstrated the fundamental utility of convolutional neural networks (CNNs) in building robust classification models for leaf diseases. Expanding on this, in [21] they introduced a comprehensive deep learning-based strategy specifically tailored for diagnosing various maize crop illnesses. Furthermore, the shift toward computational efficiency and pre-trained expertise was highlighted by [22], who endorsed the application of deep transfer learning as a superior approach for the precise classification of maize plant leaf infections. In [23], they sought to improve picture classification using the ImageNet dataset and a potent CNN. With top-5 error rates of 17.0% on ImageNet and 15.3% in the ILSVRC-2012 competition, the suggested model outperformed earlier techniques and produced toptier results. The study used dropout, data augmentation, and a very effective GPU implementation for training in order to overcome overfitting issues. In addition to discussing the possibility of additional advancement with faster GPUs and larger datasets, the research highlights the significance of depth in attaining higher performance. Mohanty et al. (2023) conducted a field study in Bangladesh using a unique, field-sourced dataset of 4,800 maize leaf images (Healthy, Common Rust, Gray Leaf Spot, and Blight) that captured real-world agricultural variability. After preprocessing with resizing, color normalization, and data augmentation, they evaluated eight deep learning architectures, with DenseNet121 achieving 99.22% accuracy and VGG19 attaining 99.44% accuracy. Their hybrid model combining ResNet50 and VGG16 features achieved exceptional performance at 99.65% accuracy, demonstrating the efficacy of leveraging complementary architectural strengths. The study underscores the feasibility of deploying such models for real-time disease detection through smartphone applications in resource-constrained agricultural settings, providing valuable validation for VGG family architectures on field-sourced data.

In [24], they explored transfer learning and data augmentation to improve corn leaf disease detection, addressing the limitations of manual observation which often leads to misdiagnosis and reduced production efficiency. Using the PlantVillage dataset, they classified images into four categories: Corn Common Rust, Northern Leaf Blight, Gray Leaf Spot, and Healthy leaves. Their optimized CNN framework refined a pre-trained GoogLeNet network by fine-tuning parameters such as learning rate and optimizer. When evaluated against other transfer learning architectures including ResNet18, VGG16, and VGG19, their model achieved an average accuracy of 97.6%, with each disease category exceeding 95% accuracy. Notably, their optimized model demonstrated a 5.9% accuracy improvement over the original GoogLeNet model. In [15], they introduced a Multi-Activation Function (MAF) module to enhance CNN performance for maize leaf disease detection, addressing limitations of single-activation-function networks. Using a field dataset of 4,428 images (sheath blight, rust, northern leaf blight, and healthy leaves) captured under diverse real-world conditions, they expanded it to 89,420 samples through geometric transforms, grayscale processing, morphological operations, Snapmix, Mosaic, and DCGAN-based augmentation. The MAF module replaces single activation layers with parallel pathways combining Tanh, ReLU, LeakyReLU, Sigmoid, and Mish functions using learnable weights that sum to unity. Evaluating multiple architectures, MAF-enhanced models consistently outperformed baselines, with ResNet50 combining Sigmoid, ReLU, and Mish achieving 97.41% accuracy a 2.33% improvement. Additional innovations included warm-up scheduling, label-smoothing, and Bi-Tempered logistic loss for robustness. An iOS-based mobile detection system validated practical applicability, demonstrating the value of activation function optimization and comprehensive augmentation strategies.

In [26], they proposed MaizeNet, a deep CNN architecture designed specifically for maize leaf disease classification, addressing challenges in capturing subtle disease features and maintaining training stability on small agricultural datasets. Using a dataset of 4,188 images across four classes (blight, common rust, gray leaf spot, and healthy), combining lab-controlled (91.6%) and real-world field images (8.4%), they implemented preprocessing including resizing to 224×224 pixels, normalization, and augmentation. The key innovation integrates channel-spatial squeeze-excite (csSE) attention mechanisms with residual learning, enabling dynamic feature recalibration both channel-wise and spatially while alleviating vanishing gradients through strided residual blocks. The architecture comprises an initial block, four attention-based residual blocks with increasing filters (32 to 256), global average pooling, and fully connected layers with dropout, totaling 2.1 million parameters. Through five-fold cross-validation with macro-averaged metrics to address class imbalance, MaizeNet achieved 95.95% accuracy with precision of 0.9528, recall of 0.9497, and F1-score of 0.9509. Ablation studies confirmed the contribution of both attention and residual components, with gray leaf spot presenting the greatest classification challenge. The research demonstrates how attention mechanisms and residual learning enhance feature extraction for agricultural classification and underscores the importance of rigorous cross-validation for developing robust disease detection systems.

### **3. Materials and Methods**

Utilizing a dataset comprising various corn leaf pathologies alongside healthy specimens, this research aims to accurately categorize leaf images into their respective health states. This study adopts a transfer learning methodology, which involves adapting generically pre-trained deep learning architectures to a specific task by modifying output layers and fine-tuning the model on target data. The fundamental principle of this approach lies in the hierarchical nature of feature detection: early layers identify universal low-level characteristics, such as textures and edges, while deeper layers capture intricate, domain-specific patterns.

By leveraging these pre-existing weights rather than initializing from scratch, transfer learning significantly accelerates the training process and reduces the volume of data required to achieve high performance. This strategy has demonstrated exceptional versatility and effectiveness across a wide range of scientific and industrial disciplines [27-30].

### 3.1 Dataset

The dataset consists of 5976 corn leaf images categorized into four distinct classes based on their health status: Gray\_Leaf\_Spot (1643 images), common rust (1047 images), blight (1303 images), and healthy specimens (1983 images). These single-leaf images are provided in JPEG format, each standardizing at a resolution of  $256 \times 256$  pixels. Data-preprocessing.

In the realm of ML and DL, data pre-processing is a vital stage that prepares the raw data to be processed by the model. This stage ensures that the data meet the quality and format required for meaningful learning and accurate predictions. Here is a detailed elaboration of the pre-processing steps we used for our study[31-34].

### 3.2 Data preprocessing

- **Image Re-sizing:** Image resizing was performed using the `load_img` function from the `keras.preprocessing.image` module, with the `target_size` parameter set to  $256 \times 256$  pixels. This function utilizes the PIL (Python Imaging Library) backend to resize all input images to a uniform spatial dimension. Standardizing image dimensions is a critical preprocessing step in convolutional neural network-based classification, as it ensures consistent input tensor shapes across the dataset, allowing the model to learn discriminative features from leaf images irrespective of their original resolutions. This uniformity is particularly essential for compatibility with the pre-trained VGG16 architecture, which requires fixed-size inputs, and facilitates efficient batch processing during both training and inference stages[35-37].
- **Color Normalization:** To address variations in illumination conditions during image acquisition, color normalization was implemented through two complementary preprocessing steps. First, pixel intensity values were rescaled to the range  $[0,1]$  by dividing by 255.0, a standard min-max normalization technique that ensures uniform input scaling across all images. Second, for compatibility with the pre-trained VGG16 architecture, channel-wise standardization was applied using the mean and standard deviation statistics derived from the ImageNet dataset. This normalization was performed via the `preprocess_input` function from `keras.applications.vgg16`, which centers the input distribution and aligns it with the statistical properties expected by the pre-trained model. The combination of rescaling and dataset-specific standardization effectively mitigates the impact of varying lighting conditions, enabling the model to prioritize disease-relevant morphological features over low-level intensity variations. This preprocessing strategy is essential for leveraging transfer learning, as it ensures consistency between the input data and the distribution on which the base model was originally trained[38-42].
- **Image Augmentation:** To enhance the diversity and robustness of the training dataset, data augmentation was applied using the `ImageDataGenerator` class from the Keras library. This technique artificially expands the training set by generating transformed versions of the original images through a series of random geometric and photometric distortions. Specifically, the augmentation pipeline included random rotations within a range of 30 degrees, width and height shifts by a factor of 0.2, and horizontal and vertical flipping. These transformations simulate the natural variability observed in real-world agricultural environments, such as changes in leaf orientation, camera perspective, and field conditions. The augmentations were applied in real-time during model training, ensuring that each epoch received slightly modified versions of the input images, which helps mitigate overfitting and improves the model's generalization to unseen data. The dataset employed in this study comprises corn leaf images collected from diverse field locations, encompassing variations in lighting, leaf positioning, and background complexity. The combination of rigorous preprocessing and real-time data augmentation ensures that the input data are both consistent and representative of the challenges inherent in automated plant disease classification under realistic conditions.

### 3.3 Deep Learning and Convolutional Neural Network (CNN)

- CNN layers explanation

As a prominent subfield of machine learning, deep learning has fundamentally transformed image classification by utilizing sophisticated neural network architectures to reach state-of-the-art performance levels. Within this paradigm, Convolutional Neural Networks (CNNs) have emerged as the leading instrument for computer vision tasks due to their unique ability to autonomously extract hierarchical spatial patterns directly from raw pixel data. The standard architectural pipeline of a CNN begins with convolutional layers, which serve as the primary engine for feature extraction by identifying increasingly complex visual structures. These are typically followed by pooling layers that perform essential dimensionality reduction, condensing the gathered information to enhance both computational efficiency and model robustness. Finally, the network concludes with fully connected layers that synthesize these high-level features to execute the definitive classification. [43-44].

---

In this study, Convolutional Neural Networks (CNNs) serve as the foundational architecture for corn leaf disease classification, enabling automated diagnosis through the learning of complex, discriminative patterns from images of both diseased and healthy leaves. However, the effective generalization of deep learning models typically necessitates large-scale annotated dataset a condition that is often difficult to satisfy in agricultural domains where data collection is constrained by environmental variability, seasonal limitations, and annotation costs. To address this challenge, the present work employs a transfer learning strategy coupled with data augmentation. Specifically, a pre-trained VGG16 model, originally trained on the large-scale ImageNet dataset, is fine-tuned on the target corn leaf dataset. This approach leverages the hierarchical features learned from natural images, allowing the model to adapt to the domain-specific task with limited training samples. Concurrently, real-time data augmentation via the Keras ImageDataGenerator is applied to introduce controlled variability through random rotations, shifts, and flips, thereby enhancing model robustness and mitigating overfitting. The combination of transfer learning and data augmentation facilitates effective adaptation of the VGG16 architecture to the corn disease classification task despite the constraints of dataset size[45-48].




Figure 2: CNN Architecture [49]

The main layers of a CNN are as follows:

The typical architecture of a CNN consists of several key layers, including the Input Layer, Convolutional Layer, RELU Layer, Pooling Layer, and Fully Connected Layer. Each of these layers plays a distinct role in transforming the input image data into meaningful features that can be used for classification.

- Convolution, Activation, pooling, normalization, dropout, dense, softmax layers

#### Convolutional Layer

The Convolutional Layer is crucial for extracting features from the input images. It applies small filters, often in stacks, to capture spatial hierarchies. For instance, a series of three connected 3x3 convolutional layers allows neurons in subsequent layers to gain increasingly larger receptive fields, enabling them to interpret a broader context of the input image [50]. Very deep convolutional networks for large-scale image recognition. This approach, which uses multiple smaller filters instead of a single large one, significantly reduces the number of parameters and enhances the efficiency of learning.

#### Activation (ReLU) and pooling Layers

Following the Convolutional Layer, the RELU (Rectified Linear Unit) Layer applies a non-linear activation function, introducing non-linearity into the model, which is essential for learning complex patterns [20]. Very deep convolutional networks for large-scale image recognition. The Pooling Layer subsequently reduces the spatial dimensions of the feature maps, preserving the most critical information while minimizing computational load.

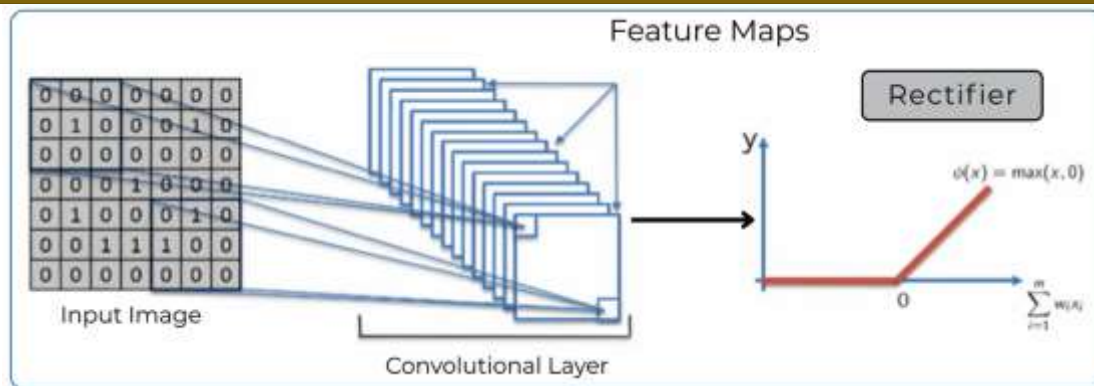


Figure 3: Activation Layer [51]

### Fully Connected Layer

Finally, the Fully Connected Layer integrates the features learned in the preceding layers to make predictions about the input data. This layer connects every neuron from the previous layer to each neuron in the current layer, culminating in the output probabilities for classification tasks. ImageNet classification with deep convolutional neural networks. [52]

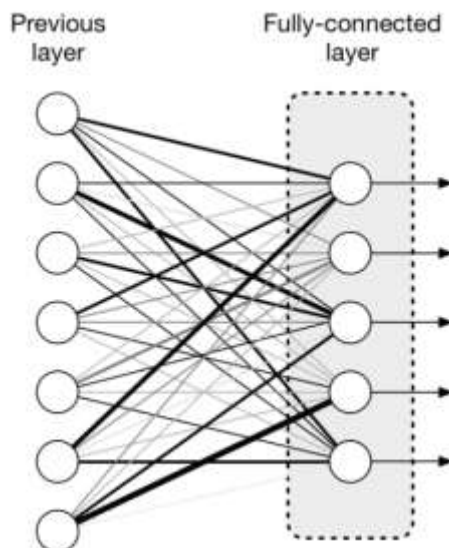


Figure 4: Fully Connected Layer. [53]

### Normalization Layer

Normalization techniques like batch and layer normalization standardize inputs or outputs to stabilize training and speed up convergence [54].

### Dropout Layer

Dropout acts as a regularization mechanism that prevents overfitting by randomly deactivating a subset of neurons throughout the training phase. By temporarily removing these connections, the technique discourages the network from developing a fragile dependency on specific individual neurons, thereby compelling the architecture to learn more robust, generalized features. [55]

### Dense Layer

Acting as the final stage of the architecture, the dense layer integrates the features identified by preceding layers to generate predictions. Positioned after the convolutional and pooling sequences, it processes flattened feature maps to map the high-level data to the final output classes. [56]

---

## Softmax Activation Function

The Softmax function transforms network outputs into a normalized probability distribution, making it essential for multiclass classification. Each value in the output vector represents the likelihood of the input belonging to a specific class, with all values summing to 1. This allows the model to identify the most probable category by establishing distinct decision boundaries for final classification. [57]

In terms of mathematics, it is represented as:

$$s(x_i) = \frac{e^{x_i}}{\sum_{j=1}^n e^{x_j}}$$

### 3.4 Transfer Learning

Transfer learning is a pivotal technique in deep learning that repurposes knowledge from pre-existing models to address new tasks, particularly when datasets are limited. By leveraging insights gained from comparable domains, this approach enhances predictive accuracy while significantly reducing the time, computational power, and high costs associated with labeling new data. Ultimately, transferring expertise from one domain to another facilitates the development of more robust and reliable machine learning systems. [58]

### 3.5 Model Architecture

The model utilizes the **VGG16** architecture as a feature extractor, initialized with weights pre-trained on the ImageNet dataset. To adapt the model for corn disease classification, the top fully connected layers were replaced with a custom classification head consisting of a **GlobalMaxPooling2D** layer followed by a **Dense** output layer with 4 neurons using the **softmax** activation function. Initially, the VGG16 base was frozen to preserve learned features, while the final stage of training involved **fine-tuning the last 8 layers** of the base model to improve specialized feature extraction. The input dimensions were standardized to **256×256×3**, ensuring compatibility with the VGG16 architecture

### 3.6 Training Process

The training was conducted in two distinct phases. Initially, the model was trained with a frozen base using the Adam optimizer at a learning rate of .0001. In the second phase, fine-tuning was performed on the top 8 layers with a reduced learning rate of 0.00001 to ensure stable convergence. Categorical cross-entropy was employed as the loss function to handle the four disease classes. batch size of 32 was used across 20 epochs for each phase. To combat overfitting and improve generalization, real-time data augmentation was implemented via ImageDataGenerator, including: Random rotations up to 30 degrees. Width and height shifts of 20%. Horizontal and vertical flipping. A 30% validation split was maintained to monitor performance, and a ModelCheckpoint was used to preserve the weights that achieved the lowest validation loss.

### 3.7 Evaluation Metrics

To provide a comprehensive performance analysis, the model was evaluated using a testing set of 1,494 images. The primary metric, Overall Accuracy, measured the percentage of correctly identified corn leaf conditions. To assess class-specific performance and account for any potential imbalances, Precision, Recall, and the F1-score were calculated for each category. A Confusion Matrix was generated to visualize specific misclassification patterns between diseases[59-65]. Furthermore, learning curves for training and validation accuracy and loss were plotted to analyze the model's convergence behavior.

## 4. Results and Analysis

The implemented VGG16 model demonstrated high efficacy in classifying corn leaf diseases. Upon evaluating the test dataset, the model achieved an overall accuracy of 95.52%, correctly identifying 1,419 out of 1,494 images. The training process showed excellent convergence, as illustrated in Figure 5.

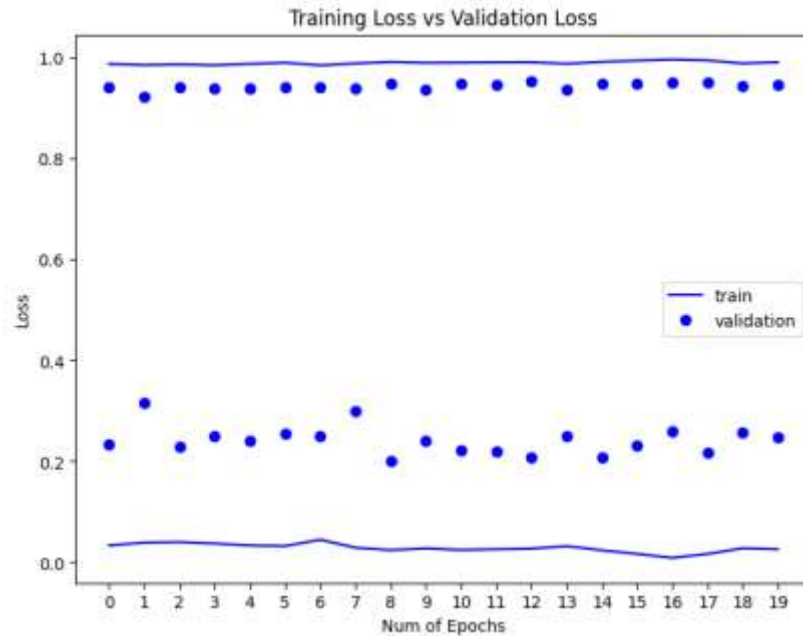


Figure 5: Training and Validation Loss Curves

The figure above demonstrates the model's stability throughout the 20-epoch fine-tuning phase. The training loss (solid line) reached near-zero levels, while the validation loss (blue dots) remained consistently low and stable, generally fluctuating between 0.20 and 0.30. This minimal gap between training and validation metrics indicates that the model generalizes well to new data and is not suffering from significant overfitting.

#### 4.1 Model Performance

The performance breakdown for the four classes is as follows:

Table 1: Model Performance metrics on the test dataset.

Class	Precision	Recall	F1-score	Correct/ Total
Blight	0.8879	0.9233	0.9053	301/ 326
Common Rust	0.9659	0.9770	0.9714	255/ 261
Gray Leaf Spot	0.9495	0.9148	0.9318	376/ 411
Healthy	1.0000	0.9980	0.9990	495/ 496

The model achieved a near-perfect result for the **Healthy** class with a precision of **100%**. **Common Rust** also showed exceptional results with an F1-score of **0.9709**. While **Blight** and **Gray Leaf Spot** had slightly more instances of confusion—as evidenced by the 38 instances where Blight was misclassified as Gray Leaf Spot—the model still maintained an F1-score above **0.89** for all categories, confirming the effectiveness of the fine-tuned VGG16 architecture.

#### 5. Conclusion

This research successfully implemented a deep transfer learning framework for the precise identification of corn leaf diseases using the VGG16 architecture. By freezing the initial layers and selectively fine-tuning the deep feature extraction layers, the model was able to bridge the gap between general ImageNet features and specialized agricultural pathology. The quantitative analysis confirms the model's high efficacy, yielding an overall accuracy of 95.52%.

The data analysis reveals that the model is highly sensitive to the Healthy and Common Rust classes, achieving near-perfect precision. The primary challenge identified was the visual overlap between Blight and Gray Leaf Spot, which accounted for the majority of misclassifications; however, the model still achieved F1-scores above 0.93 for these difficult classes. The stability of the training and validation loss curves further validates the robustness of the training strategy against overfitting.

## References

1. Abdallatif, R. F., et al. (2025). "Classification of Peppers Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 9(1): 35-41.
2. AbuEl-Reesh, J. Y. and S. S. Abu-Naser (2017). "An Expert System For Diagnosing Shortness Of Breath In Infants And Children." *International Journal of Engineering & Information Systems (IJEAIS)* 1(4): 102-115.
3. AbuJalambo, M., et al. (2026). "Spine Tumor Segmentation using Deep Learning: A Review." *Journal of Advanced Research Design* 136(1): 179-206.
4. Abu-Jamie, T. N. and S. S. Abu-Naser (2022). "Classification of Sign-Language Using Deep Learning by ResNet." *International Journal of Academic Information Systems Research (IJAISR)* 6(8): 25-34.
5. Abunasser, B. S., et al. (2025). Predictive Modeling of Underweight Malnutrition Using Neural Networks: Insights from Global Nutrition Datasets. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
6. Abunasser, B. S. and S. S. Abu-Naser (2025). Unleashing the Power of GPT-3: Revolutionary Applications in Natural Language Processing. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
7. Abu-Samra, F. Y. and S. S. Abu-Naser (2025). "Nuts Classification Using Deep Learning."
8. Al-Affifi, Y. A. and S. S. Abu-Naser (2025). "Cloud-Based Deployment of Knowledge-Based Systems: Architecture and Case Study." *International Journal of Academic Engineering Research (IAER)* 9(8): 159-165.
9. Al-Aydi, B. M. and S. S. Abu-Naser (2025). "Comparative Study of Traditional and AI-Enhanced Sorting Algorithms: QuickSort, MergeSort, HeapSort, and TimSort." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 70-79.
10. Al-Aydi, B. M. and S. S. Abu-Naser (2025). "Integrating NLP Techniques for Smarter Modern Knowledge-Based Systems." *International Journal of Academic Engineering Research (IAER)* 9(8): 95-99.
11. Albadrasawi, S. and S. S. Abu-Naser (2024). Machine and Deep Learning for Securing Traffic in Computer Networks. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
12. Albanna, R. N. and S. S. Abu-Naser (2025). "Classification of Nuts Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 9(6): 1-11.
13. Al-Bayed, M. H., et al. (2025). "Surveillance in the Age of AI: Navigating Ethical Boundaries and Human Rights."
14. Alborno, D. F., et al. (2025). "Artificial Intelligence in Drug Discovery: Unlocking New Pathways for Therapeutic Innovation."
15. AlDammagh, A. K. and S. S. Abu-Naser (2025). "Natural Language Processing in Modern Knowledge-Based Systems." *International Journal of Academic Engineering Research (IAER)* 9(8): 74-79.
16. Al-Daour, A. F., et al. (2020). "Banana Classification Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 3(12): 6-11.
17. AlDaya, D. K., et al. "Predicting Smoking-Associated Thyroid Dysfunction Using Explainable Machine Learning."
18. Aldaya, S. A. S., et al. "An Interpretable Machine Learning and Deep Learning Framework for Early Prediction of Chronic Kidney Disease Using Clinical Data."
19. Aldaya, S. A. S., et al. (2025). "Deep Learning-Based Classification of Bone Tumors Using Medical Imaging." *Interpretation* 9(12): 71-79.
20. Aldaya, S.-A. S. and S. S. Abu-Naser (2025). "Deep Learning For Grapevine Disease Detection." *International Journal of Academic Information Systems Research (IJAISR)* 9(6): 12-20.
21. Aldaya, S.-A. S. and S. S. Abu-Naser (2025). "Diagnosing Sprained Ankles Using Clips."
22. Alghalban, A. I. and S. S. Abu-Naser (2025). "Identifying Images of Chess Pieces Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 9(6): 51-55.
23. Alhaj, A. a. and S. S. Abu-Naser (2025). "Teeth Problem Diagnosis Expert System." *International Journal of Academic Engineering Research (IAER)* 9(8): 63-73.
24. AlJerjawi, N. S. and S. S. Abu-Naser (2025). "Image-Based Tomato Leaves Diseases Detection Using Deep Learning."
25. AlJerjawi, N. S. and S. S. Abu-Naser (2025). "A Rule Based System for Diagnosing Hypertension Problems." *International Journal of Academic Engineering Research (IAER)* 9(8): 129-147.
26. AlJerjawi, N. S. and S. S. Abu-Naser (2025). "AI-Assisted Multi-Criteria Sorting for Decision Support in Healthcare Systems." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 90-93.
27. AlJerjawi, N. S. and S. S. Abu-Naser (2025). "Neural Sorting Networks: Applying Deep Learning to Ranking and Sorting Tasks in NLP." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 86-89.
28. AlJerjawi, N. S. and S. S. Abu-Naser (2025). "SmartSort: An Intelligent Framework for Optimizing Sorting Efficiency Using AI." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 134-138.
29. Alkahlout, M. A. and S. S. Abu-Naser (2024). Advances in Kidney Cancer Detection: Harnessing the Power of Deep Learning for Accurate Diagnosis. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
30. Alkahlout, M. A. and S. S. Abu-Naser (2025). Thyroid Cancer Risk Classification Using Machine Learning and Deep Learning Techniques: A Comparative Study with Balanced Dataset Augmentation. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
31. Alkayyali, Z. K., et al. (2024). Classification of Cardiovascular ECGs Using MODWPT-Based Feature Extraction: A Comparative Study on Four Ailments from MIT-BIH Databases. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
32. Alkayyali, Z. K., et al. (2025). Comparative Analysis of Regressor Models for Predicting Heart Attack Risk: A Comprehensive Evaluation Using Regression Metrics and Visualization. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
33. Almoghrabi, A. and S. S. Abu-Naser (2025). "AI-Driven Adaptive Sorting Algorithms for Large-Scale Data Processing." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 155-161.
34. Almaziny, M. S. and S. S. Abu-Naser (2024). Detection and Classification of Faked and Genuine Money Using Deep Learning. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
35. Alqedra, H. I. and S. S. Abu-Naser (2025). "Knowledge Based System for Diagnosis Tomato Diseases." *International Journal of Academic Engineering Research (IAER)* 9(8): 100-110.
36. Alsaqqa, A. H. and S. S. Abu-Naser (2024). Comprehensive Analysis of Machine Learning and Deep Learning Algorithms for Phishing URL Detection. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
37. Alsaqqa, A. H. and S. S. Abu-Naser (2025). Detecting Cybersecurity Threats Using Convolutional Neural Networks and Machine Learning. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
38. AlTala, S. E. and S. S. Abu-Naser (2025). "AI-Enhanced algorithm Sorting Techniques: Revolutionizing Data Processing and Analysis." *International Journal of Academic Engineering Research (IAER)* 9(6): 44-47.
39. AlTala, S. E. and S. S. Abu-Naser (2025). "From Rules to Reasoning: Impact of NLP on Knowledge-Based Systems." *International Journal of Academic Engineering Research (IAER)* 9(8): 148-153.
40. Ashour, W. H. and S. S. Abu-Naser (2025). "Design and Development of a Clinical Diagnosis Expert System." *International Journal of Academic Engineering Research (IAER)* 9(8): 154-158.
41. Dwimah, A. and S. S. Abu-Naser (2025). "Enhancing Sorting Algorithms with Artificial Intelligence: A Hybrid Approach." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 64-69.
42. Dwimah, A. and S. S. Abu-Naser (2025). "Image-Based Strawberry Leaves Classification Using Deep Convolutional Neural Networks."
43. Dwimah, A. and S. S. Abu-Naser (2025). "Symbolic Hybrid Knowledge-Based Systems: Integrating Knowledge-Based Reasoning and Machine Learning in Explainable AI." *International Journal of Academic Engineering Research (IAER)* 9(8): 80-84.
44. Elmahmoum, A. S. A. and S. S. Abu-Naser (2025). Comparative Analysis of Data Balancing Techniques in Prostate Cancer Classification Using Machine Learning and Deep Learning. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
45. Ihlayyel, M. S. and S. S. Abu-Naser (2025). "Design and Evaluation of a Fuzzy Expert System for Early Detection of Breast Cancer." *International Journal of Academic Engineering Research (IAER)* 9(8): 116-120.
46. Ihlayyel, M. S. and S. S. Abu-Naser (2025). "Detection and Classification of Tomato Leaf Diseases Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 9(6): 21-28.
47. Kassab, M. K. I. and S. S. Abu-Naser (2025). "Image-Based Tea Leaves Diseases Detection Using Deep Learning." *International Journal of Academic Information Systems Research (IJAISR)* 9(6).
48. Kwaik, H. B. A. A. and S. S. Abu-Naser (2025). "Design and Development of a Knowledge-Based System for Medical Diagnosis." *International Journal of Academic Engineering Research (IAER)* 9(8): 111-115.
49. Massa, N. M. and S. S. Abu-Naser (2025). Predicting Breast Cancer Recurrence Using Machine Learning and Deep Learning Models: A Comparative Study. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
50. Megdad, M. M., et al. (2022). "Fraudulent Financial Transactions Detection Using Machine Learning." *International Journal of Academic Information Systems Research (IJAISR)* 6(3): 30-39.
51. Mezied, A. A. and S. S. Abu-Naser (2025). "Pepper Color Classification Using Deep Learning." *International Journal of Academic Engineering Research (IAER)* 9(8): 1-7.
52. Mezied, S. M. and S. S. Abu-Naser (2025). "Efficient Sorting of Financial Transactions Using Artificial Intelligence for Fraud Detection and Risk Assessment." *International Journal of Academic Information Systems Research (IJAISR)* 9(8): 147-154.
53. Mohaisen, B. M. and S. S. Abu-Naser (2025). "Expert System Design and Implementation for Medical Diagnostic Applications." *International Journal of Academic Engineering Research (IAER)* 9(8): 85-94.
54. Qandil, A. I., et al. (2021). "Factors Affecting Of Disputes Resolution in Workplace: UNRWA at Gaza as a Case Study." *International Journal of Academic Management Science Research (IJAMSR)* 5(2): 154-180.
55. Qandil, A. I., et al. (2021). "The level of Mediation Outcomes of Disputes Resolution in Workplace at UNRWA, Gaza." *International Journal of Academic Multidisciplinary Research (IJAMR)* 5(2): 310-327.
56. Qaoud, A. N. and S. S. Abu-Naser (2025). Deep Learning-Based Skin Cancer Classification and Localization: A Comprehensive Approach for Accurate Diagnosis and Localization of Skin Cancers. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
57. Quffa, A. and S. S. Abu-Naser (2025). "A Rule-Based Expert System for Cybersecurity Threat Detection: Evolution." Applications, and the Hybrid AI Paradigm 10.
58. Quffa, A. and S. S. Abu-Naser (2025). "A Rule-Based Expert System for Cybersecurity Threat Detection: Evolution, Applications, and the Hybrid AI Paradigm." *International Journal of Academic Engineering Research (IAER)* 9(8): 44-48.
59. Ruslan, S., et al. (2026). "Spine Tumor Segmentation Using Deep Learning: A Review." *benefits* 63(1): 271-298.
60. Salman, F. M. and S. S. Abu-Naser "Comparative Analysis of Deep Learning Architectures for Bone Fracture Detection: MobileNetV2 vs. ResNet50."
61. Samban, L., et al. (2021). "An Expert System for Knee Problems Diagnosis An Expert System for Knee Problems Diagnosis." no. August.
62. Taha, A. H. A. and S. S. Abu-Naser (2025). Predicting Loan Defaulters: A Comprehensive Analysis and Comparative Study of Machine Learning Algorithms Using a Large-Scale Loan Default Dataset. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
63. Taha, A. M., et al. (2025). Exploring Emotion Recognition Through EEG Brainwave Data: A Comparative Analysis of Machine Learning and Deep Learning Approaches. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.
64. Zarandah, Q. M., et al. (2024). Performance Evaluation of Machine Learning and Deep Learning Models for Respiratory Disease Prediction. *International Conference on Data Engineering and Communication Technology*, Springer Nature Singapore Singapore.
65. Zarandah, Q. M., et al. (2025). Efficient Respiratory Disease Classification Using Customized CNN on a Large Kaggle Dataset. *Proceedings of Eighth International Conference on Information System Design and Intelligent Applications*, Springer Nature Singapore Singapore.