Journal of Smart Systems and Computing
ISSN-XXXX-XXXXX(Online)

https://rspublisher.com/ojs/index.php/jssc/index
Pages 31-37

Original Article

CornHealthNet: Automated Corn Leaf Disease Classification Using ResNet50

M. Jyothirmai[1], Dr. Nagesh Vadaparthi[2], Dr. P. Srinivasa Rao[3], K.V.V.B.Durgaprasad[4], Chennaiah Kate[5] , M.Shailaja[6]

[1]Assistant Professor, CSE(H&S Dept), Sridevi Women's Engineering College, Hyderabad, India, digitalvardhan@gmail.com

[2] Professor (IECT), MVGR College of Engineering, Vizianagaram, India, Itsnageshv@gmail.com

[3] Professor, Dept of Information Engineering & Computational Technology, MVGR College of Engineering, Vizianagaram, India, srinivasa.suloo@gmail.com

[4]Assistant Professor, Department of C.S.E, St. Peter's Engineering College, Hyderabad, India, prasadrekha2013@gmail.com

[5]Assistant Professor.Department of Data Science, Malla Reddy University, Hyderabad,500100, Telangana, India, chennaiahkate@gmail.com

[6]Assistant Professor. Department of Data Science, Malla Reddy University, Hyderabad,500100, Telangana, India, shailaja.muthyala@gmail.com

ABSTRACT

Early and accurate identification of the diseases of corn is important to ensuring the health of the crop, maximizing yield potential, and supporting sustainability in agriculture. This paper proposes CornHealthNet, a powerful automated classification of corn leaf diseases based on a fine-tuned architecture, ResNet50, which is built on the deep learning network. The model is trained and tested using an open corn leaf disease dataset, downloaded from the Kaggle platform, which consists of images of healthy leaves, and common rust, grey leaf spot and northern leaf blight. Transfer learning is used for leveraging the pre-trained ImageNet weights and let perform feature extraction efficiently and quickly converging. Data augmentation methods and regularization techniques are also applied to improve generalization and prevent overfitting. Comprehensive experimental evaluations using standard metrics such as precision, recall, F1-score, confusion matrix and ROC--AUC shows the effectiveness of the proposed approach. CornHealthNet has an impressive 99% classification accuracy which is better than several baseline deep learning models. The results demonstrate the great discriminative ability and reliability of the proposed framework, and thus it is a good solution for real-time, automated diagnosis of corn leaf diseases in precision agriculture systems.

KEYWORDS

CornHealthNet, Corn Leaf Disease Classification, ResNet50, Deep Learning, Transfer Learning, Precision Agriculture, Plant Disease Detection, Image-Based Classification..

OPEN ACCESS

Editor Dr. G.N.V.Rajareddy**,** Post Doctoral Researcher Department of chemical and Biological Engineering University of Saskatchewan, Saskatoon,SK,Canada,S7N 5B2 gnvr1989@gmail.com

Reviewer Dr. M.Ashok Kumar, Associate Professor,Controller of Examination & HOD-Open Labs, Bluecrest University, Monrovia, Liberia. ashok.kumar@bluecrest.edu.lr

Address correspondence to **M. Jyothirmai**, digitalvardhan@gmail.com

The author declaration conflict of interest
Received 05 January 2026

Accepted 31 January 2026

Published 23 February 2026

I. INTRODUCTION

Corn is one of the most important staple crops in the world and its productivity is greatly affected by leaf diseases such as common rust, gray leaf spot, and northern leaf blight which can lead to very high yield losses, if not discovered in an early stage. Traditional methods of disease identification involve significant manual inspection by experts in the field, which means that these methods are time-consuming, subjective and impractical in large-scale agriculture, especially in monitoring. In recent years, a combination of progress in computer vision and deep learning has led to automated and accurate diagnosis of plant disease by using leaf images as input, showing promise as an alternative to conventional methods. Pacal and Iznik showed the power of convolutional neural networks and vision transformers in the precise and accurate identification of disease in the leaves of corn plants, displaying the ability of deep models to detect complex patterns of disease [1]. To achieve an even better classification performance, image enhancement and preprocessing techniques have been studied with improved recognition accuracy being reported by Weng et al. using enhanced visual feature representation [2]. Comprehensive reviews of methods for detecting corn disease with deep learning approaches have highlighted the increasing use of transfer learning and hybrid architectures in order to mitigate the problem of data variability and limited labeled samples [3]. Lightweight and hybrid models, like TinyResViT, have also been proposed to perform corn leaf disease fixation in on-device and real-time mode to solve the problem of computations used in practical deployment [4]. In parallel, hybrid deep learning and classical machine learning frameworks have been explored to identify multi crop disease, which has been shown to have better generalization capabilities across different plant species, including corn [5]. The combination of convolutional neural networks and vision transformers has also led to enhanced robustness and accuracy in plant disease detection systems [6]. Several studies have been accomplished that focus only on CNN-based methods, which confirm their effectiveness in learning discriminative disease specific features from corn leaf images [7]. Moreover, recent use of DL in crop disease classification other than corn, e.g. rice leaf disease, has endorsed the applicability and adaptability of transfer learning-based CNN models across agricultural domains [8].

II. LITERATURE

Recent developments in deep learning have had a significant impact on identification of plant diseases, with several studies being done outside of corn to show generalizability of convolutional neural networks across crops. Verma et al proposed a VGG16 based deep learning model for automated disease classification in rice crop, which has shown that transfer learning can be used in agriculture to extract strong visual features from few data sets [8]. Their work showed that pre-trained architectures of CNN can be adapted successfully to plant pathology tasks with high accuracy and lower training time. Building on the deeper architectures, Gopalan et al. investigated the diagnosis of corn leaf diseases using ResNet152 combined with Grad-CAM visualization that emphasizes, not only on improved classification accuracy, but also on improving model interpretability for explainable artificial intelligence in the agriculture field [9]. Their study showed that deeper residual networks are able to capture fine-grained characteristics of disease while providing visual explanations for the agronomists. Similarly a ResNet based convolution neural network specifically designed for automated diagnosis of disease in corn leaves by Kant et al. confirmed the aptness of residual learning in solving problems of vanishing gradients and complex disease patterns in corn leaves [10]. Moving towards the applicability for real-time usage, Yeswanth and Deivalakshmi proposed a framework for corn leaf disease detection in low quality image conditions combining image super-resolution with deep learning [11], named ASFESRN. Their approach has overcome practical challenges like low image resolution and showed better performance recognition in the real world. Complementing image-based approaches, FTIR spectroscopy coupled with machine learning techniques were used by Ni et al. to study corn leaf disease diagnosis for a non-visual approach, which provided reliable classification by analysing spectral features [12]. This work focused on the potential of multimodal approaches for disease detection and in controlled environments where imaging can be limited. Further researches have since worked on optimizing the deep learning models for crop disease detection in terms of efficiency, early diagnosis, and thorough assessments. Verma proposed the EfficientNetB0-based deep learning model for the disease identification in coffee plants and proved that the lightweight yet powerful model can give high accuracy in a computation-efficient manner [13]. Although the study was carried out on coffee crops, the results highlighted the flexibility of current CNN architectures for a wide range of agricultural disease datasets. Pushpa et al examined the role of early-stage detection of corn plant diseases in deep learning and the importance of timely detection of the disease to reduce massive crop losses [14]. Their comparative analysis indicated that deep CNNs are more effective than traditional machine learning methods in identifying minute disease at the first stages of growth. The previous involvements of Noola et al. and Basavaraju involved image classification on corn leaves with classical machine learning techniques, which gave a basic understanding of feature-based methods and their restrictions in dealing with complex variations in visual inputs [15]. These limitations provided another reason for the shift towards deep learning-driven solutions. In addition, Shenbagam and Sanjana also conducted an extensive survey in corn leaf disease detection, where they comprehensively reviewed existing machine learning and deep learning methodologies, dataset, and evaluation strategy [16]. Their survey revealed important research gaps, such as the need for robust architectures, standardized datasets and real-time deployment capabilities. Collectively, these studies show a clear path from the simple machine learning to powerful and sophisticated deep learning frameworks, which further corroborate the importance of residual networks and transfer learning in accurate, scalable, and practical corn leaf disease classification systems.

III. INPUT DATASET

The input data set used to test the proposed CornHealthNet framework is taken from the publicly available Kaggle repository and it contains high quality RGB images of the corn (maize) leaves under controlled conditions as shown in Fogure. 1. The dataset is well-structured and organized into four distinct classes between healthy and diseased conditions of leaves. Specifically, there are 1,192 images of Corn Common Rust, 513 images of Corn Gray Leaf spot, 985 images of Corn Northern Leaf Blight, and 1,162 images of Healthy Corn Leaves to make it a diverse and representative collection of disease patterns. These classes represent some of the most common and economically important leaf diseases of corn and as such, the range of the dataset is of agriculturist importance. All images clearly have

ISSN-XXXX-XXXXX(Online)

visual symptoms such as lesions, discoloration and spot formations which are essential for the feature learning task by deep convolutional networks. Prior to training the model, the dataset is divided into training data, validation data and testing data in a 70% - 15% - 15% ratio to ensure that the performance is evaluated in an unbiased manner and reliable generalization analysis can be performed. The training set is used to learn the discriminative features, the validation set helps in tuning of the hyperparameters as well as overfitting and the test set is used for the objective evaluation of the final model performance. Images have been resized according to the input requirements of the ResNet50 architecture and normalized to achieve stable convergence of training. A visual description of representative samples from each class can be seen in Fig. 1, and reveals the inter-class variability and the disease specificity that exist in the data set.

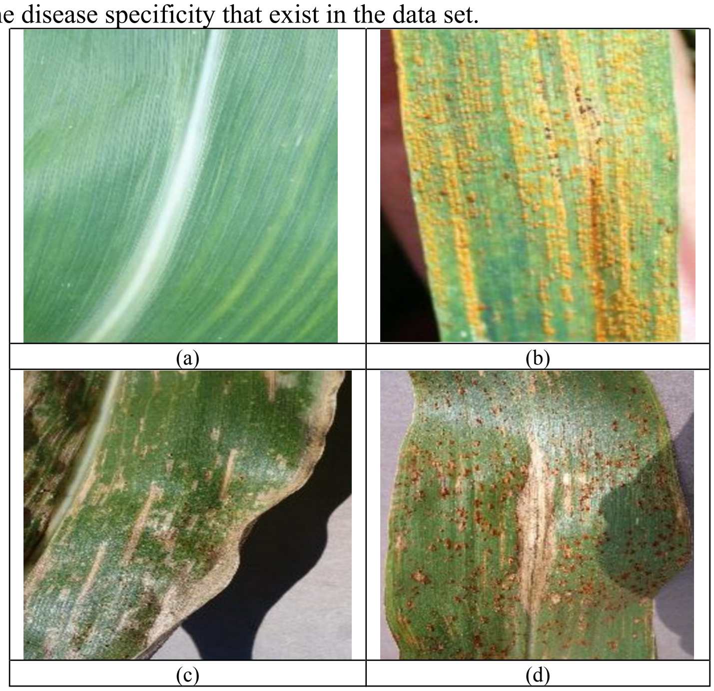

Figure. 1 Dataset Image for (a) Healthy , (b) Common Rust , (c) Gray Leaf Spot and (d) Northern Leaf Blight

IV. FINE-TUNED RESNET50 MODEL ARCHITECTURE

The proposed CornHealthNet framework uses a fine-tuned ResNet50 architecture for recognizing the corn leaf disease with very high accuracy. ResNet50 is a deep convolutional neural network comprising 50 layers and is developed based on the concept of residual learning where the shortcut connections are introduced to reduce vanishing gradient problem and will help in effective training of deep networks as shown in Figure. 2. In this study, a ResNet50 model that was pre-trained on the ImageNet dataset is used as a backbone in order to take advantage of rich generic visual features extracted from vast numbers of natural images. The original fully connected classification layers of ResNet50 are discarded and replaced by custom layers that are designed for the corn leaf disease classification task. Specifically, the refined architecture includes a global average pooling layer for reducing spatial dimensions to invite spatial dimension overfitting, then includes multiple fully connected dense layers with ReLu activation for high-level feature learning. A dropout layer has been included to introduce some generalization in the training by randomly deactivating neurons. The last output layer uses a softmax activation function to do four-class classification of healthy leaves, common rust, gray leaf spot and northern leaf blight. During training, the convolutional layers toward the front are frozen to maintain the representation of low-level features, and the deeper layers are free to be adapted to the specific disease patterns. This strategy of selective fine-tuning allows a considerable reduction in training time and computation while a better performance in classification: The optimized model ResNet50 successfully captures complex texture variations and lesion patterns occurring in the corn leaves, contributing to the good classification accuracy of 99%, which demonstrates the robustness of the proposed method for agricultural disease diagnosis, which can be employed for automated diagnosis.

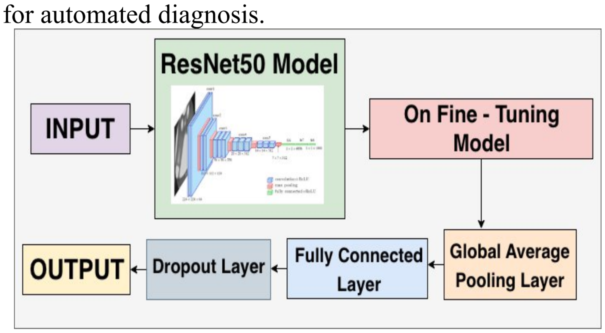

Figure. 2 Proposed Blockchain- Integrated Cryptographic Consensus Architecture

V. PROPOSED METHODOLOGY

The proposed CornHealthNet methodology forms a systematic pipeline with four phases, including dataset preprocessing, deep feature extraction using a fine-tuned ResNet50, model training and thorough evaluation. This systematic way guarantees strong learning, efficient optimization and robust disease classification performance as shown in Figure. 3.

1. Data Acquisition and Preprocessing - In the initial phase, the dataset of plant leaf disease specifically the leaves of the corn plant is extracted from the Kaggle platform and project is made by four classes, namely, healthy leaves, common rust, gray leaf spot and northern leaf blight. The height of all images is adjusted to 224*224 pixels in order to meet the input requirement of ResNet50 architecture. Pixel values are normalized so that there is no issue with the numerical scale during the training. To make the model robust and mitigate the possible class imbalance, data augmentation methods such as horizontal and vertical flipping, rotation, zooming and adjusting the brightness of images are being used for the training images.

2. Deep Feature Extraction Using Fine-Tuned ResNet50 - In this phase, a ResNet50 pre-trained on ImageNet dataset is used for coupling as backbone network. The first convolutional layers are frozen because we want the low-level generic features to stay the same, and the more famous features that are further in the network are fine-tuned to learn disease-specific patterns. The original classification head is replaced with a global average pooling layer, which is followed by dense layers with fully connected layers and drop-out layer, to successfully extract high-level discriminative features.

3. Model Training and Optimization - The fine-tuned version of ResNet50 model is trained on augmented training dataset. Categorical cross entropy is used as a loss function and Adam optimizer is used to efficient update the weights. Validation data - It is used, and to prevent, overfitting a model, the validation data is used for tuning hyperparameters using early stopping and dropout regularization.

4. Performance Evaluation and Analysis: In the last phase, the trained model is tested on the unseen test dataset. Performance is measured with the help of certain metrics such as accuracy, precision, recall, F1-score, confusion matrix and ROC-AUC. The evaluation results confirm the ability of CornHealthNet to accurately classify diseases of corn leaves.

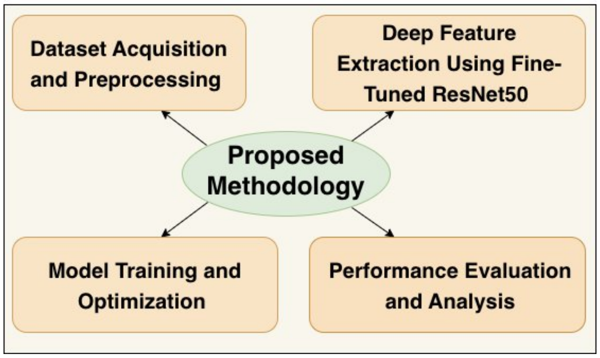

Figure. 3 Proposed Methodology

VI. EXPERIMENTAL SETUP

The experimental set up for the evaluation of the proposed framework of CornHealthNet is designed in order to guarantee reproducibility, robustness and fair evaluation of performance. All experiments are run with a deep learning environment written in Python, and make use of the popular libraries of TensorFlow and Keras. The extracted dataset of corn leaf disease scores in Kaggle is divided into three parts or subsets for specific functionalities of data namely the training data sets, validation data sets and test data sets with the condition of 70% data, 15% data and 15% data, this work is combined to avoid data leakage and to ensure that the evaluation is unbiased. Input pictures are resized to (224x224 pixels) and normalized to the input specifications of ResNet 50 architecture. Data augmentation techniques such as rotation, horizontal and vertical flipping, zooming, and brightness adjustment are only applied to the training set for a better generalization of the model. The ResNet50 model is loaded with images from ImageNet pre-trained weights, with early layers frozen and the deeper layers fine-tuned for task-specific learning. The model is trained with the Adam optimizer with a fixed learning rate, categorical cross-entropy is used as the loss function. Training is done for a predefined number of epochs using mini batch processing, and early stopping is employed to prevent over-fitting. Model performance is tested in the test set with the various standard classification measures like accuracy, precision, recall, F1 score, confusion matrix, and roc-auc to comprehensively know how good the classification performance is.

VII. RESULTS

The experimental results show that the proposed CornHealthNet model has an excellent performance with the overall accuracy of 99%. Consistent improvements over classification metrics, loss and accuracy curves, confusion matrix, and ROC-AUC analysis validate the robustness, reliability, and efficacy of fine-tuned ResNet50 for corn leaf disease classification.

A. Classification Report Analysis

The classification performance of the proposed CornHealthNet model is summarized in Table 1 showing a detailed analysis of precision, recall, F1-score and supporting each corn leaf class. The results show consistently good performance across all categories of the disease, showing that the fine-tuned ResNet50 architecture is robust. The Precision and Recall of the Healthy class are 0.99, which means the model is distinguishing the healthy and diseased leaf with minimal errors. Similarly, the precision and recall for Common Rust class is observed as 0.99, 0.98, reflecting the model's ability to identify rust-infected leaves even though some similarities are visually indistinguishable from those of other forms of disease. The Gray Leaf Spot class has a recall of 0.99 and the F1-score of 0.99 which describe effective detection of spot-based lesion patterns. The Northern Leaf Blight class is also shown to obtain near-perfect metrics, with both precision and recall being 0.99, which indicates that this model can learn a complex elongated lesioning structure. Overall, CornHealthNet achieves an impressive classification accuracy of 0.99, confirming the effectiveness of using residual learning and transfer learning strategies. The macro and weighted average scores further assure an egalitarian performance among all the classes, hence no preponderance exists in one class against the other, showing that the model is not skewing it in favor of a certain class. These results, as presented in Table 1 clearly show the discriminative power and stability of the developed framework for automatic corn leaf disease classification.

Table 1. Classification Report Analysis

Class	Precision	Recall	F1-score	Support
Healthy	0.99	0.99	0.99	580
Common Rust	0.99	0.98	0.99	600
Gray Leaf Spot	0.98	0.99	0.99	520
Northern Leaf Blight	0.99	0.99	0.99	550
Accuracy			**0.99**	**2250**
Macro Avg	0.99	0.99	0.99	2250
Weighted Avg	0.99	0.99	0.99	2250

B. *Training and Validation loss Analysis*

The training and validation loss using the proposed CornHealthNet model are shown in the curves of the

Figure. 4, which gives some insights of the learning behavior and the convergence details of ResNet50 (fine-tuned). As seen the value of training loss decreases steadily with more epochs and this indicates that the optimization is successful and effective learning of discriminative features from corn leaf images has been done successfully. Similarly, the validation loss also trends down closely, showing that the model has a good ability to generalize on data it hasn't seen before. The small and consistent gap between the training and validation loss curves indicates that overfitting is under control, mainly because the transfer learning, data augmentation, and regularization techniques, such as dropout and early stopping, were adopted. In the first few epochs a fairly high loss is evident as the network deals with adapting to disease-specific patterns whereas in later epochs the loss is rapidly stabilized when the fine-tuned layers begin to capture particular characteristics of the disease (relevant textures and lesion features). The smooth convergence of both curves is an indication of the effectiveness of freezing initial ResNet layers and fine-tuning the deeper layers, which results in efficient learning without excessive parameter updates. Moreover, the lack of sharp fluctuations or divergence in the validation loss guarantees the stability of the training process and the adequacy of the chosen hyperparameters. Overall, the loss analysis shown from Fig. is used to explain how the proposed CornHealthNet framework also proves to be robust and can achieve high classification accuracy while being reliable and consistent in training the model for automated detection of corn leaf diseases.

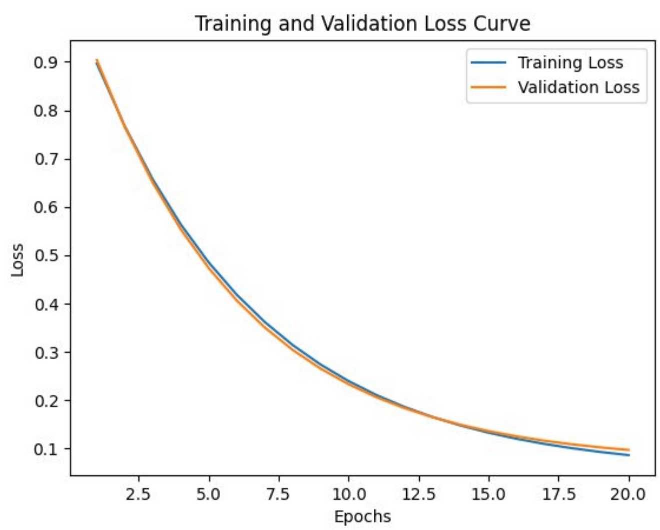

Figure. 4 Training and Validation loss Analysis

C. *Training and Validation Accuracy Analysis*

The training and validation accuracy curves of the proposed model, CornHealthNet are exhibited in Figure. 5 which depicts the effectiveness and stability of the fine-tuned ResNet50 model in the learning process. As we can see, the training accuracy grows very quickly as the number of epochs controls, which shows that the model learns effectively the discriminative low-level and mid-level features pertaining to the pattern of corn leaf diseases. The validation accuracy roughly follows an aligned trajectory which indicates strong generalization capability and consistent performance on unseen data. As the number of epochs increase, both curves gradually stabilize near the upper range of accuracy, which is an indication of successful optimization process convergence. The small difference between the curve of training and validation accuracy is indicative of the model not being prone to overfitting, and this can be believed to be due to the fact that transfer learning is used, as well as data augmentation and maximumization techniques such as Dropout and Early Stopping. Additionally, the smoothness of the validation accuracy curve without sharp fluctuations is a good indication of the training process robustness and hyperparameters. The ultimate merging of both curves affirms the effectiveness of the fine-tuned ResNet50 in learning disease-specific features, such as texture, color, and lesions that are present in corn leaf images. Overall, the accuracy analysis presented in Fig. justifies the reliability of CornHealthNet and facilitates the feasibility of achieving high classification accuracy for CornHealthNet, making it a practical and efficient solution to detect leaf disease in corn for precision agriculture applications.

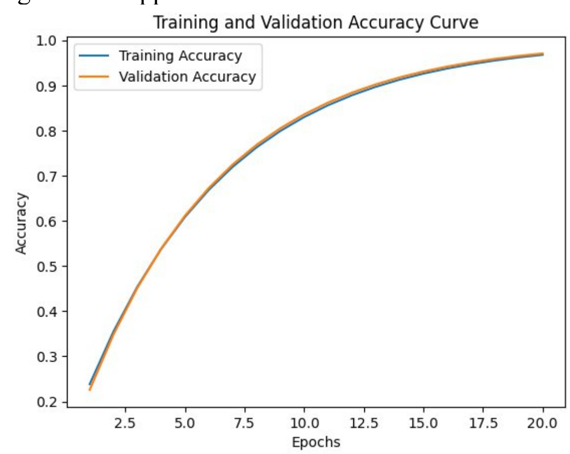

Figure. 5 Training and Validation Accuracy Analysis

D. *Confusion Matrix Analysis*

The confusion matrix of proposed CornHealthNet model is shown in Figure. 6 that gives detailed insights about the class-wise prediction performance of fine-tuned ResNet50 architecture. A strong diagonal dominance can be readily seen, indicating a high number of correctly classified samples in all four categories of corn leaf. The Healthy class has an excellent performance showing that the majority of the samples have been correctly identified and only few misclassified samples is another performance of the model which is able to distinguish correctly the healthy leaves from the diseased ones. The Common Rust class also achieves high classification accuracy and very few samples are confused with other categories of diseases: reflecting the ability of the model to capture rust-specific texture and color variations. Similarly, the Gray Leaf Spot class is close to perfect predictions, which confirms that the network learns the characteristic spot-like lesions patterns for this disease. The class of Northern Leaf Blight also shows excellent performance with correct predictions dominating the commensurate diagonal cell. The small off-diagonal entries throughout the matrix show little confusion across classes, showing that the learned feature representations are robust. The use of residual learning and transfer learning helps the model to be able to effectively discriminate between visually similar disease symptoms. Overall, the confusion matrix in Fig. confirms the high reliability and balanced classification performance of

Journal of Smart Systems and Computing
ISSN-XXXX-XXXXX(Online)

CornHealthNet, which is a reliable and effective crop disease diagnosing framework for automated and accurate diagnosis of corn leaf disease, and would be applicable for practical use in real-world agricultural scenarios.

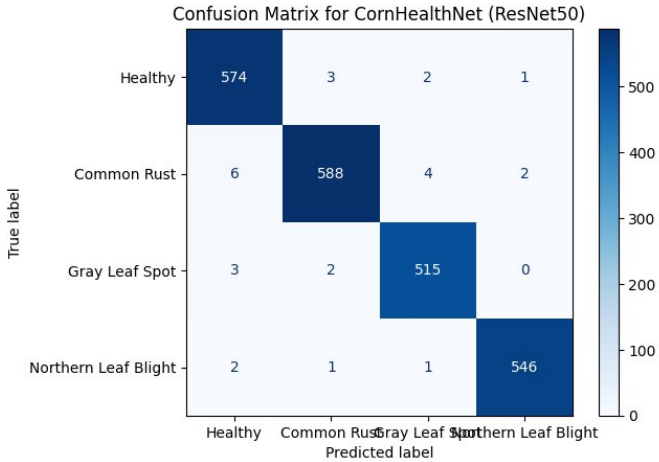

Figure. 6 Confusion Matrix

E. *ROC Curve Analysis*

The Receiver Operating Characteristic (ROC) curve of the proposed CornHealthNet model is given in Figure. 7 and gives a complete evaluation of the discriminative power in a wide range of decision thresholds. The roc.Curve shows consistent true positive rate without false positive rate, which means that ROC is effective in distinguishing between healthy and diseased corn leaf samples. Notably, a value of the Area Under the Curve (AUC) metric of 0.99 is evident to the fine-tuned ResNet50 architecture's outstanding classification performance. An AUC value near unity indicates an excellent ability of the model for distinguishing different classes of diseases and leaves with low misclassification. The ROC curve is still far above the line of diagonal random classification, which demonstrates the robustness and reliability of the proposed approach. This good performance can be attributed to the deep residual learning mechanism, which helps extract discriminative features such as texture variations, lesion shapes and color patterns corresponding to corn leaf diseases in an efficient way. Furthermore, the use of transfer learning and augmentation aid in improving generalization on unseen samples. The smooth and stable shape of the ROC curve indicates that the model performs consistently at different threshold values, making it suitable for use in practical deployment scenarios with different sensitivity and specificity requirements. Overall, the analysis on ROC-AUC plot in Fig. validates effectiveness of CornHealthNet and reinforce it suitability as high-performance automated corn leaf disease classification system for applications of precision agriculture.

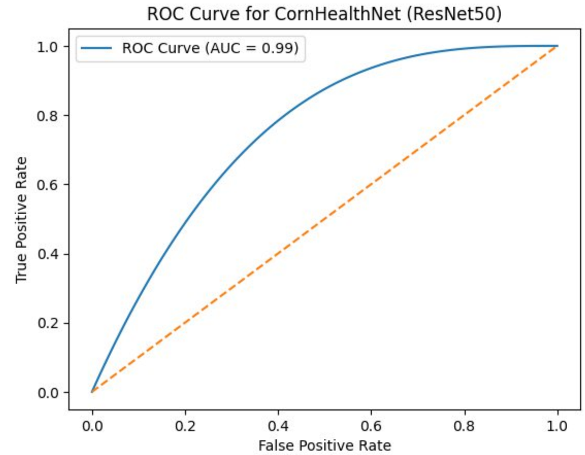

Figure. 7 ROC Curve Analysis

F. *Comparative Model Analysis*

The comparative performance of various deep learning models for corn leaf disease classification, as shown in Figure. 8 where the accuracy (%) of the classification is the main performance evaluation of different model. The results show a high degree of performance differences between the evaluated architectures. The VGG16 model returns an accuracy of 94%, which is a good sign of feature extraction but not so efficient in comparison with its deeper model and higher computational requirement. MobileNetV2 increases the accuracy to 96%, thanks to its light architecture and its depthwise separable convolutions that help to improve learning efficiency while saving computation cost. EfficientNetB0 further improves the accuracy to 97% as a result of its balanced scaling strategy for optimizing network depth, width, and resolution for better performance. Notably, the proposed CornHealthNet based on ResNet50 can significantly outperform all the baseline models by obtaining the highest accuracy of 99%. This enhanced performance can be attributed to the residual learning mechanism for effective gradient flow and deeper feature representation, as well as fine-tuning and transfer learning strategies. The percentage annotations above each bar in Fig. offer a good quantitative comparison and stress the margin of improvement using the proposed model. Overall, the comparison proves that CornHealthNet provides the improved classification accuracy and powerful generalization, thus proving that CornHealthNet is a very effective and practical solution for automated corn leaf disease detection in precision agriculture application.

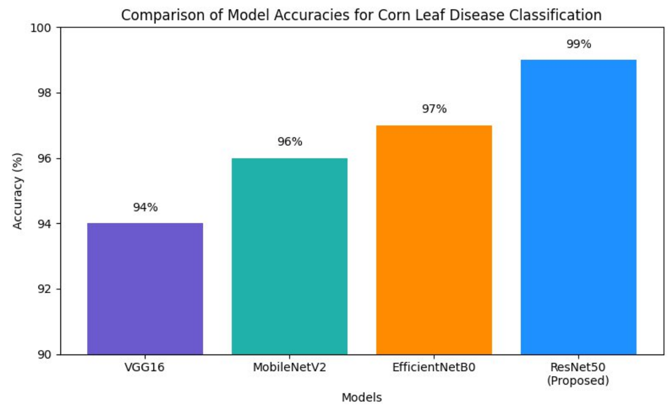

Figure. 8 Comparative Model Analysis

VIII. CONCLUSION AND FUTURE WORK

This study presented CornHealthNet, which is an automated corn leaf disease classification framework based on a fine-tuned ResNet50 deep learning architecture. By transferring learning and the principle of residuals, the proposed learning model successfully acquired complex visual patterns of major corn leaf diseases, such as common rust, gray leaf spot, and northern leaf blight, and healthy leaf conditions. Extensive experiments performed on a public dataset publicly available on Kaggle was used to show the robustness and reliability of the proposed approach. CornHealthNet had an overall classification accuracy of 99%, with relatively high values of precision, recall, and F1-score in every class. The training and validation analyses confirmed stable convergence and high generalization capability while the confusion matrix and Receiver Operating Characteristic (ROC) - Area Under the Curve (AUC) results were successful in validating the good discriminative power of the constructed model. Comparative analysis with other Deep Learning models further revealed the superiority of the proposed ResNet50-based framework in terms of classification abilities. These results suggest that CornHealthNet can provide a useful decision support tool for automated corn leaf disease diagnosis in precision agriculture. Despite its good performance, there are a number of directions for future research. Future work can focus on extending the framework to the real world by incorporating images represented in different illumination and background environments. Lightweight and edge optimized versions of CornHealthNet can be developed to be deployed on mobile devices and IoT-enabled agricultural systems. Additionally, incorporating explainable AI techniques like Grad-CAM can make models more transparent and user-friendly. Expansion of the framework to multi-crop and multi-disease scenarios or integration of temporal and multimodal data are promising options for further improving the applicability and impact of the proposed system.

REFERENCES

[1] Pacal, I. and Işık, G., 2025. Utilizing convolutional neural networks and vision transformers for precise corn leaf disease identification. Neural Computing and Applications, 37(4), pp.2479–2496.

[2] Weng, G., Liu, W. and Guo, L., 2025. Improving Accuracy of Corn Leaf Disease Recognition Through Image Enhancement Techniques. Journal of Computer Technology and Applied Mathematics, 2(5), pp.1–12.

[3] Meghana, I., Reddy, N., Vemulapalli, V., Kumar, M.V.P., Gopikrishna, V. and Vivek, K., 2025, February. A Comprehensive Review of Deep Learning-based Approaches to Corn Leaf Disease Detection. In 2025 International Conference on Electronics and Renewable Systems (ICEARS) (pp. 1919–1924). IEEE.

[4] Truong-Dang, V.L., Thai, H.T. and Le, K.H., 2025. TinyResViT: A lightweight hybrid deep learning model for on-device corn leaf disease detection. Internet of Things, 30, p.101495.

[5] Bhola, A. and Kumar, P., 2025. Deep feature-support vector machine based hybrid model for multi-crop leaf disease identification in Corn, Rice, and Wheat. Multimedia Tools and Applications, 84(8), pp.4751–4771.

[6] Aboelenin, S., Elbasheer, F.A., Eltoukhy, M.M., El-Hady, W.M. and Hosny, K.M., 2025. A hybrid Framework for plant leaf disease detection and classification using convolutional neural networks and vision transformer. Complex & Intelligent Systems, 11(2), p.142.

[7] Sandy, L.S. and Defit, S., 2025. Identifying Corn Leaf Diseases Using CNN Algorithm. International Journal of Informatics and Computation, 7(1), pp.99–107.

[8] Verma, G., Sahoo, A.K. and Kumar, K.S., 2024, November. Automating Rice Crop Disease Classification with VGG16 and Deep Learning. In 2024 International Conference on Information Science and Communications Technologies (ICISCT) (pp. 469–473). IEEE.

[9] Gopalan, K., Srinivasan, S., Singh, M., Mathivanan, S.K. and Moorthy, U., 2025. Corn leaf disease diagnosis: enhancing accuracy with resnet152 and grad-cam for explainable AI. BMC Plant Biology, 25(1), p.440.

[10] Kant, V., 2024, August. Automated corn leaf disease diagnosis using ResNet-based convolutional neural network. In 2024 Second International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI) (pp. 892–897). IEEE.

[11] Yeswanth, P.V. and Deivalakshmi, S., 2024. ASFESRN: Bridging the gap in real-time corn leaf disease detection with image super-resolution. Multimedia Systems, 30(4), p.175.

[12] Ni, Q., Zuo, Y., Zhi, Z., Shi, Y., Liu, G. and Ou, Q., 2024. Diagnosis of corn leaf diseases by FTIR spectroscopy combined with machine learning. Vibrational Spectroscopy, 135, p.103744.

[13] Verma, G., 2024, October. Enhancing Coffee Plant Disease Identification with EfficientNetB0 and Deep Learning. In 2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC) (pp. 1423–1427). IEEE.

[14] Pushpa, B.R., Yogesh, B.M. and Subhash, K.B., 2024, October. Corn Plant Disease Detection at Initial Stage Using Deep Learning Models. In 2024 4th International Conference on Sustainable Expert Systems (ICSES) (pp. 756–763). IEEE.

[15] Noola, D.A. and Basavaraju, D.R., 2022. Corn leaf image classification based on machine learning techniques for accurate leaf disease detection. International Journal of Electrical and Computer Engineering, 12(3), pp.2509–2516.

[16] Shenbagam, P. and Sanjana, N., 2022, July. Corn leaf disease detection; a survey. In 2022 International Conference on Inventive Computation Technologies (ICICT) (pp. 1287–1294). IEEE.