Wheat Plant Disease Detection Using Deep Learning Techniques

Jayasri Gollavilli
Faculty of Engineering
M.S. Ramaiah University of Applied Sciences
Bengaluru, India
gollavillijayasri8338@gmail.com

Ankitha Hudegal
Faculty of Engineering
M.S. Ramaiah University of Applied Sciences
Bengaluru, India
ankitha.1830@gmail.com

***Abstract*— Sick crops spell disaster for a farmer's existence. It is thus important to identify the problem as early as possible. However, checking the whole area of a farm is too time-consuming and costly. It is not even guaranteed that the job will be done to perfection. Why use traditional methods when the technology has the ability to access intelligent software? This process utilizes a type of AI called CNNs, which is a pattern recognition technology. This technology is able to analyze pictures of leaves and identify the signs of a problem without the need for human intervention. There is no need for the farmer to be in the fields while the machines identify the problem from the pixel level. The fields cover miles of land, and problems may be hidden from the naked eye. However, the technology ensures that things get done on time. Machines have the ability to see things that the naked eye cannot, especially over large areas. It is not just a matter of doing things fast but doing them fast and consistently. Time is of the essence when it comes to technology, but lag is minimal when it comes to accuracy. Cultivating crops is always a gamble, and the earlier you detect diseases, it works to your advantage. This is based on data and not on gut feel, which is trained to assume that there is a problem or health. It is not magic, it is advanced mathematics that works on green tones and discolorations. Plants manifest diseases differently, so it is learning by repetition. The more images it scans, the better as recognition is more pronounced.**

Recognition occurs when CNNs learn the "language" of disease. There are digital images of leaves comprising the first step toward getting this system to recognize possible diseases. These images are compiled before the assessment occurs — the images of illness and non-illness are compiled to train the system and assess results.

This training does not happen until the images are resized and distributed evenly so that the system can be trained better. The CNN is designed in tiers so that different layers compress the details and attempt to extract attributes such as pigmentation, texture, and edge contours from the leaf images. This method, developed using TensorFlow and Keras, is a nuanced system for developing deep learning systems. Once educated, the CNN is saved, reloaded and employed to assess previously unseen images of leaves. Recognition is more precise with disease determinations as it's data relative vs. hypothetical conclusions. The less human interference, the fewer errors made in determining sick vs. healthy plants. The beauty of an intelligent system is that one image can determine what's wrong with crops based upon a system constructed for time-efficient recognition and cost effectiveness. What takes mere days to notice in the field and report back to owners occurs with educated expedience for faster results with determination. Pixels become responses through layers of meticulous training that honed in on patterns that the human eye may miss. Such image-based tricks and learned machines show how technology can venture into fields just as much as it would stay in labs. Fields get an ally that learns automatically over time without any excess apparatus or need for hassle.

Keywords*— *TensorFlow, Machine Learning, Image Processing, Keras, Deep Learning, convolutional Neural Network (CNN)

I. INTRODUCTION

Growing things sustains people. It grows economies, especially in cultures where most humans are farmers. Healthy plants equate to healthy profits. When illness befalls, yield plummets. Money disappears as quickly as illness emerges without it being detected in time. Learning about these issues and how they impact longevity helps fields survive beyond their intended lifespan. Disease detection helps fields live longer than expected. Plants can become diseased by fungus, bacteria, or viruses even harsh environments can turn sickly. Symptoms usually emerge on the leaves: discoloration, markings, wilting, textural changes. Most farmers seek out their fields when symptoms develop to see what's wrong with their crops. Experts are sometimes employed. But success is not guaranteed. It depends on who's looking at the crops and what they know. The detection of crop diseases manually is labor-intensive. It's too much work. It's difficult for large farms to keep up. In rural areas, there's a lack of expertise. Sending warnings requires specialists.

The recent advancements in artificial intelligence combined with machine learning have opened up possibilities to solve some of these problems. Among these, deep learning stands out. It's especially useful with high-dimensional data, especially images. It can identify features without using written algorithms. It doesn't require hand-crafted features, starting with raw data. Since the patterns on the leaves reveal diseases in plants, using images to find these diseases is perfect. This method is currently under heavy research.

A smart program for image detection. It's called Convolutional Neural Network. It's perfect to meet the requirements of image detection. It identifies features in images in a stepwise fashion, like lines, curves, spots, patterns. The small disease spots on green leaves might not be recognized. But this network identifies them. the applicable criteria that follow. Using raw images instead of hand selected features makes this technique robust and reliable as opposed to predecessor models. The work here builds one such system that tells if a leaf looks fine or shows illness using photo examples. A collection of leaf pictures shows both sick plants and healthy ones, each marked clearly. Because accuracy matters, the group of images gets split three ways - training, checking progress, test. Six out of ten go to learning mode, two tenths help adjust settings mid-process, another two tenths stay untouched until last. Learning happens first, where patterns tied to sickness are picked up by the network slowly. While that runs, feedback comes from a smaller batch guiding small fixes. Only after everything else finishes does the untouched portion reveal how well it truly performs. That final look decides nothing earlier influenced.

Figure represents how the model is detecting the disease and sending it to the next level of wheat disease label. Getting images ready matters before using them to train deep learning models. Since pictures differ in size, clarity, brightness, or surroundings, adjustments are needed. Every image gets scaled down or up to match one standard shape. Pixel numbers then shift into a common span through rescaling. This balancing act smooths out how the system learns over time. Uniform inputs let the network focus on patterns, not noise. The model learns from consistency across its samples. This configuration uses a CNN made of several convolution layers followed by pooling layers. Based on the images of the leaves, the first layers sift out relevant features such as patterns or colored areas corresponding to certain diseases. Following this, the pooling layers compress the number of these sets to decrease the resource overheads while at the same time filtering out false patterns. The stretching and processing follow this. The distribution of the disease classes happens through the application of a softmax function to ensure that there is an even distribution in order to arrive at the outcome. The hidden layers connect everything before we end up staring at the outcome.

Based on TensorFlow and Keras, the application uses two popular frameworks that have been adopted in the industry to ensure that there is seamless neural network application development. Much versatility can be seen with the elegant handling of the architecture of the convolutional network and the ease of testing different training configurations. During training, considerable changes happen through Adam optimizer – its task is to reduce the classifying errors of data points. The tests place a lot of emphasis on accuracy scores to gauge performance. Validating data sets show their face during critical moments when the model starts to memorize instead of learning suggesting adjustments.

II. LITERATURE SURVEY

Wheat is one of the most important staple cereal crops and globally food source in agriculture, significantly contributing to food security. Wheat crop yields have also been significantly affected by many plant diseases such as leaf rust, stripe rust, powdery mildew, spot blotch, and fusarium head blight. These diseases lead to adverse yield effects on wheat crops and also the grains, causing economic losses for farmers. The conventional method for diagnosing wheat crop diseases relies on humans through visual inspection by agricultural experts or breeders. This is labor intensive, subjective and not feasible for large crop areas.
Hence, automated wheat crop disease detection using image processing techniques, machine learning and deep learning approaches have gained much attention.

The early machine learning based approaches to wheat plant disease detection mainly relied on classical image processing techniques. These techniques involve capturing images of wheat leaves, preprocessing them with denoising, contrast enhancement, and color normalization, then segmenting the diseased and healthy areas of the leaves, extracting color, texture, and shape features, and using these to train traditional classifiers like Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), Random Forests (RF), and Artificial Neural Networks (ANN). Hossain et al. [1] proved that this approach yields decent classification results in laboratory settings. However, their experiments also revealed that performance degrades rapidly when images are captured in field conditions due to the impact of uncontrolled lighting, background clutter, and variations in disease severity.

Follow-up studies applied these classical machine learning methods for the detection of specific diseases such as leaf rust and stripe rust on wheat crops. Singh and Misra [2] proposed an image-based frame- work that implements segmentation and supervised classification to identify various types of rust affecting different wheat varieties. Singh and Misra's study confirmed that machine learning models can be used to distinguish between types of disease, but they are extremely sensitive to the characteristics of the dataset used and the features selected. The authors state that handcrafted features do not capture the complexity of the patterns created by the disease when symptoms of the disease overlap or are at an early stage.

In an attempt to solve the problems presented by handcrafted features, researchers have turned to deep learning algorithms based on Convolutional Neural Networks (CNNs) that automatically learn to extract features from raw images. The first truly significant contribution was made by Lu et al.[3], who developed a weakly supervised CNN-based framework for infield wheat disease detection. Unlike other models, the framework did not require pixel level annotations from the images. It was also capable of detecting and localizing the diseases. The results of this study confirmed that CNN based models greatly outperform traditional machine learning models, a trend that suggests a promising future for using data-driven feature extraction for agricultural disease detection.

The next important innovation was Transfer Learning, which holds great promise in improving wheat disease classification when annotated image datasets are scarce. The authors of the study in question used transfer learning as developed by Kaur and Gill [4]. The researchers implemented Transfer Learning, a method that uses pretrained deep learning networks trained on large-scale image datasets, for wheat leaf disease classification. Kaur and Gill used pretrained architectures such as VGG16 and ResNet50 for their project. Their results showed remarkable accuracy and efficiency when classifying images of wheat leaf diseases. The study also established that transfer learning makes models more robust to overfitting, even with small agricultural datasets.

Comparative studies have also been performed that demonstrate the superiority of deep learning methods over traditional machine learning approaches. Sharma et al. [5] performed an extensive comparison of traditional and deep learning based techniques for wheat disease classification. However, they are not the only ones to have done so. Their findings all concluded that CNN-based models consistently had higher accuracy, precision, and recall than the SVM and Random Forest- based classifiers. The authors also reported that deep learning methods were less sensitive to the quality of background in images, which suggests robustness in real-world applications.

Not only individual studies, but also review and survey papers have been published that comprehensively assess the techniques used for wheat disease detection. The survey conducted by Patel and Shah [6] reviews only deep learning techniques for wheat leaf disease detection.

Their survey covered common datasets, CNN architectures, evaluation metrics, and deployment challenges. The authors concluded that while deep learning models are state-of-the-art in terms of accuracy, challenges such as imbalanced datasets, no benchmark datasets, and limited real-world deployment testing remain open questions.

General surveys of plant disease detection also provide clues. Kamilaris and PrenafetaBold [7] surveyed deep learning for various agricultural domains including wheat disease detection. Their survey recognized CNNs as the dominant architecture for complex visual representation capture. However, it noted some important research gaps, which included model interpretability, model complexity, and model suitability for mobile and edge devices used in precision agriculture.

Surveys have also concentrated on increasing the detection accuracy with improved architectures and hybrid approaches. Attention-based models have been proposed to improve the accuracy of the model by specifically focusing on the most affected regions of the wheat leaves. An attention-based deep learning model was proposed by Li et al. [8], which focuses on the symptomatic regions to improve the accuracy and detection capabilities of the model to recognize subtle details on the wheat leaves. The accuracy of the model can be significantly increased, and it can be confirmed that attention-based models vastly improve the accuracy of the model, especially at the early stages of the disease.

Another trend that has gained traction is the use of a hybrid approach that combines the use of deep learning and traditional classifiers. Demir and Kaya proposed a framework that combined CNN and SVM to develop a system that utilized an SVM classifier to analyze the features of the data using a CNN. This framework offered the best accuracy while being interpretable, unlike other deep learning-based models. This approach is especially important when the training data is limited.

Comparative studies on crop stress and disease detection also support the above results. Verma and Singh [10] investigated various machine learning methods for crop disease detection. They observed that all the models of deep learning methods were significantly better than other models, including those for wheat diseases.

Large-scale datasets are also increasingly being used in computer vision-based wheat research. Although not specifically targeted for disease detection, the Global Wheat Head Detection Dataset introduced by David et al. [11] has enabled robust wheat phenotyping and has helped improve the scaling of deep learning models. This dataset helps wheat disease detection research indirectly by improving model robustness, while also allowing large-scale agricultural analysis.

The literature review demonstrates a clear shift from traditional image processing and machine learning-based models to deep learning-based models for detecting wheat plant diseases. CNNs, transfer learning, attention-based models, and hybrid models have all been shown to improve accuracy and robustness for these models. Limitations related to dataset diversity, real-field deployment, efficiency of computation, and model interpretability, on the other hand, remain open questions and provide plenty of motivation for research to continue in the field.

III. PROPOSED METHODOLOGY

A. Data Set Preparation and preprocessing

Not every image set works right away. A strong CNN model needs clean, labeled pictures sorted in a clear way. Healthy wheat leaf photos sit alongside sick ones showing issues such as yellow rust, brown rust, powdery mildew, or septoria. These come from open farming databases sometimes; other times they're taken by hand out in real fields. Earlier studies used collections ranging from just hundreds up to many thousands of images.

Pictures start off resized - often to 224 by 224 or 256 by 256 pixels - so every image fits the same frame. Normalization adjusts colour intensity across photos so numbers behave consistently during learning.

B. CNN Architecture Design

Convolutions can do a lot. Let's begin with edges, then features. Features, then textures, and finally anything that can be associated with diseases. Why not? It's not only what's there, but also what's there that's different, according to something called disorders. They really stick out
what's there. Then, one at a time, they begin to make sense. Convolutions and Layers.

Let's begin with an activation function or two. Why not? Relu a curve here or there instead of a smooth transition from negative to positive numbers. Their sharp edge at zero suggests firing if required. Its simplicity enables the layering of deeper, richer feedback over time. Pooled to a smaller
size that retains important information. This permits the models to be speedy while retaining critical information. Reduction occurs here allowing lightweight processing yet with significant meaning retained. These layers translate knowledge into class probabilities. Not only transmit information but reflect deeply on what has been processed before coming to a conclusion. In the correct reflection of outcomes, all is taken into account. The Layers work to culminate in the delivery of probability distributions across classes. Probabilities are assigned here to potential diseases. Outcomes distribute to classes. Probabilities point to likely outcomes. Maximum outcome usually is correct. Totals always add up to one. Takes predictions here. Results shift when patterns change slightly. Wrong emphasis can mislead doctors. Close calls happen more than expected.

Training works better when settings such as how fast it learns, how much data it sees at once, which method updates the weights - like Adam or SGD - and how many times it runs are tuned using feedback from the validation phase. The way training is handled, along with changes to the learning
speed over time, turns out to shape both final precision and how quickly it settles on answers.

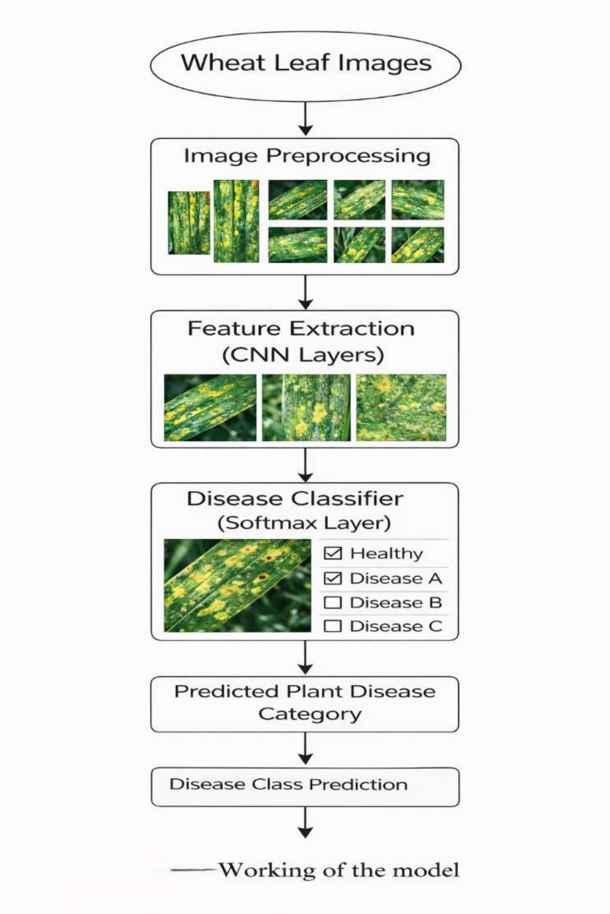

Fig-1 Model Evaluation

Fig 1 shows the complete process flow of the wheat disease detection system. Firstly, the system collects the wheat leaf images, which serve as the input data for the algorithm. Some of these images can be healthy leaf images, while others can be those suffering from various types of wheat diseases.

As shown in Fig 1, in the second step, image preprocessing is conducted. This involves the normalization, augmentation, and resizing of the images to ensure optimal image quality for the algorithm. Through image preprocessing, noise can be minimized, and changes in lighting conditions can be addressed.

Next, in the feature extraction step, the images go through the convolutional neural network layers. Here, the CNN automatically performs feature extraction from wheat leaf images. It detects various features from the images, such as texture, color, lesions, and disease spots.

The detected features are then sent to the disease classifier module. The disease classifier uses a Softmax activation layer to classify the disease into multiple classes. It estimates the probability of belonging to each class and chooses the class with the highest probability.

As a final step, the system provides the disease category of the plant as an output.

C. Transfer Learning and Model Types

Farm-focused image data might be limited, yet models trained on broader visual sets bring useful groundwork. Instead of building networks from scratch, some approaches tap into existing frameworks - ResNet or MobileNet, for instance - that already grasp complex patterns. Every class has a chance to win. Refined knowledge from ImageNet large datasets mad early performance of these models perfect. Subsequent learning speeds up due to inherited knowledge improving accuracy for deployment to crop related tasks.

There are variants: Lightweight models that require mobile devices, MobileNet brings the operations closer to the data source. MnasNet comes a close second with a mobile architecture aiming for speed and accuracy with sparse resources. These are useful for training when resources are scarce. ResNet and DenseNet use layers to squeeze all they can get. They search what others don't. They are deep enough to do so. Each learn what the other learn. They don't have shallow nets but go deep to find subtle characteristics

D. Training and Validation

Split takes place, eighty percent, ten test data, and ten validation data set. The system improves as it learns, using categorical cross entropy to minimize errors thanks to Adam and the works. Results show how often it was right how many mistakes it made, recognize patterns, a grid of results. If memorizing is excessive, the breaks hit automatically or learn rates fall.

E. Implementation

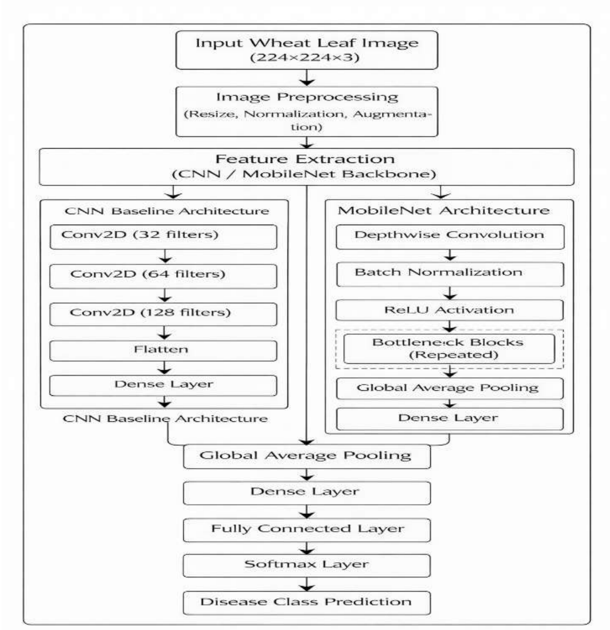

Fig-2 Overall Architecture

Fig 2 illustrates the structure of the developed framework for wheat diseases' identification based on CNN and MobileNetV2 approaches. First, the input images of size 224×224×3 are preprocessed using the methods such as resizing, normalization, and data augmentation to enhance data quality and model generalization ability. Then, these images are used as an input to the feature learning stage where both CNN and MobileNetV2 structures are applied for feature extraction.

In particular, the CNN baseline approach involves applying the Conv2D layers of different filter numbers and sizes before the flattening and dense stages for the purpose of extracting relevant information from input images. MobileNetV2 algorithm, on the other hand, uses depthwise convolution and bottleneck blocks with batch normalization, ReLU activation, and pointwise convolution. Next, the learned features undergo the global average pooling layer, and finally, dense and Softmax classifier stages are applied to identify the class of disease.

F. Deployment Considerations

Not all models need heavy computing to succeed. Some of the lightweight versions run happily on a phone. When a farmer takes pictures of a plant's leaves, the image can be instantly ingested by an app. The likely diseases are returned in short order. Leaner architectures learn better when the connectivity or computing power is poor. Recent literature shows that leaner architectures learn better in adverse conditions.

A recent survey on CNN-based approaches to detecting plant diseases shows tangible results over various crops. Not exclusive to one type of disease, the method reports results with visible success. Where traditional detections do not succeed, this model is more sensitive. Performance was boosted for unseen examples by augmenting the data with more samples. Early layers of the network had features like discoloration, surface patterns and bent leaves clearly, while later layers cut the noise and did not depend on irrelevant details for the results. Each test run reveals stronger patterns in spotting trouble early. With diverse plants involved, outcomes stay consistently useful. Instead of guesswork, clear signals emerge from image data. So far, every trial supports its growing role in farm-level monitoring.

G. Experimental Setup

A custom setup ran on Python, relying heavily on TensorFlow along with Keras tools. Performance checks unfolded across long sessions on powerful hardware fitted with a fast graphics chip. That machine helped shorten learning time dramatically. Data split carefully into three chunks - most went to practice runs, some to fine-tuning, the rest reserved strictly for final evaluation - all carved at an eighty-ten-ten cut.

H. Data preprocessing feature extraction

Preparing the data helped the model cope with variations better. By making sure it learns some extra, the accuracy of the output would be more reliable and predictable. And it would learn to habituate every corner of the image. So that it would work reluctantly. In which it tries to remove all the excessive details that produce noise and that would be depending very much on the noise.

I. Model Performance

The values are impressive, and precision and recall are high, which means most of the detections are accurate, and very few misclassifications occur. Good scores mean good results, but also depends on the level of cleanliness of the data being provided. When the results do not change with different tests, trust is formed, which shows dependability. Accuracy means an ideal condition where all the cases are correctly identified, but no false cases are identified. It is not just the results that make something successful, but also a well-rounded design decision made behind the scenes. F1-Score shows that there is no bias towards results that are favorable for diseased or healthy cases. The difficulty lies with leaves that overlap with one another and are packed very tightly. Future work could involve looking into other forms of imagery, such as beyond the spectrum of visible light, and incorporating design elements to allow for a focus. Another area of future work could involve connected devices, but it is difficult to keep it running smoothly outside.

A few look-alike plant illnesses - say, rust compared to brown spots - got mixed up more than once, showing up clearly in the error chart. Higher-quality pictures might help sort this out later on, along with smarter image focus methods drawn from newer models

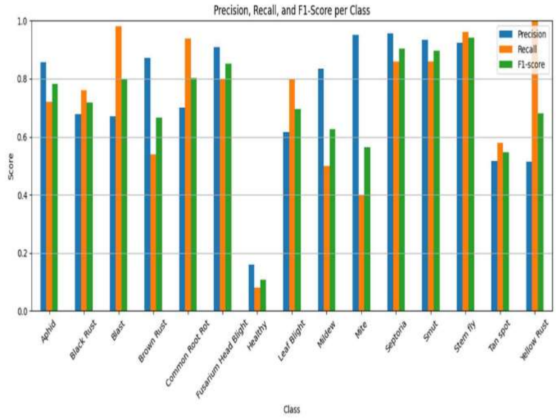

Fig 3 shows the class-wise evaluation of the developed wheat disease detection system based on the precision, recall, and F1-score evaluation metrics. It can be noted from the figure that the performance of the proposed detection approach can be seen for various wheat diseases like aphid, black rust, blast, brown rust, septoria, smut, stem fly, tan spot, and yellow rust.

The higher precision value means that the developed model can classify various disease classes in wheat crops with minimal number of false positives. On the other hand, the recall value reflects the capability of identifying samples of wheat disease classes. Finally, the F1-score takes into account both the precision and recall. It is shown that the model performs relatively good for various disease classes like septoria, smut, and stem fly, while some classes show poor performance due to class similarity, dataset imbalance, and differences in disease appearances.

J. Comparative Analysis

The project is compared between the two models in which those are CNN and Mobilenet in which those two models are widely used in image based detection projects. CNN has more standard layers which resulted in more number of parameters and complexity training time is longer and also it will be suitable to field use. Where as mobilenet light weight deep learning technique in which it specially designed for the computation. It reduces different layers with single separable convolution so gradually there will be decreasing of the accuracy. But, it fastens the speed of running time it can run efficiently on mobile with high performance.

https://doi.org/10.1051/itmconf/20268701012

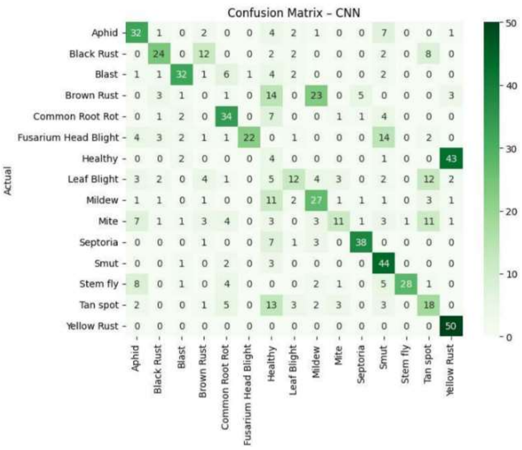

Fig-4 Confusion Matrix of CNN

Fig 4 shows the confusion matrix generated from the CNN model for wheat disease classification. The confusion matrix is used to evaluate the classification model's performance based on the comparison between actual and predicted disease classes. If the elements belong to the diagonal, then they are correctly classified. Otherwise, the elements correspond to misclassified instances between the different classes.

It can be seen that the model performed very well in terms of classification performance for some of the diseases like yellow rust, smut, and septoria where a large number of instances are accurately classified. Some of the other disease classes like leaf blight, tan spot, and mildew performed poorly due to the presence of similar symptoms in their appearance and leaf texture.

In conclusion, the confusion matrix reflects the CNN model's ability to capture discriminative disease features but still highlights some of the problems related to classification performance due to class similarity and imbalances within the data set.

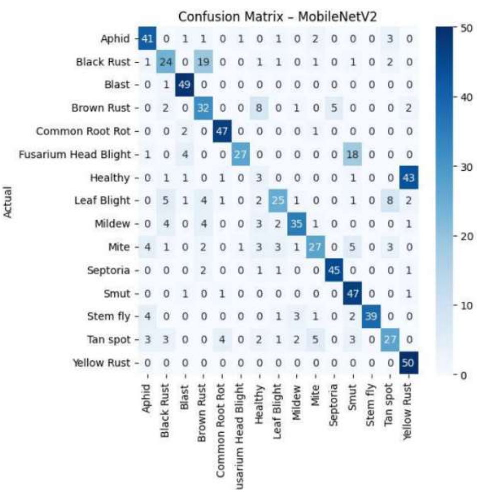

Fig-5 Confusion Matrix of MobileNetV2

Fig 5 shows the confusion matrix of the proposed classification model based on MobileNetV2. This confusion matrix reflects the relationship between actual classes and the predicted classes, where the diagonal entries show the correctly classified samples while the off-diagonal entries represent the wrongly classified samples.

The results reveal that the proposed classification model using MobileNetV2 has demonstrated superior classification capability than the traditional CNN classification model with respect to most of the disease categories. In other words, higher correct predictions have been obtained for diseases including blast, common root rot, septoria, smut, and yellow rust, which reveals the robust feature extraction capability of the proposed model MobileNetV2.

Even though some wrongly classifications exist for certain diseases, including leaf blight, mildew, and tan spot because of the similar symptoms, it can be seen that the proposed model demonstrates better generalization capability and lower classification error rate. It can be said that the superior performance of MobileNetV2 is mainly due to its lightweight depthwise convolution operation.

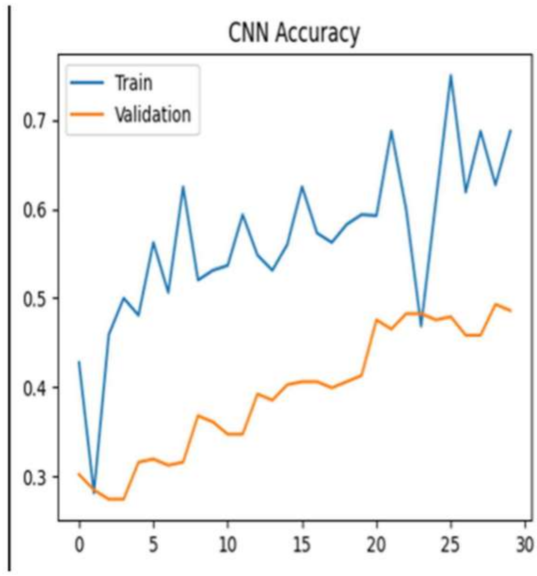

Fig-6 Accuracy of CNN

Figure 6 presents the training and validation accuracies for the proposed CNN-based wheat disease classification model over various training epochs. From the figure, the trend of the training accuracy clearly shows an improvement in the accuracy of the model during the training epochs, implying that the CNN-based model is able to learn the discriminative features from the input wheat leaf images. This improvement is reflected in the validation accuracy, which shows that the model performs well in generalization.

Fluctuations in the training and validation curves indicate difficulties in the process of learning, such as class imbalance, similar instances within different classes, and the small size of the available datasets. Overfitting can also be observed by comparing the training and validation accuracies in some epochs. However, the model performed satisfactorily in the process of classification.

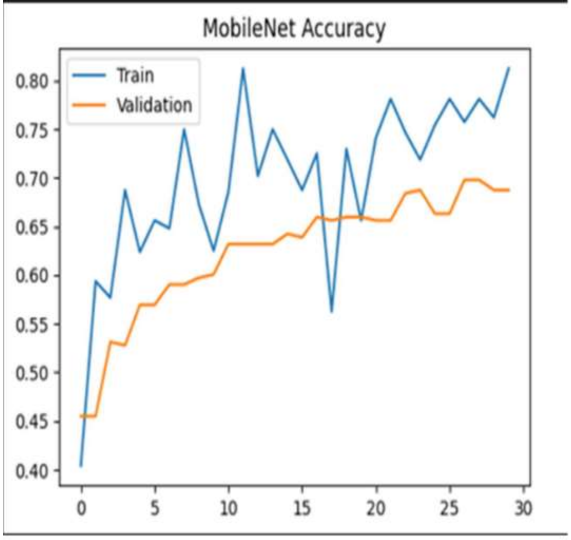

Fig-7 Accuracy of MobileNetV2

Fig. 7 depicts the training and validation accuracy of the wheat disease classification model using the architecture of MobileNetV2 for various numbers of training epochs. It can be observed that the accuracy of the proposed MobileNetV2-based model improves consistently for both training and validation sets, signifying good learning ability and enhanced feature extraction of MobileNetV2. The accuracy of the training data is greater than 80%, whereas the validation data is almost 69%–70% accurate.

This small difference between the accuracy of training and validation signifies better generalization ability and less overfitting of the model. Some minor differences in the training set could be due to imbalance in the dataset and similarities among different classes of diseases.

K. Practical Implications

Snap a leaf with your phone, get a quick read on sickness. That way, fixes start sooner, spraying stays accurate, crops grow better. Machines spotting issues help big farms keep track across wide fields. Bright sunlight might throw off the results, while shadows can confuse the system just as much.

IV. CONCLUSION

The research project focused on addressing the timely and efficient detection of diseases that affect wheat crops by proposing a classification framework that is founded on convolutional neural networks (CNN). The proposed classification framework effectively employs deep learning technology to remove the requirement of extracting numerous relevant features from images of wheat leaves. The framework is able to perform this routine task automatically. The images used to train this framework were preprocessed to account for differences in image quality as well as orientations of wheat leaves. All of this is possible because of the differences that can be encountered in real-life scenarios. The proposed classification framework has managed to conclusively prove that it is able to provide consistent results when it is used to differentiate healthy wheat leaves from other diseases. The hierarchical framework has managed to extract robust features that are associated with diseases that affect wheat crops. The performance of this framework has also managed to prove that it is a valid candidate when it comes to classifying diseases. The proposed classification framework is more accurate, efficient, and consistent when compared to traditional approaches. It is also not dependent on the presence of a trained expert.

The results obtained through this study indicate the possibility of using deep learning-based solutions for precision farming. The automated system of detecting wheat diseases can prevent farmers from incurring major losses due to crop damage and the misuse of pesticides through intelligent decisions. The methodology can be extended to other crops and diseases with just a few parameter modifications. CNN accuracy of the model is about 65% and the accuracy of the MobileNet is 81% which is of the training part.

Future research can extend the current framework to incorporate the detection of the severity of the disease, environmental factors like temperature and humidity into the model, and mobile computing applications. Additionally, the model can be extended to incorporate techniques from the domain of explanatory artificial intelligence for better trustworthiness of the model’s predictions. These can be considered as potential avenues for research to truly make the detection of the disease a key feature of smart and sustainable technology-based agriculture for healthier crops.

REFERENCES

[1] L. L. Deng et al., “Application of agricultural insect pest detection and control map based on image processing analysis,” Journal of Intelligent & Fuzzy Systems, 2020.

[2] G. G. Kennedy et al., “Pest pressure relates to similarity of crops and native plants,” Proceedings of the National Academy of Sciences of the United States of America, 2022.

[3] B. M. Atta et al., “Application of fluorescence spectroscopy in wheat crop: Early disease detection and associated molecular changes,” Journal of Fluorescence, 2020.

[4] W. X. Bao et al., “Identification of wheat leaf diseases and their severity based on elliptical-maximum margin criterion metric learning,” Sustainable Computing: Informatics and Systems, 2021.

[5] V. Devisurya et al., “Early detection of major diseases in turmeric plant using improved deep learning algorithm,” Bulletin of the Polish Academy of Sciences: Technical Sciences, 2022.

[6] Z. C. Jiang et al., “Recognition of rice leaf diseases and wheat leaf diseases based on multi-task deep transfer learning,” Computers and Electronics in Agriculture, 2021.

[7] H. Wu et al., “Identification of wheat leaf rust resistance genes in Chinese wheat cultivars and the improved germplasms,” Plant Disease, 2020.

[8] R. D. Barman et al., "Comparison of convolution neural networks for smartphone image based real time classification of citrus leaf disease," Computers and Electronics in Agriculture, 2020.

[9] K. Anitha et al., "Feature extraction and classification of plant leaf diseases using deep learning techniques," CMC—Computers, Materials & Continua, 2022.

[10] A. Tursunov, J. Y. Choeh, and S. Kwon, "Age and gender recognition using a convolutional neural network with a specially designed multi-attention module through speech spectrograms," Sensors, vol. 21, no. 17, Art. no. 5892, 2021.

[11] J. Zhao, Y. Fang, G. Chu, H. Yan, L. Hu, and L. Huang, "Identification of leaf-scale wheat powdery mildew (Blumeria graminis f. sp. Tritici) combining hyperspectral imaging and an SVM classifier," Plants, vol. 9, no. 7, Art. no. 936, 2020.

[12] K. P. Panigrahi, H. Das, A. K. Sahoo, and S. C. Moharana, "Maize leaf disease detection and classification using machine learning algorithms," in Progress in Computing, Analytics and Networking. Cham, Switzerland: Springer, 2020, pp. 659–669.

[13] K. Aurangzeb, F. Akmal, M. A. Khan, M. Sharif, and M. Y. Javed, "Advanced machine learning algorithm based system for crops leaf diseases recognition," in Proc. IEEE 6th Conf. Data Science and Machine Learning Applications (CDMA), Riyadh, Saudi Arabia, Mar. 4–5, 2020.

[14] W. Ullah, K. Muhammad, I. Ul Haq, A. Ullah, S. Ullah Khattak, and M. Sajjad, "Splicing sites prediction of human genome using machine learning techniques," Multimedia Tools and Applications, vol. 80, pp. 30439–30460, 2021.

[15] A. Paul, S. Ghosh, A. K. Das, S. Goswami, S. D. Choudhury, and S. Sen, "A review on agricultural advancement based on computer vision and machine learning," in Emerging Technology in Modelling and Graphics. Cham, Switzerland: Springer, 2020, pp. 567–581.

[16] A. Wójtowicz, J. Piekarczyk, B. Czernecki, and H. Ratajkiewicz, "A random forest model for the classification of wheat and rye leaf rust symptoms based on pure spectra at leaf scale," Journal of Photochemistry and Photobiology B: Biology, vol. 223, Art. no. 112278, 2021.

[17] L. C. Ngugi, M. Abelwahab, and M. Abo-Zahhad, "Recent advances in image processing techniques for automated leaf pest and disease recognition—A review," Information Processing in Agriculture, vol. 8, no. 1, pp. 27–51, 2021.

[18] E. Nicholls, A. Ely, L. Birkin, P. Basu, and D. Goulson, "The contribution of small-scale food production in urban areas to the sustainable development goals: A review and case study," Sustainability Science, vol. 15, no. 6, pp. 1585–1599, 2020.

[19] R. Barretto, R. M. Buenavista, J. L. Rivera, S. Wang, P. V. Prasad, and K. Siliveru, "Teff (Eragrostis tef) processing, utilization and future opportunities: A review," International Journal of Food Science & Technology, vol. 56, no. 6, pp. 3125–3137, 2020.