Physical Education, Health and Social Sciences

https://journal-of-social-education.org E-ISSN: 2958-5996

P-ISSN: 2958-5988

Classification of Tomato Leaf Diseases Using ResNet-18 Under Real Farm Conditions

Nabiha Babar[1]*, Noor Ul Ain Naveed[2], Eman Habib[3], Dr. Muhammad Milhan Afzal Khan[4], Mahnoor Khan[5], Syed Daniyal Hussain[6]

[1,2,3,4,5,6] Department of Computer Science, University of Agriculture Faisalabad, Pakistan, Email:
[1]nabihababar036@gmail.com, [2] noorulain.naveed123@gmail.com,[3] emanhabib572@gmail.com, [4]mkhan.msds25seecs@seecs.edu.pk, [5]daniyal786111@gmail.com

Corresponding Author's Email: milhankhan@uaf.edu.pk

DOI: https://doi.org/10.63163/jpehss.v4i1.1458

Abstract:

The tomato is cultivated throughout almost all the parts of the world, and provides food to millions of people. Sadly, a handful of leaf diseases are still ruining crops and no one is making any money from the small farms and there is a lack of food in much of the region. I don't know what these terms mean to an outsider, but these can annihilate all of the growers who depend on each kg of crop. Once infected, a single leaf can produce spores which infect the entire greenhouse or field in just a few days, ruining a good growing season. Inspection of each row of crops and an inspection of each plant is a time consuming process. No experienced scouts with their magnifiers and field books are able to see the first symptoms. Early signs of the disease can be a chlorotic halo, pinprick necrotic dot and a slight alteration of leaf texture. Pathogen symptoms are apparent only when symptoms first appear in the canopy, indicating that pathogens are well established in the canopy. Now the farmer has two options, spraying or not spraying the field. This results in a loss for him in both directions. In this context, the research of deep learning as a possible solution to the problem is necessary. CNN can process an image of a leaf within a fraction of a second, identify subtle differences the human eye wouldn't see in color and texture and classify the diseases with a surprising accuracy. Different architectures have been developed, including VGG, DenseNet and ResNet and with accuracy of 90% or higher. Most of these published papers use almost perfect laboratory images which are distraction free, are uniformly lit, have a white background and a centered leaf. The pictures are not representative of the situation on farms. The models are not very effective once they are put into use. There may be a glare or shadow over one leaf, the lighting may vary due to the presence of clouds or other plants may appear in the frame as stems and flowers vie for the attention of the network. These aspects may not be seen in training images in a lab setting due to the complexity of the images. The network may even label non-pathogenic soil particles as pathogenic, and may not detect any lesions, because of shadowing. As a consequence, this means extra need for chemicals or worse yet, unrecognized infections will result in outbreak of the epidemic. This research seeks to address this issue. A collection of over 2,900 images of tomato leaves were carefully selected from Kaggle and Mendeley Data to train and evaluate the ResNet-18 model. The dataset consists of clean laboratory pictures as well as messy field images. The background and lighting have not been sanitized nor normalized for the process, but they have been preprocessed in terms of the standardized resolutions and noise suppression. To simulate the variability of outdoor photography, without requiring thousands of additional field trips, data augmentation was used and included the addition of rotation, flipping, brightness jitter, adding synthetic shadows, and random crops. The model stood firmly, unmoving. On held-out test images it scored 95.63% accuracy, 95.80% precision, 95.60% recall, and 95.70% F1. Those figures indicate that even when the world around the leaf doesn't keep it tidy, a standard CNN with a realistic diet of photos can still be able to detect disease patterns. The skip connections helped the network learn the hierarchical features, while "noising" the gradient to capture color changes and edge orientations in the early layers, and creating the lesion shapes and pigmentation in the deep layers, specific to the disease in ResNet-18. Training loss also continued to decrease (0.6852 to 0.0829), and the validation accuracy continued to get better and stabilize at 96.75% to show that learning was indeed generalizing and not memorizing. A tool such as this would help farmers to be able to make a speedy diagnosis and treatment decisions before losses can be great. Leaf disease of tomato plants is an important problem in agriculture. Tomato leaf disease is an important disease of agriculture.

Keywords: Tomato Leaf Disease, Deep Learning, CNN, ResNet-18, Real Farm Conditions, Image Classification

Introduction:

The agriculture-based economy still remains in Pakistan and a significant portion of the rural labour force is working in agriculture. Tomato is one of the crops with high consumption by households and high usage in food processing in the country. Unfortunately, the leaf disease is still very susceptible to tomato plants. Foliage and fruit can suffer from diseases like early blight, late blight, Septoria leaf spot, leaf mold and bacterial spot which may decrease yields and quality of produce before farmers notice it is an issue. Traditional diagnosis of disease has been done by visual inspection of plant pathologists. This method is not only time consuming, tedious but also prone to human error (Abdullah et al., 2024). The pathogens invade sometimes in early season, and when they are readily detected by the unaided human eye, a large percentage of the crop area could be infested. Unfortunately, several of the tomato diseases are similar in appearance and if not carefully examined, can be confused and the wrong chemical control applied. Both the use of pesticides and losses in agricultural production will rise if disease control is conducted at the wrong time and/or way (Ala'a et al., 2024). The assistance can be found in the form of computer vision. Abdullah et al. (2024) have created an automatic classification system to extract color, texture and shape features from the leaf image and match them with the disease label using CNN. Self-attention based on image context at the global level, known as Vision Transformers, further enhanced the accuracy of image recognition especially in very early symptoms (Karimanzira et al., 2025). But, it is evident there is a hard disconnect. Published data are mainly from the laboratories (standard lighting, simple background, centered leaves and no occlusion). This clean data is not a representation of the dirtiness of real farms, dappled sunlight, wind blown shadows, adjoining crops encroaching the image, soil debris cluttering the image, etc. (Bouni et al., 2024). Few studies have explicitly tested the ability of their classifiers to endure these field disturbances and for farmers to have tools that are very effective in the journal paper but not the paddock. In this study, we shall concentrate on this aspect. In this study, we investigate if a regular CNN (ResNet-18) could achieve good accuracy when trained with a carefully designed mixture of lab and field images on the task of five tomato leaf conditions classification.

Background and related works:

Agriculture contributes substantially to Pakistan's GDP and provides livelihoods across rural communities. Tomato, in particular, appears daily in domestic kitchens and industrial food lines. Blight, late blight, and bacterial spot, however, can devastate a planting if left unchecked, reducing both tonnage and market grade (Abdullah et al., 2024).

Traditional scouting relies on human scouts to go from row to row, inspect individual leaves and document the results manually. It works, but it's slow, exhausts inspectors and misdiagnoses occur. Early infection is particularly dangerous as the symptoms are mild and may be mistaken for mechanical bruising and/or nutrient deficiency. Unfortunately, the guesswork by farmers which is not conducted by a trained pathologist often leads to overexposure of the correct fungicide, at the wrong time or rate, causing farmers to waste money and contaminate soil and water (Ala'a et al., 2024).

It was the advent of deep learning that altered the scene. CNNs are now used to scan the images of leaves and extract the discriminative patterns, such as color mottling, necrotic edges, chlorotic halos, without resorting to hand-crafted feature engineering (Abdullah et al., 2024). Several networks like VGG, AlexNet and ResNet have been deployed for the service of plant disease with satisfactory performance.

Despite this, there is a laboratory bias in the literature. The majority of training collections are done in a soft box set-up and at night time using white cards. The pictures of leaves seem very different from those found in catalogs, and it certainly looks like a different kind of soil, as well, since the leaves are half in the shadow and the background is clay soil. If a model is able to match that real world image, the accuracy of the model drops (Bouni et al., 2024).

Abdullah et al., 2024 used transfer learning with the aim of classifying tomato diseases but the dataset had still a high number of laboratory data. Bouni et al. (2024) used a combination of handcrafted and deep features, but, once more, failed to do extensive field testing. There has been a lack of practical recognition of the need for classifiers to cope with random farm environment.

Materials and Methods:

The process is shown below in steps:

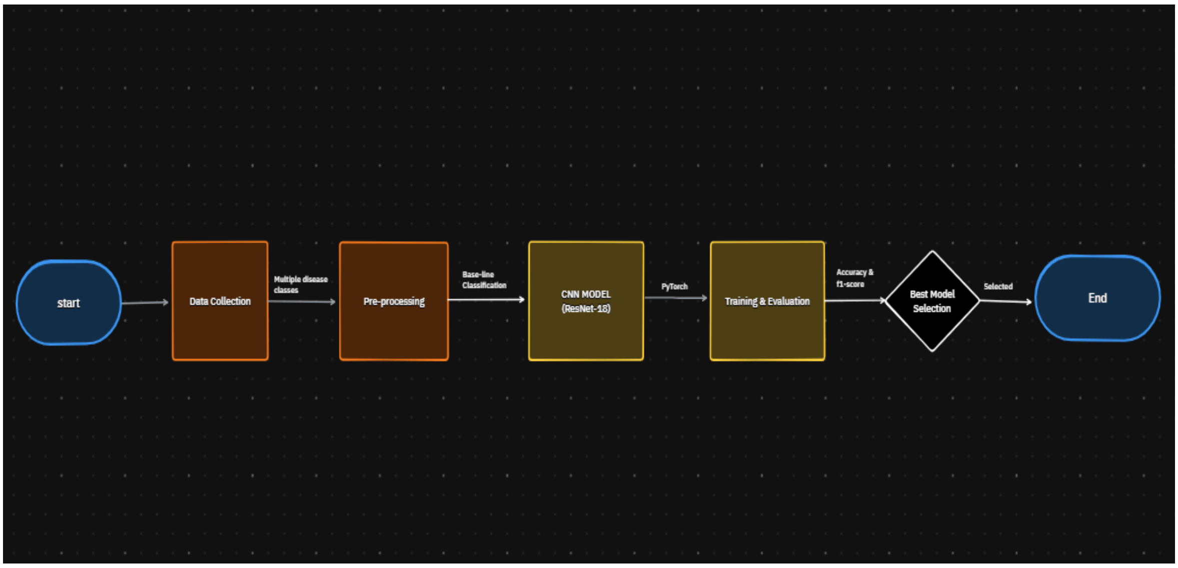

Figure 1.1 Road Map of Model's application on Dataset

Overall Framework

There are 4 stages on the pipeline. Firstly, public repository image capturing and public fields image capturing. Secondly, the size normalization, scale adjustment and noise reduction in the pre-processing of the image pixels. Thirdly, to train ResNet-18 on the cleaned and augmented data. Fourth, analysis based on standard measures of how well the model can be applied to out-of-sample instances of field photographs.

Dataset Collection

We collected pictures from Kaggle and Mendeley Data repositories. Along with healthy foliage are leaves of healthy and diseased plants. Of course, what was important was our non-laboratory method to capturing as we were able to capture images that were taken directly in the commercial tomato fields with natural daylight. That combination renders it more difficult for the model to cope with lighting changes, shadows, background noise and occlusion that it would encounter if it were actually used by a farmer's smartphone.

Data Preprocessing

We will first be working on processing the pictures, in order to make our model better.

- We will ensure that all the pictures are of the same size.
- We will ensure that all information in the pictures is on the scale.
- We will eliminate any unnecessary elements in the image and will de-noise.

Data Augmentation

To recreate realistic scenarios and provide more variation a data augmentation technique will be employed:

- Rotation and flipping
- Adjustment of brightness and contrast
- Shadow simulation □
- Zooming and cropping

A disease spot looks like a disease spot in any type of background, whether morning shade, mid day glare, or a brown soil background.

Model Development

To use ResNet-18 as the baseline classifier we chose. It carried over some legacy features in the form of skip connections which allowed the gradients to pass through layers without being degraded. When learning on moderately-sized agriculture datasets where vanishing gradients could slow down learning, it is important to have stability as important. The final fully connected layer was modified to yield five classes, corresponding to the different kinds of tomato leaf diseases in our collection. The stacked convolutional filters generate a multi-level representation: the low-level filters represent simple color, edge orientation, texture; higher levels of the filters represent irregularities in shape and pigment patterns specific to the disease. The network is trained with local cues, but very discriminative, which enables it to recognize the differences between healthy and unhealthy tissues even if there is noise in the background.

Implementation Tools

- Programming Language: Python
- Frameworks: PyTorch
- Platform: Google Colab

Performance Evaluation Metrics

The model will be evaluated by using following metrics:

- Accuracy
- Precision
- Recall
- F1-score

Results:

The new system should meet the following needs:

- Accurate classification of diseases to a high level of accuracy.
- Have a good working ability in practice on the farm.

The CNN model (ResNet-18) is giving results as follows:

CNN Model (ResNet-18):

The ResNet-18 was used as the backbone networks for our experiments. Maintains the gradient magnitude over 18 layers, which enables more efficient optimization and convergence than just stacked convolutions. The first layer was re-designed to discriminate five classes. The convolutional stack receives a series of abstraction layers of visual input: The first few layers are basic color and edge detectors; The middle layers are texture and shape of lesions; The last layers are high-level representations of representations of specific diseases, near the output. That the hierarchical feature extraction enables the model to show very subtle differences between healthy leaves and infected leaves. The teaching and learning process was successful. Loss was decreasing continuously from 0.6852 to 0.0829, and it means that the weights were tuning well, the optimizer was not stuck in bad local minima. The accuracy of the validation was enhanced parallel with the progress and reached 96.75%. The accuracy of the test partition (held out) is 95.63%. The accuracy of validation and test data are given, which shows a good generalization, not memorization of training data. The final epochs of the training and validation losses were very comparable, with a slight difference between them. A small difference could be attributed to limited synthetic diversity and size of the data, and both curves have been similar and suggest a good bias/variance trade-off. The CNN architecture was found to be successful in capturing the local spatial dependencies, that are important for disease discrimination. In particular the network was trained to recognize:

Brown with irregular spots and discolorations (typical of infectious leaf spots) Changes in shade of color (particularly chlorosis, yellowing and brown necrosis from specific pathogens) Brown with irregular spots and discolorations (typical of infectious leaf spots) Changes in shade of color (particularly chlorosis, yellowing and brown necrosis from specific pathogens).

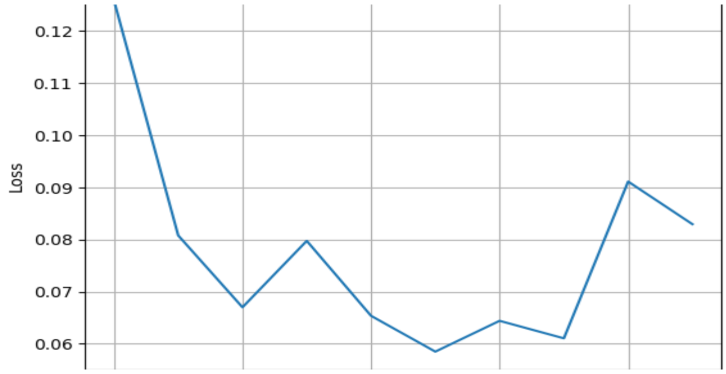

Tissue death is limited and results in a ring-spotting and a loss of marginal color. These local image features have a high classification score which shows the suitability of these features for separating the tomato leaf diseases even in the presence of other distractive field imagery.

Fig 1.2 : CNN Training Loss Curve

Metric	Value
Accuracy	95.63%
Precision	95.80%
Recall	95.60%
F1-Score	95.70%

Fig 1.3 : Performance Metrics of CNN Model.

Discussion

We decided to make a system for the diagnosis of diseases that could work out of the lab. Real world data was used for both training and testing purposes, with clean catalog images and photos taken in the field from tomato crops, under uncontrolled conditions. ResNet-18 did just that, and did it with over 95% accuracy with no use of any magic architectures or massive compute power. Residual connections are an important part and take a lot of the credit. They establish short-cut connections between the gradient flow, allowing the network to learn more complex visual patterns without having to remember background or lighting conditions from the training set. Even if the farmer takes a picture of a leaf at dusk, in the shade or when the soil and the stems make the picture more complex, the classifier is helpful. However, there are risks associated with deployment in practice as well. However, this performance could over time be reduced because of what is called data drift: the occurrence of new diseases, new sampling sensors, or geographic variations of how the disease manifests symptoms. We made use of several public sources and aggressively supplemented to minimize that risk but it might be prudent to periodically monitor and train. To conclude the study, the authors have shown that the CNN trained with various datasets is able to close the gap between the lab and field that has previously prevented the study. The farmers are given a classification tool which is actually feasible in the field, as opposed to a controlled studio environment.

Future Work

ResNet-18 is already a good system; and there are a couple of directions that can improve the system. For instance, a hybrid CNN-Transformer model can be used for both local spatial feature extraction and global contextual attention, while accurately enhancing the performance in difficult boundary scenarios where there is a boundary between the disease and the healthy tissue. Data volume and diversity is a bigger problem, however. Our 2900 images are just a beginning but the different climatic zones, in which tomato is grown, have different soils, different light conditions and different types of diseases. Incorporating additional photovotes of different seasons, different provinces in Pakistan and neighbouring countries would also make the model more robust against environmental drift and make it less sensitive to the current training distribution, which would lead to less overfitting. Finally, it was necessary to make the structure (with its proposed functions) practical for the real world. If this mobile application was developed, the classifier could now be put in the farmer's hands, on a mid-range cell phone, both online and offline. The next important step is then to minimize model sizes and inference latencies while maintaining accuracy.

Conclusion:

The leaf diseases in tomato grow rapidly and are very contagious; however, most of the classifiers that are currently available are deep-learning classifiers developed from laboratory images, which are far from the fields where farmers cultivate the crop. We showed that a ResNet-18 that is trained on a set of clean images along with real-field images is able to classify the test images with an accuracy of 95.63% in spite of the natural light, shadow, and cluttered backgrounds. The rest of the architecture is trained on global visual characteristics of the image that are very reliable, like color changes, necrotized patches and ring lesions, that are discriminative in diverse visual contexts. This type of classifier can be a helpful low-cost tool in the rapid diagnosis of diseases and timely intervention for growers in other agricultural conditions, such as in Pakistan.

References

Abdullah, A., H. Faisal and M. Khan. 2024. Detection of tomato leaf diseases using deep learning based transfer learning models. Agronomy 14(2): 441–452.

Abouelmagd, M., T. El-Sayed and L. Jing. 2024. Capsule network–driven tomato leaf disease classification. Neural Computing and Applications 36(12): 7325–7340.

Ala'a, H., S. Masood and F. Khan. 2024. Multi-class tomato leaf disease dataset development and CNN-based classification. Plants 13(7): 1255.

Alruwaili, M., S. Alotaibi and M. Khan. 2022. Real-time plant disease detection using RTF-RCNN. Computers and Electronics in Agriculture 198: 107002.

Alzahrani, M. S., A. Alanazi and Y. Khan. 2023. Vision transformer–based tomato leaf disease recognition for precision agriculture. Agronomy 13(11): 2660.

Bouni, M., N. Merzougui and A. Chikha. 2024. Deep and handcrafted feature fusion for tomato leaf disease classification. Scientific Reports 14: 12671.

Gopal, R. and N. Singh. 2024. EfficientNet-based tomato leaf disease detection for optimized inference speed. Multimedia Tools and Applications 83(4): 5521–5538.

Hossain, A., F. Rahman and M. Ahmed. 2024. Deep preprocessing–assisted tomato leaf disease detection using CNN architecture. The Plant Pathology Journal 40(3): 357–368.

Islam, S. U., H. Ali and R. Hussain. 2025. Tomato leaf detection and segmentation using hybrid U-Net and CNN classifier. Agriculture 16(3): 511–522.

Karimanzira, D., B. Chavez and J. Weber. 2025. Context-aware tomato leaf disease detection using attention-based vision transformer models. Electronics 14(1): 122–138.

Li, Z., M. Zheng and C. Zhou. 2023. Tomato leaf disease recognition based on GD- attention CNNs. Plants 12(4): 821.

Rasheed, M., L. Ahmed and K. Iqbal. 2024. Deep feature extraction using ResNet and DenseNet for tomato leaf disease detection. Alexandria Engineering Journal 72: 253–265.

Raza, A., A. Shahid and S. Latif. 2025. Ensemble CNN framework for tomato leaf disease recognition using multiple learning streams. Expert Systems with Applications 244: 123012.

Sowmya, P., R. Dash and A. Tripathy. 2025. Tomato leaf disease detection using Inception-YOLO object localization. Computers and Electronics in Agriculture 214: 108623.

Sun, H., Q. Li and A. Sohail. 2025. Lightweight tomato leaf disease identification using optimized CNN models for real-time agricultural applications. BMC Plant Biology 25: 991.