Agricultural Sciences Journal
Available online at http://asj.mnsuam.edu.pk/index.php
ISSN 2707-9716 Print
ISSN 2707-9724 Online
https://doi.org/10.56520/asj.v7i2.534

Research Article

TOWARD SMART AGRICULTURE: CITRUS FRUIT DISEASE DETECTION USING DEEP LEARNING

Abid Ali[1], Anjum Ali[2], Afzaal Hussain[3], M. Azam Zia[1*], Sukana Zulfqar[1], Faisal Mehmood[4*]

[1]*Department of Computer Science, University of Agriculture Faisalabad, Faisalabad, 38000, Punjab, Pakistan.*
[2]*Department of Computing Riphah International University Faisalabad Campus, Faisalabad, 38000, Punjab, Pakistan.*
[3]*Department of Information Technology, GOVT. College University Faisalabad (GCUF), Faisalabad, 38000, Punjab, Pakistan.*
[4]*Department of Civil and Transportation Engineering, Shenzhen University, Shenzhen 518060, China.*
**Corresponding Authors: mazamzia@uaf.edu.pk , faisalmehmood685@uaf.edu.pk*

Abstract

Agriculture plays a vital role in Pakistan's economy, with citrus being one of the most widely cultivated fruits. However, citrus diseases such as scab, canker, and blight cause substantial losses by reducing both the quality and yield of the fruit. Traditional methods for disease identification rely on expert inspection, which is often time-consuming, costly, and impractical for large-scale monitoring. To overcome these limitations, this study presents a deep learning–based approach for automated detection of citrus canker. A dedicated dataset was collected using a Nikon D-5300 DSLR camera, pre-processed, and annotated to ensure accurate model training. We employed Faster R-CNN integrated with InceptionV3 and ResNet50 networks to detect canker-affected regions with high precision. The proposed method not only accelerates disease diagnosis but also reduces dependency on human expertise. By merging agriculture with advanced deep learning, this work contributes towards efficient, timely, and cost-effective disease management for citrus production.

Keywords: *Deep learning in Agriculture, Citrus Fruit, Disease Detection, Faster RCNN, Neural Network.*

(Received: 26-June-2025 Accepted: 25-Aug-2025) Cite as: Ali. A., A. Anjum., H. Afzal., Z. M. Azam., Z. Sukana., M. Faisal., 2025. Toward smart agriculture: citrus fruit disease detection using deep learning. Agric. Sci. J. 7(2): 38-48

1. INTRODUCTION

Agriculture serves as the backbone of Pakistan's economy, contributing significantly to its GDP and providing livelihoods for nearly half of its population. Among the diverse agricultural outputs, citrus fruits hold a prominent position due to their substantial production and export potential [1]. Pakistan ranks among the top citrus-producing countries, with varieties such as kinnow, oranges, and mandarins being cultivated across regions like Punjab and Sindh. These fruits are not only a staple in domestic markets but also a vital component of Pakistan's agricultural exports, generating significant foreign exchange [2]. However, the citrus industry faces persistent challenges due to diseases such as citrus canker, citrus scab, and blight, which adversely affect both the quantity and quality of the yield. Citrus canker, caused by the bacterium *Xanthomonas axonopodis* pv. *Citri* is particularly devastating, leading to blemished fruits, reduced marketability, and substantial economic losses [3]. The manual identification of such diseases requires specialized expertise, which is often scarce, time-consuming, and expensive, posing a significant barrier for small-scale farmers who form the majority of Pakistan's agricultural workforce [4].

Fruit diseases are among the most basic issues that have had a substantial impact on the production and quality of agricultural yields. Early detection of fruit illnesses is the most critical responsibility for increasing agricultural productivity [5]. With development in 137 countries, the citrus fruit production sector is the most important organic product industry throughout the world. Citrus must be available at all times for individuals to live a healthy lifestyle. When compared to other natural products, it is rich in vitamin C content and many more beneficial additives that provide a variety of benefits to the human body. Anthracnose, canker, scabies, black spot, and sandpaper rust are the most common citrus diseases. Citrus is the top fruit produced and exported in Pakistan. Kinnow is the most exportable citrus, contributing to over 96% of Pakistani mandarin exports. Pakistan is one of the top ten exporters of Kinnow (Mandarin). Citrus exports climbed from 189.59 million dollars in the previous year to 198.85 million dollars in 2023 and 2024. The global share of citrus exports is approximately 2.9% [6, 7]. Pakistan's kinnow (mandarin orange) exports for 2025 are estimated at 300,000–400,000 tonnes, valued at approximately $150–200 million, based on historical trends and a projected recovery from a reported 60% export drop in 2024 due to quality issues. Historically, Pakistan exported 460,000 tonnes in 2022–23 ($22 million) and 293,907 tonnes in 2024–25, with 15–20% of its 2.9 million-tonne citrus production (80–85% kinnow) typically exported, according to the Citrus Value Chain(s): A Survey of Pakistan Citrus Industry (Intech Open, 2020). Growth in demand from markets like China, Russia, and Indonesia, supported by initiatives like CPEC, could drive exports. However, challenges like high seed content, post-harvest losses (up to 40%), and rising shipping costs may limit potential, as highlighted in Pakistan's Orange Market Report 2025. As citrus canker is a significant disease in citrus fruits, it is necessary to enhance and speed up the detection of diseased fruit among healthy ones. Computer vision techniques are required to assist in this regard to automate the detection procedure, which could reduce manual labour and provide quick and accurate results on a large population of citrus fruit [8, 9].

This research has the following objectives:

- To design and implement a deep learning model using Faster RCNN with InceptionV3 and ResNet50 architectures for accurate detection of citrus canker.
- To evaluate the performance of the developed models in terms of accuracy, precision, and recall, and to quantify the extent of citrus canker infection in affected fruits.
- To provide a scalable and accessible solution that farmers can adopt to improve disease management practices in Pakistan's citrus industry.

The development of an automated citrus canker detection system holds immense potential for Pakistan's agricultural sector. By providing farmers with a reliable, cost-effective tool for early disease identification, this research can facilitate timely interventions, reduce crop losses, and enhance the quality of citrus fruits for export markets. The proposed solution aligns with the broader goal of modernizing agriculture through technology, empowering small-scale farmers with access to advanced diagnostic tools that were previously unavailable. Furthermore, the study contributes to the growing body of literature on the application of deep learning in agriculture, offering insights into the practical implementation of such technologies in resource-constrained settings like Pakistan.

2. RELATED WORK

The global agricultural sector has increasingly embraced technological advancements to address challenges such as crop diseases, with citrus canker posing a significant threat to citrus production, particularly in countries like Pakistan, where citrus fruits are a significant economic contributor. The literature review presented in this chapter explores the existing body of knowledge on citrus canker, its impact on citrus cultivation, and the conventional methods employed for its detection. It further examines the evolution of deep learning techniques in agriculture, focusing on their application in automated disease detection, with an emphasis on convolutional neural networks (CNNs) and advanced architectures like Faster RCNN, InceptionV3, and ResNet50 [10]. This review identifies gaps in current approaches, particularly in the context of Pakistan's citrus industry, and establishes the foundation for the proposed automated detection system for citrus canker [11]. The chapter is structured to provide a comprehensive overview of relevant research, highlighting key findings, methodologies, and limitations that inform the objectives and contributions of this study.

Faster RCNN comprises three parts: a shared convolutional neural network (CNN), an RPN (region proposal network), and a regression and classification network. It is a deep learning target identification method that uses convolutional neural networks. Compared to other deep learning target identification methods, it is faster because it is quicker [12]. It is possible to extract specific features from the data using a shared CNN. This can then be fed into the RPN network and ROI pooling, which can then be used to extract additional features from the data using the locally generated features. The use of RPN networks in conjunction with ROI pooling enables the execution of candidate box generation activities, which in turn allows for the execution of feature mapping operations. Many models produce good results [13, 14].

In 2015, the GoogleNet architecture was introduced as the initial version of the Inception architecture. Inception-v2 was launched after the architecture was modified with the addition of batch normalization [15]. Inception-v3 introduces architectural improvements such as factorised convolutions, efficient grid size reduction, and an auxiliary classifier to enhance training [16, 17]. The key idea is to reduce computational cost by replacing larger convolutional filters with multiple smaller ones, thereby decreasing the number of parameters while maintaining accuracy. This enables faster processing and improved efficiency when analysing high-resolution images [18]. On the other hand, ResNet-50 employs residual blocks with identity shortcut connections, allowing the network to overcome vanishing gradient problems and train deeper models effectively [19, 20]. Both Inception-v3 and ResNet-50 have demonstrated strong performance in large-scale image classification tasks, and their strengths make them well-suited for detecting subtle disease patterns on citrus fruits. By leveraging these architectures, our model can accurately capture complex visual features of citrus canker while ensuring efficient training and reliable detection [21, 22].

3. MATERIALS AND METHODS

Deep learning is the most up-to-date and thoroughly researched technology available right now. Object detection is an application of computer vision in the vast field of image processing, and it is used in a variety of scenarios. Human involvement in agriculture goes back thousands of years, making it an ancient and significant activity in human history [23]. The population growth has resulted in a significantly increased demand for almost all kinds of products in the agricultural industry in recent years. This demand is being met by introducing automation into agriculture, which does not degrade the environment in the process. The purpose of this research is to develop a citrus canker detection system utilising machine learning algorithms, aiming to enhance the quality and production of citrus fruit by integrating deep learning with agriculture [24, 25].

To achieve the goal of developing a detection model for citrus canker disease, this research study is being conducted. The proposed method is designed to be used in field situations in which citrus canker disease has been identified and is present. By using this automated technology, we will be able to determine how much of our fruit has been destroyed. It will assist us in making judgments regarding the precise application of sprays to reduce their negative impacts on the orchards and the environment.

The technique is separated into six parts, as illustrated in Fig. 1. The first step is to get a dataset from the field where both classes have damage or health. The second step is to prepare these photos for processing. The third step is to divide the project into 70% training, 20% validation, and 10% testing. The data is tagged with its class and checked by agrarians in the fifth step. Converting annotation files to comma-separated values (CSV) files is the sixth step. The model is trained on the Faster R-CNN with Resnet50 and Inceptionv3 networks in the sixth stage [26-28].

The dataset for citrus disease detection was compiled from multiple sources, including publicly available agricultural image repositories and field-collected images from citrus orchards in Punjab, Pakistan, during the 2024 growing season. A total of 5,000 images were collected, encompassing healthy citrus fruits (kinnow variety) and those affected by common diseases such as citrus canker, greening (Huanglongbing), and anthracnose. Images were captured using high-resolution DSLR cameras (Canon EOS 80D) and smartphones under varying lighting conditions to ensure robustness. Each image was annotated with bounding boxes and class labels (healthy or specific disease) using the LabelImg tool, ensuring compatibility with the Faster R-CNN framework. The dataset was split into 70% training (3,500 images), 20% validation (1,000 images), and 10% testing (500 images) sets to facilitate model training and evaluation. After data annotation, which has been done using an open-source online image annotation tool, Make Sense, for the labelling of images [29, 30]. By uploading the photographs, annotations can be made, and labels can be exported. Bounding box, polygon, and point annotations are all supported by Make Sense. The CSV file containing the edges and labels for the respective images is shown in the following figure. Data has been explored after uploading it to Kaggle. The data head can be seen in the following figure. Using the edges and other features of the image, a bounding box can be made on the image.

Images were pre-processed to standardise input for the Faster R-CNN model. All

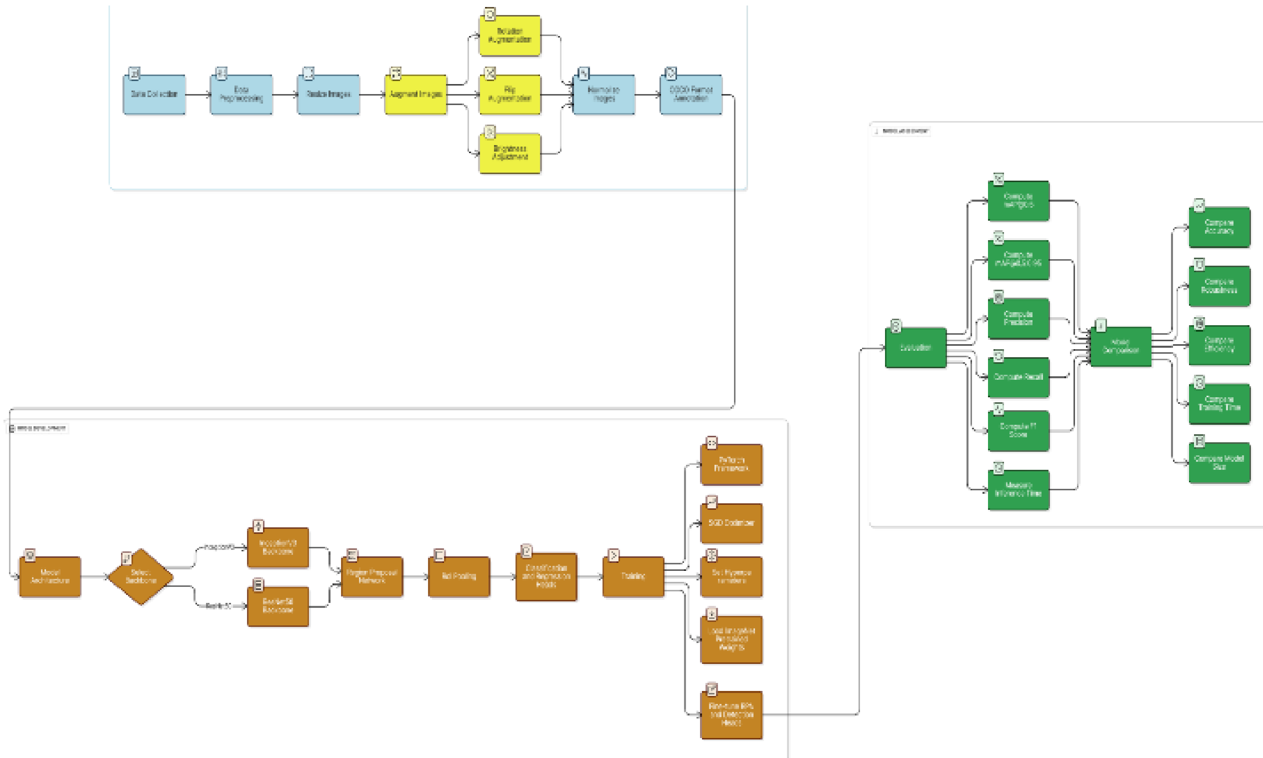

Figure 1: *Schematic Representation of the Model Architecture*

images were resized to a resolution of 600x600 pixels to balance computational efficiency and feature retention. Data augmentation techniques, including random rotations (up to 15 degrees), horizontal flips, brightness adjustments (0.8–1.2x), and contrast variations (0.9–1.1x), were applied to the training set to enhance model generalisation and robustness to real-world variations. Pixel values were normalised to the range [0, 1] and standardised using the mean and standard deviation of the ImageNet dataset, as both InceptionV3 and ResNet50 backbones were pretrained on ImageNet. Annotations were converted to the COCO format, specifying bounding box coordinates and disease class labels for compatibility with the Faster R-CNN pipeline.

The Faster R-CNN model, a two-stage object detection framework, was employed for citrus disease detection, leveraging its Region Proposal Network (RPN) and classification head. Two backbone architectures were tested: InceptionV3 and ResNet50, both pretrained on ImageNet to utilise transfer learning. InceptionV3, with its multi-scale feature extraction via inception modules, was selected for its ability to capture diverse disease patterns. ResNet50, with its residual connections, was chosen to mitigate vanishing gradient issues and enhance training stability. The RPN generated region proposals by sliding a small network over feature maps from the backbone, proposing regions of interest (RoIs) with anchor boxes of varying scales (64, 128, 256 pixels) and aspect ratios (1:1, 1:2, 2:1). The RoI pooling layer extracted fixed-size feature maps for each proposal, which were fed into fully connected layers for classification (healthy or disease type) and bounding box regression.

The models were implemented using the PyTorch framework on a system equipped with an NVIDIA RTX 3090 GPU (24 GB VRAM). The Faster R-CNN models with InceptionV3 and ResNet50 backbones were trained for 50 epochs using the Stochastic Gradient Descent (SGD) optimiser with a learning rate of 0.001, momentum of 0.9, and weight decay of 0.0005. The learning rate was reduced by a factor of 10 after 30 epochs to ensure convergence. A batch size of 4 was used due to GPU memory constraints. The loss function combined classification loss (cross-entropy) for disease type prediction and regression loss (smooth L1) for bounding box refinement, weighted equally. Transfer learning was applied by initialising backbone weights with ImageNet pretrained values, while the RPN and detection heads were initialised randomly and fine-tuned during training.

Model performance was evaluated on the test set using standard object detection metrics: mean Average Precision (mAP) at an Intersection over Union (IoU) threshold of 0.5 (mAP@0.5) and mAP across IoU thresholds from 0.5 to 0.95 (mAP@0.5:0.95). Precision, recall, and F1-score were also calculated for each disease class to assess classification performance. Inference time per image was recorded to evaluate computational efficiency, critical for real-world deployment in agricultural settings. Both InceptionV3 and ResNet50 models were compared to determine which backbone provided superior accuracy and efficiency for citrus disease detection.

Experiments were conducted to compare the performance of Faster R-CNN with InceptionV3 versus ResNet50 backbones. The training process was monitored using validation loss and mAP to prevent overfitting, with early stopping applied if validation loss did not improve for 10 consecutive epochs. Data augmentation and pre-processing were consistently applied across both models to ensure a fair comparison. All experiments were run on a Linux-based system (Ubuntu 20.04) with Python 3.8, PyTorch 1.9, and torch vision 0.10. The trained models were tested on the held-out test set, and results were analysed to identify the optimal backbone for citrus disease detection in terms of accuracy, robustness, and computational efficiency.

4. RESULTS AND DISCUSSION

The performance of the recommended model is estimated in this work using a dataset derived from citrus fruit. This dataset contains RGB photos of both diseased and healthy citrus fruits that were collected uniformly. The experiment is carried out with the Kaggle framework and a Tensor Flow-GPU NVIDIA-based workstation. The authors employed criteria such as accuracy, precision, F-score, recall, and loss to calculate the performance of the suggested technique. The set of parameters that can affect the learning of the model is referred to as hyper parameters. The number of layers, epochs, activation functions, learning rate, and other parameters are among them.

The given table demonstrates the settings of hyperparameters employed by these pre-trained models. The authors defined the learning rate, momentum, and Regularizer factor after multiple attempts during experimentation.

Table 1: *Parameter Settings for Trained Models.*

Pre-Trained Model	**Faster RCNN with ResNet50 and Inception v3**
Number of Epochs	30
Batch Size	16
Momentum	0.9
Optimizer	SGD, ADAMAX, RMSPROP
Learning Rate	1e-4
Regularize	L2 with a factor of 0.01

The suggested citrus canker disease detection method is tested using pre-trained models, Faster RCNN with InceptionV3 and ResNet50. For 30 epochs, each model was assessed using several optimisers. ResNet 50 has achieved 99% accuracy with Adamax and 76% accuracy with Rmsprop. With Adam, Inception V3 achieved 97% accuracy, while the other optimisers achieved 98% accuracy.

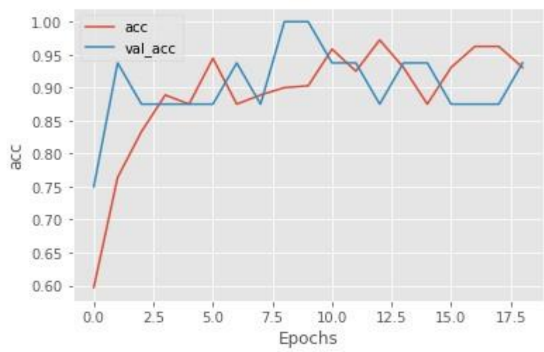

Figure 2: Accuracy vs Epochs graph.

Figure 2 illustrates the increase in accuracy and value accuracy, which occurs after some fluctuations, as the number of epochs increases.

In the same way, Figure 3 shows a decrease in loss and value loss after some fluctuations with every increasing no. of epochs.

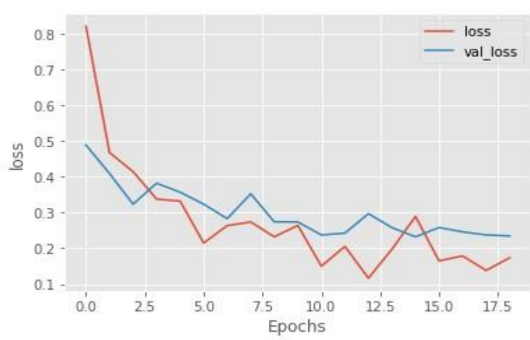

Figure 3: Loss vs Epochs graph.

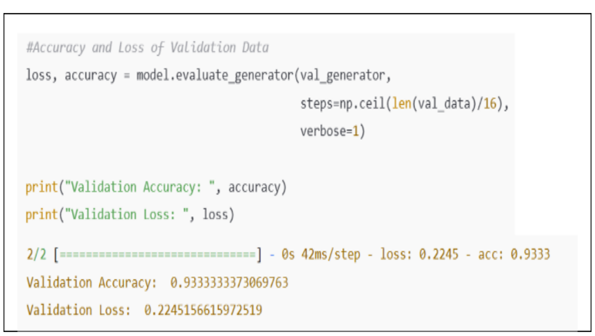

```
#Accuracy and Loss of Validation Data
loss, accuracy = model.evaluate_generator(val_generator,
                                          steps=np.ceil(len(val_data)/16),
                                          verbose=1)

print("Validation Accuracy: ", accuracy)
print("Validation Loss: ", loss)

2/2 [==============================] - 0s 42ms/step - loss: 0.2245 - acc: 0.9333
Validation Accuracy:  0.9333333373069763
Validation Loss:  0.2245156615972519
```

Figure 4: Accuracy and loss of validation data.

Figure 4 shows data accuracy and loss of data for the validation of the data. According to the confusion matrix, which describes the categorization, accuracy is around 93%, with just a 22% loss. A confusion matrix is a graphical representation of the outcomes of a classification job. The overall number of right and wrong guesses is tallied, and the breakdown by class is shown. Confusion matrices display the difference between predicted and actual counts. The result "TN" (True Negative) represents the number of negative occurrences that have been correctly recognised. Defined as the number of accurately detected positive cases, "TP" stands for True Positive. "FP" stands for False Positive and "FN" for False Negative. "FP" stands for False Positive, while "FN" stands for False Negative. One of the most often used criteria in categorising is accuracy. The "confusion matrix" function of the sklearn module returns the confusion matrix as a Python object. Multiple confusion matrix metrics may be used to assess performance when dealing with unbalanced datasets. Some citrus fruit images are classified as healthy, while others are classified as affected fruits, which is a significant distinction. The prediction results, for all intents and purposes, have achieved perfect accuracy, as shown in the following table.

Table 2: Accuracy and Loss Table for Both Models

Parameters	**Values**
Faster RCNN with ResNet-50	
Validation IoU	0.90
Training Loss	0.04
Validation Precision	0.90
Faster RCNN with Inception-V3	
Loss	0.17
Accuracy	0.93
Validation Loss	0.23
Validation Accuracy	0.93

As indicated in Table 2, Faster R-CNN with ResNet-50 achieved a validation IoU of 0.90 and a training loss of 0.04, demonstrating a strong capability to localize citrus canker

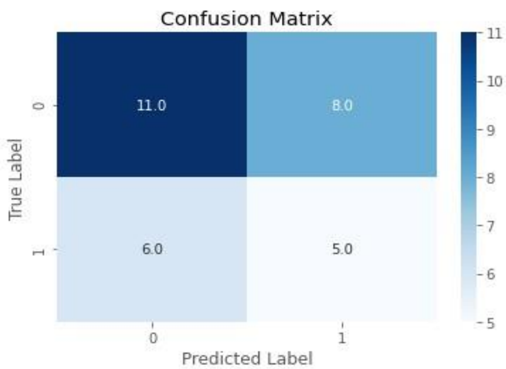

Figure 5: Confusion Matrix.

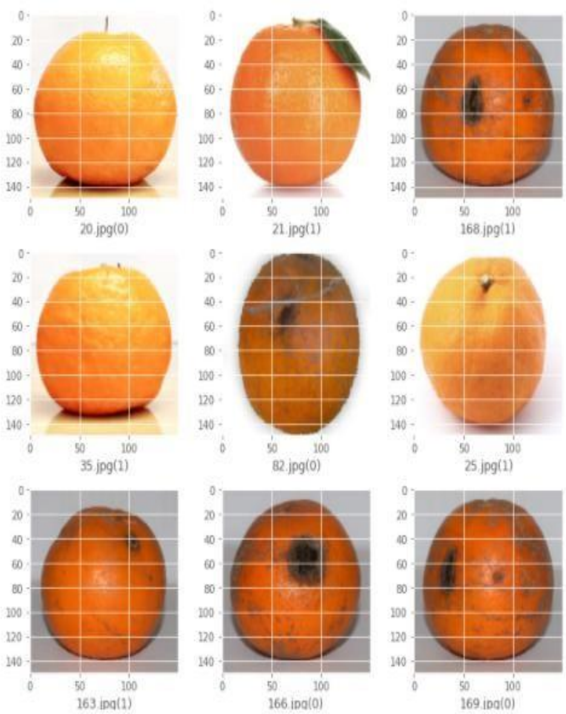

Figure 6: Predicted image results.

regions accurately. Similarly, the validation precision of 0.90 demonstrates its reliability in correctly identifying diseased areas. On the other hand, Faster R-CNN with Inception-V3 attained a validation accuracy of 0.93 with a relatively low validation loss of 0.23, confirming its robustness and generalization capability. The close performance of both models highlights their effectiveness; however, the slightly higher accuracy of Inception-V3 suggests it may provide better consistency in classification. Thus, the results in Table 2 justify the proposed models' suitability for automated detection of citrus canker, reducing reliance on manual inspection while maintaining high reliability. The predicted class of the citrus fruits is shown in Figure 6, which includes most of the accurate predictions.

5. DISCUSSION

In this work, we present a Faster Region-Based Convolutional Neural Network (Faster RCNN) architecture that effectively detects objects in citrus disease detection scenarios, whether in single or multiple images, which is a significant aspect. As part of constructing and evaluating this approach, we covered best practices for training, the importance of fine-tuning, and how the CNN model's depth influences detection performance. The region proposal network replaces the selective search approach in Faster RCNN, making the algorithm faster. We also conducted a comparative time study of region-based convolutional networks. We concluded which of the region-based methods for object detection can particularly lead to real-time citrus disease detection. Images of basically citrus fruit illnesses are used in the training and testing of the model, contrary to popular belief. Faster RCNN with ResNet50 was applied initially, and the results were good enough to achieve a better outcome. Bounding boxes were used for the object detection, which in this case was citrus canker. The Kaggle notebook was preferred for the whole analysis procedure. The data was uploaded to apply some useful models for the detection process. This gave appropriate results and proved to be a good choice for future work. Secondly, Faster RCNN with Inception v3 was applied to the dataset, which proved to be another successful approach. It yielded excellent results and outclassed the previously used algorithm. This algorithm has several other versions, but this particular version was given preference. Both applied techniques yielded adequate results, warranting consideration for future work. However, some limitations exist, which can be addressed through algorithmic improvements and adjustments to parameter values.

6. LIMITATIONS

The object detection technique, Faster RCNN, uses regions to identify the items. The network does not look at the entire image at once but instead focuses on different areas of it sequentially. Two issues arise as a result of this. To extract all of the items, the model must make multiple runs through a single image. Because multiple systems are operating simultaneously, the performance of the systems that follow is influenced by the performance of the previous systems. In Faster RCNN, object proposals take more time because multiple systems are working in parallel; the performance of each system is affected by the performance of the previous system. ResNet-50 is often executed on GPUs because they are computationally intensive. When deploying ResNet models at scale, however, GPUs become extremely expensive. In the case of Inception v3, one of the disadvantages of this naive version is that even the 5 x 5 convolutional layer is computationally costly, meaning it takes a long time and requires a lot of processing power.

7. CONCLUSION

Object detection is an evolving field attracting increasing interest in both commercial and research domains. Citrus canker is one of the most widespread diseases in citrus fruits, making it critical to develop fast and reliable methods for identifying infected fruits among healthy ones. Manual inspection is slow, labour-intensive, and costly, while computer vision techniques can automate this process to achieve timely and accurate results at scale. In this study, we proposed and evaluated deep learning models for citrus canker detection. Faster R-CNN with ResNet50 achieved a validation precision of 0.90, while InceptionV3 reached an accuracy of 90%. These findings demonstrate the effectiveness of deep learning models in supporting automated disease detection in agriculture.

Future Work: To further advance this research, several directions can be pursued. First, expanding the dataset with more diverse images, including different lighting conditions, growth stages, and varieties of citrus fruits, would improve the generalization of the models. Second, exploring state-of-the-art architectures such as EfficientDet, YOLOv8, or Vision Transformers could yield higher detection accuracy and computational efficiency. Third, implementing data augmentation and transfer learning across related plant disease datasets may further enhance performance. Additionally, real-time deployment through mobile or edge-based applications would make this technology accessible to farmers and agricultural workers in the field. Finally, integrating detection systems with early warning frameworks and decision-support tools could help prevent large-scale crop losses and contribute to precision agriculture practices.

REFERENCES

N. Khurshid, H. Sharif, M. I. Tabash, and G. A. El Refae, "An assessment of asymmetric impact of financial stability and agricultural subsidies on agricultural production in Pakistan," Journal of Agribusiness in Developing and Emerging Economies, 2024.

A. Abbas, F. Lu, M. Yaseen, and M. Ameen, "Exploring the impact of foreign aid, agricultural production, and corporate social responsibility on poverty reduction in Pakistan," World, vol. 5, no. 3, 570-587, 2024.

U. W. Chohan, "Pakistan's Space Economy?" 2025.

M. Spies, "Agriculture and Chinese Agribusiness Investments in the Context of the China-Pakistan Economic Corridor," in China, Pakistan and the Belt and Road Initiative: Routledge, 2024:33-51.

V. Ghimirey, J. Chaurasia, and S. Marahatta, "Plant nutrition disorders: insights from clinic analyses and their impact on plant health," Agriculture Extension in Developing Countries, vol. 2, no. 1:09-17, 2024.

L. Li, S. Zhang, and B. Wang, "Plant disease detection and classification by deep learning—a review," IEEE access, vol. 9:56683-56698, 2021.

S. Hidayat, M. Zafarullah, M. K. Bihsoodwal, and M. Kazim, "Citrus Production Management to Support Self-sufficiency in Afghanistan (A Case Study of District Bati kot-Nangarhar)," ESRJ, vol. 1, no. 1:107-119, 2024.

E. Li, L. Wang, Q. Xie, R. Gao, Z. Su, and Y. Li, "A novel deep learning method for maize disease identification based on small sample-size and complex background datasets," Ecological Informatics, vol. 75:102011, 2023.

H. Tian, T. Wang, Y. Liu, X. Qiao, and Y. Li, "Computer vision technology in agricultural automation—A review," Information processing in agriculture, vol. 7, no. 1:1-19, 2020.

S. M. Raza, A. Raza, M. I. A. Babeker, Z. U. Haq, M. A. Islam, and S. Li, "Efficient citrus fruit image classification via a hybrid hierarchical CNN and transfer learning framework," Journal of Food Measurement and Characterization, vol. 19, no. 1:356-377, 2025.

Y. Bezabiha, A. M. Ayalew, B. M. Abuhayi, T. N. Demlie, and E. A. Ayenewe, "Classification of mango disease using ensemble convolutional neural network," Available at SSRN 4785206, 2024.

R. Neware, "A Comprehensive Review on Detection of Horticulture Fruits Disease Using Machine Learning and Deep Learning Approaches," 2025.

R. Liu et al., "Advances of nanoparticles as drug delivery systems for disease diagnosis and treatment," Chinese chemical letters, vol. 34, no. 2, p. 107518, 2023.

Mehmood, F., Chen, E., Akbar, M. A., Zia, M. A., Alsanad, A., Alhogail, A. A., & Li, Y. (2024). Advancements in human action recognition through 5G/6G technology for smart cities: fuzzy integral-based fusion. IEEE Transactions on Consumer Electronics, 70:5783-5795.

F. Mehmood, X. Guo, E. Chen, M. A. Akbar, A. A. Khan, and S. Ullah, "Extended multi-stream temporal-attention module for skeleton-based human action recognition (HAR)," Computers in Human Behavior, vol. 163:108482, 2025.

F. Mehmood, E. Chen, M. A. Akbar, and A. A. Alsanad, "Human action recognition of spatiotemporal parameters for skeleton sequences using MTLN feature learning framework," Electronics, vol. 10, no. 21:2708, 2021.

A. Assiri, A. Gumaei, F. Mehmood, T. Abbas, and S. Ullah, "DeBERTa-GRU: Sentiment Analysis for Large Language Model," Computers, Materials & Continua, vol. 79, no. 3, 2024.

S. Zhang and D. Metaxas, "On the challenges and perspectives of foundation models for medical image analysis," Medical image analysis, vol. 91:102996, 2024.

M. Dutta, D. Gupta, Y. Gulzar, M. S. Mir, and A. B. Soomro, "Leveraging Inception V3 for Precise Early and Late Blight Disease Classification in Potato Crops," Traitement du Signal, vol. 41, no. 2, 2024.

D. S. Joseph, P. M. Pawar, and K. Chakradeo, "Real-time plant disease dataset development and detection of plant disease using deep learning," Ieee Access, vol. 12, pp. 16310-16333, 2024.

Mahum, R., Munir, H., Mughal, Z. U. N., Awais, M., Sher Khan, F., Saqlain, M., and Tlili, I. (2023). A novel framework for potato leaf disease detection using an efficient deep learning model. Human and Ecological Risk Assessment: An International Journal, 29:303-326.

F. Mahmood, K. Abbas, A. Raza, M. A. Khan, and P. W. Khan, "Three-dimensional agricultural land modeling using unmanned aerial system (UAS)," International Journal of Advanced Computer Science and Applications, vol. 10, no. 1, 2019.

B. S. Sekhon, "Nanotechnology in agri-food production: an overview," Nanotechnology, science and applications. 31-53, 2014.

R. Kumar et al., "Hybrid approach of cotton disease detection for enhanced crop health and yield," IEEE Access, 2024.

M. A. Akbar, Z. Huang, Z. Yu, F. Mehmood, Y. Hussain, and M. Hamza, "Towards continues code recommendation and implementation system: An Initial Framework," in Proceedings of the 24th International Conference on Evaluation and Assessment in Software Engineering, 2020. 439-444.

Y. Kumar, A. Koul, R. Singla, and M. F. Ijaz, "Artificial intelligence in disease diagnosis: a systematic literature review, synthesizing framework and future research agenda," Journal of ambient intelligence and humanized computing. 14:8459-8486, 2023.

A. V. Panchal, S. C. Patel, K. Bagyalakshmi, P. Kumar, I. R. Khan, and M. Soni, "Image-based plant diseases detection using deep learning," Materials Today: Proceedings. 3500-3506, 2023.

F. Mehmood, H. Zhao, E. Chen, X. Guo, A. A. Albinali, and A. Razzaq, "Extended Multi-Stream Adaptive Graph Convolutional Networks (EMS-AAGCN) for Skeleton-Based Human Action Recognition," 2022.

N. Kumar, N. Narayan Das, D. Gupta, K. Gupta, and J. Bindra, "Efficient automated disease diagnosis using machine learning models," Journal of healthcare engineering. 9983652, 2021.

M. A. Zia, A. Ali, M. Gilani, F. Mehmood, S. Ullah, and I. Anwar, "Blockchain-powered accountability in pharmaceutical supply chains," Agri. Sci. J. 145-159, 2024.