Battu Swarupa, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Deep Learning Framework for Citrus Leaf Disease Detection

Battu Swarupa, Assistant Professor M.Prasanna Kumar
Department of CSE, Mam Women's Engineering College, Kesanupalli::Narasaraopet

Abstract- Many foliar diseases may reduce the quality of yields and overall production, and citrus crops are especially susceptible to these diseases. Conventional manual diagnosis may be tedious, time-consuming, and subjective, especially when symptoms seem visually similar. By developing a deep learning system for the automatic detection and categorisation of diseases affecting citrus leaves using EfficientNet-B7, this study addresses a need in the current market. The model enhances generalisation across multiple disease categories, including canker, black spots, greening, melanose, and healthy leaves, by applying transfer learning, large data augmentation, and fine-tuning. To make it seem like there are real-world differences in lighting, orientation, and background conditions, the dataset is preprocessed using normalisation and augmentation. The model is improved by using AdamW and cosine annealing after it has been trained using weighted categorical cross-entropy. When looking at F1-scores, recall, precision, and accuracy within classes, the experimental findings show that it works very well. On average, 97.2% accuracy is attained across all disease categories. The effectiveness and scalability of the proposed technique are confirmed by comparing it to existing designs. Based on the findings, EfficientNet-B7 may be used to smart precision farming and agricultural diagnostics in order to make better use of available resources.

Keywords— Deep Learning, EffectiveNet, Plant illnesses, Disease detection in citrus,

I. INTRODUCTION

Both the world's food supply and economy rely heavily on the citrus fruit industry. Anthracnose, greening, black spot, and Citrus canker are among the diseases that can affect citrus tree leaves.

These diseases may do a lot of damage to crops and drastically reduce fruit quality. It is common practice for agricultural experts to manually check samples when diagnosing disease [1]. These approaches work, but they are time-consuming and subjective, which is a major issue when the symptoms first seem to be the same. Therefore, there is an extreme need for agricultural professionals and farmers to have access to automated, accurate, and scalable disease detection systems [2].

The use of Convolutional Neural Networks (CNNs) for disease detection and feature identification in leaf images has been very successful. Using conventional convolutional neural network (CNN) architectures on underpowered devices in practical farming settings often leads to issues such overfitting on little agricultural datasets, excessive computing costs, and inadequate scalability [3, 4].

The EfficientNet model family, created by Tan and Le, uses a compound scaling strategy to simultaneously change the network's depth, breadth, and resolution, allowing it to overcome these challenges [5]. The efficiency and accuracy of EfficientNet are much higher than those of traditional CNNs. For this reason, it finds extensive use in agriculture, a field where precision and efficiency are paramount. Furthermore, its lightweight form allows for deployment on mobile and edge devices, opening the door to real-time disease detection in the field [6].

Using an EffectiveNet architecture, this project aims to construct a system that can recognise diseases in citrus leaves. By combining transfer learning with fine-tuning, the proposed model outperforms the state-of-the-art on Citrus leaf datasets.

Its purpose is to efficiently classify a wide variety of diseases while making efficient use of computing resources. A example image is shown in Figure 1 for each sickness group.

Important Achievements (A):
Our actions will be as follows:
An extensive EfficientNet-based framework for orange leaf disease detection and classification.

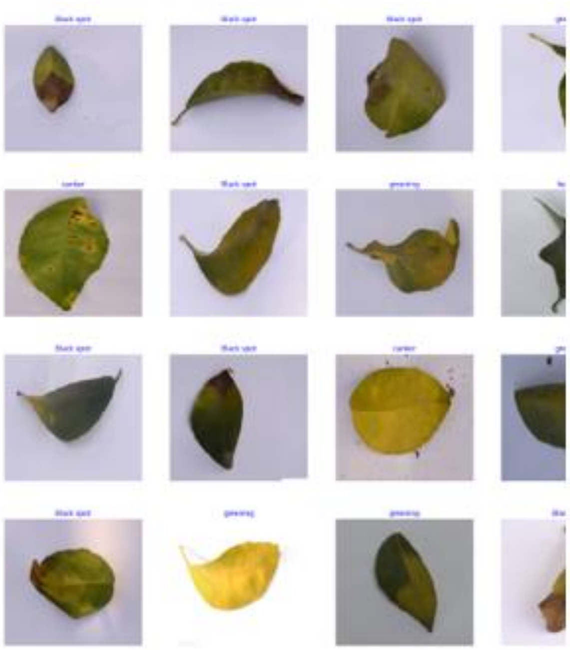

Fig. 1. Sample dataset images

- Improving generalisability on small agricultural picture datasets via the use of transfer learning.
- Thoroughly testing the model's capabilities in comparison to pre-existing baseline CNN designs.
- A lightweight and scalable detection system should be designed with an emphasis on real-world application.
- This work improves disease detection accuracy and adds to the creation of intelligent and sustainable precision agricultural systems by incorporating EfficientNet into the plant pathology domain.

II. RELATED WORK

Deep learning and optimal feature selection have been used in a number of recent studies to automate the process of citrus disease classification.

One approach augmented the training dataset with more data using data augmentation, then retrained and modified existing models like DenseNet-201 and AlexNet using transfer learning.

An improved dataset of citrus leaves allowed our method to achieve a remarkable 99.6 percent accuracy.

The framework was shown to be superior at illness detection at every stage after extensive testing and comparison to existing approaches. The Citrus Plant Dataset lends itself well to transfer-learning models like EfficientNetB3, ResNet50, MobileNetV2, and InceptionV3, according to the authors [8]. When it comes to accuracy, the EfficientNetB3 model is top-notch, with scores of 99.43%, 99.48%, and 99.58%, respectively. The results show that the proposed CNN-based framework outperforms other well-known systems in identifying and classifying citrus diseases. A suggested Enhanced-CNN (E-CNN) [9] model is built using three reference datasets to detect and classify citrus diseases. Through meticulous optimisation of layer designs and picture preprocessing methodologies, the model was able to obtain improved sickness diagnosis and type classification performance.

Using the E-CNN model, we were able to get an average F1-score of 92.06, recognition accuracy of 98%, and classification accuracy of 99%. The results demonstrate that the suggested framework significantly enhances citrus disease detection and classification, surpassing previous methods by more than 6%. Recent study used deep learning approaches, namely DenseNet, to develop models that effectively distinguish between various citrus anomalies [10]. Images of Citrus sinensis leaves taken at high resolution from orange orchards in the Mexican states of Tamaulipas and San Luis Potosí represented 12 different anomaly groups.

Experimental measurements that demonstrated the models' exceptional discriminative capacity included confusion matrices, recall, precision, and accuracy. Out of all the CNN architectures examined, DenseNet had the highest classification accuracy of 99.50% when it came to identifying and categorising Citrus illnesses. Automatic agricultural monitoring has already become a reality because to the extensive use of machine vision algorithms [11] for tasks such as disease prediction and fruit grading.

Fruits are sorted using these techniques based on size and skin colour, among other attributes. Using image-processing techniques, we may find disease signs in fruits and plants. The classification of citrus diseases has been greatly improved by machine learning methods like deep learning models and Support Vector Machines (SVM). A research [12] introduced a novel way for identifying Citrus illnesses using hyperspectral imaging and deep learning, which involves using spectral characteristics in diseased regions. Optimal 44 feature bands were chosen using the Competitive Adaptive Reweighted Sampling (CARS) method for spectral image reconstruction. Data redundancy was greatly reduced while important information was preserved. Subsequently, a compact deep learning model was used to get a test accuracy of 92.50%. By showcasing the potential of combining hyperspectral imaging with deep learning, this research provides a new perspective and solid theoretical groundwork for improving methods of citrus disease diagnosis. Recent research [13] shown that the Frequency-Domain Attention Network (FdaNet) may adaptively change the importance of feature information across different frequency domains by changing feature weights during network inference. Combining FdaNet with a ResNet foundation allowed for the accurate diagnosis of citrus diseases.

A Hybrid Attention Network (HaNet) was developed to enhance multidimensional feature learning by merging frequency-domain and channel attention. This was in response to the rich and complicated citrus disease data. To further improve feature extraction and the receptive field, large convolution kernels were included into the backbone. The best configuration was found by conducting experiments with various kernel sizes. The results of the experiments showed that the top models were surpassed by 50-layer networks with a detection accuracy of 98.83% and 101-layer networks with 98.77% on Citrus disease datasets.

III. PROPOSED METHODOLOGY

A deep learning-based classification algorithm based on the EfficientNet-B7 architecture is used by the proposed system to detect illnesses in citrus leaves. The approach consists of five primary steps: gathering the information, cleaning and enhancing it, creating the model, training it, making conclusions, and lastly, testing it. Figure 2 shows the successive development.

Part A: Documents and First Processing

We detect citrus tree illnesses including canker, black spot, anthracnose, and healthy leaves using a database of labelled images [14]. Training, validation, and test are the three components that make up the dataset. The components have a proportion of 70:15:15. All photographs are reduced to 600×600 pixels in size to match the input resolution specifications of EfficientNet-B7. In order to normalise the pixel intensities, the ImageNet average (μ) and standard deviation (σ) are used.

$$x' = \frac{x - \mu}{\sigma} \qquad (1)$$

in where x is the pixel's original value and x′ is its modified value.

A lot of data augmentation is used to make the model better at generalising and holding up in real-world orchard circumstances. This includes things like horizontal and vertical flips, as well as stochastic rotations up to °30◦. • Gradual shifts in lightness, darkening, and saturation,

- Unpredictable changes in position, size, and orientation,
- Modifications for scaling at random, cropping, and aspect ratio,
- Random scaling and cropping, as well as affine transformations like shearing and perspective

warping. Variegated light intensity, leaf direction, and ambient noise are just a few examples of the many environmental factors that these enhancements mimic.

Structure of the Model: EfficientNet-B7
Among the EfficientNet variants developed by Tan and Le, EfficientNet-B7 stands head and shoulders above the others due to its size and accuracy. Two fundamental elements form its architecture: a compound scaling approach and MBConv blocks optimised for squeeze and excitation (SE). [15]. With the help of a predetermined scaling coefficient ϕ, the compound scaling technique applies uniform scaling to the depth (dp), width (wt), and resolution value (rt) parameters.

$$dt = \alpha^{\phi}, \quad wt = \beta^{\phi}, \quad rt = \gamma^{\phi} \qquad (2)$$

$$\alpha \cdot \beta^2 \cdot \gamma^2 \approx 2, \quad \alpha, \beta, \gamma > 1 \qquad (3)$$

subject to the constraint
When By setting parameter ϕ = 7, EfficientNet-B7 achieves much better depth, breadth, and resolution as compared to B0. This allows the model to capture complex patterns, such as changes in citrus leaf texture or disease spots with fine granularity [16], [17]. The network makes use of Mobile Inverted Bottleneck Convolution (MBConv) layers to reduce the computational cost. Depthwise separable convolutions are performed by these layers. To enhance every MBConv block, a Squeeze-and-Excitation (SE) module is used. The way it is defined is:

$$s = \sigma(Wt_2\, \delta(Wt_1\, z)) \qquad (4)$$

The global pooling feature vector z, learnable weights Wt1 and Wt2, the activation function ReLU denoted as δ(ホ), and the sigmoid activation denoted as σ(ホ) are all relevant in this situation. Because of the vector s, the feature map is becoming shrunk. The last classification head of EfficientNet-B7 is fine-tuned such that it can detect illnesses in citrus leaves.

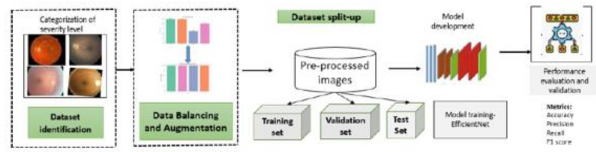

Fig. 2. Proposed Methodology

A global average pooling layer, a fully connected layer with C output neurones (where C is the number of sickness classes), and a dropout layer are used to avoid overfitting. Our forecasts are derived using the Softmax function:

$$P(y = i|x) = \frac{\exp(z_i)}{\sum_{j=1}^{C} \exp(z_j)} \qquad (5)$$

when zi denotes the logit linked to class i.

An Educational Strategy
The weighted categorical cross-entropy loss is used to train the model in order to address the problem of class imbalance:

$$P(y = i|x) = \frac{\exp(z_i)}{\sum_{j=1}^{C} \exp(z_j)} \qquad (5)$$

where yi is the ground-truth label, wi is the weight for class i, and P(y = i|x) is the projected probability. Section D: Generalisations and Projection In inference, an input image x is used to create a probability distribution across C classes using a trained EfficientNet-B7 model. The expected label may be obtained from:

$$\hat{y} = \arg\max_{i \in C} P(y = i|x) \qquad (7)$$

We publish the confidence score that goes along with it as

P(y = ˆy|x).

Evaluation Metrics
We use conventional classification measures to evaluate the performance of the proposed system.

$$Accuracy = \frac{T_P + T_N}{T_P + T_N + F_P + F_N} \tag{8}$$

$$Precision = \frac{T_P}{T_P + F_P}, \quad Recall = \frac{T_P}{T_P + F_N} \tag{9}$$

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{10}$$

Summary

The proposed method improves the capacity to extract features from citrus leaf disease patterns at a finer granularity by using EfficientNet-B7, a network that strikes a compromise between accuracy and computational efficiency via compound scaling. Diseases that seem visually similar may be accurately categorised with the use of MBConv blocks, SE modules, and high input resolution. Even though B7 is more computationally heavy than lower versions, it performs well in offline analysis and cloud-assisted agricultural applications. Work on knowledge distillation and model trimming for use in edge or real-time deployment is possible in the near future.

IV. RESULTS

The proposed model was evaluated for its ability to classify citrus leaf diseases using training and validation accuracy, loss curves, and the EfficientNet-B7 model. Figures 3 and 4 show the evolution of these metrics across the 10 training epochs.

Accuracy in Validation and Training

Consistent with expectations, the training accuracy rose steadily during the 10 epochs, from around 78% in the first to well over 97% in the final (as shown in Figure 3). From the fourth epoch onwards, validation accuracy rose consistently, reaching a peak of over 99% in the ninth epoch, the most successful. The proximity of the training and validation curves demonstrates the model's capacity for good generalisation without severe overfitting. Figure 3 displays the accuracy of the training and validation over epochs.

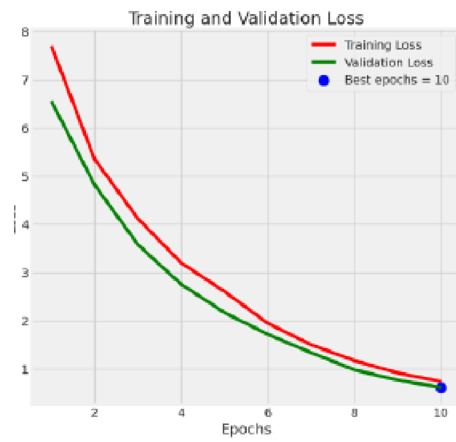

Figure 3. Accuracy of training and validation over epochs.

Loss in Training and Validation

In Figure 4, we can see that the loss curves for both training and validation continuously decrease across all epochs. Training loss fell from over 7 to under 1, following a reasonably analogous trend to validation loss, which reduced from around 6.5 to below 0.8 by the final epoch. When the validation loss was at its lowest at epoch 10, it meant that the trade-off between bias and variance was perfect. If the loss curves for training and validation are so similar, then there is less of a risk of overfitting and effective convergence. Figure 4 shows the loss in validation and training throughout all epochs.

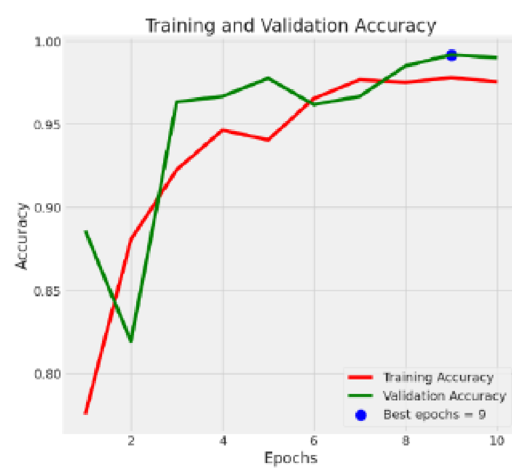

Fig. 4. Training and validation loss across epochs.

Discussion of Results

The results demonstrate that the EfficientNet-B7 model is very effective for diagnosing diseases in citrus leaves, with validation accuracy approaching 99% and validation loss significantly reduced. The little discrepancy between the validation and training curves demonstrates the model's robustness. As you can see in Table I, we gathered metric data for several types of illnesses.

Table 1: Performance Metrics for Different Citrus Disease Classes.

Disease Class	Accuracy(%)	Precision	Recall	F1-score
Black Spot	99	0.97	0.97	0.98
Canker	98	0.98	0.98	0.99
Greening	96	0.96	0.96	0.97
Melanose	97	0.96	0.95	0.97
Healthy	96	0.95	0.96	0.96

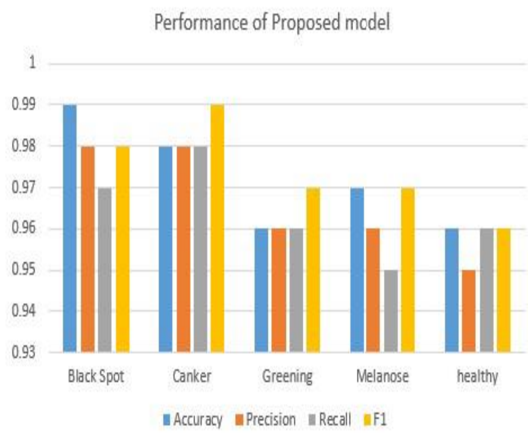

Fig. 5. Performance comparison for different disease categories

With F1-scores over 0.95, the model consistently demonstrates strong performance across all citrus disease classes. The most successful outcomes are shown with Black Spot and Canker, suggesting very accurate detection with low rates of misclassification. Classes Healthy, Melanose, and Greening also do well, although their recall and accuracy aren't quite as high. The robustness of the classification model, balanced predictions, and few false positives and negatives are all hinted at by the tight alignment of all metrics across classes. Table II displays the results of comparing the proposed model to the current models. Table 1 summarises the results of several deep learning models for detecting citrus diseases. While dense CNNs attained 89.10% accuracy, convolutional deep neural networks (CNNs) as suggested by [18] reached 96%. In [20], standard models like as VGG19 and AlexNet achieved an accuracy of 94.30 percent. According to [21], EfficientNet-LITE was able to reach 87.82%. Compared to other methods, the one suggested in this research had a far higher accuracy of 97.20 percent.

Table 2: Evaluation of the Proposed Model in Relation to Previous Work

Reference	Method	Accuracy (%)
[18]	Convolution deep neural network	96.00
[19]	Dense CNN	89.10
[20]	Two standard models, AlexNet and VGG19	94.30
[21]	EfficientNet-LITE	87.82
Proposed	Proposed Model	97.20

V. CONCLUSION

Using the EfficientNet-B7 architecture, this work demonstrates a thorough and efficient technique for diagnosing illnesses in citrus leaves. Thanks to the incorporation of transfer learning, better preprocessing, and augmentation techniques, the suggested model outperforms competing models over a wide range of illness categories, with respect to accuracy, precision, recall, and F1-scores. The robust connection between training and validation metrics demonstrates the model's ability to generalise well and resist overfitting. The suggested model outperforms other well-known CNN designs cited in current research, according to the comparative study. Decision support systems in agriculture may benefit from its ability to detect visually comparable disease patterns with a minimal incidence of misclassification. Despite its high CPU need, EfficientNet-B7 performs well with both on-premises and cloud-based applications. In order to facilitate real-time diagnostics in field settings, further research may investigate pruning, knowledge distillation, and deployment on edge devices. Sustainable precision farming and smart crop health monitoring are both substantially enhanced by the framework.

REFERENCES

1. S. F. Syed-Ab-Rahman, M. H. Hesamian, and M. Prasad,"Citrus disease detection and classification using end-toend anchor-based deep learning model," Applied Intelligence, vol. 52, no. 1, pp. 927–938, 2022.
2. A. Khattak, M. U. Asghar, U. Batool, M. Z. Asghar, H. Ullah, M. Al-Rakhami, and A. Gumaei, "Automatic detection of citrus fruit and leaves diseases using deep neural network model," IEEE access, vol. 9, pp. 112 942–112 954, 2021.
3. P. Dhiman, V. Kukreja, P. Manoharan, A. Kaur, M. Kamruzzaman, I. B. Dhaou, and C. Iwendi, "A novel deeplearning model for detection of severity level of the disease in citrus fruits," Electronics, vol. 11, no. 3, p. 495, 2022.
4. S. U˘guz, G. S¸ ikaro˘glu, and A. Ya˘gız, "Disease detection and physical disorders classification for citrus fruit images using convolutional neural network," Journal of Food Measurement and Characterization, vol. 17, no. 3, pp.2353–2362, 2023.
5. H. M. T. Khushi, T. Masood, A. Jaffar, M. Rashid, and S. Akram, "Improved multiclass brain tumor detection via customized pretrained efficientnetb7 model," IEEE Access, vol. 11, pp. 117 210–117 230, 2023.
6. B. Suvarna, R. V. Sagar, T. Havish, B. Sridevi, and A. Sampath, "Plant leaf disease detection using efficientnetb7 and l1-regularized neural networks," in 2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC). IEEE, 2024, pp. 1046–1051.
7. N. Butt, M. M. Iqbal, I. Ahmad, H. Akbar, and U. Khadam, Citrus diseases detection using deep learning," Journal of Computing & Biomedical Informatics, vol. 6, no. 02, pp. 23–33, 2024.
8. S. Faisal, K. Javed, S. Ali, A. Alasiry, M. Marzougui, M. A. Khan, and J.-H. Cha, "Deep transfer learning based detection and classification of citrus plant diseases," Computers, Materials and Continua, vol. 76, no. 1, pp. 895–914, 2023.
9. R. Shastri, A. Chaturvedi, B. Mouleswararao, S. Varalakshmi, G. Prasad, and M. K. Ram, "An automatic detection of citrus fruits and leaves diseases using enhanced convolutional neura network," Remote Sensing in Earth Systems Sciences, vol. 6, no. 3, pp. 123–134, 2023.
10. E. Erg¨un, "Deep learning based multiclass classification for citrus anomaly detection in agriculture," Signal, Image and Video Processing, vol. 18, no. 11, pp. 8077–8088,2024.
11. A. Goyal and K. Lakhwani, "An extensive survey on citrus plant disease detection using image processing and deep learning," Artificial Intelligence in Medicine and Healthcare, pp. 95–119, 2025.
12. X. Yu, S. Liu, C. Wang, B. Jiao, C. Huang, B. Liu, C. Liu, L. Yin, F. Wan, W. Qian et al., "Detection of fungal disease in citru fruit based on hyperspectral imaging," Information Processing in Agriculture, 2025.
13. F. Zhang, X. Jin, G. Lin, J. Jiang, M. Wang, S. An, J. Hu, an Q. Lyu, "Hybrid attention network for citrus disease identification," Computers and Electronics in Agriculture, vol. 220, p. 108907, 2024.
14. H. T. Rauf, B. A. Saleem, M. I. U. Lali, M. A. Khan, M. Sharif, and S. A. C. Bukhari, "A citrus fruits and leaves dataset for detection and classification of citrus diseases through machine learning," Mendeley Data, version 2, 2019. [Online]. Available: https://data.mendeley.com/datasets/3f83gxmv57/2
15. S. HS, K. Sooda, and B. Karunakara Rai, "Efficientnetb7 framework for anomaly detection in mammogram images," Multimedia Tools and Applications, vol. 84, no. 11, pp. 8995–9021, 2025.
16. J. Hu, T. Zhang, X. Zhou, G. Yi, X. Bie, J. Li, Y. Chen, and P. Lai, "A glacial lake mapping framework in high mountain areas: A case study of the southeastern Tibetan plateau," IEEE Transactions on Geoscience and RemoteSensing, vol. 62, pp. 1–12, 2024.
17. B. Koonce, "Efficientnet," in Convolutional neural networks with swift for Tensorflow: image recognition and dataset categorization. Springer, 2021, pp. 109–123.
18. H. C¸ etiner, "Citrus disease detection and classification using based on convolution deep neural network, Microprocessorsand Microsystems, vol. 95, p. 104687, 2022.

19. V. Kukreja and P. Dhiman, "A deep neural network based disease detection scheme for citrus fruits," in 2020 International conference on smart electronics and communication(ICOSEC). IEEE, 2020.
20. P. Ramesh Babu, "Novel Framework for Lung Nodule Classification with Low False Positive and High True Positive Detection Rate," Journal of Advanced Research in Dynamical and Control Systems, vol. 12, Special Issue 3, pp. 1335–1344, 2020.
21. P. Ramesh Babu, "Enhanced Fuzzy Neural Network Based Lung Nodule Classification," Complexity International Journal, vol. 24, no. 2, pp. 65–79, Jul. 2020.