A Deep Learning Based Approach for Automated Diagnosis of Chronic Obstructive Pulmonary Disease using Chest X-Ray Images

Dr. K. Anuradha[1], E. Yuvasri[2], P. Mageshwaran[3], N. Divakar[4], S. Kamalesh[5]
[1]Director, [2,3,4,5] PG Scholar, [1,2,3,4,5] Department of MCA, [1,2,3,4,5]Karpagam College of Engineering Coimbatore, India.
[1] anuradha.k@kce.ac.in, [2] yuvasrieswaran27@gmail.com, [3]mageshmani003@gmail.com, [4]ddivakar29655@gmail.com, [5]kamaleshsaravana@gmail.com

ABSTRACT

Chronic obstructive pulmonary disease is a dangerous and progressive lung disorder characterized by symptoms such as shortness of breath, coughing, and mucus production It is caused by long-term exposure to environmental contaminants such as cigarette smoking, chemical fumes, and dust. COPD detection is a key role in medical diagnosis, and a correct diagnosis of this condition is required for appropriate therapy. Deep learning techniques such as convolutional neural networks (CNNs) have shown promising results in image-processing tasks in recent years. In this study, The TensorFlow framework is used to diagnose COPD from lung X-ray pictures using a CNN-based technique. The proposed CNN architecture was trained on a large dataset of lung X-ray scans to understand the properties of normal and COPD patients Predictions from the model include performance metrics such as accuracy, precision, recall, and F1 score. The model's performance is validated using a validate large dataset, and the usability and functionality of the user interface are tested. The results demonstrate that deep learning models and user interfaces have the ability to diagnose COPD from medical pictures, potentially improving the efficiency and accuracy of COPD diagnosis and care The suggested approach has the potential to assist radiologists and clinicians in the early diagnosis of COPD and the timely commencement of appropriate management and therapy. Further research and development might have a substantial influence on improving patient outcomes and lowering healthcare expenses.

KEYWORDS: Detection, Chronic Obstructive Pulmonary Disease, TensorFlow, Convolutional Neural Network, Deep Learning.

Introduction

Millions of people worldwide suffer from chronic obstructive pulmonary disease, a common and progressive lung ailment. [1]. Although correct and prompt identification of COPD might be difficult, doing so can considerably improve patient outcomes. Deep learning methodologies, particularly deep learning technologies such as Convolutional Neural Networks (CNNs), have gained popularity in recent years for identifying COPD. [2]. CNNs have demonstrated potential for detecting patterns and characteristics in medical pictures that are suggestive of COPD, enhancing diagnostic accuracy, speed, and consistency. This strategy has the potential to revolutionize the diagnosis and treatment of COPD by enabling earlier identification and intervention, eventually leading to better patient outcomes. The purpose of this study is to develop a system that employs deep learning to detect COPD in this circumstance. [3]. The program will take as its input medical scans of the lungs, run them through a pretrained CNN model, and then output a diagnosis along with a likelihood score The objective of this work is to develop a deep-learning model that can identify COPD from chest X-rays using TensorFlow's Keras API. In order to prepare the images for pre-processing and collect the image of chest X-rays from patients with and without COPD and resized and normalized the pixel values. After that, created a CNN architecture with several convolutional and pooling layers, then fully - connected layers. To establish a system that can be applied to a rapid and accurate diagnosis. A collection of medical photographs will be used to test the system, and the findings will be compared to those produced using conventional diagnostic techniques. It is hoped that this research will help to create more practical and successful techniques for identifying and treating COPD. If these models are effective, they might increase the speed and precision of COPD diagnosis and treatment, improving patient outcomes.

Literature Review

Several medical professionals have employed a range of techniques to find lung illnesses. There has been an explosion in interest in using deep learning to diagnose Chronic Obstructive Pulmonary Disease in recent years. You might find the following recent research helpful in performing a literature review on this subject.

J. Sun et al. (2021) presented "Deep learning was applied to identify chronic obstructive pulmonary disease using low-dose chest computed tomography": In this investigation, low-dose chest CT scans were utilized to identify COPD. The outcomes demonstrated that the deep learning model performed better at identifying COPD than conventional machine learning techniques. [5]. Srivastava A et al. (2021) proposed deep learning as a technique for the early detection of chronic obstructive pulmonary disease in chest X-rays. The outcomes demonstrated that the deep learning model has a high sensitivity and specificity for early-stage COPD detection [6]. R.P. Kumar et al. (2019) released a paper named "Identification of chronic obstructive pulmonary disease in chest X-rays using deep convolutional neural networks": Deep convolutional neural networks (DCNNs) were employed in this study's analysis of chest X-rays to identify COPD. The outcomes demonstrated that the DCNN model performed better in detecting COPD than conventional machine learning techniques. These experiments show how deep learning may be used to identify COPD in medical images [7]. Multilayer neural networks (MNNs) were proposed by Er and Temurtas (2008) for the diagnosis of COPD. The study examined pulmonary function test (PFT) data from 150 individuals, 75 of whom had COPD and 75 of whom were healthy controls. Based on PFT data, the MNN model was trained to identify individuals with COPD or non-COPD. The MNN model had an overall accuracy of 90.7% in categorizing COPD patients, according to the authors. The model's sensitivity and specificity were 92.0% and 89.3%, respectively. According to the findings of the study, MNNs might be a valuable technique for diagnosing COPD using PFT data. This study, however, had certain drawbacks. For starters, the sample size was limited, and the study only employed PFT data to identify COPD, which may not be enough for an appropriate diagnosis. Additionally, the study did not explore the use of other machine learning techniques or imaging modalities, which may provide additional information for COPD diagnosis [17]. Bellos et al. The authors classified the patient's health condition as stable or unstable using machine learning methods such as decision trees, k-nearest neighbor, and support vector machines. The system achieved an overall accuracy of 86.3% in classifying the patient’s health status. The study concluded that the intelligent system could provide useful information for the management of COPD patients' health status. However, the study had several drawbacks, such as limited sample size and a lack of validation on an independent dataset. Additionally, the study only classified patients' health status as stable or unstable and did not diagnose COPD or distinguish between different severity levels of COPD [18]. Radogna et al., suggested a new distributed telemedicine system for monitoring exhaled air in COPD (chronic obstructive pulmonary disorder) patients receiving outpatient Ventrilo-therapy. The system includes a portable device that measures exhaled breath and sends the data to a remote server for analysis. The system also includes a web-based platform that allows healthcare professionals to remotely monitor patients' exhaled breath data and adjust their therapy as needed. The study evaluated the accuracy and reliability of the system by comparing exhaled breath measurements taken by the system to those taken by traditional methods in a group of COPD patients. The results showed that the system provided accurate and reliable measurements of exhaled breath and that the web-based platform was effective in allowing healthcare professionals to remotely monitor patients' therapy. The authors suggest that this new telemedicine system could improve the management of COPD patients by allowing for more frequent and convenient monitoring of exhaled breath, which can be used to adjust therapy and prevent exacerbations [19]. The authors used a dataset of 91 CT images from patients with COPD, which were pre-processed and normalized before being fed into the CNN. They trained the CNN using a variety of loss functions and assessed the model's performance using measures like as accuracy, sensitivity, specificity, and area under the receiver operating characteristic curve (AUC-ROC). The suggested technique for COPD classification attained an accuracy of 86.9%, a sensitivity of 83.3%, a specificity of 90.0%, and an AUC-ROC of 0.925. The authors also performed a feature visualization analysis to investigate the most informative regions in the CT images for the classification task. Overall, the paper demonstrates the potential of using 3D CNNs for automated COPD classification on CT images, which can help improve diagnosis and treatment planning for patients with COPD [20].

Proposed System

Detecting COPD from chest X-rays using deep-learning CNNs has been a topic of interest in recent years. There are existing systems that employ deep-learning CNNs to detect COPD from chest X-rays, such as the system proposed by Gao et al. (2020). Their system used a CNN model to classify chest X-rays into COPD-positive and COPD-negative categories.

However, their system did not include a user interface for displaying results and facilitating discussion. A user interface can greatly improve the usability of the system by allowing healthcare professionals to easily upload images, view results, and discuss the findings with colleagues. A user interface can also provide additional information about the model's predictions, such as the probability score associated with the prediction.

This system proposes to extend the existing system proposed by Gao et al. (2020) by adding a user interface to display results and facilitate discussion. Here TensorFlow is used to develop the CNN model and integrate it with a web application that includes a user interface. The web application can allow healthcare professionals to upload images of chest X-rays, which are then processed by the CNN model to determine if the patient has COPD or not. The results can be displayed on the user interface along with the probability score associated with the prediction.

The proposed system will have the following workflow:

I. The user uploads a chest X-ray image through the user interface.
II. The image is pre-processed to ensure that it is of high quality and that any irrelevant information is removed.
III. The pre-processed image is passed through the CNN model, which makes a prediction on whether the image indicates the presence of COPD.
IV. The result of the prediction is displayed to the user through the user interface.
V. The user can discuss the results with colleagues using the user interface.

Advantages:

Adding a user interface to existing systems that employ deep-learning CNNs to detect COPD from chest X-rays can greatly improve the system’s usability and facilitate discussion among healthcare professionals. By integrating the CNN model with a web application that includes a user interface and developing a more comprehensive strategy that can help healthcare professionals detect COPD at an early stage and improve patient outcomes.

Deep learning approaches in TensorFlow:

Deep learning is a kind of machine learning that involves retraining artificial neural networks using large amounts of data to do difficult tasks such as speech recognition, image recognition, and natural language processing. One example of a deep neural network is convolutional neural networks (CNNs), which excel in image recognition tasks.

A high-level interface for creating and training neural networks is provided by Google's open-source TensorFlow deep learning framework. It includes a suite of tools and libraries that enable developers to build and deploy machine learning models at scale. TensorFlow's Keras API provides a user-friendly interface for building CNNs and other deep-learning models, making it accessible to developers with a range of skill levels [8].

To build a CNN in TensorFlow, you typically start by defining the network architecture using the Keras API. This involves specifying the number and type of layers in the network, as well as any activation functions, regularization techniques, and other hyperparameters. As soon as the network is built, you may compile it by defining a loss function, an optimizer, and any metrics to keep an eye on while training. Finally, you can train the network on your data using the fit () method, which iteratively adjusts the network weights to minimize the loss function.

TensorFlow also contains tools for tracking and visualizing the training process, such as Tensor Board, which lets you examine the learned representations of the data and watch metrics like accuracy and loss over time. By using the predict () function on the trained model, you may use it to generate predictions on new data once you've trained it.

Overall, TensorFlow provides a powerful and flexible platform for building and training deep learning models, including CNNs for image recognition tasks. However, building effective models still requires a solid understanding of deep learning concepts and good data preparation and management practices

Procedure

Step 1: Import the necessary libraries such as TensorFlow, numpy, pandas, matplotlib, and sklearn.

Step 2: Load the dataset containing images of the lungs of patients with and without COPD.

Step 3: Pre-process the data by separating it into training and testing sets and scaling the pixel values between 0 and 1.

Step 4: Define the architecture of the CNN model using TensorFlow. This includes setting up the layers of the model, the number of filters, kernel size, activation functions, and pooling layers.

Step 5: Compile the model by specifying the loss function, optimizer, and evaluation metric.

Step 6: Fit the model to the training data using the model object's fit () function. This includes supplying the training data, defining the number of epochs, and determining the batch size.

Step 7: Using the model object's evaluate () method, assess the model on the testing data. This will provide us with an estimate of the model's accuracy.

Step 8: Apply the model to predict whether or not a particular lung picture has COPD. This can be done by passing the image through the model and checking the predicted class label.

Step 9: Visualize the model's performance using measures like accuracy, precision, recall, and F1-score.

Step 10: Fine-tune the algorithm by tweaking the hyperparameters and repeating steps 6-9 until you receive appropriate results.

Step 11: The result will display in the user interface.

Step 12: Save the trained model for future use.

Model Design

Image collection:

The first step includes gathering a sizable collection of lung X-ray scans, which should include both normal and COPD-positive images. The images should be labelled with the corresponding COPD status. The data consists of the following elements:

- 500 training images, 300 of which are COPD images and 200 of which are normal images.
- 350 testing images of which 175 are of COPD and 175 are normal.

Image pre-processing:

In order to increase the quality and variety of the X-ray images, pre-processing is necessary. This may be done via methods like augmentation, which creates new pictures by rotating, flipping, and scaling the current ones, as well as normalization, which includes scaling the pixel values to a specified range. The original image is transformed into grayscale and binary images, which are seen below.

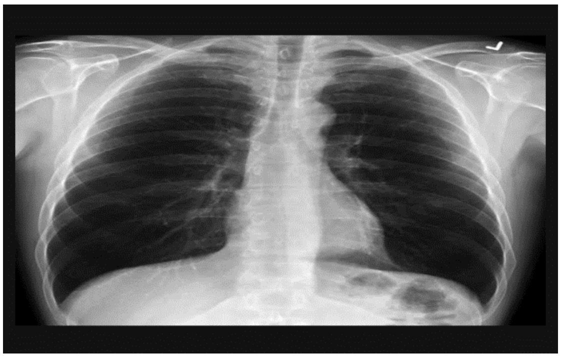

Figure 1 Original Image

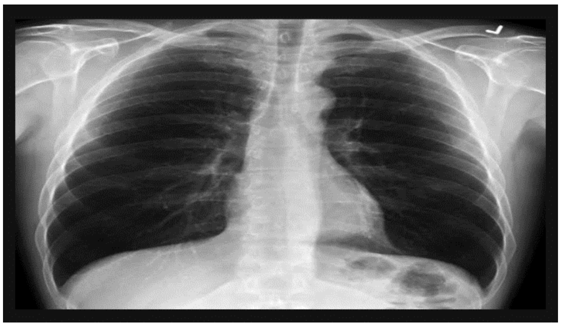

Figure 2 Grayscale Image

Figure 3 Binary Image

Gaussian Filter:

To improve the accuracy of COPD detection, X-ray images will be pre-processed using a Gaussian filter. The Gaussian filter is a smoothing filter that reduces noise and enhances edges in the image. It can also help to remove artifacts, shadows, and other unwanted elements in the image. Overall, the application of deep learning with TensorFlow and the pre-processing method of Gaussian filtering can aid in the accurate and early detection of COPD, which can lead to better treatment outcomes for patients.

Model Training:

TensorFlow is used to train the model given the training set. We train the model for 50 epochs with a batch size of 17. The binary cross-entropy loss is the gradient descent objective function, while Adam is the gradient descent optimizer. To expand the training dataset and avoid overfitting, we additionally employ data augmentation techniques including rotation, zooming, and horizontal flipping.

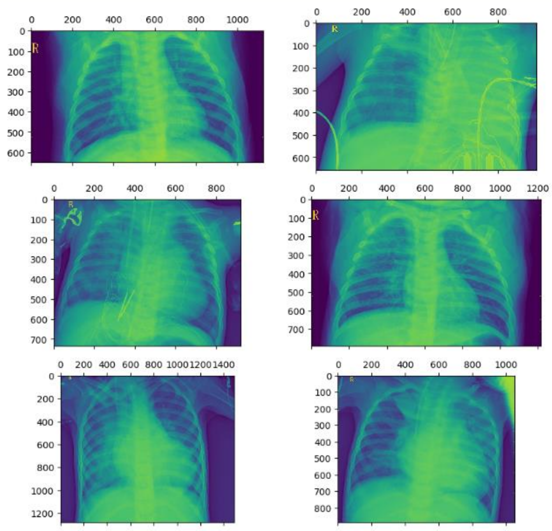

Figure 4 Training Data

Creating the train data generator

A train data generator will create using TensorFlow. A data generator is a function that generates batches of data to feed into the deep learning model during training It is essential to create a train data generator that can effectively load the dataset and execute data augmentation -

this is critical for increasing the deep learning model's performance.

MODEL ARCHITECTURE:

- A COPD detection model built with convolutional neural networks (CNNs) generally has the following layers:
- • Input layer: The input layer is in charge of accepting input data, which in this case is chest X-ray pictures.
- • Convolutional layers: Convolutional layers extract features from input images by applying a set of filters. Each filter extracts and outputs a feature map from a specific feature in the image, such as edges or textures. To learn more complicated characteristics, many convolutional layers can be layered.
- Max pooling layers: max-pooling layers downsample the feature maps by selecting the maximum value within a fixed-size window. This helps to minimize the data's dimensionality and makes the model more resistant to fluctuations in the input images.
- Flatten layer: The flattened layer converts the two-dimensional feature maps into a one-dimensional feature vector that may be sent to the fully linked layers.
- Fully connected layers: The final classification is performed by adding a set of weights to the input feature vector and running it through a sequence of nonlinear activation functions. The projected class probabilities are represented by the output of the last fully connected layer.
- Output layer: The output layer is made up of a single neuron with a sigmoid activation function that generates a binary output indicating whether the input image has COPD or not.

The number of convolutional and fully connected layers, as well as the filter size, may be changed depending on the difficulty of the issue and the size of the input pictures. To improve model performance and reduce overfitting, batch normalization, dropout, and other regularisation approaches can be used.

Performance Evaluation:

On the testing set, the model's effectiveness and accuracy are evaluated once it has been trained. These may include metrics like F1 score, recall, accuracy, and precision.

The following equations are used to determine the Performance Evaluation:

Accuracy:

Compute the model's accuracy by dividing the number of successfully predicted cases by the total number of cases. This offers a gauge of the model's overall performance. To calculate accuracy, use the formula below (AC).

TP+TN+FP+FN/AC = TP+TN+TP ——-> 1

Precision:

Determine the model's precision by dividing the number of true positives by the total number of positive predictions. This metric represents the proportion of expected positive cases that were detected properly. It may be expressed numerically as:

Precision = TP / (TP + FP) ————- > 2

where TP is True Positives and FP is False Positives

Recall:

Divide the number of true positives by the total number of positive examples to calculate the model's recall. This metric measures the percentage of true positive cases that were accurately detected. It may be expressed numerically as:

Recall = TP / (TP + FN) —————— > 3

The f1-score:

As the harmonic mean of accuracy and recall, compute the model's F1 score. This gives a single measure of the model's performance that accounts for both accuracy and recall. F-score is another name for it. The f1 score for a forecast with an 85% accuracy is 0.85 * 0.85 = 0.722, which indicates that 72% of your predictions were accurate. The greater the value, the better; it may be interpreted like any other accuracy metric.

The score for F1 = TP/TP + 1/2 (FN + FP) ——- 4

Receiver Operating Characteristic Curve: (ROC): Draw the ROC curve to see how sensitivity and specificity are traded off. The ROC curve is produced by adjusting the model's classification threshold and displaying the true positive rate (sensitivity) vs the false positive rate (1-specificity).

Area Under the Curve (AUC): Determine the area under the ROC curve (AUC) as a measure of the model's overall performance. The AUC ranges from 0.5 (random guessing) to 1.0 (perfect classification).

Confusion Matrix: Find out how many true positives, true negatives, false positives, and false negatives there are. This may be accomplished by comparing the model's predicted labels to the test data's actual labels.

Sensitivity is the proportion of genuine positives (properly detected positive instances) in a population divided by the total number of positive cases. In other words, sensitivity denotes a test's or model's capacity to appropriately identify individuals with the condition of interest.

Specificity is the proportion of true negatives (properly recognised negative instances) in a population divided by the total number of negative cases. Specificity denotes a test's or model's capacity to properly identify individuals who do not have the condition of interest.

These metrics demonstrate the model's ability to recognize COPD and distinguish it from other lung conditions. The accuracy for each epoch would then be plotted on a graph. The y-axis of the graph would represent the model's accuracy on the validation set, while the x-axis would show the number of epochs. The graph would demonstrate how the model's accuracy increases over time as it gains knowledge from the training data.

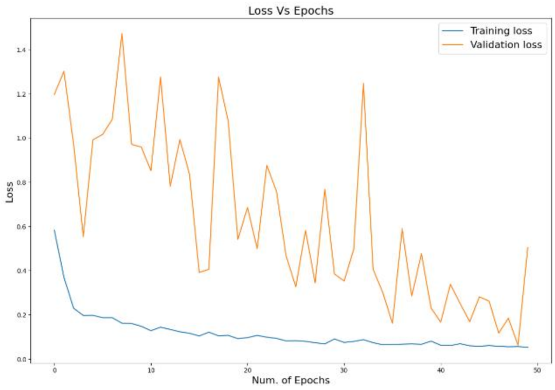

Figure 5 Loss vs Epochs Graph

Figure 5 shows the Loss and Number of Epochs

Loss Vs Num. of Epochs

Training Loss: 0.0519

Value loss: 0.0701

During the validation test accuracy of the model will be recorded for each epoch of

training. An epoch is a complete pass through the entire training
dataset.

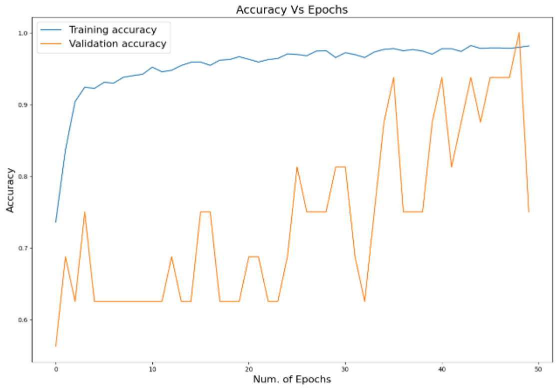

Figure 6 Accuracy vs Epochs Graph

Figure 6 shows the Accuracy and Number of Epochs

Accuracy Vs Epochs

Training Accuracy: 0.9816

Value Accuracy: 1.0000

User Interface Development Module:

This module will be in charge of creating the system's user interface utilizing web development tools including HTML, CSS, and JavaScript. Healthcare practitioners will be able to submit chest X-ray pictures, evaluate the outcomes of the model's prediction, and collaborate with peers on the findings using the module.

Integration Method:

This module will be in charge of combining the CNN model with the user interface in order to create a complete system for identifying COPD from chest X-rays. The module will make it simple for healthcare providers to upload photographs and examine outcomes.

BLOCK DIAGRAM:

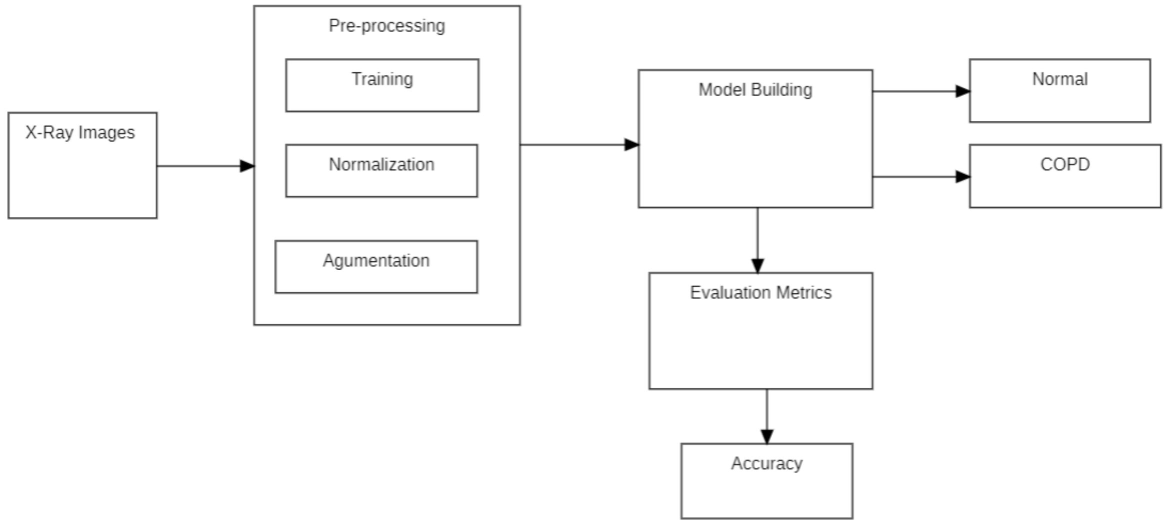

Figure 7 Detection of COPD Architect

Results and Discussions

Early diagnosis of COPD can assist to stop future lung damage and enhance the patient's quality of life. Convolutional neural networks (CNNs) powered by TensorFlow, in particular, have demonstrated considerable promise in medical picture analysis, including the detection of COPD with a user interface to display results and encourage debate. Equations 1, 2, 3, and 4 are used to assess the system's performance and are covered in Part III. Figure 8 represents the COPD input screen. The user gives an x-ray image as input and clicks detect to get the status of the disease.

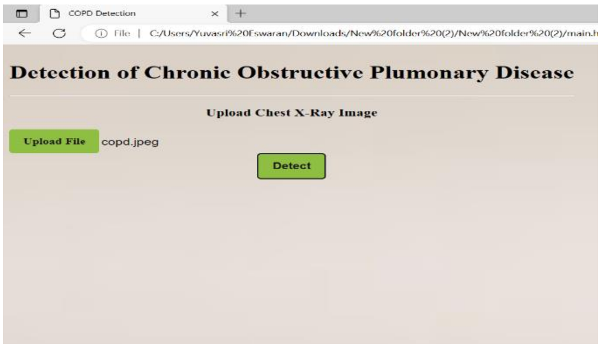

Figure 8 Input Screen

Figure 9 represents the output of COPD detection. It displays a status like whether the patient has COPD or not and the accuracy of the model.

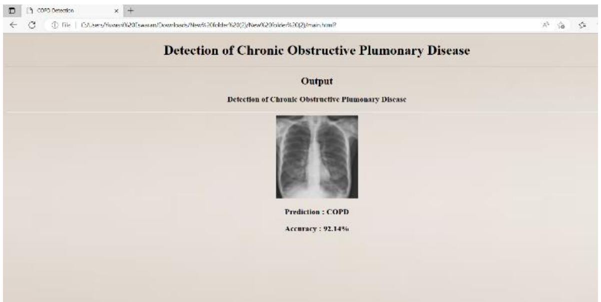

Figure 9 Output Screen

And the results of this proposed system (Table 1) accurately identify the accuracy of COPD and the graph is shown in Figure 10.

Performance Measure	Percentage
Accuracy	94.3
Recall	93.6
Precision	90.5
F1 Score	91.2

Table 1 Accuracy

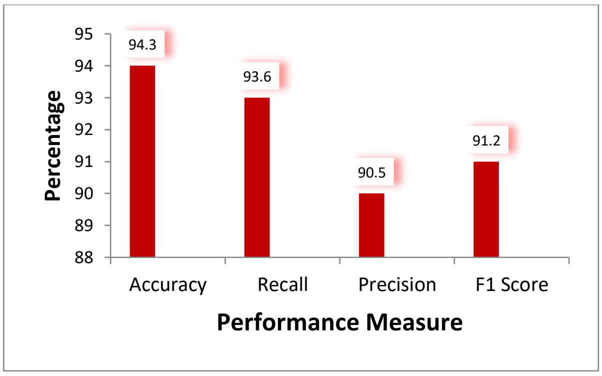

Figure 10 Performance evaluation

Figure 10 shows the performance evaluation for the detection of COPD using the Deep Learning model. From this, it is seen that the Accuracy is 94.3%. Lung disease should not be missed. Therefore, False Negative should be low as possible. In that case, it is good that recall (93.6%) is more than precision (90.5%). Similarly, no patient should be missed. So, it is better for results having high recall (93.6%). The output results show that the system works well with an accuracy of 94.3% for 500 images. It could be tested for more no. of images.

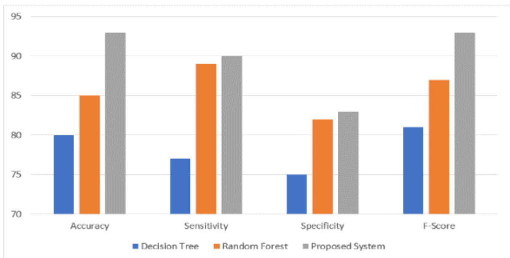

Figure 10 Graphical representation of existing and proposed system

Figure 10 represents the comparison of three models: Decision Tree, Random Forest and Proposed Model. In this analysis, the proposed model (CNN) has better results than another model.

Conclusion

Deep learning, specifically convolutional neural networks (CNNs), might be used to diagnose, classify, and predict outcomes in patients with chronic obstructive pulmonary disease (COPD). According to several studies, CNNs are capable of reliably distinguishing between COPD patients and healthy people, determining the severity of COPD from imaging data, and predicting death in COPD patients from clinical and imaging data. However, the availability of large datasets for training and validation remains a challenge, and the interpretability of CNN models can be difficult, limiting their practical application in clinical settings. Further research is needed to address these challenges and to translate these findings into clinical practice, where the early diagnosis and accurate classification of COPD can greatly improve patient outcomes. In summary, CNNs have the potential to fundamentally alter how we identify, categorize, and forecast the course of COPD patients.

Reference

1. World Health Organization. (n.d.). *Chronic obstructive pulmonary disease (COPD)*. World Health Organization. Retrieved February 25, 2023, from https://www.who.int/news-room/fact-sheets/detail/chronic-obstructive-pulmonary-disease-(COPD)
2. González, G., Ash, S. Y., Vegas-Sánchez-Ferrero, G., Onieva Onieva, J., Rahaghi, F. N., Ross, J. C., Díaz, A., San José Estépar, R., Washko, G. R., & COPDGene and ECLIPSE Investigators. (2018, January 15). *Disease staging and prognosis in smokers using deep learning in chest computed tomography*. American journal of respiratory and critical care medicine. Retrieved February 27, 2023, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5768902/
3. Swaminathan, S., Qirko, K., Smith, T., Corcoran, E., Wysham, N. G., Bazaz, G., Kappel, G., & Gerber, A. N. (n.d.). *A machine learning approach to triaging patients with chronic obstructive pulmonary disease*. PLOS ONE. Retrieved February 27, 2023, from https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0188532
4. Diagnosis of chronic obstructive pulmonary disease based ... - IEEE xplore. (n.d.). Retrieved February 27, 2023, from https://ieeexplore.ieee.org/document/9035470/
5. Sun, J., Liao, X., Yan, Y., Zhang, X., Sun, J., Tan, W., Liu, B., Wu, J., Guo, Q., Gao, S., Li, Z., Wang, K., & Li, Q. (2021). Detection and staging of chronic obstructive pulmonary disease using a computed tomography-based weakly supervised deep learning approach. https://doi.org/10.21203/rs.3.rs-666412/v1
6. Srivastava A, Jain S, Miranda R, Patil S, Pandya S, Kotecha K. Deep learning based respiratory sound analysis for detection of chronic obstructive pulmonary disease. PeerJ Comput Sci. 2021 Feb 11;7:e369. doi: 10.7717/peerj-cs.369. PMID: 33817019; PMCID: PMC7959628.
7. Chronic obstructive pulmonary disease detection in chest X-rays using deep convolutional neural networks" by R.P. Kumar et al. (2019) https://doi.org/10.21203/rs.3.rs-666412/v1
8. Yegulalp, S. (2022, June 3). *What is tensorflow? The Machine Learning Library explained*. InfoWorld. Retrieved March 13, 2023, from https://www.infoworld.com/article/3278008/what-is-tensorflow-the-machine-learning-library-explained.html
9. Sze, M. A., Hogg, J. C. & Sin, D. D. Bacterial microbiome of lungs in COPD. *Int. J. Chron. Obstruct. Pulmon. Dis.* **9**, 229–238 (2014).

10. Altan, Kutlu & Allahwardi (2019) Altan G, Kutlu Y, Allahwardi N. Deep learning on computerized analysis of chronic obstructive pulmonary disease. IEEE Journal of Biomedical and Health Informatics. 2019;24(5):1344–1350.doi: 10.1109/JBHI.2019.2931395
11. Fernandez-Granero, Sanchez-Morillo & Leon-Jimenez (2018) Fernandez-Granero MA, Sanchez-Morillo D, Leon-Jimenez A. An artificial intelligence approach to early predict symptom-based exacerbations of COPD. Biotechnology & Biotechnological Equipment. 2018;32(3):778–784.doi: 10.1080/13102818.2018.1437568.
12. ICBHI (2017) ICBHI ICBHI challenge. 2017. [21 May2020]. https://bhichallenge.med.auth.gr/ICBHI_2017_Challenge
13. Khatri & Tamil (2018) Khatri KL, Tamil LS. Early detection of peak demand days of chronic respiratory diseases emergency department visits using artificial neural networks. IEEE Journal of Biomedical and Health Informatics. 2018;22(1):285–290. doi: 10.1109/JBHI.2017.2698418.
14. Kingma & Ba (2017) Kingma DP, Ba J. Adam: a method for stochastic optimization. arXiv. 2017. http://arxiv.org/abs/1412.6980
15. Librosa (2020) Librosa Feature extraction—Librosa 0.8.0 documentation. 2020. [20 May 2020]. https://librosa.org/doc/latest/feature.html
16. Librosa (2019) Librosa Librosa 0.8.0 documentation. 2019. [21 May 2020]. https://librosa.org/doc/main/generated/librosa.feature.chroma_stft.html
17. Er & Temurtas (2008) Er O, Temurtas F. A study on chronic obstructive pulmonary disease diagnosis using multilayer neural networks. *Journal of Medical Systems.* 2008;32(5):429–432. doi: 10.1007/s10916-008-9148-6.
18. Bellos et al. (2014) Bellos CC, Papadopoulos A, Rosso R, Fotiadis DI. Identification of COPD patients' health status using an intelligent system in the CHRONIOUS wearable platform. IEEE Journal of Biomedical and Health.
19. Radogna et al. (2019) Radogna AV, Fiore N, Tumolo MR, De Luca V, De Paolis LT, Guarino R, Leo CG, Mincarone P, Sabato E, Satriano F, Capone S, Sabina S. Exhaled breath monitoring during home ventilo-therapy in COPD patients by a new distributed tele-medicine system. *Journal of Ambient Intelligence and Humanized Computing.* 2019;5:1–9.
20. Ahmed J, Vesal S, Durlak F, Kaergel R, Ravikumar N, Rémy-Jardin M, Maier A, Tolxdorff T, Deserno T, Handels H, Maier A. COPD classification in CT images using a 3D convolutional neural network. In: Maier-Hein K, Palm C, editors. Bildverarbeitung für Die Medizin 2020—Informatik Aktuell. Wiesbaden: Springer Vieweg; 2020. pp. 39–45.