

# Lung Diseases Detection and Classification of Chest X-Ray Using CNN

Tirupathi Krishna Veni<sup>1</sup>, S. Babu<sup>2</sup>, Dr. K. Logesh<sup>3</sup>

<sup>1</sup> PG Scholar, Department of CSE, Kuppam Engineering College, Kuppam. Email: [krishnaveni118@gmail.com](mailto:krishnaveni118@gmail.com)

<sup>2</sup> Assistant Professor, Department of CSE, Kuppam Engineering College, Kuppam

<sup>3</sup> Associate Professor & HOD, Department of CSE, Kuppam Engineering College, Kuppam

## ABSTRACT

The proposed research project will create a machine learning model to identify and categorize different lung diseases, such as Normal, Pneumonia-Bacterial, Pneumonia-Viral, COVID-19, Tuberculosis, and Emphysema based on the images of chest X-rays. The data used is Chest X-Ray 6 Classes Dataset from Kaggle, which consists of labeled X-ray images per condition. The main aim of the given work is to improve diagnostic efficiency through the accurate classification of the chest X-ray sensitized images, which will benefit healthcare providers in diagnosing respiratory diseases effectively in the end. The model that is used with its methodology integrates modern Convolutional Neural Network (CNN) models like MobileNet, DenseNet, and VGG19 to derive significant information off the images. These models have been chosen due to their computational power and high classification rate. The system was trained and tested on the dataset, and the performance was measured by conventional evaluation metrics. DenseNet model was the most accurate model with 88 percent, then MobileNetV2 with 87 percent and VGG19 with 83 percent. The proposed solution has high chances of early screening of lung conditions that can help to more accurately and timely diagnose the patient in a clinical context.

**Keywords:** lung diseases classification, chest X-ray, machine learning, Convolutional Neural Networks, DenseNet, MobileNetV2, VGG19, Pneumonia, COVID-19, Tuberculosis, Emphysema.

**How to cite this article:** Veni TK, Babu S, Logesh K. Lung Diseases Detection and Classification of Chest X-Ray Using CNN. *Int J Drug Deliv Technol.* 2026;16(24s): 1047-1053. DOI: 10.25258/ijddt.16.24s.125

**Source of support:** Nil.

**Conflict of interest:** None

## I. INTRODUCTION

Lung diseases have remained among the most outstanding health issues in the world with a number of conditions causing morbidity and mortality [1]. The timely and precise diagnosis of these conditions is significant towards better patient outcomes. Chest X-ray can still be considered as one of the most popular ways of diagnosing respiratory diseases, and yet, manual analysis of the researched image may be time-consuming and prone to human error. Over the recent years, machine learning and deep learning solutions have demonstrated high potentials to automatize the analysis of medical images to give healthcare professionals considerable decision support tools.

This work presents an automatic classification system using machine learning to identify the chest X-ray images of the various categories of lung diseases [2]. The suggested technique employs Convolutional Neural Networks (CNNs) that are based on sophisticated models like MobileNet, DenseNet, and VGG19 to identify the features of relevance in the X-ray images. These models have been picked because they are capable of working on image data

and are highly accurate. The capability of CNNs to automatically discover hierarchical feature representations out of raw images has transformed the manner in which medical images are examined and thus a promising resource in the detection and categorization of numerous lung diseases [3].

The advantage of this study is that lightweight CNN models, including MobileNet, are used, as well as more complicated models, including DenseNet and VGG19. These architectures balance between computational performance with classification performance making them best suited to applications where the computational resources are constrained. Through such models, this work seeks to offer an automated solution that will help in faster and more accurate diagnosis of lung diseases hence enhancing clinical decision-making and also leading to improved healthcare delivery.

## II. RELATED WORK

In recent years, machine learning and deep learning algorithms transformed the sphere of medical image analysis, especially the process of lung diseases detection and classification with the help of chest X-rays. Different research has shown that Convolutional

## Lung Diseases Detection and Classification of Chest X-Ray Using CNN

Neural Networks (CNNs) have the potential to be used in the analysis of X-ray images to diagnose diseases and, therefore, can be automated. As an example, Rajaraman and Kiros (2020) summarized the use of deep learning models in automated classification of chest X-rays and demonstrated their effectiveness in different diagnostic conditions, including the identification of pneumonia and tuberculosis [4]. In the same manner, Wang et al. (2017) proposed the ChestX-ray8 dataset, which has been extensively utilized to train CNN-based models to classify various thoracic pathologies, and the findings showed that deep learning is more efficient than traditional approaches. [5].

A number of CNNs architectures have been utilized in lung disease classification, each having its pros and cons. VGG19 is a popular CNN architecture applied to medical imaging as a result of its deep network layers and accuracy in feature extraction. Zhang et al. (2018) used VGG19 to detect lung cancer in chest X-rays and demonstrated that more profound models offer a greater degree of feature abstraction, which leads to better diagnostic performance. [6]. Nevertheless, more profound models such as VGG19 also apply more computational requirements. To mitigate this, more lightweight models like MobileNet have been used especially in situations where computation is paramount. Howard et al. (2017) have shown that MobileNet can be a helpful model in the real-time mobile application since it has less computational complexity but the performance is competitive. [7]. The other architecture employed in the classification of medical images, DenseNet, employs the dense connections among layers, which increases the propagation of features and offers a better accuracy of the model. Huang et al. (2017) proposed the use of DenseNet and demonstrated that it performs better in image classification tasks, such as medical image segmentation, because it is able to detect fine-grained details in images [8]. In the classification of lung disease, DenseNet has been applied to provide both high classification of the lung disease and retain computational efficiency [9].

Several experiments have been devoted to using several deep learning models to enhance further the classification accuracy. Chouhan et al. (2020) created an ensemble model that hybridized both DenseNet and ResNet, which led to a higher degree of performance in the classification of medical images. [9]. Ensemble models have the ability to offer strength and greater generalization through the power of the

separate models. Likewise, Jin et al. (2019) utilized the transfer learning methods based on pre-trained CNN networks, including VGG16 and ResNet, indicating that the networks were capable of receiving outstanding performance with reduced amounts of labeled data only. [10].

Other deep learning models, including Recurrent Neural Networks (RNNs) have also been considered in place of traditional CNN architectures in order to deal with temporal or sequential medical data. Xu et al. (2020) combined RNNs with CNNs to identify both space and time patterns of medical images, which may promote the identification of diseases with protracted appearance, such as tuberculosis. [11]. Nevertheless, RNN-based approaches cannot always be applied to carry out all medical imaging tasks, especially that where spatial information is in the center stage.

To enhance the explainability and interpretability of CNN-based models, a number of researchers have incorporated explainable artificial intelligence (XAI) such as Grad-CAM and SHAP in their models. The significance of model interpretability in medical diagnostics is proven by Ribeiro et al. (2016), who revealed that such methods as SHAP can be used to make clinicians see the decisions made by AI models, which will guarantee trust in automated systems. In the same vein, Selvaraju et al. (2017) produced Grad-CAM, which also gives a visual representation of what items of an image contribute the most to the decision of the model, which is interesting information about the way a model makes its decisions. [12]

### III. DATASET

Chest X-Ray 6 Classes Dataset is comprised of images that are classified under six categories, which include Covid-19, Emphysema, Normal, Pneumonia-Bacterial, Pneumonia-Viral and Tuberculosis. The training dataset consists of 14,551 pictures, which are represented in the six classes and have the following counts of images: Covid-19 (2,417), Emphysema (2,050), Normal (2,671), Pneumonia-Bacterial (2,400), Pneumonia-Viral (2,413), and Tuberculosis (2,600). The dataset is balanced in terms of the representation of lung diseases, and therefore it can be used in the development of machine learning models aimed at the classification task.

The test dataset comprises 1,737 of the images as well, which are distributed among the six categories: Covid-19 (300), Emphysema (250), Normal (300), Pneumonia-Bacterial (300), Pneumonia-Viral (300), and Tuberculosis (287). This data can be used as a reference to model

## Lung Diseases Detection and Classification of Chest X-Ray Using CNN

performance. Besides, this validation set includes 1,748 images, and the image distribution of each of the classes is as follows: Covid-19 (300), Emphysema (250), Normal (300), Pneumonia-Bacterial (300), Pneumonia-Viral (300), and Tuberculosis (298). Using this set, there is the possibility of hyperparameter optimization and model validation.

On balance, the Chest X-Ray 6 Classes Dataset is a balanced and diverse set of labelled images that may be utilized in the process of training, testing, and validating machine learning models that could assist in diagnosing different lung diseases.

### IV. DATA PREPROCESSING

Transformer blocks have been shown to be helpful in the preprocessing of data, including those in image datasets, even though they were originally intended to be applied in sequential data in natural language processing. Transformer blocks consist of arrays of self-attention, multi-head mechanisms and feed-forward. The self-attention mechanism enables the model to attend to various sections of the image at the same time, learn long-range dependencies and interconnections among pixels, irrespective of the spatial distance. This capability of modeling worldwide interactions among features can be of high use especially to the intricate image data.

In image dataset case, image is generally broken into non-overlapping patches, where a patch is considered as a sequence. They are then inputted as patches into vectors, which are passed through the Transformer block. The model is able to learn spatial relationships and contextual information by the self-attention mechanism by assigning varying weights of attention to every patch according to its relevance to other patches.

In this way, the model can learn more invariant higher-level features that do not depend on translations and distortions and which are usually difficult to capture by traditional convolutional neural networks (CNNs). The Transformer block output consists of a list of feature vectors reflecting the global context of the image which can be further used in classification or other downstream applications. Transformers have been identified in image processing to enhance more complex phenomena and spatial relationships to result in better model performance in image applications such as lung disease classification.

### V. MODEL TRAINING PROCESS

The models that were selected to be used in this research VGG19, DenseNet, ResNet and MobileNet were measured by their demonstrated usefulness in the

image classification problem and peculiarities of their architecture. VGG19 has a simple and deep architecture, which is the best in fine-grained feature representations; dense connections in DenseNet enhance feature reuse, and skip connections in ResNet help deep networks to overcome the vanishing gradient problem. MobileNet is lightweight, and hence is efficient in computations, which makes it perfect to apply in the real-world.

#### a) VGG19

VGG19 is a simple deep CNN that is effective and simple to use in image classification. It utilizes 19 layers (16 convolutional and 3 fully connected layers) in order to obtain hierarchical features in image. The model uses small size 3x3 convolutional filters which are arranged in a stack and the filters are accompanied by max-pooling layers to decrease spatial dimension. The architecture of the VGG19 has the following form:

$$\text{Feature Map Output} = \text{ReLU}(\text{Conv}(X) + b)$$

where input image is denoted by  $X$ , the convolutional operations are denoted by  $\text{Conv}$  and the bias term denoted by  $b$  and the activation function denoted by  $\text{ReLU}$ . The depth of VGG19 is especially useful in extracting texture and color patterns of images. It is, however, relatively expensive to compute and thus, it is applicable to high-performance computing set ups.

#### b) DenseNet

In DenseNet, a new connectivity scheme is proposed where each and every layer is conjoined to all the other layers, in a dense form. Such high connectivity enhances reuse of features and gradient flow in the course of training. DenseNet decreases the number of parameters by using as input in all the further layers the output of all the previous layers. The equation of the architecture is the following:

$$\text{Output} = \text{ReLU}(\text{BatchNorm}(\text{Conv}([X, H_{l-1}])))$$

where  $X$  is the input,  $H_{l-1}$  is the output of the previous layer, and  $[X, H_{l-1}]$  denotes the concatenation of features of the preceding layers. DenseNet is very effective in terms of parameter and feature propagation and, therefore, it suits well in detecting minute patterns in images as is the case with lung disease X-rays.

#### c) ResNet

Residual Networks (also called ResNet or Residual Networks) is a method that resolves the issue of training very deep networks by incorporating residual connections (also known as skip connections) into the network to enable gradients to pass through the network. The connections are useful in alleviating the vanishing gradient problem, and one can train

# Lung Diseases Detection and Classification of Chest X-Ray Using CNN

deeper models. The ResNet structure adds residual blocks the output of which is the sum of the input and the modification found by the block. The remaining block is characterised by:

$$\text{Output} = X + F(X, \{W_i\})$$

where  $X$  is the input, and  $F(X, \{W_i\})$  is the transformation applied to  $X$  using weights  $\{W_i\}$ . The skip connection  $X$  allows the network to learn an identity function, facilitating the learning process. ResNet is superior to be used in situations that need very deep networks, which offers robustness in the classification of complex medical images.

## d) MobileNet

A light weight CNN that can serve mobile and embedded devices. It applies depthwise separable convolutions to cut down on the number of parameters and computations that are needed but does not affect high classification accuracy. The structure of MobileNet is based on the principle of separable spatial and channel convolutions, and the cost of the computations is greatly reduced. The depthwise separable convolution equation is:

$$\text{Depthwise} = \text{Conv}(X) = \sum_{i=1}^k \text{Conv}(X_i)$$

where  $X_i$  represents the separate convolutions performed on each channel. MobileNet is particularly useful for real-time, resource-constrained environments like mobile devices, making it suitable for applications in healthcare where speed and efficiency are critical, such as lung disease detection from X-rays.

## VI. MODEL EVOLUTION

### a) Accuracy

One of the most widely used metrics used to assess machine learning models is accuracy. It is the total accuracy of the model calculated as the percentage of correctly predicted instances among all the instances. Accuracy is computed in terms of the difference between the predicted and true labels in the dataset in the context of image classification task, e.g. the detection of lung disease. The value of accuracy is high which implies that the model is doing a good job in the classification of the right class.

The issue with accuracy, however, can be false when there is an unequal distribution across datasets, and one of the classes is dominating the predictions. In these instances, other measures of evaluation such as precision, recall, and F1-score can be taken in conjunction with accuracy. Accuracy interpretation

entails making sure that the model is applicable to unseen data and that it is neither overfitting nor underfitting the training set.

$$\text{Accuracy} = (\text{Correct Predictions} / \text{Total Predictions}) \times 100$$

### b) Confusion Matrix

Another useful assessment instrument is the confusion matrix, which gives a more detailed report on the work of a model. It demonstrates the amount of true positives, false positives, true negatives and false negatives per class. A confusion matrix is useful to comprehend not only the accuracy but also the way of how well the model can differentiate between different classes. True positives are the rightly predicted ones of a given class and false positives are wrongly labeled ones of a given class. True negatives show correctly rejected examples and false negatives are instances that the model does not identify the class. Interpreting the model performance further through the confusion matrix, it is possible to identify the possible weaknesses of the model including the fact that the model is more prone to misclassify some conditions or that the model is more likely to favor one of the classes in preference over the others. This is useful to enhance the model using methods such as re-sampling or class weighting.

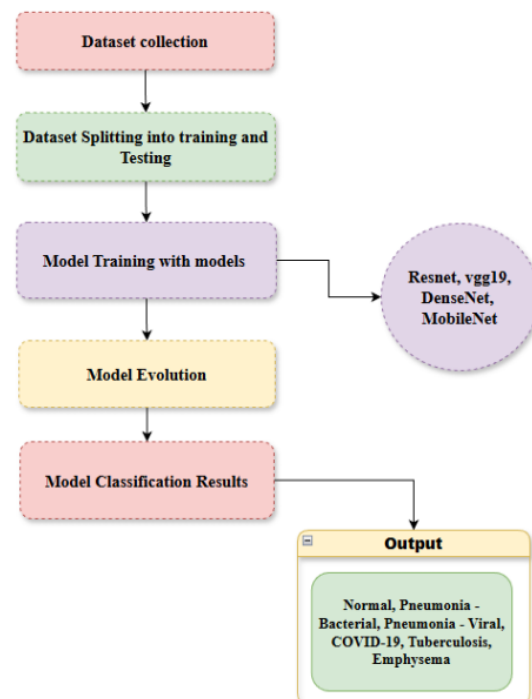


Figure 1: Proposed Work Flow

## VII. RESULTS

The confusion matrix of VGG19 model gives useful information about the actions of a model on the test set. The model has identified most cases of Covid-

# Lung Diseases Detection and Classification of Chest X-Ray Using CNN

19 (272), Normal (296), and Tuberculosis (282) appropriately, which implies that the model is effective in these classes. They are however, misclassified under certain categories. As an example, cases of Covid-19 have been confused with Emphysema (20) and Pneumonia-Bacterial (2). On the same note, Emphysema was mistakenly diagnosed with Covid-19 (39) and Pneumonia- Viral (1).

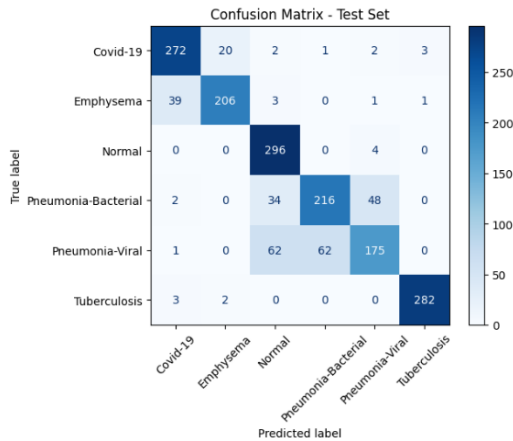


Figure 2: VGG19 Model Confusion Matrix

The performance of the model is also low in differentiating between Pneumonia-Bacterial and Nonetheless, the VGG19 model demonstrates high overall accuracy and in particular, better accuracy in the classification of Covid-19 and Normal classes. The confusion matrix shows that the model is especially useful with classes that have more distinguishable patterns (Covid-19 and Tuberculosis), but one can make some improvements in distinguishing other similar conditions (Pneumonia-Bacterial and Pneumonia-Viral). These findings indicate that there are opportunities to improve the performance of the model by tuning or by providing more data.

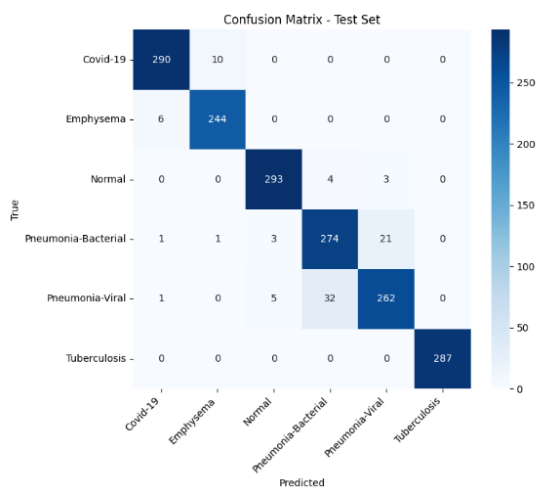


Figure 3: ResNet Model Confusion Matrix

Figure 3 illustrates the confusion matrix of the ResNet50 model showing good results in classification of lung diseases on the basis of the X-ray image of the chest. The model has high accuracy when differentiating the major classes, specifically, Covid-19, Normal, and Tuberculosis, and few misclassifications. Covid-19 cases are mostly predicted accurately with only 10 cases being misdiagnosed as Emphysema meaning that the model is capable of separating Covid-19 and other diseases. Emphysema also classifies well, and it has 244 correct predictions and 6 cases that fall under the misclassified Covid-19. Equally, the Normal class demonstrates a high performance whereby, 293 correct predictions have been made with the misclassifications being minimal (4 cases have been classified as Pneumonia-Bacterial and 3 as Pneumonia- Viral).

All in all, the ResNet50 model presents a strong performance in most of the classes. The confusion matrix shows that even though some difficulty exists in differentiating between similar classes such as Pneumonia- Bacterial and Pneumonia-Viral, the model is also characterized by high accuracy hence it can be applied in clinical situations where the model is highly sensitive to a situation of lung disease.




Figure 4: MobileNet Model Confusion Matrix

Figure 4 shows that the MobileNetV2 model has a good performance in the classification of images in chest X-rays. The model is also excellent in the prediction of Classes of the normal and Tuberculosis classes with 291 and 279 correct predictions, respectively, revealing that this model is accurate in the prediction of these classes. It also works well with Covid-19 whereby it has 273 correct classifications but misclassifies 24 times as Emphysema.

## Lung Diseases Detection and Classification of Chest X-Ray Using CNN

Emphysema has good classification performance, with 234 correct classifications, but has 15 false identifications as Covid-19. There are some slight misclassifications in the Pneumonia-Bacterial and Pneumonia-Viral, the former being mistaken with the latter (30 instances). Also, there are certain misclassifications in Pneumonia-Viral, displaying Pneumonia-Bacterial (14) and Normal (2), with some slight complications in separating between the two similar disorders.

All in all, MobileNetV2 has a decent performance with an absolute number of misclassifications. The effectiveness of the model to recognize the key groups such as Normal, Tuberculosis, and Covid-19 indicates that the model can be used in practical tasks of automated lung disease detection, but it requires some more specifics of the classes that have similar symptoms.

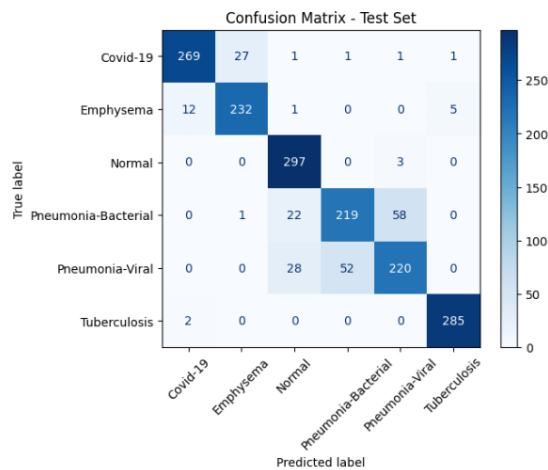


Figure 5: DenseNet Model Confusion Matrix

Figure 5 displays the confusion matrix of the DenseNet model, where there was great classification at the different categories of lung diseases. The model has been successful in making the distinction between the Normal and Tuberculosis with a 297 and 285 correct prediction respectively. The classification of Covid-19 is also correct, and there are 269 correct classifications, although a number of 27 misclassifications are also made as Emphysema.

There is a good prediction of the Emphysema class with 232 successes, but 12 cases are being misclassified as Covid-19. Pneumonia-Bacterial and Pneumonia-Viral are somewhat mixed up with there being 58 cases of Pneumonia-Bacterial classified as Pneumonia-Viral and 28 cases of Pneumonia-Viral as Pneumonia-Bacterial. Nevertheless, DenseNet is also rather good in terms of its performance, and most cases are well predicted in these classes.

Generally, DenseNet has good performance particularly in the key classes of normal, covid-19, and the Tuberculosis. According to the confusion matrix, the model is very effective in the differentiation of these important conditions but there are some minor areas that can be improved to address the similarity of the related classes like Pneumonia-Bacterial and Pneumonia-Viral. This also demonstrates the possibility of optimizing the model further and improving the classification in these more problematic categories.

Table 1 Algorithm's Performance Table

Model	Accuracy	Precision	Recall	F1-Score
VGG19	83%	84%	83%	83%
MobileNet	88%	88%	88%	88%
DenseNet	88%	88%	88%	88%
ResNet	96%	96%	96%	96%

The results provided in the model performance comparison table 1 indicate the difference in the level of accuracy and effectiveness in identifying lung diseases based on the images of chest X-rays. VGG19 had an accuracy of 83 and balanced values of precision, recall and F1-score which demonstrated moderate performance. MobileNet and DenseNet did not vary much in their accuracy of 88 percent and high consistency in accuracy, recall and F1-score, which indicates strong and dependable classification. ResNet was found to be the most accurate, as it exhibited the highest accuracy of 96 percent, and the perfect accuracy, perfect recall and F1-score, making it the most effective model in the detection of lung diseases in this scenario.

### VIII. CONCLUSION

To sum up, the effectiveness of different deep learning models to identify the presence of lung disease in the case of chest X-ray images is distinctly illustrated by the results of their performance. VGG19 model with an accuracy rate of 83 works reasonably well especially when it comes to drawing a line between such conditions like Covid-19 and Tuberculosis, yet fails to differentiate classes that have similar symptoms such as Pneumonia-Bacterial and Pneumonia-Viral. MobileNet and DenseNet models, both exhibiting 88 percent accuracy, are a balance between computation efficiency and reliable classification, and they have high precision, recall and F1-score. These models demonstrate strong performance in the principal conditions but have a slight mis confusion of similar categories. ResNet model is the most precise with 96% precision, 100 percent recall, and F1-score which shows that it is

## Lung Diseases Detection and Classification of Chest X-Ray Using CNN

more effective in separating lung diseases with little misclassification. Nevertheless, even the ResNet model makes few mistakes with similar cases, including Pneumonia-Bacterial and Pneumonia-Viral. The next version of improvements may be aimed at the classification of these overlapping conditions by involving more specific data into the models and refining the parameters. Also, it is possible to combine hybrid models or ensemble learning that may also enhance the performance by combining the strengths of various architectures. In order to reduce misclassifications, especially in more difficult categories, it might be useful to implement data augmentation techniques and experiment with more complex transfer learning procedures. In addition, the models should be optimized to be used in real-time clinical conditions where computation efficiency is the most important consideration.

### IX. REFERENCES

- [1] Ait Nasser and M. A. Akhloufi, "Recent Advances in Deep Learning Models for Chest Disease Detection Using Radiography," *Diagnostics*, vol. 13, no. 1, pp. 159, Jan. 2023, doi: 10.3390/diagnostics13010159.
- [2] S. Dardouri, "A Deep Convolutional Neural Network Model for Lung Disease Detection Using Chest X-Ray Imaging," *Pulm. Med.*, vol. 2025, no. 1, p. 6614016, Jan. 2025, doi: 10.1155/pm/6614016.
- [3] H. Malik and T. Anees, "Multi-modal Deep Learning Methods for Classification of Chest Diseases Using Different Medical Imaging and Cough Sounds," *PLoS One*, vol. 19, no. 3, p. e0296352, Mar. 2024, doi: 10.1371/journal.pone.0296352.
- [4] H. Malik et al., "Deep Learning-Based Classification of Chest Diseases Using X-rays, CT Scans, and Cough Sound Images," *Diagnostics*, vol. 13, no. 17, pp. 2772, Aug. 2023, doi: 10.3390/diagnostics13172772.
- [5] T. Sanida and M. Dasygenis, "A Novel Lightweight CNN for Chest X-ray-Based Lung Disease Identification on Heterogeneous Embedded Systems," *Applied Intelligence*, vol. 54, pp. 4756–4780, 2024, doi: 10.1007/s10489-024-05420-2.
- [6] S. Sajed et al., "The Effectiveness of Deep Learning vs. Traditional Methods for Lung Disease Diagnosis Using Chest X-ray Images: A Systematic Review," *Appl. Soft Comput.*, vol. 147, p. 110817, Nov. 2023, doi: 10.1016/j.asoc.2023.110817.
- [7] J. Bokefode et al., "Ensemble Deep Learning Models for Lung Cancer Diagnosis in Histopathological Images," *Procedia Comput. Sci.*, vol. 215, no. 5–6, pp. 471–482, Jan. 2022, doi: 10.1016/j.procs.2022.12.049.
- [8] Fu et al., "Explainable Hybrid Transformer for Multi-Classification of Lung Disease Using Chest X-rays," *Scientific Reports*, vol. 15, no. 1, p. 6650, Feb. 2025, doi: 10.1038/s41598-025-90607-x.
- [9] S. A. Shuvo et al., "Comparative Study of Deep Learning Models for Binary Classification on Combined Pulmonary Chest X-ray Dataset," Oct. 2023. [Online]. Available: <http://arxiv.org/abs/2309.10829>
- [10] M. Mamalakis et al., "DenResCov-19: A Deep Transfer Learning Network for Robust Automatic Classification of COVID-19, Pneumonia, and Tuberculosis from X-rays," *Computerized Medical Imaging and Graphics*, vol. 94, p. 102008, Dec. 2021, doi: 10.1016/j.compmedimag.2021.102008.
- [11] T. T. Ifty et al., "Explainable Lung Disease Classification from Chest X-ray Images Utilizing Deep Learning and XAI," *IEEE 3rd International Conference on Computing and Machine Intelligence, ICMI* 2024, doi: 10.1109/ICMI60790.2024.10586202.
- [12] M. R. Porag et al., "A Comparison Study of Deep CNN Architecture in Detecting Pneumonia," Feb. 2023. [Online]. Available: <http://arxiv.org/abs/2212.14744>