International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

Cnn-Based Classification of Chest X-Rays for the Diagnosis of Lung Diseases

M. Venkataramana, K. Hemalatha, K. Sudeshna, K. Viswa Subhakar, Mr. R.S. Ravindra Babu
M. Tech, Assistant Professor, Department of CSE
Department of Computer Science & Engineering, LENDI INSTITUTE OF ENGINEERING & TECHNOLOGY

Date of Submission: 15-06-2024 Date of Acceptance: 25-06-2024

ABSTRACT
Lung diseases are a major global health concern that poses serious threats to human beings. These diseases not only cause discomfort and hinder daily activities but can also lead to long-term complications, including respiratory failure and even death. Various factors contribute to the development of lung diseases, including environmental pollution, smoking, genetics, and infections. Timely identification of these diseases is crucial in providing early intervention, which can ultimately save lives. Convolutional Neural Networks (CNNs) have become increasingly significant in the healthcare industry for the early detection of chest diseases. They rely heavily on X-rays, which are a critical diagnostic tool in accurately identifying lung diseases. The primary objective of this research is to utilize X-ray images to predict the occurrence of lung diseases such as Pneumonia, Coronavirus disease (COVID-19), and Tuberculosis (TB). The CNN architecture is employed to predict these diseases, and the model's performance metrics are subsequently presented. The trained model achieved F1 scores of 0.9895, 0.9828, and 0.9857 for Pneumonia, COVID-19, and Tuberculosis disease classes respectively, while it achieved 0.9713 for the Normal class.
Keywords: Convolutional Neural Networks, Pneumonia, COVID-19, Tuberculosis, X-ray.

I. INTRODUCTION

Lung diseases are a diverse group of conditions that affect the respiratory system, which includes the lungs and the airways that lead to and from them. These diseases can range from mild and easily treatable to chronic and life-threatening. Asthma, Chronic Obstructive Pulmonary Disease (COPD), Pneumonia, Tuberculosis, and Pulmonary Fibrosis are among the most prevalent lung diseases.

This study focuses on three significant lung diseases: pneumonia, COVID-19, and tuberculosis.

- **Pneumonia** is a lung infection that is both widespread and severe, affecting individuals of all ages. However, it poses a more significant risk to young children, older adults, and individuals with compromised immune systems. As per the World Health Organization (WHO), pneumonia is the primary cause of mortality in children under the age of 5 years worldwide, accounting for around 15% of all deaths in this age category. In 2019, pneumonia accounted for an estimated 2.5 million deaths globally. In low- and middle-income countries, pneumonia is often caused by bacterial infections, and access to effective antibiotics can be limited. Prevention efforts, including vaccination against pneumococcal disease and Haemophilus influenzae type b (Hib), can help reduce the burden of pneumonia.

- **COVID-19** is a contagious and infectious disease caused by the SARS-CoV-2 virus that primarily attacks the respiratory system, causing symptoms such as cough, fever, and breathing difficulties. As of March 2023, the World Health Organization reports that the COVID-19 pandemic has had a substantial impact on global health, with over 400 million confirmed cases and over 6 million deaths worldwide. The pandemic has also had significant social and economic consequences.

- **Tuberculosis (TB)** is a bacterial infection that mainly targets the lungs and spreads through the air when an infected individual coughs or sneezes. As per the World Health Organization, Tuberculosis (TB) is among the

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

top 10 causes of death globally, and it is the primary cause of mortality resulting from a single infectious agent. Globally, around 10 million cases of TB were reported in the year 2020.

X-rays are a common diagnostic tool used to detect and diagnose various lung diseases, including Pneumonia, COVID-19, and Tuberculosis. X-rays are a form of electromagnetic radiation that can penetrate the body, generating images of internal structures. X-rays are particularly useful in detecting Pneumonia, as they can show the presence of fluid or pus in the lungs, which is a hallmark of the disease. According to the American Lung Association, among the frequently used imaging tests for diagnosing Pneumonia, a chest X-ray is one of the most common. X-rays are also useful in diagnosing tuberculosis, as they can reveal characteristic patterns of lung damage, such as cavities and nodules. Moreover, X-rays can be utilized to track the advancement of COVID-19 in patients displaying severe respiratory symptoms. Although low levels of radiation exposure are involved in X-rays, the advantages of identifying and treating lung ailments at an early stage are much greater than the possible hazards. According to the American College of Radiology, the radiation exposure from a chest X-ray is relatively low and equivalent to the amount of radiation a person receives from natural sources over several months.

The growing challenges and prevalence of lung diseases have spurred researchers to explore effective methods for accurate disease identification and treatment. Artificial Intelligence (AI) has become a very promising instrument, particularly the employment of Machine Learning (ML) and Deep Learning (DL) algorithms. Using chest X-ray images, these algorithms have demonstrated encouraging outcomes in precisely detecting different lung diseases, such as Pneumonia, Tuberculosis, and COVID-19. In our research, we employed deep Convolutional Neural Networks (CNNs) as feature extractors, to obtain robust and unique characteristics from chest X-ray images, which facilitated the accurate recognition of infected and uninfected images with high precision.

II. LITERATURE SURVEY

The most prevalent and cost-efficient approach to diagnose lung diseases, such as Pneumonia, Tuberculosis, and COVID-19, is through the utilization of chest X-rays. This review of the literature presents an outline of recent investigations concerning CNN-based classification of lung ailments.

Zhang et al. (2020) introduced an innovative deep-learning framework for the categorization of COVID-19 and pneumonia using chest X-rays. The suggested approach employed a deep CNN to extract image features, along with a global average pooling layer and a fully connected layer for classification. The outcomes revealed that the proposed method obtained a classification accuracy of 93.62% for COVID-19 and pneumonia.

In 2020, Narin et al. proposed a deep learning technique for identifying COVID-19 through chest X-rays. The suggested method involved a pre-trained ResNet50 CNN model and was able to detect COVID-19 cases with 98.08% accuracy.

Oh et al. (2020) suggested a multi-task learning model to classify COVID-19, pneumonia, and tuberculosis from chest X-rays. Their approach utilized a multi-scale CNN architecture and obtained an overall accuracy of 89.5% for the classification of these diseases.

In 2020, Islam et al. suggested a deep learning method for categorizing COVID-19, pneumonia, and normal cases through chest X-rays. The proposed approach incorporated a ResNet50 CNN model and obtained an accuracy of 94.5% for the classification of these cases.

Lately, the classification of chest X-rays based on CNNs has demonstrated significant potential in identifying respiratory ailments, including Pneumonia, Tuberculosis, and COVID-19. The research articles reviewed in this survey emphasize the efficacy of deep learning-based techniques in categorizing chest X-rays for the diagnosis of lung diseases.

Upon reviewing the above mentioned studies, it was observed that the majority of them primarily focused on the binary classification of either Pneumonia or COVID-19 or other diseases, with a few on focussing on multi- class classification. While some of the studies that employed multi-class classification achieved good accuracy rates, our proposed model differs from them in terms of accuracy, utilization of "Stratified Shuffle Split" technique and using class weights for dealing with class imbalance. We combined the datasets and trained a CNN model to diagnose Pneumonia, Tuberculosis, and COVID-19 through chest X-ray classification. Our proposed model achieved a maximum accuracy rate of 98.46% on the test set.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

III. DATASET

In our work, we employed a dataset named "**Chest X-Ray (Pneumonia, Covid-19, Tuberculosis)**" available on Kaggle's website that contained 7135 chest X-ray images. The dataset consists of three primary folders (train, test, val) and contains subfolders for each image type (Normal, Pneumonia, Covid-19, Tuberculosis).

The 'train' folder has a total of 6326 images, with 1341, 3875, 460, and 650 images belonging to the Normal, Pneumonia, Covid-19, and Tuberculosis classes, respectively. In the 'test' folder, there are 771 images with 234, 390, 106, and 41 images of Normal, Pneumonia, Covid-19, and Tuberculosis classes, respectively. Lastly, the 'val' folder has 38 images, with 8, 8, 10, and 12 images belonging to the Normal, Pneumonia, Covid-19, and Tuberculosis classes, respectively.

The process of creating the above dataset involved combining images from three different Kaggle datasets through collaboration:

1. **Chest X-Ray Images (Pneumonia)**: The dataset comprises chest X-ray images of patients with Pneumonia and those without it. The chosen images were obtained from a retrospective analysis of chest X-ray scans of pediatric patients aged between one to five years, who had undergone routine clinical care and imaging at Guangzhou Women and Children's Medical Centre in Guangzhou. To ensure the quality of the images used in the analysis, all low-quality or unreadable scans were removed during an initial screening process. Two experienced physicians evaluated the diagnoses of the remaining images before allowing their use in training the AI system. Furthermore, to eliminate any potential grading errors, a third expert reviewed the evaluation set.
2. **Tuberculosis (TB) Chest X-ray Database**: A team of scientists from Qatar University, the University of Dhaka, Malaysia, along with medical professionals from Hamad Medical Corporation and Bangladesh collaborated to develop a chest X-ray image database comprising both TB-positive cases and normal images. The database consists of 700 TB images accessible to the public, 2800 TB images that require signing an agreement to download from the NIAID TB portal, and 3500 normal images.
3. **Chest X-ray (Covid-19 & Pneumonia):** The dataset comprises chest X-ray images of individuals diagnosed with either COVID-19 or Pneumonia, along with individuals not affected by any of these respiratory ailments. The images were sourced from various publicly accessible resources.

IV. METHODOLOGY

1. PREPROCESSING

Preprocessing is a critical step in preparing data for training a Convolutional Neural Network (CNN). Preprocessing aims to convert the raw data into a more appropriate format for analysis, with the purpose of enhancing the CNN's performance. Without proper pre-processing, training learning algorithms on raw image data can lead to poor classification performance. A proposed approach is shown in Figure 1, which outlines how the image dataset is processed for analysis. We have preprocessed the dataset in the following ways:

- **MERGING:** To merge the images from the 'train', 'test', and 'val' folders into a single dataframe, we iterated through all the images in the directories and added their paths to the 'path' column and their respective labels to the 'label' column. By doing so, we were able to combine the data from different folders into a single cohesive dataset.
- **SPLITTING:** The distribution of the 4 classes: Pneumonia, Normal, Tuberculosis, and COVID-19 is in the ratio 60:22:10:8. The dataset is imbalanced as illustrated by fig. 1 . We partitioned our dataset into training and testing sets using the Stratified Shuffle Split method, with an 80:20 ratio. We further split the training set using the StratifiedShuffleSplit() method from the sklearn library to obtain a validation set, again in an 80:20 ratio. By doing so, we ensured that the class distribution was maintained in each set, resulting in the same 60:22:10:8 ratio of classes. We ultimately obtained a training set of 4566 images, a test set of 1427 images, and a validation set of 1142 images. The distribution of classes across the training and validation sets is illustrated in Fig. 2.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

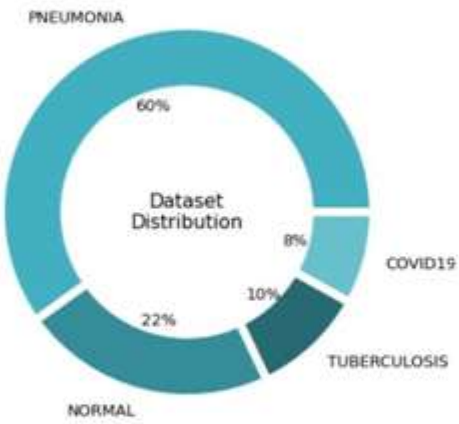

Fig. 1. Distribution of classes in the original dataset

The Stratified Shuffle Split is a machine learning cross-validation method utilized to partition a dataset into training and testing subsets while maintaining the proportion of class distribution in the target variable. In this method, the dataset is first partitioned into groups based on the values of the target variable. Next, the samples are shuffled randomly in each group and divided into training and testing sets. By preserving the class proportions in the target variable for both the training and testing sets, this technique is particularly valuable when working with imbalanced datasets. By utilizing a stratified shuffle split, the model can be trained on a representative subset of the complete dataset. This technique ensures that the model can accurately categorize new samples from each class. Python libraries such as Scikit-Learn and Keras offer this technique as a standard feature.

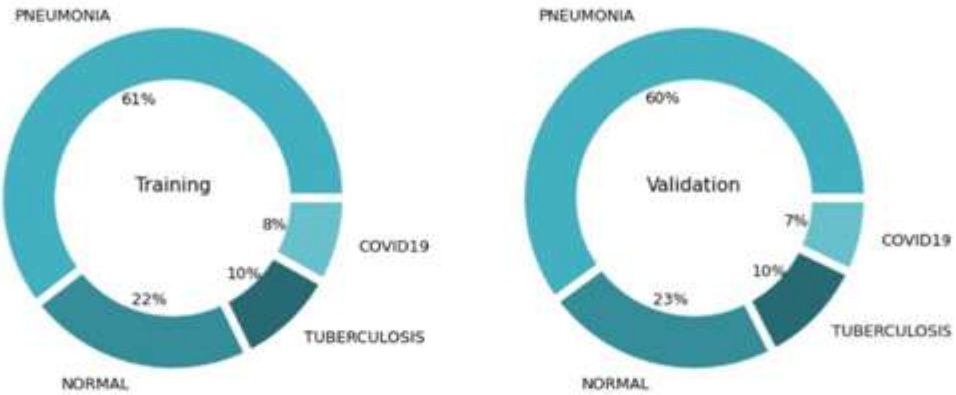

Fig.2. Distribution of classes in Training and Validation sets after Stratified Shuffle Split

- **RESIZING:** To standardize the images in our dataset, we resized them all to a uniform size of 128x128. This was necessary because the images originally had varying sizes, and resizing them to a fixed size ensures that they can be processed consistently and efficiently by our model.
- **RESCALING:** After resizing all images in the dataset to a uniform size of 128x128, we rescaled them using the maximum pixel value in each image. This step is crucial to keep the pixel values in a consistent range, leading to improved model performance during training and inference. By scaling pixel values to a range of [0, 1] using the maximum pixel value in each image, we ensure numerical stability and prevent issues like vanishing gradients.
- **ADJUSTMENTS:** We adjusted the brightness and contrast of the images within the range of 0.64 to

1.37 in the training set, resulting in their enhancement. The methods ImageEnhance.Brightness() and ImageEnhance.Contrast() were used to adjust the brightness and contrast of the images, respectively. The brightness adjustment can lighten or darken the overall image and bring out hidden details. By adjusting the contrast, the distinction between the brightest and darkest pixels is intensified, resulting in a more vivid image. By randomly selecting factors within a given range and applying them to the image, we can simulate different lighting and contrast conditions that may occur in real-world scenarios. By doing so, we can prepare the model to be more robust to these variations during training and inference.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

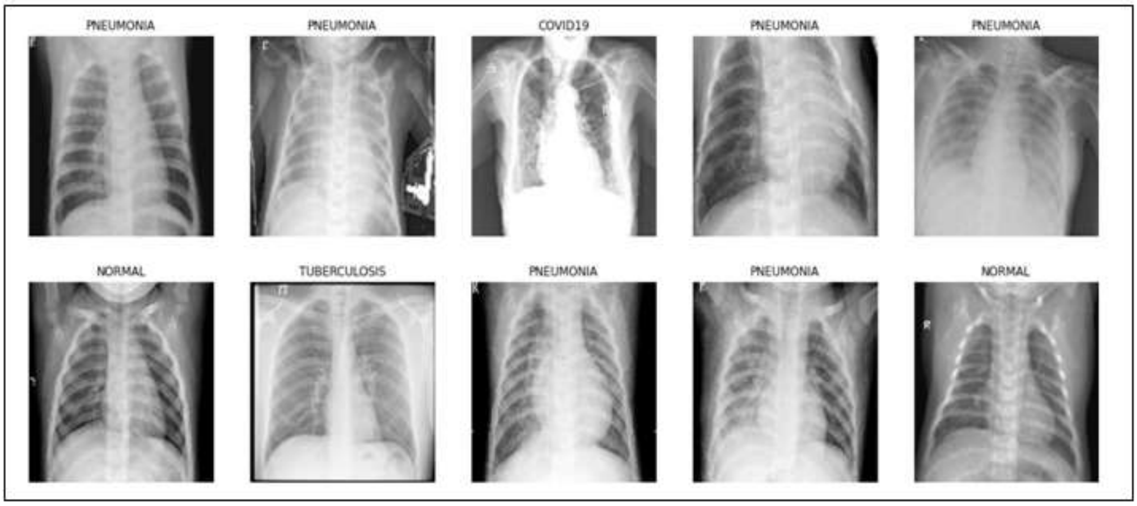

Fig.3. X-ray images of various classes in the preprocessed dataset

2. PROPOSED CNN ARCHITECTURE

Our proposed CNN architecture as shown in fig.4 involves a sequential model that integrates several layers, such as Convolutional, Max-Pooling, Flatten, Dropout, and Dense layers, to facilitate classification tasks. The Convolutional layer utilizes the 'ReLU' activation function, while the output block uses the 'Softmax' activation function.

The convolutional layers are accountable for the primary extraction of features and the identification of patterns in the input data.

The role of Max-Pooling layers is to down-sample the feature maps created by the convolutional layers by selecting the maximum value in each pooling window.

Flatten layer reshapes the output of the previous layers into a 1D vector to be fed to the next layer.

During the training phase, the dropout layer is utilized to avoid overfitting by randomly removing some neurons.

Dense layers are fully connected layers that take input from the previous layer and produce the final output. These layers play an important role in the classification task as they provide the necessary computations to generate the output labels.

The 'ReLU' (rectified linear unit) activation function is used in each of the convolutional layers in this CNN model for feature extraction. The Rectified Linear Unit (ReLU) function is a prevalent activation function in deep learning because of its efficiency and simplicity. By setting all negative values to zero and leaving positive values unaltered, it accelerates training and enhances the overall performance of the model.

For classification purposes, the final output layer typically utilizes the 'softmax' activation function, which is a popular choice for solving multi-class classification tasks. By compressing the output values from the preceding layer into a probability distribution across classes, this step ensures that the total of the output values equals 1. This allows the model to make a probabilistic prediction of the class for a given input image, with the highest probability corresponding to the predicted class.

In our work, a convolutional neural network (CNN) model is defined using the Sequential model from the Keras library.

Fig.5 shows the summary of our CNN model for image classification. The model consists of multiple layers, including convolutional layers (Conv2D), max pooling layers (MaxPooling2D), and a flatten layer (Flatten).

- The model receives an input image of 128x128 pixels with 3 channels (RGB). It begins with a Conv2D layer having 32 filters, a kernel size of 3x3, and an activation function of Rectified Linear Unit (ReLU). Next, a second Conv2D layer with 32 filters and the same kernel size and activation function as the first layer follows. The third layer, a max pooling layer with a pool size of 2x2, decreases the feature maps' spatial dimensions by a factor of 2.
- The fourth and fifth layers are similar to the first and second layers, respectively, with 64 filters in each layer. Another max pooling layer follows the fifth layer. The sixth and seventh layers are again similar to the previous ones but have 128 filters each. A final max pooling

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

layer follows the seventh layer.

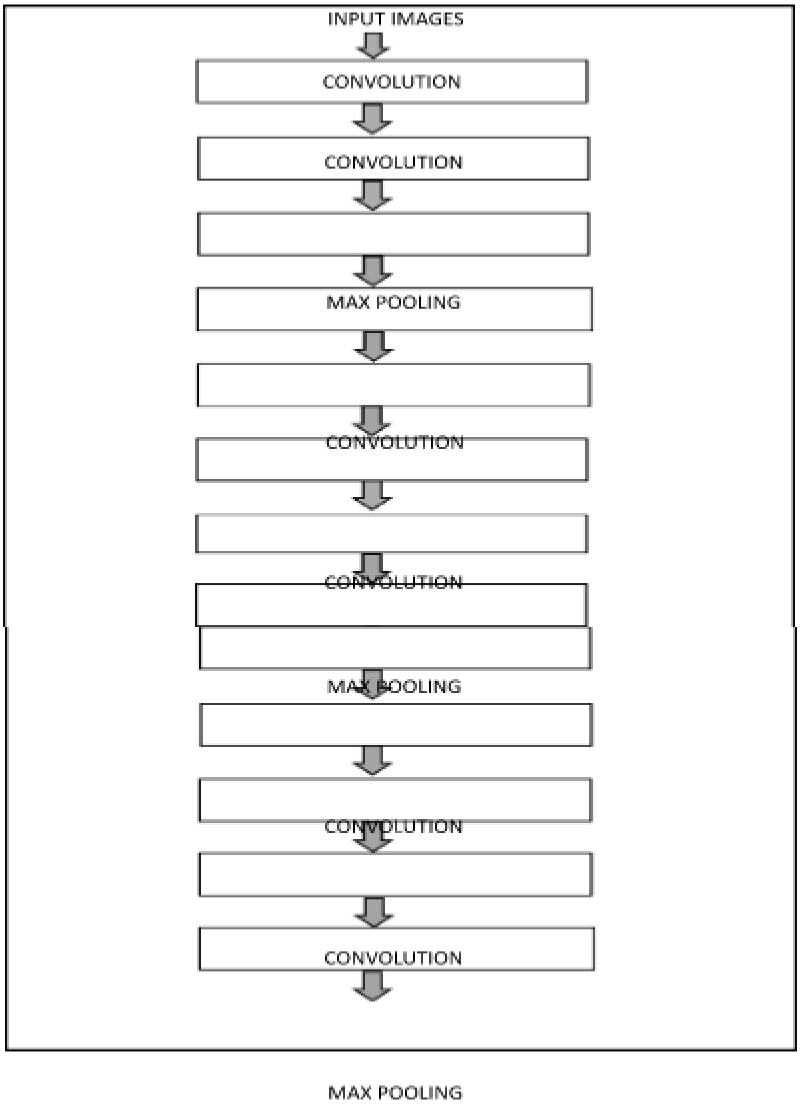

Fig. 4. Proposed CNN Structure

- A Flatten layer is utilized to convert the output of the final max pooling layer into a 1D vector. The next layer is a dropout layer (Dropout) with a dropout rate of 0.3, which randomly drops out 30% of the connections to prevent overfitting.
- The subsequent layer is composed of a fully connected (Dense) layer that comprises 128 neurons and employs a ReLU activation function. The ultimate layer is also a fully connected layer with four neurons, each representing a different class, and utilizes a softmax activation function to generate a probability distribution across the classes.
- The summary also shows the number of parameters in each layer, the total number of trainable and non-trainable parameters in the model, and the output shape of each layer.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

Layer (type)	Output Shape	Param #
block1_conv1 (Conv2D)	(None, 128, 128, 32)	896
block1_conv2 (Conv2D)	(None, 128, 128, 32)	9248
pool1 (MaxPooling2D)	(None, 64, 64, 32)	0
block2_conv1 (Conv2D)	(None, 64, 64, 64)	18496
block2_conv2 (Conv2D)	(None, 64, 64, 64)	36928
pool2 (MaxPooling2D)	(None, 32, 32, 64)	0
block3_conv1 (Conv2D)	(None, 32, 32, 128)	73856
block3_conv2 (Conv2D)	(None, 32, 32, 128)	147584
pool3 (MaxPooling2D)	(None, 16, 16, 128)	0
flatten (Flatten)	(None, 32768)	0
dropout1 (Dropout)	(None, 32768)	0
dense1 (Dense)	(None, 128)	4194432
final (Dense)	(None, 4)	516

Total params: 4,481,956
Trainable params: 4,481,956
Non-trainable params: 0

Fig. 5. Summary of the model

A similar CNN architecture with sequential model using Convolutional, Max-Pooling, Flatten, Dropout, and Dense layers for classification can be found in the paper "Deep learning-based classification of breast tumours using transfer learning" by J. H. Kim et al. For classification purposes, the authors employed a CNN model featuring three convolutional blocks, where each block was succeeded by max pooling layers, along with two fully connected (dense) layers. Dropout was also used to prevent overfitting during training. The model achieved an accuracy of 88.8% on a dataset of breast tumour images.

3. MODEL TRAINING

Upon completion of building our CNN model, we proceeded to train it for image classification. The process of training a CNN involves various steps, such as compiling the model, training the model, evaluating the model, and making predictions.

- The initial stage in training an image classification model through Convolutional Neural Networks (CNN) is to compile the model, which involves specifying the optimizer, loss function, and metrics. The optimizer upgrades the model's weights by reducing the loss function, whereas the metrics assess the model's performance. For our study, we utilized the Adam optimizer with a learning rate of 0.00008, sparse categorical cross-entropy loss function, and accuracy as the evaluation metric.
- The next step is to train the compiled model on a set of training images using the fit() method. During the training process, the model learns to classify images by adjusting the weights of its layers based on the input images and their corresponding labels. In our case, the fit method is used along with callbacks to improve the training process. Class weights are calculated to account for class imbalance, early stopping is used to stop the training process if validation loss does not improve for 5 epochs, plateau is used to decrease the learning rate if validation loss remains the same for 2 epochs, and Tensorboard is used to monitor the training process and save the logs.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

- Once the model is trained, it is evaluated on a validation set to measure its accuracy and performance. This is done using the evaluate() method.
- Finally, the trained model can be used to make predictions on new unseen images using the predict() method, which obtains the model's predicted class labels for each input image.

4. FEATURE EXTRACTION AND CLASSIFICATION OF DISEASES

CNNs are used to extract both low-level and high-level features from chest X-rays for identifying medical conditions like pneumonia, tuberculosis, and COVID-19. The low-level features include edges, corners, and textures, while high-level features consist of anatomical structures' shapes. These features are automatically learned by CNNs during the training process, where the convolutional layers of the network process the input chest X-ray images using learnable filters or kernels. These filters convolve over the input image, extracting different patterns and features at varying levels of abstraction. For instance, the initial convolutional layers learn basic features like edges, corners, and textures that capture the image's fundamental structure. In contrast, the subsequent convolutional layers learn more complex features, such as the shapes of anatomical structures, which are crucial for identifying specific medical conditions.

Our approach involves loading preprocessed images from a dataset and feeding them into our model as a numpy array with multiple channels. The model consists of several convolutional layers, with each layer responsible for extracting distinct features from the input image. The first two layers have 32 filters with a size of 3x3, while the second two have 64 filters with the same size. The third set of layers consists of 128 filters with a size of 3x3.

The convolutional layers are then followed by max-pooling layers that help reduce the dimensions of the image while preserving its essential features. The max-pooling layers have a pool size of 2x2. Following the convolutional and pooling layers, the output is flattened and forwarded through a dropout layer, which randomly omits certain neurons during training to avoid overfitting.

The flattened output is then passed through a fully connected layer with 128 neurons and a ReLU activation function. This layer is specifically designed to capture high-level features of the image, serving as a feature extractor.

The final step involves passing the output through a softmax layer comprising of four neurons. This layer computes the likelihood of the input image belonging to each of the four classes. During training, the model learns the optimal weights for each layer by using a training dataset of chest X-ray images and their corresponding labels. During prediction, the model takes in a new chest X-ray image and produces a probability distribution over the four classes, allowing it to predict the most likely class for the input image. Fig. 6 shows the probability distribution of a chest X-ray image and the prediction of disease.

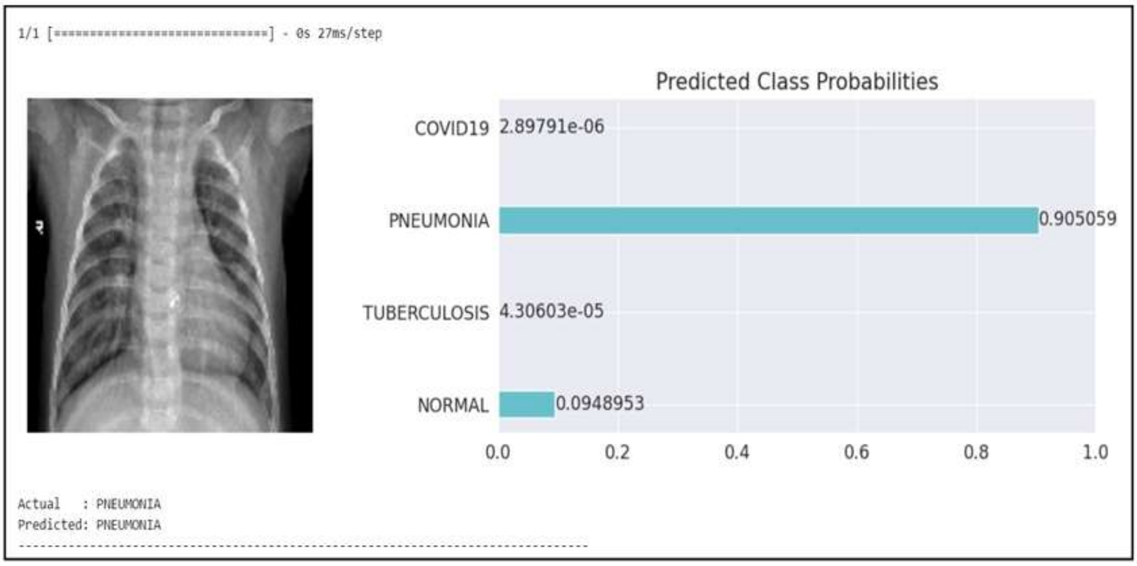

Fig. 6. Prediction of a chest X-ray image after training

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

5. EVALUATION METRICS

Various evaluation metrics are used to assess the performance of models, with some being more suitable for regression models and others for classification models. In this particular research, the performance of the model was evaluated using accuracy, recall, precision, and F1 score. Figures 7, 8, and 9 show the evaluation results of our model.

- **Accuracy** is a commonly used evaluation metric that measures the percentage of correctly classified instances. It is important to note that accuracy can be misleading if the dataset is unbalanced. Therefore, accuracy should be used with caution when evaluating the performance of a model on unbalanced data. It can be calculated as shown in the formula below:

Accuracy = (TP+TN) / (TP+TN+FP+FN)

- **Recall** is a significant evaluation metric that quantifies the proportion of positive instances that the model correctly identifies. Its computation involves dividing the true positive class by the sum of the true positive class and the false-negative class.

Recall = TP / (TP+FN)

- **Precision** is a common evaluation measure utilized in conjunction with recall to evaluate the efficacy of classification algorithms. It gauges the proportion of positive predictions that accurately correspond to positive instances. Precision is calculated by dividing the true positive class by the sum of the true positive and false positive classes.

Precision = TP / (TP+FP)

- **F1 score** is a combined measure of precision and recall and provides an overall measure of a model's performance. It is obtained by multiplying the product of precision and recall by 2, and then dividing the result by the sum of precision and recall evaluation metrics. This metric is valuable in scenarios where the dataset is imbalanced or where both precision and recall are equally significant.

F1 Score = (2 x Precision x Recall) / (Precision + Recall)

V. RESULTS

The following graph illustrates the accuracy of our model in training and validation.

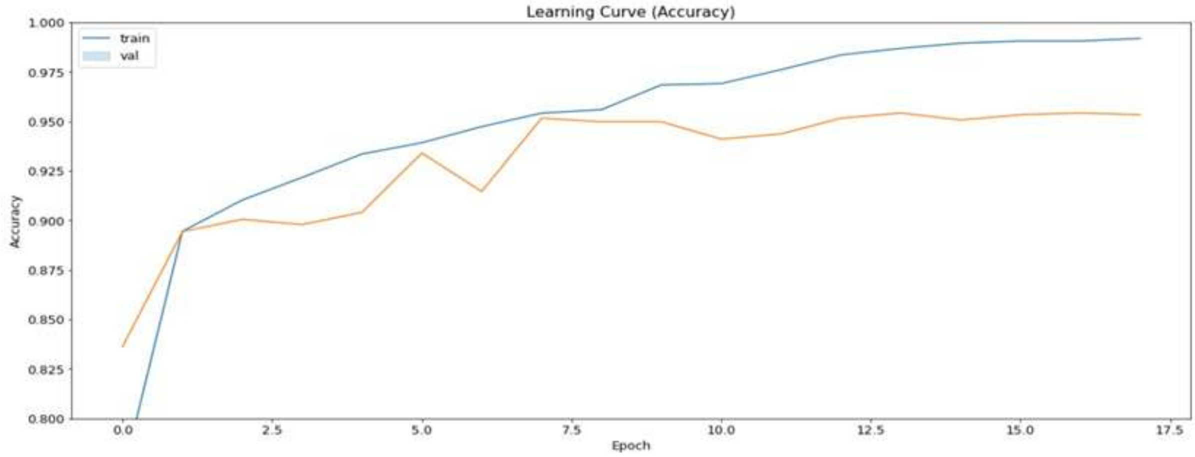

Fig. 7. Model accuracy per epoch

The following graph illustrates the loss of our model in training and validation.

International Journal of Advances in Engineering and Management (IJAEM)
Volume 6, Issue 06 June 2024, pp: 878-888 www.ijaem.net ISSN: 2395-5252

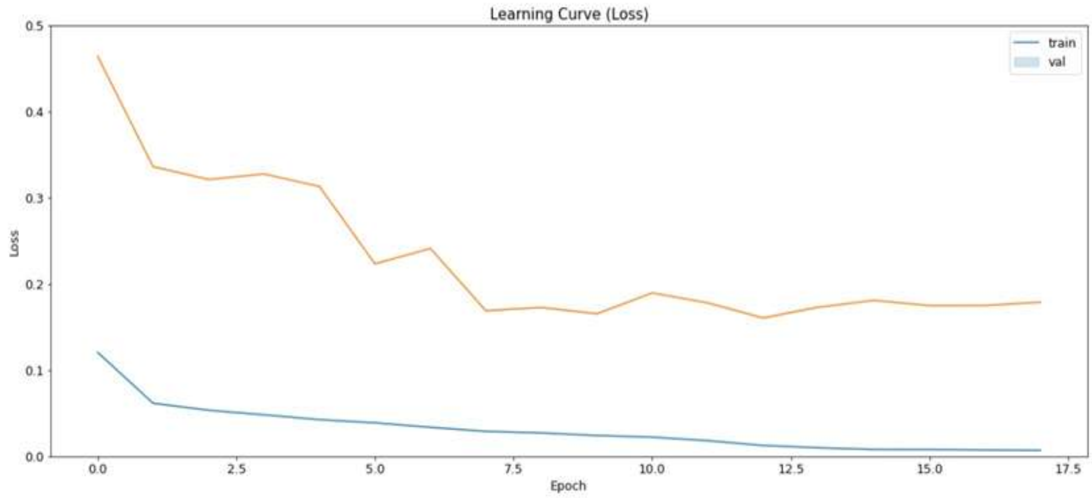

Fig. 8. Model loss per epoch

The following figure shows the model's performance on the test set.

	precision	recall	f1-score	support
COVID19	0.9744	0.9913	0.9828	115
NORMAL	0.9807	0.9621	0.9713	317
PNEUMONIA	0.9860	0.9930	0.9895	854
TUBERCULOSIS	0.9928	0.9787	0.9857	141
accuracy			0.9846	1427
macro avg	0.9835	0.9813	0.9823	1427
weighted avg	0.9846	0.9846	0.9845	1427

Fig. 9. Model test results

VI. CONCLUSION

Our goal was to create a convolutional neural network (CNN) model to assist in detecting lung diseases, including Pneumonia, Tuberculosis, and COVID-19, by analyzing chest X-ray images. We first preprocessed the dataset and then trained the CNN model for 4-class classification, which included Pneumonia, Tuberculosis, COVID-19, and Normal classes. After the implementation of the model, we obtained F1 scores of 0.9895, 0.9857, 0.9828, and 0.9713 for Pneumonia, Tuberculosis, COVID-19, and Normal, respectively. Our model achieved an overall maximum accuracy of 98.46% on the test set, demonstrating its efficacy in diagnosing lung diseases such as Pneumonia, Tuberculosis, and COVID-19.

REFERENCES

[1]. The World Health Organization has provided information on pneumonia at https://www.who.int/health-topics/pneumonia.

[2]. Information on the Coronavirus disease (COVID-19) accessed from the WHO https://www.who.int/emergencies/disease-outbreak-news/item/2022-DON305.
[3]. The World Health Organization offers information on tuberculosis at https://www.who.int/health-topics/tuberculosis.
[4]. American Lung Association. Pneumonia Diagnosis https://www.lung.org/lung-health-diseases/lung-disease-lookup/pneumonia/diagnosin-pneumonia.
[5]. American College of Radiology. Radiation Safety in X-ray, Interventional Radiology and Nuclear Medicine Procedures https://www.acr.org/-/media/ACR/Files/Radiology-Safety/Radiation-Safety-X-ray-Interventional- Radiology-Nuclear-Medicine.pdf
[6]. Lin Li,Lixin Qin,Zeguo Xu, Youbing Yin, Xin Wang, Bin Kong, Junjie Bai, Yi Lu, Zhenghan Fang, Qi Song, Kunlin Cao, Daliang Liu, Guisheng Wang, Qizhong Xu, Xisheng Fang, Shiqin Zhang, Juan Xia, and Jun Xia. "Artificial Intelligence Distinguishes COVID-19 from Community Acquired Pneumonia on Chest CT" https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7233473/
[7]. Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, Andrew Y Ng.
[8]. "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning"https://arxiv.org/abs/1711.05225
[9]. Chest X-Ray (Pneumonia,Covid-19,Tuberculosis) Dataset https://www.kaggle.com/datasets/jtiptj/chest-xray-pneumoniacovid19tuberculosis
[10]. Pneumonia X-ray images references https://data.mendeley.com/datasets/rscbjbr9sj/2 https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia
[11]. "Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning" https://www.cell.com/cell/fulltext/S0092-8674(18)30154-5