Research Article

Training AI to Detect Tuberculosis in Chest X-Rays

Khushi Agarwal

Greenwood High International, Varthur Sarjapur Road, Bengaluru, Karnataka, 560087, India

ABSTRACT

Tuberculosis (TB) remains a significant global health challenge, leading to over 1.3 million deaths in 2022. Early and accurate detection of TB is crucial for effective treatment. Shortage of expert radiologists, especially in resource-constrained settings is one of the most significant impediments in detection of TB. This study explores the application of artificial intelligence in detecting TB features in chest X-rays. Utilizing a data set of 4,200 labeled chest X-ray images, a convolutional neural network (CNN) was developed to identify key radiological features associated with TB, including infiltrates, cavities, pleural effusion, enlarged lymph nodes, and miliary patterns. The model achieved an accuracy of 97.98%, demonstrating its potential as a supplementary tool for healthcare professionals. However, limitations such as reliance on a homogeneous dataset and lack of integration of patient medical history suggest further enhancements are necessary. Future work should focus on diversifying the dataset and incorporating comprehensive diagnostic elements to improve the model's applicability in varied clinical scenarios.

Keywords: Tuberculosis; AI; Detection; X-rays; Radiologists; Artificial; Intelligence

INTRODUCTION

Tuberculosis is the second leading cause of death amongst all infectious diseases after Covid-19, responsible for 1.3 million deaths in 2022, despite being a curable disease (1). Caused by the bacterium *Mycobacterium tuberculosis*, this transmissible disease primarily affects the lungs but can spread to the rest of the body as well. (2). Comorbidity between TB and other diseases like diabetes, HIV, AIDS, malnutrition, and cavitary diseases have been observed (3). TB can easily be treated by a 6-12 month course of antibiotics (4).

Chest X-rays are a widely available and commonly used tool for TB screening. While X-rays help detect and locate the disease, definitive diagnosis still requires the identification of the *M. tuberculosis bacteria* (5). X-rays can often help differentiate TB from other conditions since they can show certain features specifically related to TB. Typically, to diagnose TB using an X-ray, a clear scan of the thorax is required.

The limited availability of expert radiologists, especially in resource-constrained regions (6), has highlighted the need for artificial intelligence based solutions to assist in the interpretation of medical images. This study provides

Corresponding author: Khushi Agarwal, E-mail: thisiskhushi.agarwal@gmail.com.

Received November 23, 2024**; Accepted** December 15, 2024
https://doi.org/10.70251/HYJR2348.24175182

a method to aid healthcare practitioners in improving the efficiency, accuracy, and timeliness of TB detection, which can then be confirmed using blood tests or skin tests. The purpose of this study is to look into how AI can be trained to look for these specific features in an X-ray to detect TB.

LITERATURE REVIEW

One of the major factors contributing to TB's pervasiveness is that it shares symptoms with many common ailments and hence goes undetected (7). The symptoms of TB include prolonged cough, chest pain, weakness, fatigue, fever, weight loss, and night sweats (1). Early and accurate diagnosis is crucial for effective treatment and disease control. Chest X-rays are a useful tool to diagnose TB since specific features present in an X-ray can differentiate TB from other possible diseases (8). The features this paper will be looking at to make its diagnoses include infiltrates, cavities, pleural effusion, enlarged lymph nodes, and miliary pattern.

Infiltrates are areas of increased opacity, often found in the upper lobes of the lungs in cases of TB (9). When particles of the bacteria are inhaled and are small enough to by-pass the upper respiratory body defenses, they often lodge deep inside the lungs, leading to infection (5). Inflammatory cells and fluid accumulate as the body tries to fight off this infection, leading to the white or cloudy appearance of infiltrates on X-rays (10). Generally, centrilobular nodules around 2-4mm with branching opacities are observed (11). The location of these infiltrates can help determine the stage and form of TB (5).

Cavities are round, hollow spaces within infiltrates caused due to lung tissue destruction (9). This destruction is a result of the immune response to the infection and is recognised in advanced cases of TB. The number and size of these cavities affect the extensiveness of the disease, with more severe cases requiring more aggressive treatment. A larger number of cavities can carry a greater number of TB bacteria, and are hence associated with a higher bacterial load (5). A cavity of approximately 2 cm can hold up to 100 million bacilli. Cavities can potentially lead to many complications, like the risk of bronchiectasis, hemoptysis, and superinfection (12).

Pleural effusion refers to the accumulation of fluid within the pleura layers of the lung (13). The fluid is predominantly lymphocytic, occasionally containing other organisms. It represents a delayed reaction to the presence of the TB bacteria in the pleura and pleural space, and generally clears out within a few weeks (14). The detection of pleural effusion can indicate a more extensive form of the infection, which means it might require additional diagnostic procedures like thoracentesis to confirm diagnoses (13). Studies have reported that pleural thickening $\geq$10 mm can indicate a significant symptom of TB in patients (14).

Enlarged lymph nodes are the result of inflammation and swelling caused when the bacteria is trapped in the nodes and attacked by the immune cells. The bacilli reproduce inside the immune cells if they are not killed, ultimately killing the host immune cells (macrophages) (5). This leads to inflammatory cells being attracted to the area. After causing damage in the lungs, the bacilli move to the lymph nodes and hence lead to them becoming enlarged (2). Enlarged lymph nodes with a short axis dimension >1 cm suggest active disease (11). The presence of these nodes can help differentiate TB from other lung diseases like pneumonia, and is a critical feature in the detection of TB (13).

A miliary pattern is visible in scans when the bacteria spread to the bloodstream, infecting other organs in the body. Detecting a miliary pattern is crucial since it represents a condition that requires immediate treatment. As it suggests that other organs are affected, a more aggressive and comprehensive treatment plan might be required (13). In the beginning stages, nodules appear as 1 mm nodules, both sharply and poorly defined, randomly spread throughout both lungs. If left untreated, the nodules might reach 3-5mm, resembling millet seeds (11).

An accurate and early diagnosis is crucial for effective treatment of TB. Models for this purpose have been developed before, like the one developed by Khan et al (22). This neural network model focused on differentiating between positive and negative cases of TB on a large dataset of more than 12,600 patient records. Another technique has been proposed by Chandra et al for a similar purpose, where the model focused on detecting abnormal chest X-rays by extracting a specific feature from an X-ray such as those described above (23). This study aims to develop a similar model that can help reduce the diagnostic burden on healthcare professionals by classifying X-rays to be either Normal or Tuberculosis, acting as a quicker tool for diagnoses.

MATERIALS AND METHODS

The methodology of this study focuses on developing and applying a machine learning algorithm to analyze chest X-rays to detect specific radiological features indicative of TB. The dataset chosen for training and

testing the neural network is from Kaggle (15). All the images of chest X-rays were already labeled as either "Normal" or "Tuberculosis". The TB database is collected from the following sources:

1. NLM dataset: National Library of Medicine (NLM) in the U.S. has made two lung X-ray datasets publicly available: the Montgomery and Shenzhen datasets.
2. Belarus dataset: Belarus Set was collected for a drug resistance study initiated by the National Institute of Allergy and Infectious Diseases, Ministry of Health, Republic of Belarus.
3. NIAID TB dataset: NIAID TB portal program dataset, which contains about 3000 TB positive CXR images from about 3087 cases. -Note: Due to the data-sharing restriction, we have to direct the potential user to NIAID website where you can get a data-sharing agreement signing option and you can get DICOM images from there easily.
4. RSNA CXR dataset: RSNA pneumonia detection challenge dataset, which comprises about 30,000 chest X-ray images, where 10,000 images are normal and others are abnormal and lung opacity images.

This database has been used in the paper titled "Reliable Tuberculosis Detection using Chest X-ray with Deep Learning, Segmentation and Visualization" published in IEEE Access in 2020. Besides the Kaggle dataset, the neural network in this study was developed using Tensorflow in Google Colab, which is a hosted Jupyter Notebook service (16). To begin with, the 4200 images of the dataset needed to be split into a training set, testing set, and validation set. The testing dataset was made using 20% of the dataset (840 images), 20% of the remaining made the validation set (672 images), and the remaining 80% completed the training set (2688 images).

Figure 1 shows a sample of 9 randomly selected images in a 3 × 3 matrix. The purpose of this Figure is to provide a visualization of the data this study worked with. As shown in Figure 1, four out of these nine images are classified as 'Tuberculosis', while the rest are in the 'Normal' class. The cloudy appearance of the TB X-rays is expected due to pleural effusion, as the literature review explains. All the images from the dataset used in this study are of the dimensions 512 by 512, as Figure 2 shows.

The next step in building the neural network is configuring the dataset for performance. For this, the dataset needs to be dynamically tuned in randomized batches. Next the data needs to be rescaled. Since each

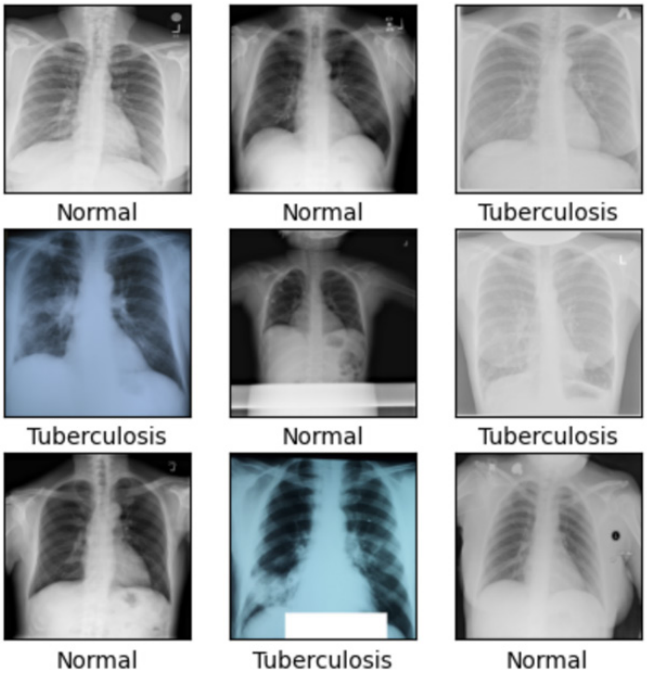

Figure 1. Examples of normal and TB-affected chest X-rays. This Figure illustrates a set of chest X-rays depicting both normal lungs and lungs affected by TB. The normal X-rays show clear lung fields without significant abnormalities, while the TB X-rays display signs such as infiltrates or areas of increased opacity. These images serve to highlight the visual differences between healthy and TB-infected lungs.

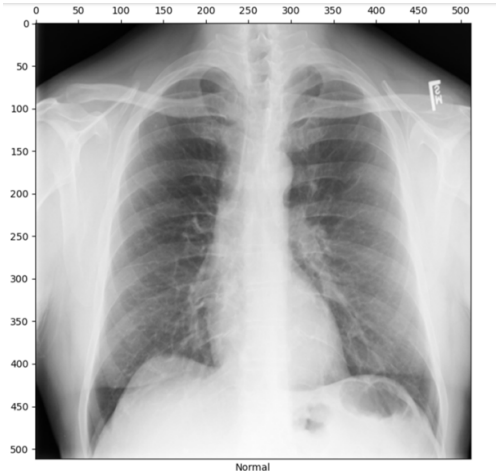

Figure 2. Chest X-ray image with pixel values. This Figure represents how AI converts an image into pixels to process it. The image is converted into a 512 × 512 grid with each pixel corresponding to a specific location of the image. This allows the image to be processed as numerical data, enabling the model to analyze patterns and identify relevant features.

pixel represents 3 possible values; red, green, and blue, the RGB channel values range from 0 to 255, which is not ideal for a neural network. Therefore, the input data is rescaled to range between 0 and 1.

After this, the different filters need to be applied. The first layer is a convolutional layer that extracts the necessary features from the input data, which is the image and creates a feature map. The second filter is the max pooling filter, added to reduce the dimensionality of each feature, which reduces the number of parameters the model needs to learn, hence shortening the training time for the model. Next, a dropout layer is added, which randomly sets a fraction of input units to zero during training. This means that during each training step, some neurons are randomly dropped out of the network, which forces the network to learn redundant representations, thereby preventing overfitting (17). The number of neurons to be dropped is determined by the dropout rate, which was 15% in this study. Adding all these filters results in a 3D tensor which must be flattened to a vector before being fed to the Dense layers. The dense layers help map the derived features to the required class: Tuberculosis and Normal. The final layer is the softmax activation which calculates the output based on probabilities. In other words, the probability each class is assigned with the higher probability being the model's output for the given input. Figure 3 provides a summary of the model structure, showing all the filters that have been added with their parameters.

```
Model: "sequential_4"
_________________________________________________________________
 Layer (type)                Output Shape              Param #
=================================================================
 rescaling_4 (Rescaling)     (None, 512, 512, 3)       0

 conv2d_4 (Conv2D)           (None, 510, 510, 32)      896

 max_pooling2d_4 (MaxPoolin  (None, 255, 255, 32)      0
 g2D)

 dropout_9 (Dropout)         (None, 255, 255, 32)      0

 flatten_4 (Flatten)         (None, 2080800)           0

 dense_8 (Dense)             (None, 128)               266342528

 dense_9 (Dense)             (None, 2)                 258

=================================================================
Total params: 266343682 (1016.02 MB)
Trainable params: 266343682 (1016.02 MB)
Non-trainable params: 0 (0.00 Byte)
_________________________________________________________________
None
```

Figure 3. CNN model summary for TB detection. This Figure displays the structure of the Convolutional Neural Network (CNN) used for TB detection from chest X-ray images. The model includes a rescaling layer, followed by a convolutional layer (Conv2D), a max pooling layer (MaxPooling2D), a dropout layer, a flattening layer, and two dense layers. The final dense layer outputs predictions for two classes: Tuberculosis and Normal. The model has a total of 26.6 million trainable parameters.

The model is now built. This involved compiling the model, adding a loss function and an optimizer function.

After this, the model can be trained. Here is where a certain number of epochs need to be run, which was five for this paper. Epochs refer to the iterations run over the entire training data. Five epochs resulted in a training accuracy of 98.59% and a validation accuracy of 98.21%. While running more epochs could lead to a higher accuracy, the time each additional epoch was taking was considerable. With the default CPU processor, each epoch took approximately 30 minutes, so the TPU processor available on Google Colab was used to speed up the results.

The model then needed to be tested, which was carried out in the same manner and led to an accuracy of 97.98%. This meant that out of the 840 validation images, 17 were incorrectly identified.

RESULTS

Before running the model, the dataset was already classified with an accuracy of 98.3% by Kaggle (15). Out of 4200 images, 3360 were 'Normal' cases, and 840 were 'Tuberculosis' positive cases.

After the dataset was split as described in the methodology, it was fed into the model. As the model was being trained, its accuracy and loss were plotted to visualize the process. In Figure 4, the training accuracy of the model is shown by the blue line and the validation accuracy is shown by the orange line. The graph shows a steep increase of training accuracy in the beginning till the second epoch, after which the rate of increasing accuracy reduces. The model begins with an accuracy of 88.32%, which reaches 96.54% by the second epoch, and a final accuracy of 98.59%. On the other hand, the validation accuracy increases at a lower rate when compared to the training accuracy. The validation accuracy also fluctuates more, not reaching a point where it stabilises. This accuracy begins with a higher value of 95.54%, which increases to 96.28% by the third epoch, and reaches a final value of 98.21%. Despite beginning with a value significantly higher than that of the training accuracy, the value of the validation accuracy is very close to that of the training accuracy.

Along with this, the training and validation loss was also plotted. Figure 5 depicts the training loss using the blue line and the validation loss using the orange line. The loss here refers to the difference between the actual value and the predicted value. The testing loss began with a value of 953.2%, which was significantly reduced by the second epoch to 10.03% and reached a final value of 4.95%. The negative relationship between the accuracy and loss of a model is shown by these results, since a steep increase in training accuracy leads to a steep decline in loss. Furthermore, the validation loss began with a value of 11.22%, which reduced to 3.83% by the fourth epoch but increased to 5.19% by the end of the fifth epoch. The sudden change in validation loss also correlates to the sudden fall in accuracy shown in Figure 4.

Figure 5 shows how the two lines intersect and then diverge, only to repeat the same pattern again, leaving the two lines diverging: an indication of overfitting.

To gain a better understanding of the images incorrectly classified by the model, they were displayed as shown by Figure 6. Out of the 840 validation images, 17 were incorrectly identified. The fourth image is predicted to have TB, but is labeled as normal. The chest X-ray shows no clear signs of opacity or infiltrates, which indicates a possibility that the image has been incorrectly labelled. However, there is also a possibility of the image not being clear enough for the model to correctly predict its class. This prediction of a false positive can be easily verified using a TB skin test or blood test, which is why it is fairly inconsequential. A positive prediction alerts doctors, which acts as a preventative measure in case it is simply misclassified. The rest of the images in Figure 6 are examples of false negatives, where the prediction is normal but the actual label is Tuberculosis, which would have greater consequences during clinical applications. The reason behind these misdiagnoses can be attributed to many factors. The blurriness seen in the majority of the X-rays in Figure 6 is caused due to an advanced stage of TB, where the opacity is visible throughout the lungs. It is possible that the model misinterpreted these images because the dataset it was trained with did not consist of enough images with advanced cases of TB, and so the lack of clarity was ruled out to be bad quality rather than a progressed feature of TB. The purpose of such an algorithm is to reduce a radiologist's work by limiting the number of chest X-rays they need to look at themselves. So, if an X-ray is classified as negative by the algorithm, a radiologist has no reason to verify that result. A doctor would believe they can rule out a diagnosis of TB as a possibility, since the symptoms of TB are very similar to those of other respiratory diseases. This misdiagnosis by the AI leads to the bacteria spreading further while the patient reaches a more severe stage, where they require more aggressive treatment.

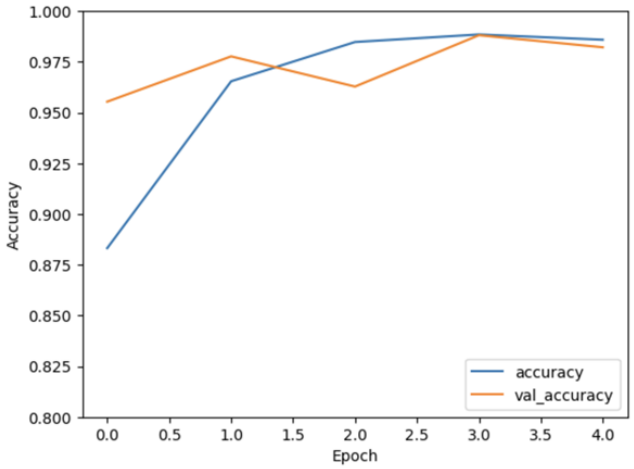

Figure 4. Model Accuracy Across Epochs. The plot shows the training accuracy (blue line) and validation accuracy (orange line) of the model over four epochs. As training progresses, both training and validation accuracy increase, indicating the model's learning performance. The final training accuracy achieved is approximately 98.59%, and the final validation accuracy is approximately 98.21%.

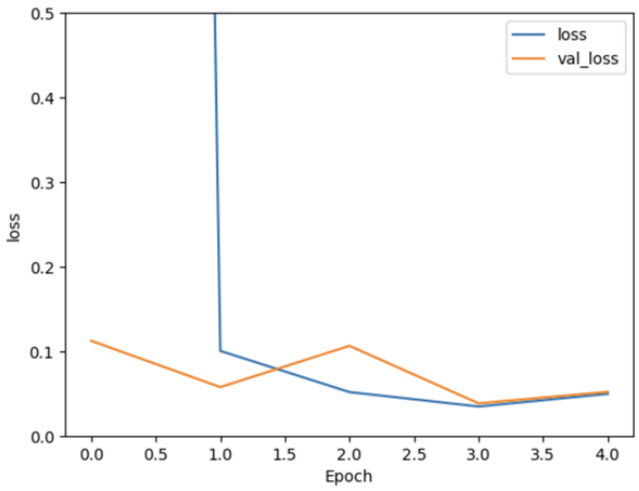

Figure 5. Model Loss Across Epochs. This plot illustrates the training loss (blue line) and validation loss (orange line) over four epochs. The training loss drops sharply after the first epoch, indicating rapid learning, and stabilizes near zero as training progresses.

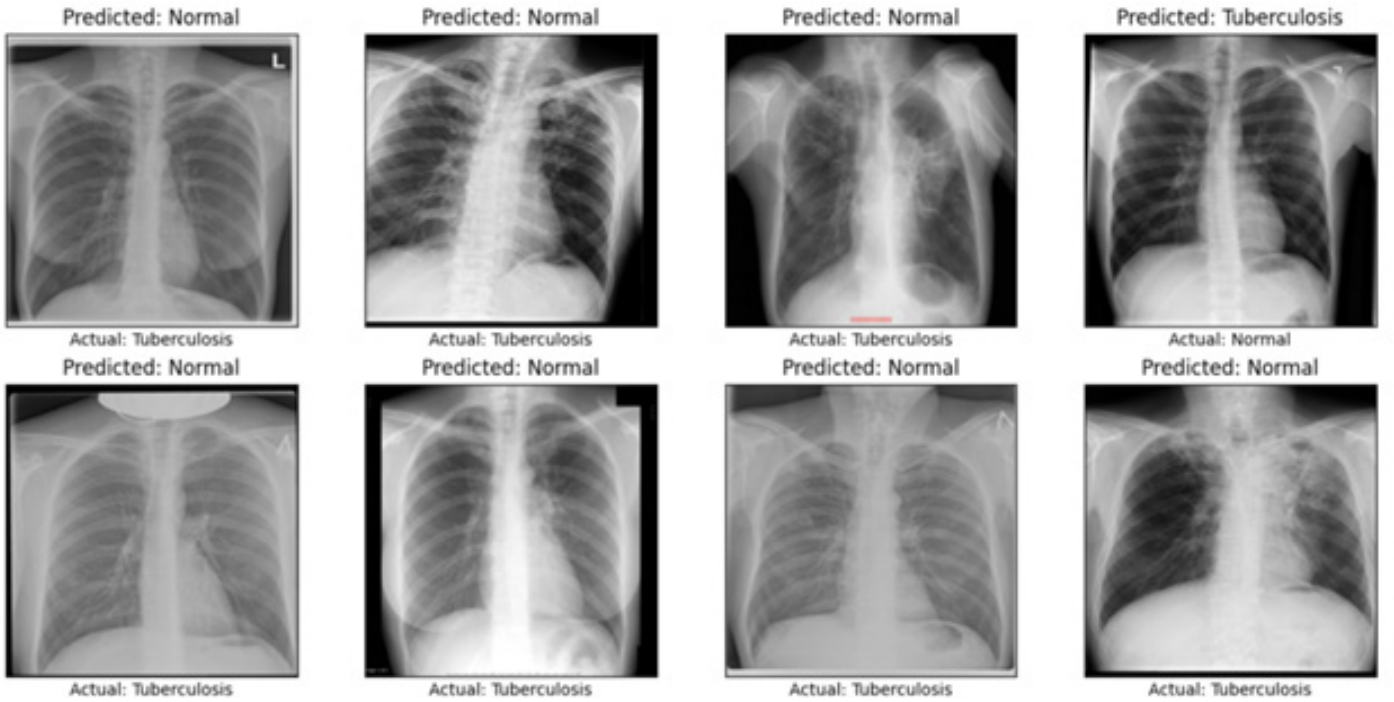

Figure 6. Misclassified Chest X-ray Images by the Model. This is a sample of 8 images out of the 17 incorrectly classified images. The model mainly made mistakes identifying TB positive cases, incorrectly predicting 'Normal' for 7 of the 8 images shown here. Out of the 17 misclassified images, 13 were predicted to be 'Normal', while the rest were predicted to be 'Tuberculosis'.

DISCUSSION

This study demonstrated the potential of AI-based models to detect TB by utilizing chest X-rays as an input. This model's high final accuracy of 97.98% can be partly attributed to data homogeneity, since only one dataset was used, but a high accuracy is seen in other related models too. The AI model developed in this paper differs from other studies since it worked with a limited dataset of 4200 homogeneous images, but did result in accuracy similar to that of models trained using more diverse datasets. The accuracy of the model developed by Liu et al, for example, ranged between 89.0% to 96.1% when tested on 3 different datasets (19). Similarly, in a study conducted by Lakhani and Sundaram, an ensemble of AlexNet and GoogleLeNet was used, which resulted in an accuracy of 98.7% (20).

Since the accuracy of this model has not been compared with actual human diagnoses, its true clinical utility is still uncertain. The algorithm used in this study was trained and validated using the same pre-labelled dataset, which was split to differentiate between the images used for training and validating but still affects the chances of overfitting. In a similar study, chest X-rays were sourced from real cases of TB and possible TB in different hospitals, and then classified by a team of expert radiologists. Their classification was compared with the classification of 3 pre-trained commercial models to verify their actual accuracy in the context of the real world, which provided a better understanding about the clinical usability of such a model (18).

Furthermore, this and other similar studies are limited to the diagnosis of TB, but can be extended to detection of early TB (21). This model's applications are also narrowed since it cannot focus on specific areas of an X-ray and make a prediction with the same reliability. This means that the user's X-ray input needs to match the same resolution and magnification as that of the Kaggle dataset to garner results, which is a challenge in resource-constrained areas. This model only works with X-ray images, even though TB can be detected using other imaging scans like MRI and CT scans. The model has also not been trained using a diverse dataset, preventing it from being accessible to different regions and age groups. As previously mentioned, this study's model is at risk of overfitting, where it can accurately predict the presence of TB in the used dataset but not necessarily in new, unseen

cases. In addition, the model also relies solely on images to make its diagnoses, not taking the patient's medical history or environmental factors into consideration. A limitation of this model is its binary nature, where it can only determine if a chest X-ray has features of TB or not.

Despite this model's limitations, it can still be used in the development of new models. Doctors and radiologists keep a large list of possible diagnoses in mind based on the patient's history and other factors, something this model has not been trained to do. To realistically imitate the human diagnostic process, a model should be able to work with different types of imaging scans and diseases. A more accurate model should be able to work through a database of possible diseases and using any imaging scan as an input, tell the user the probability of each disease. To generate this database, the model should also be trained to consider a patient's medical history and environmental factors, the same way a human diagnosis is made. This would make for a better, more accessible tool for doctors and patients in resource-constrained regions. Once the limitations of this model are addressed and improved, this model can also help confirm human diagnoses, increasing the overall accuracy of a diagnosis. Other than its application in diagnosing, AI like this can also be used to keep track of a patient's progress with treatment, using serial images to predict possible treatment. It can also be expanded to include patient history and lab results into its predictions, which can potentially increase its accuracy too.

CONCLUSION

This study aimed to look into how effective AI can be when trained to look for specific features in a chest X-ray and accordingly classify it as either normal or Tuberculosis. By building a model as described by the methodology of this paper, a testing accuracy of 97.98% was achieved. While a high accuracy is crucial in the context of the real world, the accuracy of this model was due to the limitations of this paper.

The high accuracy was mainly due to the use of a clean and well-curated dataset for training and testing the model. However, since the model has only been exposed to this single dataset, its generalizability is significantly reduced. This limitation can lead to overfitting, where the model effectively "memorizes" the training data, but cannot make accurate predictions on new, unseen data. This poses a serious concern in practical applications, as X-ray images can vary considerably between different hospitals and clinical settings, potentially resulting in misdiagnoses and reducing the reliability of the AI system.

These limitations can be improved by using more diverse datasets with X-rays of varying quality to train the neural network. This model can also be made more accessible for resource-constrained areas by training it to work with other types of imaging scans.

DECLARATION OF CONFLICT OF INTERESTS

The author declares that there are no conflicts of interest regarding the publication of this article.

REFERENCES

1. World Health Organization. Tuberculosis. World Health Organization. 2024. Available from: https://www.who.int/news-room/fact-sheets/detail/tuberculosis (Accessed on 2024-09-10)
2. American Lung Association. Learn about Tuberculosis [Internet]. www.lung.org. American Lung Association; 2020. Available from: https://www.lung.org/lung-health-diseases/lung-disease-lookup/tuberculosis/learn-about-tuberculosis (Accessed on 2024-09-10).
3. Co-morbidities [Internet]. www.who.int. Available from: https://www.who.int/southeastasia/activities/co-morbidities-tb (Accessed on 2024-09-13).
4. American Lung Association. Treating and Managing Tuberculosis | American Lung Association [Internet]. www.lung.org. 2020 [cited 2024 Sep 13]. Available from: https://www.lung.org/lung-health-diseases/lung-disease-lookup/tuberculosis/treating-and-managing (Accessed on 2024-09-13).
5. Nardell EA. Tuberculosis (TB) - Infectious Diseases. MSD Manual Professional Edition. 2022 Available from: https://www.msdmanuals.com/en-in/professional/infectious-diseases/mycobacteria/tuberculosis-tb (Accessed on 2024-09-12).
6. Henderson M. Radiology Facing a Global Shortage. www.rsna.org. 2022. Available from: https://www.rsna.org/news/2022/may/Global-Radiologist-Shortage (Accessed on 2024-09-11).
7. Bloom BR, Atun R, Cohen T, Dye C, Fraser H, Gomez GB, et al. Tuberculosis [Internet]. 3rd ed. Holmes KK, Bertozzi S, Bloom BR, Jha P, editors. PubMed. Washington (DC): The International Bank for Reconstruction and Development / The World Bank; 2017. Available from: https://www.ncbi.nlm.nih.gov/books/NBK525174/ (Accessed on 2024-09-103).
8. Chest X-Ray [Internet]. www.cdc.gov. 2021. Available from: https://www.cdc.gov/tb/webcourses/tb101/page16966.html (Accessed on 2024-10-26).
9. Sherrel Z. Chest X-ray for tuberculosis (TB): What to expect,

results, and more [Internet]. www.medicalnewstoday.com. 2023. Available from: https://www.medicalnewstoday.com/articles/tuberculosis-x-ray (Accessed on 2024-09-10).

10. Zhai W, Wu F, Zhang Y, Fu Y, Liu Z. The Immune Escape Mechanisms of Mycobacterium Tuberculosis. *International Journal of Molecular Sciences*. 2019; 20 (2): 340. https://doi.org/10.3390/ijms20020340.
11. Bhalla AS, Goyal A, Guleria R, Gupta AK. Chest tuberculosis: Radiological review and imaging recommendations. *Indian Journal of Radiology and Imaging*. 2015; 25 (3): 213. https://doi.org/10.4103/0971-3026.161431.
12. Gadkowski LB, Stout JE. Cavitary Pulmonary Disease. *Clinical Microbiology Reviews*. 2008; 21 (2): 305–333. https://doi.org/10.1128/CMR.00060-07.
13. Krishna R, Rudrappa M, Antoine MH, Alahmadi MH. Pleural effusion [Internet]. Nih.gov. StatPearls Publishing; 2024. Available from: https://www.ncbi.nlm.nih.gov/books/NBK448189/
14. Vorster MJ, Allwood BW, Diacon AH, Koegelenberg CFN. Tuberculous pleural effusions: advances and controversies. *Journal of Thoracic Disease*. 2015 Jun 1 [cited 2024 Sep 11]; 7 (6): 981–991.
15. Rahman T, Khandakar A, Kadir MA, Islam KR, Islam KF, Mahbub ZB, et al. Tuberculosis (TB) Chest X-ray Database [Internet]. www.kaggle.com. 2020. Available from: https://www.kaggle.com/datasets/tawsifurrahman/tuberculosis-tb-chest-xray-dataset (Accessed on 2024-09-10)
16. Google Colab. Available from: https://research.google.com/colaboratory/faq.html (Accessed on 2024-10-10)
17. GeeksforGeeks. Implementing Dropout in TensorFlow [Internet]. GeeksforGeeks. 2024. Available from: https://www.geeksforgeeks.org/implementing-dropout-in-tensorflow/ (Accessed on 2024-10-27)
18. Qin ZZ, Sander MS, Rai B, Titahong CN, Sudrungrot S, Laah SN, et al. Using artificial intelligence to read chest radiographs for tuberculosis detection: A multi-site evaluation of the diagnostic accuracy of three deep learning systems. *Scientific Reports*. 2019 Oct 18; 9 (1). https://doi.org/10.1038/s41598-019-51503-3.
19. Liu J, Liu Y, Wang C, Li A, Meng B, Chai X, Zuo P. An Original Neural Network for Pulmonary Tuberculosis Diagnosis in Radiographs. In: Kůrková Věra, Manolopoulos Yannis, Hammer Barbara, Iliadis Lazaros, Maglogiannis Ilias., editors. Artificial Neural Networks and Machine Learning – ICANN 2018. ICANN 2018 (Vol 11140) Cham, Switzerland: Springer; 2018; pp. 158–166. https://doi.org/10.1007/978-3-030-01421-6_16.
20. Lakhani P, Sundaram B. Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. *Radiology*. 2017 Aug; 284 (2): 574–582. https://doi.org/10.1148/radiol.2017162326.
21. Hansun S, Argha A, Liaw ST, Celler BG, Marks GB. Machine and Deep Learning for Tuberculosis Detection on Chest X-Rays: Systematic Literature Review. *Journal of Medical Internet Research [Internet]*. 2023; 25: e43154. doi: 10.2196/43154. https://doi.org/10.2196/43154.
22. Khan MT, Kaushik AC, Ji L, Malik SI, Ali S, Wei DQ. Artificial neural networks for prediction of tuberculosis disease. *Frontiers in microbiology*. 2019; 10: 395. doi: 10.3389/fmicb.2019.00395. https://doi.org/10.3389/fmicb.2019.00395.
23. Chandra TB, Verma K, Singh BK, Jain D, Netam SS. Automatic detection of tuberculosis related abnormalities in Chest X-ray images using hierarchical feature extraction scheme. *Expert Systems with Applications*. 2020; 158: 113514. https://doi.org/10.1016/j.eswa.2020.113514.