



Diagnosis Of Tuberculosis Using Deep Learning

D N Keerthana ^{1*}, D Jayasree ², V Jyothirmayee ³,
C Jyoshna Reddy ⁴, S Ankitha ⁵

¹⁻⁵Department of Computer Science and Engineering,
Aditya College of Engineering , Madanapalle , Andhra Pradesh , India;

* Corresponding Author : D. N. Keerthana; keerthanajoy@gmail.com

Abstract: Preventing the spread of tuberculosis (TB) and enhancing patient outcomes are largely dependent on early detection. This presents a new method for detecting tuberculosis by utilizing convolutional neural networks to examine X-ray pictures of the chest. To extract and learn hierarchical information from x-ray images- a critical component for effective predictions the CNN's architecture has been meticulously built. By examining the interpretability of the CNN model, the study sheds light on the characteristics and patterns that the network considers important when forming predictions. For medical professionals to trust the model and use it as a useful diagnostic tool, this feature is essential. Firstly, it outlines the limitations of conventional TB diagnostic methods, emphasizing the need for more accurate, efficient, and accessible solutions.

Keywords: Tuberculosis Detection, Chest X-Rays, Convolutional Neural Networks, Medical Imaging, Deep Learning.

1. Introduction

Tuberculosis (TB) remains a significant global health challenge, caused by the infectious bacterium *Mycobacterium tuberculosis*. With millions of new cases reported each year, TB primarily attacks the respiratory system, posing a threat to individuals worldwide. This introduction sets the stage for a focused exploration into the symptoms, affected body parts, and innovative approaches, such as Convolutional Neural Networks, to enhance the detection and understanding of TB through the analysis of chest X-rays. Tuberculosis (TB) is an infectious disease caused by *Mycobacterium tuberculosis*. Global health concern: Millions of new cases are reported annually. Airborne transmission [1]. Highlights the contagious nature of the disease. Impact on public health: Emphasizes the significant burden TB places on healthcare systems worldwide. Project Focus: Utilizing Convolutional Neural Networks for enhanced TB detection through chest X-ray analysis. How *Mycobacterium tuberculosis* infects the lungs: *Mycobacterium tuberculosis* is transmitted through the air via respiratory droplets, Inhaled bacteria reach the lungs and may form small, granulomatous lesions, over time, these lesions can lead to more extensive lung damage. Our project aims to revolutionize tuberculosis (TB) detection through the application of Convolutional Neural

Networks (CNN) on chest X-ray images. TB remains a major global health concern, with early detection being crucial for effective treatment and prevention of transmission [2]. By leveraging deep learning techniques, specifically CNNs, we can automate the process of TB detection, reducing the reliance on highly skilled radiologists and potentially speeding up diagnosis. Our CNN model is trained on a large dataset of chest X-ray images, distinguishing between normal and TB-infected lungs with high accuracy.

The CNN architecture learns intricate patterns and features from the X-ray images, enabling it to detect subtle abnormalities indicative of TB infection. Through extensive experimentation and validation, we have achieved promising results, demonstrating the efficacy and reliability of our approach [3]. This project not only showcases the power of deep learning in medical image analysis but also holds immense potential for real-world application in healthcare settings, particularly in regions with limited access to skilled radiologists. By providing a rapid and accurate TB screening tool, our work contributes to the global efforts in combating this infectious disease and improving public health outcomes



1.1. Utilizing GANs for Dataset Augmentation

Generative Adversarial Networks (GANs) are a class of artificial intelligence algorithms used in unsupervised machine learning, introduced by Ian Goodfellow and his colleagues in 2014. They are composed of two neural networks, the generator and the discriminator, which are trained simultaneously through a competitive process. The fundamental goal of a GAN is to generate new data that is indistinguishable from real data.

1.2. Employing an Ensemble of CNNs for Classification

An ensemble approach involves using multiple CNN models to classify the same input and then aggregating their predictions to make a final decision. This strategy leverages the strength of multiple learning models to achieve better performance than any single model could on its own. Key benefits include:

Increased Accuracy: Different CNN architectures may capture different aspects of the data. Combining these models can lead to higher accuracy and sensitivity in TB detection.

Reduced Overfitting: Ensemble methods can reduce the risk of overfitting by averaging out biases and variances across multiple models, leading to more generalizable results

2. Related Work

Previously discussed, deep Convolutional Neural Networks (CNNs) have gained popularity for their enhanced ability in image classification tasks. The convolutional layers, equipped with various filters, are adept at identifying both spatial and temporal characteristics within images. Transfer learning proves to be particularly beneficial for CNN applications where the available dataset is small [4]. It has been effectively applied across various sectors, including manufacturing, healthcare, and security screening, eliminating the need for large datasets, and significantly shortening the typically extensive training durations required for developing deep learning algorithms from the ground up.

For the detection of tuberculosis (TB), nine well-known pre-trained deep learning CNN models were utilized, namely ResNet18, ResNet50, ResNet101, DenseNet201, ChexNet, SqueezeNet, InceptionV3, VGG19, and MobileNetV2, with ChexNet being the exception, all initially trained on the ImageNet database [5]. The Residual Network, or ResNet, designed to address the issues of vanishing gradient and degradation, comes in several variants (ResNet18, ResNet50, ResNet101, ResNet152) distinguished by the number of layers.

ResNet, which focuses on learning residuals rather than features, has shown efficacy in biomedical image classification through transfer learn Dense Convolutional Network, or DenseNet, is noted for requiring fewer parameters than traditional CNNs by avoiding redundant feature map learning. Its architecture includes narrow layers that contribute a small set of new feature maps, offering variants like DenseNet121, DenseNet169, DenseNet201, and DenseNet264. Each DenseNet layer has direct access to the input image and loss function gradients, thus significantly lowering computational costs and making it an optimal choice for image classification tasks. The ChexNet model, a specialized variant of DenseNet121, is trained extensively on chest X-ray images. SqueezeNet and MobileNetV2 stand out for their compact design. SqueezeNet is based on a fire module consisting of a Squeeze Layer with 1×1 filters leading to an Expand Layer that combines 1×1 and 3×3 filters [6]. VGG highlights the importance of network depth, utilizing small receptive fields in its convolutional layers and 1×1 filters for linear transformation, followed by a ReLU layer. Its design maintains spatial resolution post convolution. VGG variants include VGG16 and VGG19. MobileNet, on the other hand, is constructed with depth-wise separable convolutions, except for the initial full convolution layer, leading to a streamlined structure capped by a Softmax layer for classification after a final average pooling step. Incorporating 28 layers when counting both depth-wise and pointwise convolutions, MobileNet emphasizes efficiency. Inception modules in CNNs promote more efficient computation and allow for deeper networks by using 1×1 convolutions for dimensionality reduction, addressing computational cost and overfitting among other challenge.

Identifying tuberculosis (TB) from chest X-ray (CXR) images presents a challenge due to the variety of manifestations like cavities, varying sizes of opacities, consolidation, focal lesions, and nodules. Traditionally, research in this area has focused on extracting handcrafted features from CXR images and using a classifier to distinguish between them. Particularly, textural, and geometrical characteristics are crucial for identifying specific patterns on CXRs. One of the pioneering efforts in automating TB detection by Ginneken et al. involved analysing local textures in CXRs to identify various abnormalities. Their approach segmented lung fields into overlapping areas, from which texture features were extracted. A k-nearest neighbour classifier then assessed each area, combining the scores from all areas into a final score using a weighted multiplier, achieving an area under the curve (AUC) of 0.82 on a TB dataset of 388 images and 0.98 on a 200-image dataset of interstitial lung disease (ILD). This demonstrated the effectiveness of integrating scores from different regions to enhance system performance. Hogeweg et al.

proposed a similar approach by merging scores from different detection systems. Their method involved dividing the CXR image into small circular patches to extract features and calculate a texture score with an LDA (Linear Discriminant Analysis) classifier. Combining this score with a clavicle detection module helped eliminate false positives, and a shape abnormality score was determined using the Mahalanobis distance at the image level. These detection systems could be combined either serially or in parallel to improve TB detection accuracy.

Further advancements by Hogeweg et al. included a more refined TB detection strategy that separately analysed textural, focal, and shape abnormalities before merging them into a comprehensive TB score, utilizing commercial software for focal analysis.

Jaeger et al. introduced using intensity mask, lung model mask, and Log Gabor mask for lung segmentation, followed by employing various shape and texture descriptors to identify pathological patterns. They utilized histogram bins of each descriptor as features, with a linear support vector machine (SVM) classifying the CXR images as normal or abnormal. They also explored a method using two distinct feature sets for automated detection, leading to AUC values of 0.87 and 0.90 for object detection and CBIR feature vectors, respectively.

Some studies have focused on detecting specific TB manifestations. For instance, Shen et al. developed a hybrid method for identifying TB cavities using adaptive thresholding and an active contour model (ACM), with a Bayesian classifier confirming the detections. However, this method struggled with cavity detection if the initial contour placement failed.

Xu et al. introduced a dual-scale approach, first identifying cavity candidates through template matching and feature extraction, then refining these detections with active contour segmentation and SVM classification for false positive reduction, achieving an 82.8% accuracy on a 35 CXR dataset.

Karargyris et al. demonstrated the use of shape and texture features for detecting pulmonary and pleural abnormalities, finding that separate processing of these features not only sped up the process but also improved accuracy.

Maduskar et al. developed an algorithm for detecting pleural effusion, leveraging anatomic landmarks and lung segmentation refinement to assess the severity of pleural effusion with an AUC of 0.87, though its effectiveness was limited in cases where chest wall or pulmonary fissure issues obscured the costophrenic recess.

Recent shifts towards deep learning for TB detection have shown promising results. Hwang et al.'s work with the Alex net network and transfer learning marked the first deep CNN-based TB detection method, illustrating the enhanced performance of networks utilizing transfer learning compared to those without.

3. Experimental Method/Procedure/Design

The application of Generative Adversarial Networks (GANs) for dataset augmentation, combined with the utilization of U-Net for lung segmentation, represents a sophisticated approach in the effort to enhance the detection of tuberculosis (TB) from chest X-rays (CXR). This methodology addresses several of the limitations inherent in deep learning models, primarily those related to dataset quality and quantity, while also refining the focus of the analysis on the most relevant parts of the image.

3.1. Dataset Augmentation with GANs: GANs, with their dual architecture comprising a generator and a discriminator, have shown remarkable capability in generating realistic images. In the context of TB detection, GANs are used to augment existing CXR datasets, generating new images that replicate the diverse manifestations of TB. This augmentation is crucial, especially in scenarios where the available real-world datasets are imbalanced, with a predominance of either TB-negative or TB-positive cases. By synthesizing high-quality, diverse CXR images that exhibit or do not exhibit TB characteristics, GANs help in creating a more balanced and comprehensive dataset. This enriched dataset aids in training more robust deep learning models, capable of recognizing a wide array of TB presentations, thus improving model generalization and reducing bias.

3.2. Lung Segmentation with U-Net: Following dataset augmentation, the U-Net architecture plays a pivotal role in precisely segmenting lung regions from the augmented CXR images. U-Net, known for its efficiency in medical image segmentation, enables the model to focus solely on the lung areas, which are most relevant for TB detection. This segmentation step is crucial as it eliminates background noise and focuses the model's attention, reducing the computational complexity and potentially increasing the accuracy of TB detection. By providing a clear demarcation of lung boundaries, U-Net ensures that subsequent analysis is concentrated on lung textures, patterns, and anomalies indicative of TB, thereby enhancing the sensitivity and specificity of the detection process.

3.3 Integrated Approach for TB Detection: Integrating GANs for dataset augmentation and U-Net for lung segmentation before applying a deep

learning classifier for TB detection forms a powerful triad that leverages the strengths of each approach. The GAN-augmented dataset ensures that the model is trained on a varied and comprehensive set of images, addressing the challenge of data scarcity and imbalance. Meanwhile, segmentation by U-Net sharpens the focus of the analysis, ensuring that the classifier's predictions are based on relevant lung features rather than extraneous information. This integrated approach not only improves the model's performance in detecting TB but also contributes to the development of a more reliable and efficient diagnostic tool that can be deployed even in resource-constrained settings.

3.4.AlexNet: AlexNet's success demonstrated the power of deep learning methods in computer vision and led to a surge of interest in the field. It also popularized the use of GPUs for accelerating the training of deep neural networks. AlexNet is a classic convolutional neural network. It consists of convolutions, max pooling and dense layers as the basic building blocks. Grouped convolutions are used in order to fit the model across two GPUs Alexnet has revolutionized the field of Deep learning.

4. Results and Discussion

CNN : A convolutional neural network (CNN or ConvNet) is a network architecture for deep learning that learns directly from data. CNNs are particularly useful for finding patterns in images to recognize objects, classes, categories. They can also be quite effective for classifying audio, time-series, and signal data.

- Accuracy: 98.9%
- Recall:95.0 %
- Precision: 98%
- F1-Score: 97%

Here we have the graphs from the training history signifying how the models have improved over time with each training.



Figure.1 Plot of loss vs validation loss while training

Accuracy : Accuracy is the most common metric to be used in everyday talk. Accuracy answers the question "Out of

All the predictions we made how many were true" As we will see later, accuracy is a blunt measure and can sometimes be misleading.

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \dots\dots\dots (4.1)$$




Fig 2. Plot of accuracy vs validation accuracy while training

Precision : Precision is a metric that gives you the proportion of true positives to the amount of total positives that the model predicts. It answers the question "Out of all the positive predictions we made, how many were true"

$$\text{Precision} = \frac{TP}{TP+FP} \dots\dots\dots (4.2)$$

Recall : Recall focuses on how good the model is at finding all the positives. Recall is also called true positive rate and answers the question "Out of all the data points that should be predicted as true, how many did we correctly predict as true" As you can see from the definitions of precision and recall they are tightly connected.

$$\text{Recall} = \frac{TP}{TP+FN} \dots\dots\dots (4.3)$$

F1 Score : F1 Score is a measure that combines recall and precision. As we have seen there is a trade-off between precision and recall, F1 can therefore be used to measure how effectively our models make that trade-off. One important feature of the F1 score is that the result is zero if any of the components (precision or recall) fall to zero. Thereby it penalizes extreme negative values of either component

$$\text{F1 Score} = 2 \cdot \frac{\text{Precision} \cdot \text{recall}}{\text{Precision} + \text{recall}} \dots\dots\dots (4.4)$$

Table. 1 Comparison between different pre- defined networks

Architecture	Number of Parameters	Top 5 Accuracy	Top 1 Accuracy
AlexNet	82,379,342	98.0%	84.0%
VGG16	139,347,555	91.0%	74.0%
GoogLeNet	24,000,000	92.0%	74.5%

Table. 2 Results comparison of different algorithms

CLASSIFIERS	ACCURACY%	PRECISION	
		0	1
CNN	98.9	98	100
ANN	88	86	89
RNN	91	86	87

[7] **How To Cite:**

5. Conclusion and Future Scope

The project aimed at harnessing deep learning technologies, specifically through the use of Generative Adversarial Networks (GANs) for dataset augmentation and U-Net for precise lung segmentation, marks a significant stride towards improving the detection and diagnosis of tuberculosis (TB) from chest X-rays (CXR). By leveraging these advanced AI methodologies, the project has demonstrated potential in enhancing diagnostic accuracy, reducing false negatives, and improving the speed of TB detection compared to traditional methods. The use of GANs for generating additional training images addresses the challenge of limited datasets, which often hampers the training of robust deep learning models. Meanwhile, the application of U-Net for segmentation ensures that the focus is maintained on relevant areas of the CXR images, thereby improving the model's ability to recognize TB indicators. Expansion to Multimodal Imaging: Exploring the use of deep learning for analysing additional types of medical imaging (e.g., CT scans) alongside CXRs could provide a more comprehensive view of lung health, aiding in the detection of TB and other pulmonary conditions. Improvement of Model Interpretability:

References

- [1]. Global Tuberculosis Report 2019, WHO, Geneva, Switzerland, 2019.
- [2]. SK. Sharma and A. Mohan, "Tuberculosis: From an incurable scourge to a curable disease journey over a millennium," Indian J. Med. Res., vol. 137, no. 3, p. 455, 2013
- [3]. C. Silverman, "An appraisal of the contribution of mass radiography in the discovery of pulmonary tuberculosis," Amer. Tuberculosis vol.6 0, no. 4, pp. 466–482
- [4]. A. H. van't Hoog, H. K. Meme, K. F. Laserson, J. A. Agaya, B. Screening strategies tuberculosis prevalence surveys: The value of chest radiography and symptoms," PLoS ONE, vol. 7, no. 7 Jul. 2012, 'Art. no. e386
- [5]. A. P. Brady, "Error and discrepancy in radiology: Inevitable or avoidable?" Imag., vol. 8, no. 1, pp 171–182, Feb.

- [6]. A.J.Degnan, E. H.Ghobadi, P.Hardy, E.Krupinski, E. P. Scali, L. Strat A. Ulano, E. Walker, A. P.Wasnik, and W. F. Auffermann 'Perceptual and interpretive error in diagnostic radiology Causes and potential solutions, Academic Radiol., vol. 26, no. 6, pp. 833–845, Jun. 2

How To Cite:

D N Keerthana , D Jayasree , V Jyothirmayee , C Jyoshna Reddy , S Ankitha , " Diagnosis of Tuberculosis using Deep Learning ", International Journal of Computational Sciences and Engineering Research (IJCSEr) , Volume 1, Issue 2 , pp 17-21 , May 2024. CrossRef DOI: <https://doi.org/10.63328/ijcser-v1ri2p4>

Declaration

Conflicts of Interest: The authors declare no conflict of interest.

Author Contribution: All authors wrote the main manuscript text and also consent to the submission.

Ethical approval: Not applicable.

Consent to Participate: All authors consent to participate.

Funding: Not applicable, and No funding was received

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Personal Statement: We declare with our best of knowledge that this research work is purely Original Work and No third party material used in this article drafting. If any such kind material found in further online publication, we are responsible only for any judicial and copyright issues.

Generative AI Statement: The authors declare that generative AI was not used in the preparation or creation of this manuscript. The alternative text (alt text) accompanying the figures in this article was generated by Frontiers with the assistance of artificial intelligence. Reasonable efforts have been made to ensure the accuracy and appropriateness of the generated alt text, including review by the authors wherever possible. If you identify any inaccuracies or issues, please contact the editorial team.

Publisher's Note: The statements, opinions, and claims presented in this article are exclusively those of the authors and do not necessarily reflect the views of their affiliated institutions, the publisher, editors, or reviewers. Any products discussed or evaluated, as well as any claims made by their manufacturers, are not warranted, approved, or endorsed by the publisher.

Acknowledgements

We thank everyone who inspired our work.