




Automated Classification of Tuberculosis from Chest X-Ray Images Using a Custom CNN

Meetali Verma ^{1*}, Vinod Sharma ²

Department of Computer Science & IT, University of Jammu, Jammu, India

* Corresponding Author: **Meetali Verma**

Article Info

P-ISSN: 3051-3502

E-ISSN: 3051-3510

Volume: 07

Issue: 02

Received: 08-07-2026

Accepted: 06-08-2026

Published: 04-09-2026

Page No: 88-94

Abstract

Tuberculosis (TB) is still a leading global infectious disease, therefore early and accurate diagnosis is critical for the effective treatment and control of the disease. The paper presents a deep learning based approach for automatic classification of chest X-ray images into Normal and Tuberculosis (TB) classes. The experiments were performed on a publicly available Kaggle Tuberculosis Chest X-ray Dataset consisting of 3,500 Normal images and 700 TB images. We developed a custom Convolutional Neural Network (CNN) to learn discriminative features from chest X-ray images automatically. The preprocessing pipeline consisted of resizing images to 224×224 pixels, normalizing the pixels, and noise reduction. We applied data augmentation techniques to help the model generalize better and to mitigate the problem of class imbalance. The CNN was trained with Adam optimizer, binary cross entropy as loss function, learning rate scheduling and early stopping. Experimental results demonstrate that the proposed CNN achieved a classification accuracy of 96.07% on the test dataset. The findings demonstrate the potential of the proposed CNN-based approach for automated TB screening from chest X-ray images and provide a foundation for developing computer-aided diagnostic systems, particularly in resource-constrained healthcare environments.

DOI: <https://doi.org/10.54660/IJMER.2026.7.2.88-94>

Keywords: Tuberculosis (TB), Chest X-ray, Convolutional Neural Network (CNN), Deep Learning, Medical Image Classification

1. Introduction

Tuberculosis (TB) is a serious infectious disease caused by *Mycobacterium tuberculosis* and remains a significant public health challenge globally ^[1]. Although TB diagnosis and treatment have advanced, TB is especially common in low- and middle-income countries, where limited healthcare infrastructure and shortages of trained medical professionals may delay timely diagnosis. Prompt diagnosis of TB is essential for initiating adequate treatment and reducing transmission of disease ^[2].

Pulmonary TB is primarily a lung disease and can create different types of abnormalities which are visible in chest X-ray (CXR) images. CXR is one of the most common screening methods because it is relatively cheap, quick and more accessible than some laboratory based diagnostic techniques. Interpretation of chest radiographs can be difficult, however, because manifestations associated with TB can vary widely between patients. The diagnostic process may also be influenced by image quality and the experience of the interpreting radiologist, creating a need for reliable computer-assisted screening approaches ^[3].

The rapid development of artificial intelligence (AI), especially deep learning, has opened up new opportunities for automated analysis of medical images ^[4]. CNNs have shown great capabilities in image classification among deep learning techniques by automatically learning relevant spatial and visual features directly from input images. Compared to traditional machine learning techniques that often require manually engineered features, CNNs can learn hierarchical representations automatically through training, which makes them good candidates for medical image analysis ^[5].

In this work we propose a custom CNN based framework for automatic binary classification of chest X-ray images into Normal and Tuberculosis classes. Publicly, a Kaggle TB chest X-ray dataset is available with 3,500 Normal and 700 TB images. Image preprocessing and data augmentation are incorporated in the proposed framework in this paper to improve the quality and generalization of the model. The CNN is trained with the binary cross-entropy loss and the Adam optimization algorithm and evaluated on an independent test set.

The proposed CNN achieved 96.07% classification accuracy on the test set. We showed that a custom-designed CNN can be used for automated TB screening using chest X-ray images, and it can be a computer-aided support tool, especially in healthcare settings where there is limited access to specialized diagnostic expertise.

The rest of this paper is organized as follows. Section 2 discuss about related work. In Section 3, we describe the dataset and the preprocessing procedures. Section 4 describes the research methodology and the proposed CNN methodology. Section 5 is about experimental results and discussion while at the end section 6 concludes the paper.

2. Literature Review

Deep learning based analysis of chest X-ray images has emerged as a major research direction for automated screening of tuberculosis (TB). In previous studies, convolutional neural networks (CNNs) have been shown to learn clinically relevant patterns directly from chest radiographs and to assist in identifying TB-related abnormalities.

Lakhani and Sundaram^[6] explored deep CNNs for automated TB detection by training AlexNet and GoogLeNet on chest radiographs. In their study, they achieved AUC values close to 0.99 and provided early evidence that deep learning could discriminate TB-related radiographic patterns effectively. Later, Stirenko *et al.*^[7] added lung-region segmentation with CNN-based classification and showed that isolating the relevant anatomical region can improve the detection process. Hwang *et al.*^[8] developed a large-scale deep learning system with about 2.5 million chest X-ray images from the Korean National Health Insurance database. Their independent evaluation showed a sensitivity of 90.9% and specificity of 96.4%, suggesting the promise of deep learning in large-scale TB screening.

With the increasing availability of pretrained networks, transfer learning has also been widely used for chest x-ray classification. Melendez *et al.*^[9] investigated VGG-16-based transfer learning for TB detection, and discovered model performance differences between different patient subgroups. They stated that the correct training data and model adaptation are important factors in deploying deep learning systems for different populations. Subsequently, several researchers have investigated architectural modifications to improve feature extraction.

Liu *et al.*^[10] presented a multi-scale attention-based CNN for TB detection, obtaining an AUC of 0.978 on a diverse multi-site dataset. The study showed that attention mechanisms can help CNNs focus on image regions containing potentially relevant pathological information. A comparative study of existing architectures such as CheXNet, VGG-16, ResNet-50, and Inception-V3 with Shenzhen and Montgomery chest X-ray datasets was done by Rajaraman *et al.*^[11]. Their ensemble approach achieved an AUC of 0.99, demonstrating

the effectiveness of deep CNN architectures for automated TB classification.

Lightweight and computationally-efficient networks have also attracted attention as TB screening systems may have to operate in resource constrained environments. Xu *et al.*^[12] presented a MobileNetV2 based framework and obtained 96.8% accuracy, demonstrating the potential of lightweight CNN architectures for automated screening. Pasa *et al.*^[13] proposed an efficient deep learning approach for large-scale TB screening, leveraging curriculum learning to address challenges arising from noisy labels in public datasets. Similarly, Yadav *et al.*^[14] proposed a DenseNet based architecture using Squeeze-and-Excitation modules on the Montgomery County Chest X-ray dataset and showed a sensitivity of 97.2%. Rahman *et al.*^[15] worked on a Kaggle TB dataset using different transfer learning architectures and found the highest accuracy of 97.37% with InceptionV3 and contrast enhancement using CLAHE.

While these studies show the effectiveness of different deep learning architectures, a direct comparison of reported results is hampered by differences in datasets, preprocessing procedures, training strategies and evaluation protocols. Additionally, relatively few studies have been devoted to the development and evaluation of a custom CNN architecture for binary TB classification while using a relatively simple and reproducible training framework. Hence, in this work, we propose a custom CNN for classification of chest X-ray images into the Normal and TB classes. The aim is to evaluate if a specific CNN architecture is able to produce robust classification results while keeping a relatively simple model design that can be applied to computer-assisted TB screening. The main contributions of this paper are as follows:

- **Custom CNN Framework:** A custom Convolutional Neural Network (CNN) is developed for automatic binary classification of chest X-ray images into Normal and Tuberculosis classes.
- **Image Preprocessing and Augmentation:** To enhance the quality of images and the generalization ability of the models, a preprocessing pipeline including image resizing, normalization, noise reduction and data augmentation is utilized.
- **Automated TB Classification:** Instead of using manually engineered features, the proposed CNN learns the relevant visual features straight from the chest X-ray images.
- **Performance Evaluation:** The proposed model is evaluated rigorously in terms of accuracy, precision, recall, F1-score, sensitivity and specificity to assess its classification performance.
- **Good Classification Performance:** The proposed CNN achieves 96.07% test accuracy, which shows its potential for automated TB screening using chest X-ray images.

3. Materials and Methods

This section presents the dataset, preprocessing procedures, and data augmentation techniques used for developing the proposed CNN-based tuberculosis classification model.

3.1. Dataset Description

The research work used the Kaggle Tuberculosis Chest X-ray Dataset which is publicly available. Dataset contains 2 classes i.e. Normal and Tuberculosis (TB). The dataset contains total 4200 chest X-ray images. The dataset contains

3500 Normal images and 700 TB images which leads to an imbalanced class distribution.

The dataset was split into training, validation and test sets with a ratio of 70:10:20. The same class distribution was kept

the same across the three subsets. The training set consisted of 2,940 images, the validation set of 420 images and the independent test set of 840 images. The full distribution is shown in Table 1.

Table 1: Distribution of the Tuberculosis Chest X-ray Dataset

Class	Total Images	Training (70%)	Validation (10%)	Test (20%)
Normal	3,500	2,450	350	700
Tuberculosis	700	490	70	140
Total	4,200	2,940	420	840

The test set therefore consisted of 700 Normal and 140 TB images, which were kept completely separate from the training process.

3.2. Preprocessing

To get a uniform input format and to enhance the quality of the images before CNN training, image preprocessing was performed. All chest X-ray images were resized to 224×224 pixel. This ensured that each image had the same spatial dimensions as required by the proposed CNN.

Gaussian filtering was applied to reduce image noise and minor acquisition related artifacts. Then, the filtered images were normalized by scaling the pixel intensity values from the original range of 0–255 to 0–1. These pre-processing steps allowed a uniform representation of the input images and helped maintain stable model training.

3.3. Data Augmentation

Data augmentation^[16] was applied so as to improve the generalization ability of the CNN and reduce the overfitting. This was particularly important as the dataset is much less populated with TB images than Normal images. Only the training set was augmented, and the validation and test set were kept the same.

During training, we used seven augmentation operations: rotation, horizontal flipping, zooming, width shifting, height shifting, brightness adjustment, and shearing. Rotation caused slight changes in the orientation of the image,

horizontal flipping and shifting simulated changes in the position of the image. Zooming and shearing also added to the spatial variation, and the model was able to cope with the differences in image brightness by adjusting the brightness. The use of these transformations increased the diversity of training samples without altering the original validation and test distributions. Consequently, the model was trained on more varied image representations while its final performance was evaluated using unmodified test images.

3.4. Research Methodology

The overall research methodology followed in this study is illustrated in Figure 1. The proposed framework consists of five main stages: dataset preparation, image preprocessing, data augmentation, CNN model development and training, and performance evaluation.

The Chest X-ray dataset was first collected and classified into Normal and Tuberculosis classes. The entire dataset was partitioned into training, validation, and testing sets with a ratio of 70:10:20. The training images were also pre-processed for resizing (224×224 pixels), noise reduction and normalization of pixel values.

In the next step, data augmentation was applied to the training images by applying different transformations such as rotation, horizontal flipping, zooming, shifting, brightness adjustment, and shearing. This step was taken to increase the diversity of training samples and to enhance the generalization ability of the model to unseen images.




Fig 1: Research methodology

Following preprocessing and augmentation, a custom Convolutional Neural Network (CNN) was developed for binary classification. The CNN automatically learns relevant

visual features from the input chest X-ray images through successive convolutional and pooling operations. The extracted features are then passed to fully connected layers

for final classification into the Normal and TB categories. The model was trained with the Adam optimizer and binary cross-entropy loss. We trained the network with a learning rate of 0.0001. The model was trained for a maximum of 50 epochs, with the learning-rate adjustment and early stopping to improve training stability and reduce overfitting. We tracked model performance on the validation set during training and the test set was completely unseen until the final evaluation. The trained CNN was finally tested on the independent test set in terms of accuracy, precision, recall, F1-score, sensitivity, and specificity.

3.5. Proposed CNN Architecture

A custom Convolutional Neural Network (CNN) was developed for the binary classification of chest X-ray images into Normal and Tuberculosis (TB) classes. Figure 2 and Table 2 show complete architecture of CNN. The input images (size 224×224 pixels) are fed to four convolutional blocks and a classification head.

Each convolutional block is composed of a Conv2D layer, then Batch Normalization, ReLU activation and MaxPooling layer. The convolutional blocks learn features of increasing complexity. The early blocks learn low level features such as edges and textures whereas the deeper blocks learn more complex and high level features related to pulmonary abnormalities.

The first convolutional block has 32 filters, the second block has 64 filters, the third block has 128 filters and the fourth block has 256 filters. After each block, MaxPooling is used to reduce the spatial dimensions of the feature maps but keep the important information.

The resulting feature maps are then passed to the classification head, where they are flattened and connected to a dense layer with 128 neurons with ReLU activation. To prevent over fitting, a Dropout layer with rate 0.5 is added. And finally, a single neuron with Sigmoid activation gives the probability of input image belonging to the TB class.




Fig 2: presents the complete architecture of the proposed CNN.

Table 2: Architecture of the Proposed CNN

Layer/Block	Main Components	Filters/Units
Input	Image	224×224
Block 1	Conv2D + BatchNorm + ReLU + MaxPool	32
Block 2	Conv2D + BatchNorm + ReLU + MaxPool	64
Block 3	Conv2D + BatchNorm + ReLU + MaxPool	128
Block 4	Conv2D + BatchNorm + ReLU + MaxPool	256
Classification Head	Flatten + Dense + ReLU	128
Regularization	Dropout	0.5
Output	Dense + Sigmoid	1

4. Experimental Setup

We have implemented the proposed CNN model using a deep learning framework for automatic classification of chest X-ray images into Normal and Tuberculosis classes. The preprocessed images were fed to the CNN with the input size of 224×224 pixels. The model was trained with Adam optimizer with learning rate of 0.0001, and binary cross entropy as loss function.

The training process was done for a maximum of 50 epochs. We employed early stopping to avoid excessive training and overfitting, and ReduceLROnPlateau to automatically decrease the learning rate when the validation performance plateaued. The model was monitored for its performance on the validation dataset during training and the test dataset was not seen until the final evaluation.

The performance of the trained CNN was evaluated using accuracy, precision, recall, F1-score, sensitivity and

specificity. A confusion matrix was also generated to evaluate the classification of Normal and TB images and to identify the correctly and incorrectly classified samples.

5. Results and Discussion

The proposed CNN was evaluated on the independent test set containing 840 chest X-ray images, including 700 Normal and 140 Tuberculosis images. The performance of the model was evaluated in terms of accuracy, precision, recall, F1-score, sensitivity and specificity. The proposed CNN achieved an overall accuracy of 96.07% which shows the ability of the CNN to distinguish between Normal and Tuberculosis chest X-ray images. The full classification performance is shown in Table 3. The model correctly predicted a large proportion of TB positive cases in the test set with a sensitivity of 94.29%. The specificity of 96.43% shows that the majority of Normal chest X-ray images were

correctly identified. The achieved F1-score of 88.89% shows the balance between precision and recall.

Table 3: Performance of the Proposed CNN

Metric	Performance (%)
Accuracy	96.07
Precision	84.08
Recall	94.29
F1-Score	88.89
Sensitivity	94.29
Specificity	96.43

The training and validation accuracy and loss curves were analyzed to examine the learning behavior of the CNN during training shown in Figure 3 and Figure 4 respectively. These curves provide an indication of whether the model learned effectively and whether significant overfitting occurred during the training process. A confusion matrix was generated using the independent test set to visualize the number of correctly and incorrectly classified Normal and TB images. It provides a class-wise view of the model's predictions and helps identify false-positive and false-negative classifications.




Fig 3: Training and validation accuracy curve.




Fig 4: Training and validation loss curve.




Fig 5: Confusion matrix of the proposed CNN on the test dataset.

5.1. ROC Curve Analysis

The Receiver Operating Characteristic (ROC) curve was used to evaluate the ability of the proposed CNN to distinguish between Normal and Tuberculosis chest X-ray images across different classification thresholds. The ROC curve is plotted by considering the True Positive Rate (Sensitivity) against

the False Positive Rate (1 – Specificity).

The Area Under the Curve (AUC) provides a threshold-independent measure of the model's discriminatory capability. A higher AUC indicates better separation between the Normal and TB classes. The loss values obtained during the experiment are summarized in Table 4.




Fig 6: ROC curve of the proposed CNN for Tuberculosis classification.

Table 4: ROC-AUC and Loss Performance of the Proposed CNN

Metric	Value
ROC-AUC	0.9891
Training Loss	0.1021
Validation Loss	0.1284
Test Loss	0.1317

The training loss of 0.1021 decreased during model optimization, while the validation loss reached 0.1284. The final test loss was 0.1317, showing a relatively small difference between the validation and test losses. Together with the high ROC-AUC, these results indicate that the CNN learned useful discriminative features and maintained good performance on unseen test images.

5.2. Discussion

The experimental results demonstrate that the proposed custom CNN provides effective performance for binary classification of chest X-ray images. The model achieved a test accuracy of 96.07%, indicating that the majority of Normal and Tuberculosis images were correctly classified. The obtained 94.29% sensitivity demonstrates that the model was able to correctly identify a high proportion of TB-positive cases, which is particularly important for a screening-oriented application.

The model also achieved a specificity of 96.43%, indicating reliable identification of Normal chest X-ray images. The precision and F1-score were 84.08% and 88.89%, respectively. In addition, the ROC-AUC value of 0.9891 indicates strong class-discrimination capability across different classification thresholds.

The training, validation, and test losses were 0.1021, 0.1284, and 0.1317, respectively. The relatively close validation and test loss values suggest that the model maintained consistent performance on unseen data. Overall, the results indicate that the custom CNN can effectively learn relevant patterns from chest X-ray images and can be considered a potential approach for automated TB screening.

6. Conclusion and Future Work

This study presented a custom Convolutional Neural Network (CNN) for the automated classification of chest X-ray images into Normal and Tuberculosis classes. The proposed approach incorporated image preprocessing and

data augmentation to improve the consistency and generalization of the model. The CNN was evaluated on an independent test set of 840 images and achieved a classification accuracy of 96.07% and a ROC-AUC of 0.9891. The model also obtained a sensitivity of 94.29% and specificity of 96.43%, demonstrating its ability to distinguish between TB and Normal chest X-ray images.

The results indicate that a custom CNN can effectively learn discriminative patterns from chest radiographs and can serve as a potential computer-aided screening approach for TB, particularly in resource-constrained settings. However, the study is limited by the relatively small and imbalanced dataset and the use of a single publicly available dataset. Further research can focus on evaluating the model using larger and multi-center datasets, improving class balance, incorporating explainable AI techniques such as Grad-CAM, and validating the model on real-world clinical data. Future work may also investigate lightweight deployment of the model for use in low-resource healthcare environments.

Author Contributions

Author-1 (Meetali Verma): Data Curation, Conceptualization, Methodology, Writing-Original Draft Preparation, Result Analysis.

Author-2 (Prof. Vinod Sharma): Validation, Formal Analysis, Supervision.

Funding: This investigation acknowledged no precise support from public, commercial, or non-profit funding organizations.

Availability of Data and Material:

<https://www.kaggle.com/datasets/tawsifurrahman/tuberculosis-is-tb-chest-xray-dataset>

Declarations

Ethics approval and consent to participate: Not applicable. This study does not involve human participants, animals, or personal data.

Consent to publish: Not applicable.

Competing interests: The authors declare that there are no conflict of interest is involved in this paper.

References

1. World Health Organization. Global tuberculosis report 2013 [Internet]. Geneva: World Health Organization; 2013 [cited 2026 Jul 2³]. Available from: https://www.google.co.in/books/edition/Global_Tuberculosis_Report_2013/1rQXDAAAQBAJ
2. Silva DO, *et al.* Impact on the ability of healthcare professionals to correctly identify patient-ventilator asynchronies of the simultaneous visualization of estimated muscle pressure curves on the ventilator display: a randomized study (Pmus study). *Crit Care*. 2023;27(1):128. doi:10.1186/S13054-023-04414-9.
3. Rana G. Technical adequacy of chest X-ray. In: *Essential radiology: a question and answer guide*. 2019. p. 29-31. doi:10.1007/978-3-030-26044-6_7.
4. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436-444. doi:10.1038/nature14539.
5. Haber E, Ruthotto L, Holtham E, Jun SH. Learning across scales—multiscale methods for convolutional neural networks. *Proc AAAI Conf Artif Intell*. 2018;32(1):3142-3148. doi:10.1609/aaai.v32i1.11680.
6. Lakhani P, Sundaram B. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology*. 2017;284(2):574-582.
7. Stirenko S, *et al.* Chest X-ray analysis of tuberculosis by deep learning with segmentation and augmentation. In: *Proceedings of the IEEE 38th International Conference on Electronics and Nanotechnology (ELNANO)*; 2018; Kyiv, Ukraine. p. 422-428.
8. Hwang EJ, *et al.* Development and validation of a deep learning-based automated detection algorithm for major thoracic diseases on chest radiographs. *JAMA Netw Open*. 2019;2(3):e191095.
9. Melendez J, *et al.* A novel multiple-instance learning-based approach to computer-aided detection of tuberculosis on chest X-rays. *IEEE J Biomed Health Inform*. 2015;19(1):142-150.
10. Liu Y, Wu Y, Ban Y, Wang H, Cheng Q. Rethinking computer-aided tuberculosis diagnosis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*; 2020. p. 2646-2655.
11. Rajaraman S, *et al.* Iteratively pruned deep learning ensembles for COVID-19 detection in chest X-rays. *IEEE Access*. 2020;8:115041-115050.
12. Xu C, Zheng F, Wang L, Qian X. Efficient deep learning model for tuberculosis detection using MobileNetV2. *Front Med (Lausanne)*. 2021;8:712327.
13. Pasa F, Golkov V, Pfeiffer F, Cremers D, Pfeiffer D. Efficient deep network architectures for fast chest X-ray tuberculosis screening and visualization. *Sci Rep*. 2019;9(1):6268. doi:10.1038/s41598-019-42557-4.
14. Yadav O, Gupta P, Kumar R. Pulmonary tuberculosis identification from chest X-rays using DenseNet with SE attention. *Comput Med Imaging Graph*. 2021;92:101953.
15. Rahman T, *et al.* Reliable tuberculosis detection using chest X-ray with deep learning, segmentation and visualization. *IEEE Access*. 2020;8:191586-191601.
16. Perez L, Wang J. The effectiveness of data augmentation in image classification using deep learning [Internet]. arXiv:1712.04621. 2017. Available from: <https://arxiv.org/abs/1712.04621>

How to Cite This Article

Asiquzzaman M, Zaman E, Kamal S, Rahim A, Mahi KS. Hemoptysis and severe thrombocytopenia in adult dengue fever: spontaneous platelet recovery with conservative management: a case report. *Int J Multidiscip Evol Res*. 2026;7(2):84-87. doi:10.54660/IJMER.2026.7.2.84-87.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.