Tuberculosis Detection Model using Machine Learning

Ayman Husain

May 2025

1 Abstract

[This project proposes the development of a supervised deep learning model to detect tuberculosis (TB) from chest X-ray images by classifying them as either TB-positive or TB-negative. Using publicly available labeled datasets, the model will be trained and evaluated for performance in real-world diagnostic contexts. The core focus is on leveraging convolutional neural networks (CNNs), such as GoogLeNet or TB-Net, and exploring interpretability techniques like Grad-CAM to visualize the model's decision-making process. The final goal is to build a diagnostic tool that is both accurate and computationally efficient, with potential for deployment in low-resource healthcare settings. Alongside model development, this project will include a critical evaluation of different architectural choices and preprocessing strategies, with attention to medical relevance and ethical considerations.

2 Introduction

Tuberculosis (TB) remains one of the deadliest diseases in the world, with the WHO reporting A total of 1.25 million people died from tuberculosis (TB) in 2023 alone. Furthermore, they have set out on a mission to end the TB epidemic by the year 2030 [1]. As a result, there is a current rise in solutions to early and accurate detection of this disease so that it can be treated quickly and/or prevent transmissions. Traditional methods for testing for TB consists popularly of sputum microscopy and culture tests, which have proven to be time-consuming and an inefficient use of resources, especially for those in less fortunate environments. In this vacuum of cheaper and more efficient solutions, TB detection models using deep learning to analyze chest X-rays, has arisen. Below, I have reviewed several pieces of literature, all using similar methodologies of building a deep convolutional neural network (DNN) to achieve this feat, while also assessing their rooms of improvement that I could possibly work on.

3 Current Approaches

It is worth discussing briefly, the present state of tuberculosis detection in terms of methods and their limitations. Commonly, early detection of pulmonary TB is done through a process of microscopic examination of sputum cultures and chest X-rays[2]. There also exist strains that are drug-resistant, requiring a drug susceptibility test (DST), as the culturing will not work on them. As a result, the processes have proven to be tedious and time-consuming, often taking a longer time for pre-diagnosis processes and fail to provide reports at a fast enough rate, which is critical to early detection. In response to this, many immunoassay techniques have been developed by the medical industry, which take less time and have a higher sensitivity. However, they are faced with the other side of the coin's issues, having a high cost and need for well-established infrastructures. While these are simplifications of the different TB detection processes that exist now, a more detailed table overview can be found in M. Singh's paper.

They go on to discuss newer methods that are being introduced to combat this issue of early TB detection, mostly focusing on different AI techniques that are still being experimented with. The foundations for these different convolutional neural networks that are suggested are classified into three different methods, being: unsupervised learning, semi-supervised learning and supervised learning. For this project, I intend to create a deep learning CNN model trained by supervised learning. This means that I will be feeding my model labelled data as input and outputs to find common patterns, as opposed to having one or both sets of unlabeled data like in the other methods. The author also mentions the use of transfer learning,that is for when I am trying to transfer knowledge bases between different models. Although, I do not believe I will be using this model unless I want to use a pre-made CNN like GoogLe Net, and just transfer my databases to compare it with the model I will be building.

In most cases, they have applied different versions or models of CNN to detect TB using CXR images. Wong et al. [3] proposes a deep CNN model with a self-attention mechanism called TB-Net. This model showed much better performance compared to standard CNNs, showing results in accuracy, sensitivity and specificity of 99.86%, 100.0% and 99.71%. being particularly useful due to its ability to capture spatial dependencies in medical images, making it more efficient at classifying the test result as negative or positive for TB. However, it is worth mentioning that this model is almost too powerful or power hungry, in that it is not suitable for resource-constrained settings. With all that being said, my project is heavily influenced by this research model, due to its high figures in results and due to its self-attention layer which will hopefully allow it to better locate the tuberculosis affected areas and learn more efficiently from the practice datasets.

Sharma et al. [4] focused on integrating deep learning with visualization techniques such as Grad-CAM to highlight key regions in CXR images contributing to TB classification. This method enhances model interpretability, which is critical for medical applications where decision transparency is essential. Another significant approach is presented by Jaeger et al. [5], who employed a UNet-based segmentation model for TB detection. While UNet is effective for biomedical image segmentation, its computational complexity and reliance on large training datasets pose challenges for real-world implementation. On the other hand, Sharma aimed their paper at being more about visualization techniques relating to deep learning. For example, the most important technique they refer to is Grad-CAM to highlight key regions in CXR images, allowing it to better classify an input as TB positive or negative. However, a significant downside of this approach is that it relies on high-resolution images to be doing this, which again, may not be feasible in an underprivileged setting.

A commonality between all the papers is their emphasis on preprocessing techniques, but there is one paper that puts it on a pedestal more than the others, being Norval[6]. They mention methods such as lung ROI extraction and contrast enhancement that basically isolate the different regions of the lung from the X-ray and then improve the image to improve visibility of key features or points of interest. They also introduce a unique hybrid method which combines traditional computer-aided detection (CAD) with deep learning and has shown a model accuracy of 92.54%. It does however, add a whole new level of complexity to the model by adding this much pre-processing and could slow the system down.

4 Preliminary Design

The system I am proposing will follow a supervised learning approach, making use of labeled chest X-ray (CXR) datasets where each image is annotated as TB positive or TB negative. The workflow will begin with data preprocessing, including resizing, normalization, and optional region-of-interest (ROI) extraction or lung segmentation to isolate key features. This step will ensure that input data is standardized and optimized for training. The model architecture will be centered on convolutional neural networks. In parallel, the project will experiment with a TB-Net-inspired model that incorporates self-attention layers, which may offer improved localization of pathological features.

These models will be trained and evaluated using PyTorch or TensorFlow, depending on compatibility with available resources and pretrained weights. Visualization and interpretability will be incorporated through tools like Grad-CAM, enabling a visual overlay of attention maps that show the regions influencing the model's decisions. This component is essential for validating the model's clinical relevance and transparency. Should performance bottlenecks or generalization issues arise, fallback plans include simplifying the model architecture or using image augmentation to diversify the dataset.

5 Preliminary Evaluation Plan

To assess the model's success, a combination of quantitative and qualitative evaluation methods will be employed. The quantitative metrics will include accuracy, precision, recall, F1-score, and ROC-AUC, providing a comprehensive view of performance across different decision thresholds. A confusion matrix will also be used to understand the balance of false positives and false negatives. These metrics will be calculated on a held-out test set using cross-validation to mitigate overfitting.

Qualitative evaluation will focus on the interpretability of the model's predictions. Grad-CAM visualizations will be generated to ensure the model is focusing on clinically relevant regions of the lungs when making predictions. Experiments will be conducted to compare the performance of the GoogLeNet-based model against a more domain-specific TB-Net architecture and a custom CNN model. Additional evaluation will include ablation studies to analyze the impact of preprocessing steps such as ROI extraction and contrast enhancement.

6 Risk Analysis

Several technical and practical risks have been identified. One key risk is label quality within the dataset, as mislabeled or inconsistent images can misguide the training process. This will be mitigated through manual inspection and by using class weighting or data cleaning techniques. Another risk is limited generalizability due to dataset homogeneity; this will be addressed through image augmentation and cross-validation. Overfitting is a common concern in medical imaging and will be handled using regularization techniques such as dropout, early stopping, and reduced model complexity.

There is also a risk that Grad-CAM or other interpretability tools may not produce informative or medically meaningful visualizations. To address this, the project will prepare alternative methods such as integrated gradients or LIME. Finally, hardware availability may be a constraint. In the event that access to a GPU is limited, training will be staggered over time or performed using cloud platforms like Google Colab Pro, which supports GPU acceleration.

7 Conclusion

Deep learning has shown immense potential for TB detection using X-ray images, with CNN-based architectures demonstrating high accuracy at a cheaper cost to both the users and the moderators of the programs. The integration of self-attention mechanisms, preprocessing techniques, and visualization methods enhances performance and interpretability. However, challenges such as data availability, model transparency, and generalizability need to be addressed. Future research should focus on improving model explainability, developing lightweight architectures for real-time applications, and ensuring cross-population generalization. I hope to make the necessary changes to my program by learning from these previously explored methods and come out with an early TB detection model that can build on the work of my predecessors.

]

8 Timeline of Events

Week — Milestone / Deliverable

Week 1 — Finalize dataset sources and begin data exploration and cleaning
Week 2 — Baseline model training (simple CNN) and validation split setup
Week 3 — Implement and train core model (e.g., GoogLeNet or TB-Net)
Week 4 — Run performance evaluations and log baseline results
Week 5 — Integrate Grad-CAM visualizations and begin writing analysis
Week 6 — Early Semester Break (no major deliverables – light debugging or writing)
Week 7 — Mid-project presentation and feedback integration
Week 8 — Begin drafting final paper and evaluate alternative architectures
Week 9 — Conduct interpretability experiments, finalize visuals
Week 10 — Submit full first draft of final report for review
Week 11 — Final debugging, polishing visualizations, finalize all figures
Week 12 — Final presentation preparation and optional auxiliary features
Week 13 — Submit final paper and present to department/class
Week 14 — Thanksgiving Break (buffer time)
Week 15 — Final archive, documentation, and project reflections

References

[1] World Health Organization. Tuberculosis fact sheet, n.d.

[2] M. Singh, G. V. Pujar, S. A. Kumar, M. Bhagyalalitha, H. S. Akshatha, B. Abuhaija, A. R. Alsoud, L. Abualigah, N. M. Beeraka, and A. H. Gandomi. Evolution of machine learning in tuberculosis diagnosis: A review of deep learning-based medical applications. *Electronics*, 11(17):2634, 2022.

[3] A. Wong, J. R. H. Lee, H. Rahmat-Khah, A. Sabri, and A. Alaref. Tb-net: A tailored, self-attention deep convolutional neural network design for detection of tuberculosis cases from chest x-ray images. arXiv preprint arXiv:2104.03165, 2021.

[4] V. Sharma, S. Nillmani, S. K. Gupta, and K. K. Shukla. Deep learning models for tuberculosis detection and infected region visualization in chest x-ray images. *Intelligent Medicine*, 4(2):104–113, 2024.

[5] S. Jaeger, A. Karargyris, S. Candemir, L. Folio, J. Siegelman, F. Callaghan, X. Zhiyun, K. Palaniappan, R. K. Singh, S. Antani, G. Thoma, W. Yi-Xiang, L. Pu-Xuan, and C. J. McDonald. Automatic tuberculosis screening using chest radiographs. *IEEE Transactions on Medical Imaging*, 33(2):233–245, 2014.

[6] M. Norval, Z. Wang, and Y. Sun. Pulmonary tuberculosis detection using deep learning convolutional neural networks. In *Proceedings of the 2019 3rd International Conference on Video and Image Processing, ICVIP 2019*, ACM International Conference Proceeding Series, pages 47–51. Association for Computing Machinery, 2019.

[7] C. J. Liu, C. C. Tsai, L. C. Kuo, et al. A deep learning model using chest x-ray for identifying tb and ntm-ld patients: A cross-sectional study. *Insights into Imaging*, 14:67, 2023.

[8] S. Kant and M. M. Srivastava. Towards automated tuberculosis detection using deep learning. In *Conference Paper*, n.d.

[9] S. Hansun, A. Argha, S. T. Liaw, B. G. Celler, and G. B. Marks. Machine and deep learning for tuberculosis detection on chest x-rays: Systematic literature review. *Journal of Medical Internet Research*, 25:e43154, 2023.

[10] L. Lin-Sheng, L. Yang, L. Zhuang, Y. Zhao-Yang, Z. Wei-Guo, and W. P. Gong. From immunology to artificial intelligence: Revolutionizing latent tuberculosis infection diagnosis with machine learning. *Military Medical Research*, 10:1–37, 2023.

[11] Kaggle. Chest x-ray lungs segmentation, n.d.

[12] Kaggle. Tuberculosis (tb) chest x-ray dataset, n.d.