

Efficient Tuberculosis Diagnosis Using Deep Learning based Object Detection Method

Pranjali Kasture^{1*}, Neha Kapadia², Pratibha Prasad³, Kajal Patel⁴

^{1,2,3,4} Assistant professor, Department of Information Technology, Thakur College of Engineering and Technology, Mumbai, India

* Corresponding author: kasturepranjali@gmail.com

ARTICLE INFO	ABSTRACT
Received: 26 Dec 2024	Tuberculosis is a lung disease that poses a threat to life leading to causes of death so early and precise identification is important. The purpose of this work is to progress the field of tuberculosis (TB) detection by implementing deep learning techniques. This entails developing a model that can handle a variety of instances, including differences in ill lungs and the presence of tuberculosis bacteria. The work offers a complete methodology for tuberculosis detection based on deep learning. A diversified chest X-ray image dataset is collected and pre-processed. The proposed model, a combination of ResNet50 and Fast R-CNN (region-based convolutional neural network), creates a powerful synergy that provides a reliable method for precisely detecting and localizing objects within images. By utilizing ResNet50 as a backbone, Fast R-CNN benefits from the rich feature maps generated by the deep network. It facilitates precise region proposal generation and improves the overall accuracy of object localization. Alternative object detection algorithms, such as RetinaNet and SSD, are compared with the proposed model. We achieved a notable accuracy of 96.7% by using the Fast R-CNN approach with a ResNet50 backbone. This accomplishment highlights the effectiveness of our selected strategy by outperforming other algorithms.
Revised: 12 Feb 2025	
Accepted: 22 Feb 2025	
Keywords: Tuberculosis, Deep learning, CNN, Backbone, Object detection	

INTRODUCTION

Mycobacterium tuberculosis (MTB) is an infectious disease that poses a threat to life, with a 25% infection rate estimated. This is the most common infectious disease-related cause of mortality [1]. Current advancements have introduced multiple diagnostic techniques that rely on molecular analysis and bacterial culturing. As per a World Health Organization (WHO) report, it was estimated that 9.9 million individuals would experience the impact of this disease in the year 2020 [2]. Significantly, tuberculosis CXR images are frequently mistaken for other diseases with comparable radiological patterns [3] potentially resulting in incorrect treatment for patients and consequently exacerbating their health condition. Lately, Deep learning has demonstrated impressive results in picture recognition and categorization using the convolutional neural network (CNNs) model [4]. This work focuses on how the preprocessing stage affects the efficacy of the deep learning technique and it is not significantly addressed in the previous research.

The availability of extensive, well-labeled CNNs with deep convolutional layers and datasets has resulted in significant achievements in image identification. CNNs enable the acquisition of extremely descriptive, data-driven, features of a hierarchical image through an adequate training set. In addition, acquiring medical imaging datasets with annotations as comprehensive as those within ImageNet remains a challenging task. As per Koo and Cha [5], Pre-processing of images is essential to the training process of CNN models, significantly improving their classification performance. Image enhancement is a vital component of this pre-processing. Therefore, investigating the connection between image enhancement and the CNN model holds great significance. Deep learning algorithms for machine learning have been utilized within tuberculosis identifying through adjusting the deep-layered CNN parameters in a range of studies [6]-[10]. Additionally, the concept of transferable knowledge within the framework was used to apply deep learning to detect tuberculosis by employing pre-trained models and combining them into ensembles in other studies. Hooda and companions [11] introduced a Deep learning method for categorizing CXR pictures into groups for TB and non-TB, achieving a precision of 82.09%. In a separate study, Evangelista, as well as Guedes, presented an automated system that relied on excellent pattern recognition through CNNs to detect TB with an accuracy of 88.76%. In clinical settings, enhancing the precision of tuberculosis detection from chest X-rays with a strong and adaptable approach may enhance the dependability of computer-aided diagnostic tools. To improve

classification accuracy, one can explore different deep learning algorithms, adapt currently available, efficient algorithms, or combine several top-executing algorithms into an ensemble model. However, X-ray images encompass not only the lungs but also other areas of the thorax, even though diseases like TB primarily affect the lung region. Therefore, concentrating on the lung region when training and classifying X-ray images can significantly enhance TB detection performance. As far as our information indicates, no work has been done on the use of deep-learning networks to segment lung sections for tuberculosis diagnosis. This paper's primary focus is on TB detection, employing transfer knowledge-based CNN techniques in both unaltered and partitioned lung regions within X-ray images.

LITERATURE REVIEW

Current literature on TB detection through X-ray image analysis enhances pre-trained models and considers ensemble models to improve detection performance. It integrates conventional radiologic assessments and transfer learning. Transparency is continuously achieved by using visualization approaches, and common research foci include consideration of availability of data, cost, and global access. The overall goal of this body of work is to improve patient outcomes and support public health efforts by providing a more accurate, inexpensive, and easily assessed method for TB detection.

Lakhani et al., 2017 demonstrated the potential of Convolutional Neural Networks (CNNs) for accurate detection of TB (Tuberculosis), specifically noting the requirement for large, diverse datasets to train a model [12]. Jaeger et al. in 2014 explored the application of CNNs for automatic identification of TB, showing positive results in identifying TB indicators and highlighted the scalability of a TB screening solution [13]. Hwang et al., 2016 demonstrated the effectiveness of CNNs in accurately discerning TB from chest X-ray images highlighted the identification of features to support the diagnosis of TB from Chest X-ray images and the potential for deep learning models to assist radiologists in creating a diagnosis [14].

A evaluation of different CNN designs and their suitability for medical image analysis, including TB diagnosis, was presented by Shin et al. in 2016. The significance of transfer learning and the essential requirement for diverse datasets to ensure resilient model training were discussed [15]. Hosny et al., in 2018 emphasized the application of computer-aided detection (CAD) systems powered by deep learning for large-scale validation in TB detection. The study demonstrated the potential of these systems in assisting radiologists with accurate TB diagnosis from chest X-rays [16]. Lopes et al., in 2017 highlighted the use of a deep CNN for TB detection, achieving high accuracy rates comparable to expert radiologists. The authors discussed deep learning in TB diagnosis, addressing the limitations of traditional detection methods [17][18].

METHODOLOGY

DATASET DESCRIPTION

The dataset used for the proposed research is Tuberculosis X-ray (TBX11K) dataset. It is a comprehensive collection of 1,200 chest X-ray (CXR) images. Figure 1 exhibits this dataset which was created to fill the gap in training data for computer-aided tuberculosis diagnosis (CTD). With this dataset, it becomes possible to train advanced detectors for more accuracy.



Figure 1 (a). A normal chest X-ray



Figure 1 (b). A chest X-ray exhibiting tuberculosis-related diseases

The TBX11K dataset outperforms other tuberculosis datasets in several notable aspects, making it conducive for training deep Convolutional Neural Networks (CNNs). Table 1 lists the other tuberculosis datasets with the number of samples. Firstly, unlike prior datasets with only a limited number of X-ray images, TBX11K boasts a substantial compilation of 11,200 images. This size advantage enables the training of exceptionally deep CNNs. Table 1 provides an overview of some datasets of chest X-rays used for the classification of TB.

Table 1: List of Tuberculosis Datasets with the number of samples

Dataset	Class and number of images	Total Sample
Montgomery County (MC)	NORMAL:80 ABNORMAL:58	138
Shenzhen	NORMAL:326 TB:336	662
DA	TB:78 Non-TB:78	156
DB	TB:75 Non-TB:75	150
TBX11k	TB:1200 Non-TB:10000	11,200

DATA PREPROCESSING

The Convolutional Neural Network (CNN) takes image-based input, allowing for the encoding of specific properties within the network. In our study, as we deal with grayscale images, the pixel array is structured as height $\times$ width $\times$ depth, such as $224 \times 224 \times 1$. To extract valuable features from the image data, we adjusted certain parameters in the Conv layers during training. Simultaneously, the acquired knowledge of parameters within the context of Fully Connected (FC) layers is responsible for categorizing the extracted features into target classes, distinguishing between normal and TB images. The Conv layers progressively capture visual features from the raw input images hierarchically. The images are represented as vectors, with binary labels indicating '1' for TB images and '0' for other images in our case.

The experiment involved the utilization of transfer learning with the ResNet-50 models. ResNet is a neural network that incorporates skip connections to bypass multiple layers. This leap over layers serves the purpose of mitigating the vanishing gradient problem by reusing activations from the previous layer before the adjacent layer has its weights initialized. ResNet comprises multiple layers, with each layer being frequently replicated.

MODEL FRAMEWORK

Diagnosing tuberculosis (TB) can be difficult because there are so many different signs of TB that could present on medical imaging, especially chest X-rays (CXR). There are many signs, including nodules, large opacities, conglomerates, localized lesions, cavities, and sometimes microscopic opacities, that complicate a straightforward and consistent diagnosis. Conventional diagnostic techniques are hindered by these symptoms, and deep learning models even face difficulties. It takes large and varied datasets that cover these variances to properly train CNNs to identify these many TB symptoms. To create robust models, a broad range of manifestations must be presented to the system. This enables the model to learn and distinguish between patterns connected to tuberculosis and those resulting from other illnesses or anomalies. The intricacy of tuberculosis (TB) detection in medical imaging calls for the development of deep learning models, coupled with the presence of varied and extensive datasets, ongoing algorithmic refinement, and cooperation between medical professionals and AI researchers to increase the precision and dependability of TB diagnosis using CXR images. They used CXR pictures to build a single system for tuberculosis detection. Their system is made up of numerous modules that must follow certain instructions to classify the input image. An extraction function comes after a pre-processing module based on denoising the dataset.

CNN ARCHITECTURE

CNNs are popularly used for a lot of computer vision problems, especially classification, detection, and recognition problems. Usually, a CNN model comprises three layers: a convolutional layer, a pooling layer, and a fully connected (FC) layer. Figure 2 represents the architecture of the CNN model. Each layer is connected to the one before it using kernels that have a predetermined, fixed-size receptive field. To capture global or local image properties, the CNN model learns to configure its hyperparameters from an extensive dataset. Compared to software that was created by humans, the model architecture's various layer types and activation functions demonstrate superior representational features. This adaptability is followed in the model's training and fine-tuning. After being adjusted for tuberculosis detection, the model is trained on a ready-made dataset.

To identify and distinguish TB-related manifestations shown in the chest X-ray pictures, the pre-trained CNN's weights and parameters are changed. To enable the model to learn from the training data while guaranteeing its

generalizability through performance validation on distinct data, typically, the dataset is partitioned into training and validation sets. To maximize the model's capacity to precisely identify TB-related patterns, training, validation, and adjustment cycles are used to continuously improve the model's performance. After training and validation, the model may be used to assess fresh chest X-ray pictures, providing predictions and insights into possible tuberculosis infections.




Figure 2. Architecture of CNN

BACKBONE ARCHITECTURE

The choice of a suitable backbone architecture plays a crucial task in influencing the performance of a model. The utilization of ResNet50 as backbone architecture can be seen in Figure 3. It brings distinctive strengths to the method. ResNet50 is a very effective convolutional neural network (CNN) architecture that has become a foundation for computer vision tasks. The architecture's major best feature is the introduction of residual learning. This innovation uses skip connections to mitigate the vanishing gradient problem when training deep networks. The layers in ResNet50 add up to 50 layers, and ResNet50 is well-known for its depth and ability to learn very complicated features in images.

The use of residual connections enables the model to learn residual mappings. This is significant in developing the capability of the model to learn and represent complex patterns and hierarchical features in the input data. Additionally, the ResNet50 structure makes it easy to reuse features at different levels through skip connections, allowing for information to pass through to different layers directly. This helps reduce information loss and helps learn more abstract representations efficiently, which makes it ideal for exploring tasks. When used as a backbone architecture in a deep learning model, ResNet50 provides a strong base for meaningfully extracting features from images. It combines its depth, residual connections, and feature reuse to optimize performance in distinguishing and classifying deep representations of visual data. In images or more complex computer vision applications, the capabilities of ResNet50 make it an optimal addition to improving the effectiveness of neural networks overall.




Figure 3. ResNet50 architecture.

OBJECT DETECTION METHOD

Choosing Fast R-CNN for tuberculosis chest X-ray classification is an appropriate approach for capitalizing on its fast object detectors to solve a problem stemming from a set of complications associated with medical imaging. Fast R-CNN is specifically designed for object localization and accurately detects and locates various forms of anomalies (e.g., nodules or opacities) in chest X-images (e.g., opacities) as an integral part of the diagnostic process for TB. By using a backbone network like ResNet50, the model is capable of extracting hierarchical and semantically rich features and can spatially capture subtle patterns of manifestations of TB, which are important for the classification system. The framework's ability to accept different images sizes, combine with pre-trained models, and use multiclass classification is another reason it was considered viable in providing accurate and detailed diagnostic information. Overall, Fast R-CNN is a useful model that meets the identification and classification requirements of tuberculosis chest X-ray classification.

The combination of ResNet50 and Fast R-CNN (Region-based Convolutional Neural Network) produces a powerful combination in object detection, offering an effective way to correctly classify and locate objects in an image. We will discuss how the addition of ResNet50 improves the efficiency of Fast R-CNN. ResNet50's deep architecture and effective feature extraction makes it an ideal backbone network for Fast R-CNN. In the context of object detection, the goal is not only to classify objects, but also to accurately locate objects within an image. The depth and feature capturing capabilities provided by ResNet50 is advantageous in this objective. The hierarchical representations learned through ResNet50 allow Fast R-CNN to identify complex patterns and contextual information to be more accurate and contextually aware in object detection.

Fast R-CNN, as a region-based approach, operates in two stages: generation of region proposals and detection of objects. In the initial phase, a region proposal network (RPN) is utilized to produce probable candidate regions to encompass objects. By utilizing ResNet50 as a backbone, Fast R-CNN benefits from the rich feature maps generated by the deep network, facilitating precise region proposal generation and improving the overall accuracy of object

localization. It can be briefly seen in Figure 4.




Figure 4. Methodological framework

The combination of ResNet50 and Fast R-CNN overcomes the problem of feature extraction and provides strong support for detecting objects in complex scenes. ResNet50's depth and hierarchical feature representations help the model capture the subtleties of visual information for more reliable and accurate object detection results. This combination is a sophisticated solution for object detection, demonstrating the interaction between advanced backbone architectures and new state-of-the-art detection frameworks.

Other object detection algorithms, including RetinaNet, and SSD (Single Shot Multibox Detector), offer options, and both algorithms have unique benefits. RetinaNet is characterized by its fast-RCNN identification technique known as focal loss and a feature pyramid network. The focal loss separates example locations (background and abnormalities) by identifying the location of the abnormality within the dataset. In medical imaging where there are limited samples of positive examples, it aims to handle dealing with these dataset imbalances. In addition, the feature pyramid network enhances multi-scale and multi-sizability detection, which is important when detecting various sizes of lesions in X-ray images of tuberculosis, as the detection algorithm may encounter different lesions due to the disease. In that respect, SSD also has great potential in object detection. More specifically, SSD conducts detection of abnormalities by using features captured from various scales and uses predefined aspect ratios to produce bounding boxes around detections. This multi-scale mapping model, along with the use of multi-ids for lesions related to tuberculosis in chest X-rays, can make the SSD applicable. In terms of models selected for lesion detection, Algorithm 1 depicts the stepwise approach of pre-trained ResNet-50, followed by fine-tuning and lastly employing the applied model.

Algorithm 1 for Object Detection using Fast R-CNN

Input: Image dataset containing objects to be detected

Output: Detected objects

1. **procedure** OBJECT_DETECTION_FAST_RCNN
2. Load pre-trained ResNet-50 model as the backbone of Fast R-CNN
 - a. Download pre-trained ResNet-50 weights
 - b. Initialize Fast R-CNN with ResNet-50 backbone
3. Fine-tune the Fast R-CNN on the specific object detection dataset
 - a. Replace the classification head with a Region Proposal Network (RPN)
 - b. Train the model to predict class labels
4. Apply the trained Fast R-CNN on the input images
 - a. For each image, use the RPN to propose regions likely to contain objects.

b. Classify the refined regions to obtain the final object detection results.

5. Output the detected objects

6. **end procedure**

MATHEMATICAL FORMULATION

When using a Convolutional Neural Network, an image is fed through a series of convolutional layers. These layers use biases b_l and filters designated as W_l . By applying these filters to the input image, the convolutional process helps the network learn and extract hierarchical features, which improves the network's overall comprehension and representation of the input data.

$$F_l(x) = \max(0, W_l * x + b_l) \quad (1)$$

where $*$ denotes convolution and $\max(0, x)$ is the rectified linear unit (ReLU) activation function.

As a result, feature map F is created. To generate important outputs, the Region Proposal Network (RPN) processes the feature map F through a series of tiny convolutional layers. Among these outputs are the object scores, which determine the likelihood that an object will be present at every point on the feature map. These scores, which provide a quantitative estimate of the probability that an object will be present at different positions in the feature map, are produced by applying a sigmoid function. By suggesting candidate regions for additional examination, the RPN aids in the succeeding steps of the object detection process.

$$p_{obj}(x) = \text{sigmoid}(W_{obj} * F(x) + b_{obj}) \quad (2)$$

where $p_{obj}(x)$ is the output probability indicating the presence of an object at position (x) , $\text{sigmoid}(\cdot)$ is the sigmoid activation function, which normalizes the input between 0 and 1, W_{obj} is the weight matrix applied to the features extracted by $F(x)$, $F(x)$ is the output from a preceding layer, capturing spatial relationships through convolution, $*$ The convolution operation, represents the spatial filtering of the input, b_{obj} is the bias term added to the linear transformation of the convolutional output, (x) is the position in the input space where the object probability is evaluated.

Bounding box offsets allow these predict adjustments to pre-defined boxes using linear regression.

$$t_x, t_y, t_w, t_h = W_{reg} * F(x) + b_{reg} \quad (3)$$

Here, t_x indicates the predicted shift in the x-direction (horizontal), t_y indicates the predicted shift in the y-direction (vertical), t_w indicates the predicted log of the width scale factor, t_h indicates the predicted log of the height scale factor, W_{reg} indicates the weight matrix applied to the features obtained from $F(x)$ which is the output a previous hidden layer that learns spatial relationships via convolutional operation, and b_{reg} indicates the bias addition to the linear transformed convolution output.

Region of Interest (ROI) Pooling is a critical step in object detection when using convolutional neural networks. This involves extracting fixed-size feature vectors from the proposed regions (ROIs) in the feature map F . In order to form consistent representations of the ROIs (regardless of their original shape), a strategy such as max-pooling is implemented. By using these pooling techniques on the ROIs, the network makes sure that the extracted features stay the same size, which makes it easier for the network to identify and categorize objects in these designated regions of interest and makes it easier for the next steps of the object identification process to proceed.

The collected feature vectors are passed through fully connected layers in the classification and bounding box regression process when it comes to object detection. These layers forecast the likelihood that every Region of Interest (ROI) will belong to a particular object class in the categorization aspect. A SoftMax function is used to normalize the output probabilities and provide a probability distribution over the different object classes, enabling this prediction. Bounding Box Regression optimizes the localization accuracy of the recognized items by simultaneously refining the predicted bounding box coordinates for each ROI. When combined, these procedures aid in the thorough comprehension and accurate location of items in the input image.

$$p(c_i|x) = \text{softmax}(W_c * f(x) + b_c) \quad (4)$$

where f is the extracted feature vector and c_i is the class label.

Similar to the Region Proposal Network (RPN), bounding box regression uses linear regression techniques to forecast modifications or enhancements to the bounding boxes. The goal of this procedure is to improve the bounding box predictions made in the initial phases of object detection. The model adjusts the bounding box's dimensions and coordinates by using linear regression, guaranteeing a more accurate alignment with the real limits of the items it has spotted. Bounding box regression, in the end, refines the spatial information offered by the original region proposals and adds to the overall improvement in the localization precision of items inside the image.

RESULTS & DISCUSSION

The research outcomes unveil the effectiveness of cutting-edge object detection algorithms, each leveraging distinct backbone architectures. In employing ResNet-50 as the backbone for Fast R-CNN, RetinaNet, and SSD, the study delves into the nuanced performance of these models through a meticulous evaluation of key metrics. The performance metrics include Accuracy, Area Under the Curve (AUC), Average Precision (Ave. Prec.), and Average Recall (Ave. Rec.).

Table 2. Performance comparison of object detection methods on tuberculosis chest X-ray classification using ResNet-50 as the backbone

Method	Backbone	Accuracy	AUC	Ave. Prec.	Ave. Rec.
RetinaNet	ResNet-50	90.6	90.2	84.8	86.8
SSD	ResNet-50	90.3	91.8	86.6	89.2
Fast R-CNN	ResNet-50	96.7	93.6	87.7	90.5

In Table 2 presented below, the Fast R-CNN method utilized ResNet-50 Backbone is obtaining the best results with an accuracy of 96.7%, AUC of 93.6%, average precision of 87.7% and average recall of 90.5%. Overall, with respect to the Fast R-CNN method, in terms of accuracy RetinaNet shows a difference of 6.1%, difference of AUC of 3.4%, average precision difference of 2.9% and average recall difference of 3.7%. A similar comparison is made for the SSD method: Accuracy difference of 6.4%, an AUC difference of 1.8%, average precision difference of 1.1% and average recall difference of 1.3%. TABLE 1(RESNET -50)




Figure 5. Key Results for Backbone ResNet-50

Results are also visualized in Figures 5 and 6 interpreting better performance of Fast R-CNN than RetinaNet and SSD.




Figure 6. AUC For Backbone ResNet-50

Table 3 provides an overview of the performance metrics for three object detection methods employing the VGG16 backbone. As shown in the following Table 3, the Fast R-CNN method using VGG16 Backbone achieves the highest results, producing an accuracy of 93.6%, a score of 93.2% for AUC, an average precision of 87.7%, and an average recall of 90.5%. When comparing these results to the RetinaNet method, there is a difference in accuracy of 4.2%, AUC difference of 4%, average precision difference of 6.9%, and average recall difference of 6.3%. The SSD method produced an accuracy difference of 3%, AUC difference of 1.1%, average precision difference of 2.6%, and an average recall difference of 0.1%.

Table 3. Performance comparison of object detection methods on tuberculosis chest X-ray classification using VGG16 as the backbone

Method	Backbone	Accuracy	AUC	Ave. Prec.	Ave. Rec.
RetinaNet	VGG16	89.4	89.2	80.8	84.2
SSD	VGG16	90.6	92.1	85.1	90.4
Fast R-CNN	VGG16	93.6	93.2	87.7	90.5




Figure 7. Key Results for Backbone VGG16

Results are also visualized in Figures 7 and 8 interpreting better performance of Fast R-CNN than RetinaNet and SSD.




Figure 8. AUC For Backbone VGG16

CONCLUSION

In the battle against tuberculosis, the application of deep learning algorithms for TB detection has greatly increased the precision and effectiveness of diagnostic procedures, providing a glimmer of hope. With the help of a robust synergy between ResNet50 and Fast R-CNN (Region-based Convolutional Neural Network), the suggested model offers a dependable way to accurately recognize and localize objects within images. With ResNet50 serving as its backbone, Fast R-CNN leverages the rich feature maps produced by the deep network to produce precise region proposals and enhance object localization accuracy overall. The suggested approach is contrasted with alternative object detection algorithms like SSD and RetinaNet. We utilized a ResNet50 backbone and the Fast R-CNN technique to obtain an astounding accuracy of 96.7%. This achievement, which outperforms competing algorithms, demonstrates the efficacy of our chosen approach. As AI advances, we may expect more precise and real-time object identification capabilities, which will be addressed in subsequent developments.

REFERENCES

- [1] Catrin Sohrabi, Zaid Alsafi, Niamh O'Neill, Mehdi Khan, Ahmed Kerwan, Ahmed Al-Jabir, Christos Iosifidis, Riaz Agha, World Health Organization declares global emergency: A review of the 2019 novel coronavirus (COVID-19), *International Journal of Surgery*, 2020
- [2] Chih-Cheng Lai, Tzu-Ping Shih, Wen-Chien Ko, Hung-Jen Tang, Po-Ren Hsueh, Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) and coronavirus disease-2019 (COVID-19): The epidemic and the challenges, *International Journal of Antimicrobial Agents*, 2020
- [3] M. van Cleeff, L. Kivihya-Ndugga, H. Meme, J. Odhiambo, and P. Klatser, "The role and performance of chest X-ray for the diagnosis of tuberculosis: A cost-effectiveness analysis in Nairobi, Kenya," *BMC Infectious Diseases*, vol. 5, no. 1, p. 111, Dec. 2005.
- [4] S. Graham, K. D. Gupta, J. R. Hidvegi, R. Hanson, J. Kosiuk, K. A. Zahrani, and D. Menzies, "Chest radiograph abnormalities associated with tuberculosis: Reproducibility and yield of active cases," *Int. J. Tuberculosis Lung Disease*, vol. 6, no. 2, pp. 137–142, 2002.
- [5] K.-M. Koo and E.-Y. Cha, "Image recognition performance enhancements using image normalization," *Hum-centric Comput. Inf. Sci.*, vol. 7, no. 1, pp. 1–11, Nov. 2017
- [6] L. G. C. Evalgelista and E. B. Guedes, "Computer-aided tuberculosis detection from chest X-ray images with convolutional neural networks," in *Proc. Anais do XV Encontro Nacional de Inteligência Artif. e Computacional (ENIAC)*, Oct. 2018, pp. 518–527.
- [7] F. Pasa, V. Golkov, F. Pfeiffer, D. Cremers, and D. Pfeiffer, "Efficient deep network architectures for fast chest X-ray tuberculosis screening and visualization," *Sci. Rep.*, vol. 9, no. 1, pp. 1–9, Dec. 2019.
- [8] Q. H. Nguyen, B. P. Nguyen, S. D. Dao, B. Unnikrishnan, R. Dhingra, S. R. Ravichandran, S. Satpathy, P. N. Raja, and M. C. H. Chua, "Deep learning models for tuberculosis detection from chest X-ray images," in *Proc. 26th Int. Conf. Telecommun. (ICT)*, Apr. 2019, pp. 381–386.
- [9] Hernández, Á. Panizo, and D. Camacho, "An ensemble algorithm based on deep learning for tuberculosis

- classification,” in Proc. Int. Conf. Intell. Data Eng. Automated Learn., Manchester, U.K., 2019, pp. 145–154.
- [10] Rohilla, R. Hooda, and A. Mittal, “Tb detection in chest radiograph using deep learning architecture,” in Proc. ICETETSM, Aug. 2017, vol. 6, no. 8, pp. 1073–1085
- [11] R. Hooda, S. Sofat, S. Kaur, A. Mittal, and F. Meriaudeau, “Deep-learning: A potential method for tuberculosis detection using chest radiography,” in Proc. IEEE Int. Conf. Signal Image Process. Appl. (ICSIPA), Kuching, Malaysia, Sep. 2017, pp. 497–502.
- [12] Lakhani and B, Subdram. "Deep learning and its application to chest radiograph interpretation: A review of the literature." *Journal of Thoracic Imaging*, 32(2), 2017.
- [13] S. Jaeger, S. Antani, S. Candemir "Two-tiered unsupervised local alignment for lung segmentation in chest radiographs." *IEEE Transactions on Medical Imaging*, 33(2), 2014.
- [14] S. Hwang, J. Jeong and H.J Kim "Deep learning for lung cancer detection: Tackling the Kaggle Data Science Bowl 2017 challenge." *Journal of Digital Imaging*, 29(4), 2016.
- [15] Shin, H.E Kim "Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics, and transfer learning." *IEEE Transactions on Medical Imaging*, 35(5), 2016.
- [16] Hosny, A. Brock "Deep learning for lung cancer prognostication: A retrospective multi-cohort radiomics study." *PLoS Medicine*, 15(11), 2018.
- [17] U. Lopes and J.F Valiati "A deep convolutional neural network approach for detection of tuberculosis from chest X-ray images." *Computers in Biology and Medicine*, 88, 2017.
- [18] Kasture, P., & Shirsath, K. (2023). Recent Studies on Deep Learning Techniques for Stock Market Forecasting. *2023 International Conference on Sustainable Computing and Smart Systems (ICSCSS)*, 57–62. <https://doi.org/10.1109/ICSCSS57650.2023.10169585>