RetinaNet Localization of Radiographically Suspicious Pulmonary Tuberculosis on Chest X-Ray: A Transfer Learning Approach

Jack W. Luo, MD, McGill University; Chang Heng Mo, DCS; Tiffany Fung, MD; Jason Chong, MD, MHI, FRCPC

Introduction

Tuberculosis is a leading cause of death worldwide and despite improved screening methods and treatments, its burden remains a significant Canadian and global issue. Limited resources available to interpret radiographs pose key barriers to timely diagnosis and isolation. The integration of artificial intelligence in the diagnostic process can reduce the burden of interpretation and improve screening efficacy.

We propose a RetinaNet based convolutional neural network (CNN) pipeline to automatically detect, localize and sub-classify suspected tuberculosis on chest radiographs. Given the importance of timely detection and isolation of active TB cases, such an approach could reduce the time to interpretation and warn both clinicians and radiologists to this critical finding.

Hypothesis

A RetinaNet based CNN model trained on chest radiographs suspect for tuberculosis should be able to adequately localize and classify TB under its three main pulmonary presentations (cavitary, infiltrate and miliary). In addition, transfer learning on a prior pneumonia dataset, the *RSNA Pneumonia Detection Challenge*, should greatly increase model performance, in particular for pulmonary infiltrate localization.

Methods

We extract 175 cases of TB and 8,933 controls using a keyphrase search approach on PACS, with studies from 2006 to 2017. Search queries were created to identify radiographically suspicious TB chest radiographs. Each positive case was manually annotated for opacity, cavitary or miliary regions of interest with rectangular bounding boxes. Among the 175 frontal images of radiographically suspicious tuberculosis were identified, 156 were positive for infiltrate, 89 for cavitary, and 19 for miliary, with all cavitary cases also having infiltrates.

We employed a RetinaNet with an underlying ResNet101 convolutional neural network (CNN) architecture in order to perform localization of TB. Pretraining was performed on the RSNA Pneumonia Challenge, a dataset comprising of over 2,500 chest radiographs positive for infiltration alongside bounding boxes. Further training was then performed on our local dataset for 40 epochs, and standard RetinaNet parameters. The final trained networks were evaluated using mean

Average Precision (mAP) for localization and using AUC for classification.

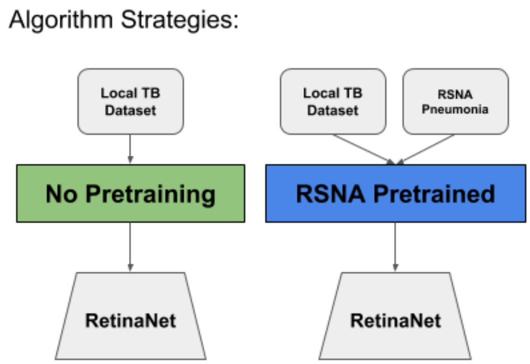

The study evaluated the effect of augmenting a RetinaNet based model of radiographically suspicious TB, using solely a local TB dataset, versus the effect of utilizing the RSNA Pneumonia Detection Challenge as a baseline for pre-training.

FIGURE 1: Methodology of our experimental design

Results

Overall performance on all three pulmonary presentations for TB performs at an mAP[0.50-0.95] of 0.320 on the test dataset. For infiltrate cases, the network performs at an mAP[0.50-0.95] of 0.384 with pre-training and 0.305 without. For cavitary cases, the network performs at an mAP[0.50-0.95] of 0.253 with pre-training, and 0.169 without. For miliary cases, the network performs at an mAP[0.50-0.95] of 0.346 with pre-training, and 0.335 without. The overall classification AUC of our network is 0.897.

Given the significant improvement in localization performance with RSNA Pneumonia Challenge cases, we believe our network effectively transfers latent pulmonary infiltrate representations to TB. For comparison, the best Kaggle leaderboard score on the *RSNA Pneumonia Challenge* was 0.246 (mAP[0.50-0.95]).

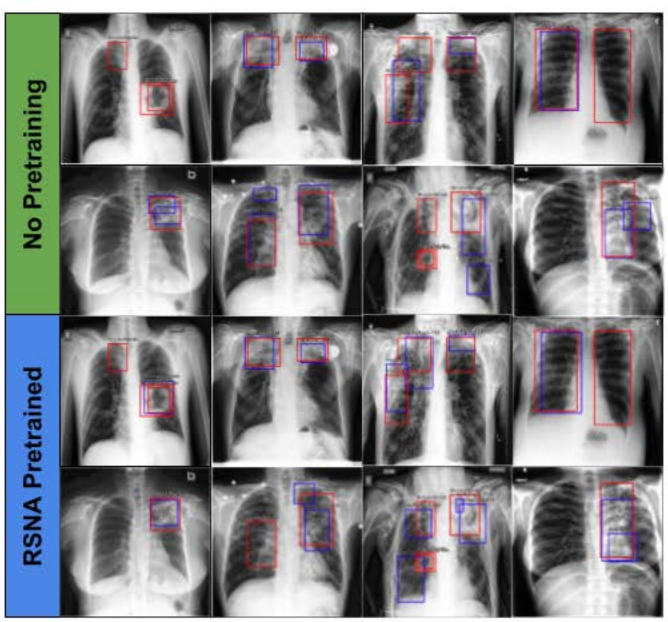

Random sampling of test-group predicted bounding boxes, without pretraining, and with RSNA pretraining. Images in the same relative position are identical cases. Ground truth [red box] and Predictions [blue box] are depicted per case.

FIGURE 2: Random sampling of test-group predicted bounding boxes with and without pre-training

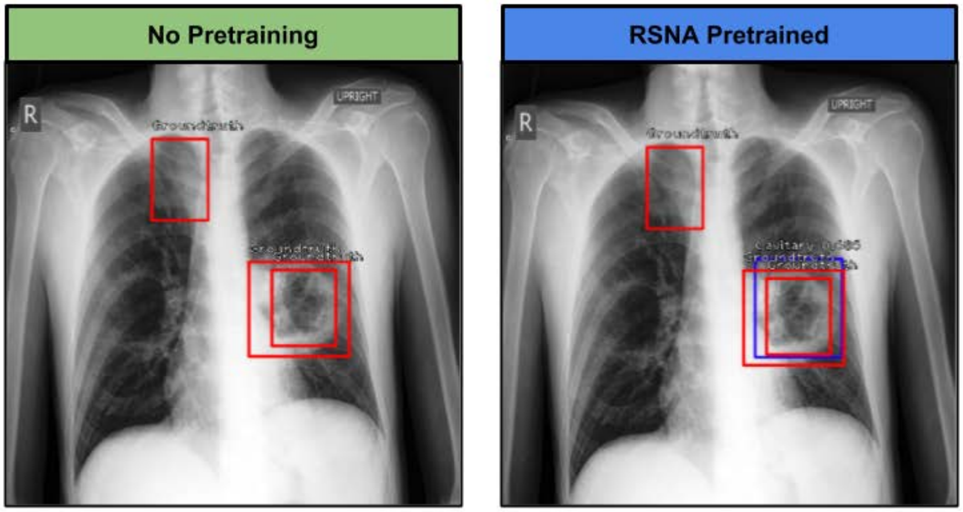

Comparison of predicted bounding boxes, without and with pre-training. Despite access to the same cavitary training examples, the DCNN failed to identify a conspicuous left mid lung zone cavitary lesion without pretraining. With the addition of RSNA pre-training, the cavitation was successfully identified, however with no improvement of the right apical infiltrate.

FIGURE 3: Comparison of predicted bounding boxes with and without pre-training for a cavitary lesion

Conclusion

Utilizing a RetinaNet CNN transfer learning approach is able to identify automatically identify, localize, and sub-classify findings of radiographically suspicious tuberculosis on chest radiographs, with maximum performance for lung infiltrates.

Statement of Impact

The performance of our tuberculosis localization/classification model could greatly reduce the time to interpretation of this critical finding.

Keywords

chest radiographs, tuberculosis, CNN, object localization, transfer learning