Abstract

In the face of an ever increasing patient burden and limited medical expertise, developing nations require effective screening strategies to deal with the high mortality rates associated with infectious diseases such as pulmonary TB and pneumonia. The recent use of deep convolutional neural network (DCNN) models to perform automated chest X-ray screening presents a viable solution to this need. However, in order to operate with high accuracy, these deep models are often required to detect subtle cases of pathology under weak training supervision.

This dissertation addresses the challenge of improving DCNN-based pathology classification under a weakly-supervised setting, by developing a multi-instance learning (MIL) method to detect local discriminative information within weakly-labelled frontal chest X-ray images.

The developed MIL method extends from previous works in the literature, and consists of two training stages which can be fully automated to enable end-to-end learning. The first stage sensitises a model to local sources of discriminative information, while the second stage boosts this model to recognise non-discriminative sources.

Validation of the method is initially performed on two synthetic datasets, after which it is experimentally tested across four chest X-ray datasets containing pathological findings for TB and pneumonia. The overall results show that the MIL method enables the detection of small and subtle findings for pathology, outperforming conventional weakly-supervised classification for pulmonary TB detection, while producing poorer performances in detecting pathologies with larger spatial extents. Consideration of these findings motivates for the combined use of the MIL method along with other weakly supervised techniques, such holistic classification, in order to improve overall pathology classification.