Characterization of Pulmonary Fibrosis on HRCT Images Using Deep Learning.

Xavier Ignacio Gonzalez[1,*], Diego Llarrull[1,*], Mirabela Rusu[2#], Ansaf Salleb-Aouissi[1#]
[1]Columbia University. Data Science Institute, New York City, NY, USA. 10025
[2]General Electric, Global Research, Niskayuna, NY, USA 12309
* Contributed equally to the work. # Contributed equally to the work.

Introduction: Interstitial lung diseases (ILDs) are relatively rare conditions but include a broad spectrum of diseases resulting from the development of diffuse abnormalities in the lung. With more than 150 histological diagnoses, these diseases have very different prognosis. Due to this variability, it is essential to accurately identify the type of ILD, to allow the administration of the correct treatment. High-resolution computed tomography (HRCT) imaging has a central role in this aim, as the appearance and quantification of lung tissue patterns are very informative for establishing the differential diagnosis. The goal of this study is to investigate the utility of Deep Learning in identifying pulmonary fibrosis in HRCT images. Previous approaches predict pulmonary fibrosis based on HRCT using a combination of feature extraction and traditional machine learning classifiers. Yet, to the best of our knowledge, the utility of deep learning for this task has not yet been investigated.

Dataset and Methodology: The dataset used to build the model is described in [1] and contains HRCT scan of 101 subjects with confirmed pathologic diagnosis of diffuse disease, including pulmonary fibrosis or normal regions. Two expert radiologists outlined region of diffused diseases for 13 conditions. Our approach focused on distinguishing pulmonary fibrosis from the other diffuse diseases (aggregated into a non-fibrosis group). The framework had several steps. First, we pre-processed the data by normalizing the image intensity, to unify the grayscale distribution, and by correcting for the different resolution. Second, each image with radiologist outlines was decomposed in 28x28 pixel subframes. Third, to correct the imbalance in the dataset, the positive cases (i.e. subframes with fibrosis label) were oversampled by adding their reflections. Specifically, each positive subframe resulted in 3 additional positive subframes obtained by mirroring and reflection along the horizontal and vertical. The oversampling resulted in 2036 positive subframes, and 2886 negative subframes. These resulting subframes were input into the convolutional neural network CNN with two convolutional and two subsampling layers.

Results and Discussion: Two sets of experiments were conducted. First, a 5-fold cross validation was executed to obtain the optimal CNN configuration and optimal set of hyperparameters such as pixel size pattern, number of kernels in the layers, kernel sizes, subsampling scale, etc. These experiments returned a cross-validation error ranging between 16 and 20%. The second experiment consisted into a leave-one-out per patient validation technique. All images of one patient were removed from the training dataset and were used for testing, while the remaining images were used in the training of the CNN. This second experiment used the optimal hyperparameters identified in the 5-fold cross validation and returned a median error rate of 13,5%. Moreover, our approach generated prediction maps to facilitate interpretability (see example in Figure 1).

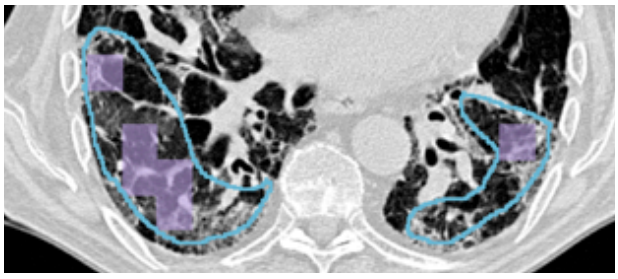

Figure 1. Lung image with expert radiologist annotation (light blue contour), and with automatic prediction maps (purple rectangles).

Conclusions: This study shows the utility of CNNs to classify tissue patterns into lung images. Our approach provided encouraging results compared to previous works. Future work pointed to more sophisticated configurations of CNN may provide even better performance.

References:

[1] Depeursinge, Adrien, et al. "Building a reference multimedia database for interstitial lung diseases." Computerized medical imaging and graphics 36.3 (2012): 227-238.