

Pneumonia Detection on Chest X-Rays using Transfer Learning

William G Morgan
3/18/2019

I. Definition

Project Overview

Every year, in the United States alone, more than 1 million adults are hospitalized with pneumonia and roughly 5% or 50k die from the disease (CDC, 2017, <https://www.cdc.gov/nchs/fastats/pneumonia.htm>) . That's 15 to 1 per 100,000 population that end in death.

Currently, a chest X-Ray is one of the best available methods for diagnosing pneumonia (<https://www.radiologyinfo.org/en/info.cfm?pg=pneumonia>). However, identifying or detecting pneumonia in chest X-Ray is very challenging and relies on the expertise of a radiologists.

Problem Statement

The workload of clinical radiologists continues to increase year over year, which places pressure on radiological services to become more efficient while maintaining and improving the quality of both patient care and radiology outcomes.

The purpose of this research is to develop a tool that can reasonably identify pneumonia by use of machine learning while taking advantage of Domain Adaptation thru Transfer learning.

We will demonstrate the performance of the models by statistical metrics like accuracy, recall, precision, and F1 score and comparing the metrics of these models to those of clinical radiologists.

The accuracy of most radiologists ranges from 88.5% to less than 85.2% depending on workloads ([https://www.jacr.org/article/S1546-1440\(10\)00134-1/abstract](https://www.jacr.org/article/S1546-1440(10)00134-1/abstract)).

Therefore, a machine learning algorithm that can accurately predict at 88.5% or better consistently, thus will out perform a human. There are documented studies as referencing the loss in accuracy as fatigue sets in

(<http://www.diagnosticimaging.com/practice-management/radiologists%E2%80%99-diagnostic-accuracy-drops-day-goes>).

The below outline provides the workflow of the approach used in solving the identified problem:

1. Explore the Data
 - a. Review training data
 - b. Review validation data
 - c. Review testing data
 - d. Dimensionality of the data
 - e. Statistical summary
2. Data Preprocessing/Augmentation
 - a. Apply transformations
 - i. Random scaling
 - ii. Cropping
 - iii. Rotation
 - b. Normalize the means and standard deviations of the images to what the network expects
3. Build a new feed forward classifier
 - a. ReLU
 - b. Dropout
 - c. Backpropagation
 - d. Tune hyperparameters
 - e. Features of pretrained model [Transfer Learning]
4. Model Tuning
5. Predict, summarize results

The expectations are that we will successfully be able to train a model using a convolutional neural network [CNN] with one of two architectures; Vgg13 or Alexnet, that will exceed the 77% metric as shown in the CheXnet model being used as our benchmark metric.