Pneumonia Diagnosis with Convolutional Neural Networks

Jim Jiayi Xu, Dylan Vizcarra, Apoorva Srivastava, Denana Vehab
jjx002@eng.ucsd.edu, dvizcarr@ucsd.edu, a1srivas@eng.ucsd.edu, dvehab@eng.ucsd.edu

Predicting

Using machine learning, we aim to diagnose patients' chest x-rays (CXRs) to reduce specialists' workload and increase rate of accurate diagnosis. We did this using convolutional neural networks, and overall, our machine learning pneumonia classifier has 81% accuracy, 60% precision, and 80% recall on the testing set.

Data and Features

We use the Chest X-Ray Images (Pneumonia) dataset from Kaggle [1]. It comes with labelled chest x-ray jpegs with and without pneumonia. We performed no manual feature engineering. Instead, we allocated this task to the convolutional layers in our neural networks

No Pneumonia

Pneumonia

Preprocessing:

We performed histogram equalization on all images to enhance contrast.

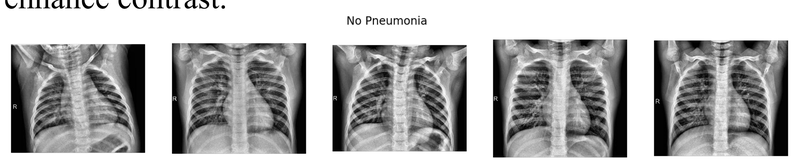

Models

We used VGG16 both as a template for our custom neural network and transfer learning experiments

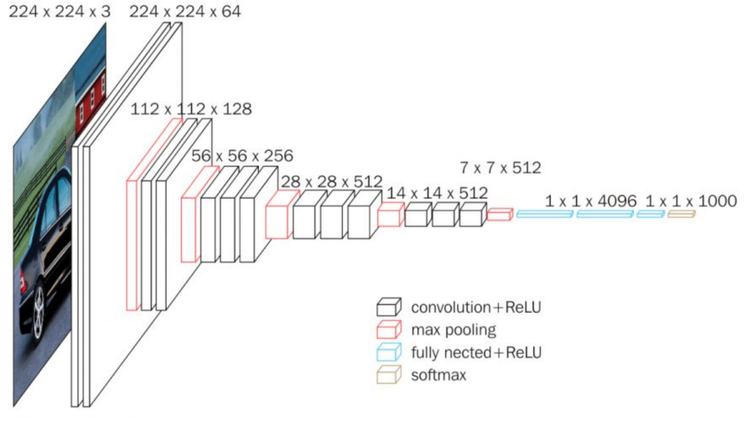

Results

Accuracy Loss

Custom NN:

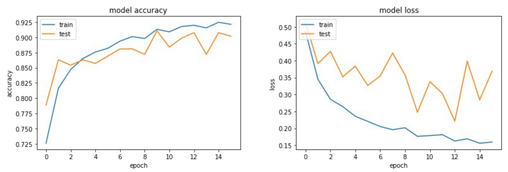

Pretrained VGG:

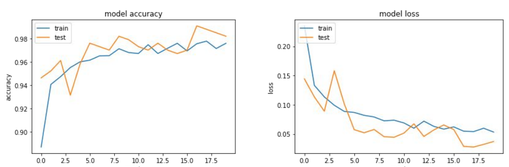

Final Results:

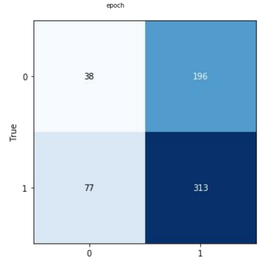

Discussion

The final test dataset accuracy is noticeably lower than the validation accuracy. We tuned our model and hyperparameters on the validation dataset, and are surprised that performance is so different when using the test dataset. Following standard procedure when presenting machine learning performance, we did not tune our machine learning systems on the testing dataset. Although our model did not perform as well on the testing dataset as we have hoped, we still present this result as an honest representation of our machine learning model.

Future

We recommend adding more data samples to the relatively small dataset and trying other successful CNN architectures beside VGG16. Additionally, we would try to fine tune the pretrained VGG16 by unfreezing the last convolutional layers and jointly training with the classifier. After implementing those two changes, and obtaining improved results, we would try to localize the pneumonia using an adjusted CNN architectures such as mask RCNNs.

References

[1] "RSNA Pneumonia Detection Challenge," Kaggle. [Online]. Available: https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data. [Accessed: 04-Jun-2019].