Classifying Pneumonia in Chest X-Rays

Benjamin Antin, Joshua Kravitz, Emil Martayan
{bantin, kravitzj, emilmar}@stanford.edu

Problem

Detecting Pneumonia from chest X-rays is a challenging task, even for experienced radiologists. With computer-aided diagnosis, physicians can make chest X-ray diagnoses more quickly and accurately. We aim to **train a model to help physicians in making chest X-ray diagnoses**.

Data

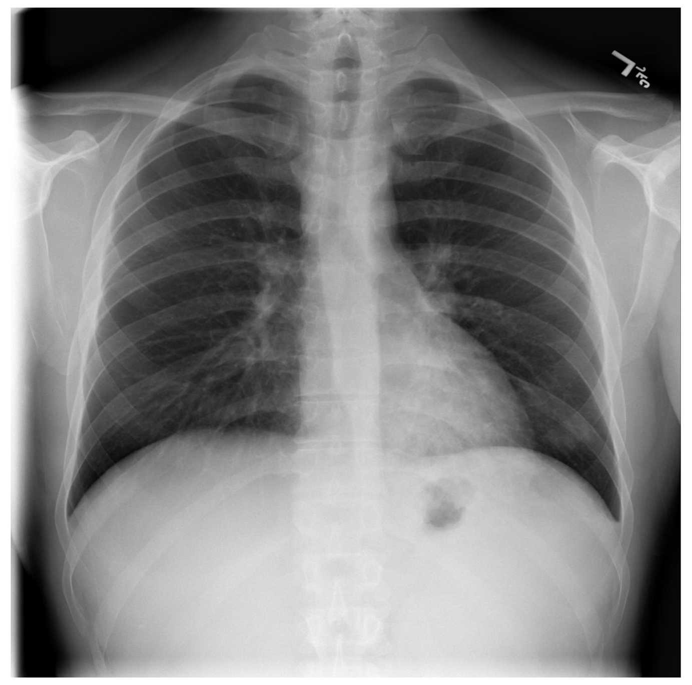

Pneumonia

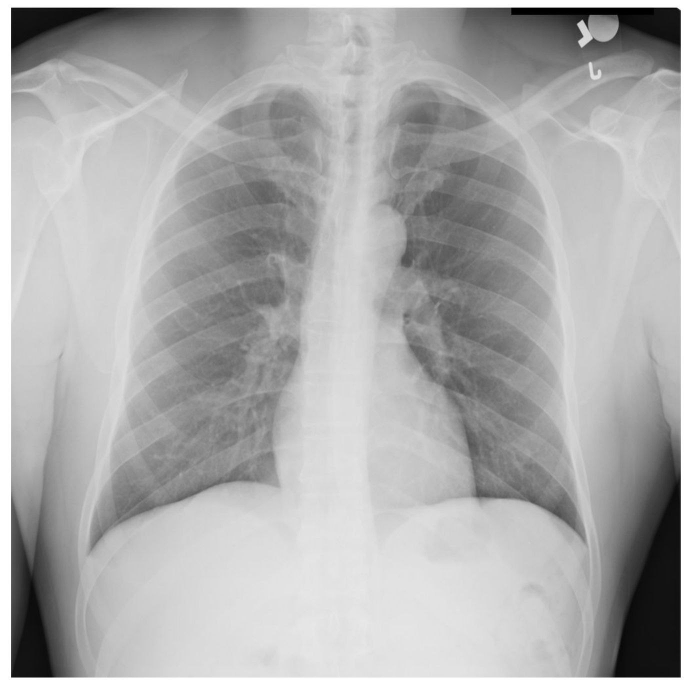

No Pneumonia

- 112,000 labeled grey-scale chest X-ray images by ***Wang, et al***[1]
- For our project, we focus on classifying pneumonia vs. no pneumonia, though the dataset has labels for other diseases

Pre-Processing

- Downsample 1028x1028 pixel images to 224x224 (and smaller sizes for logistic regression).
- Normalize and randomly flip the images (CNN only).

Baseline: Logistic Regression

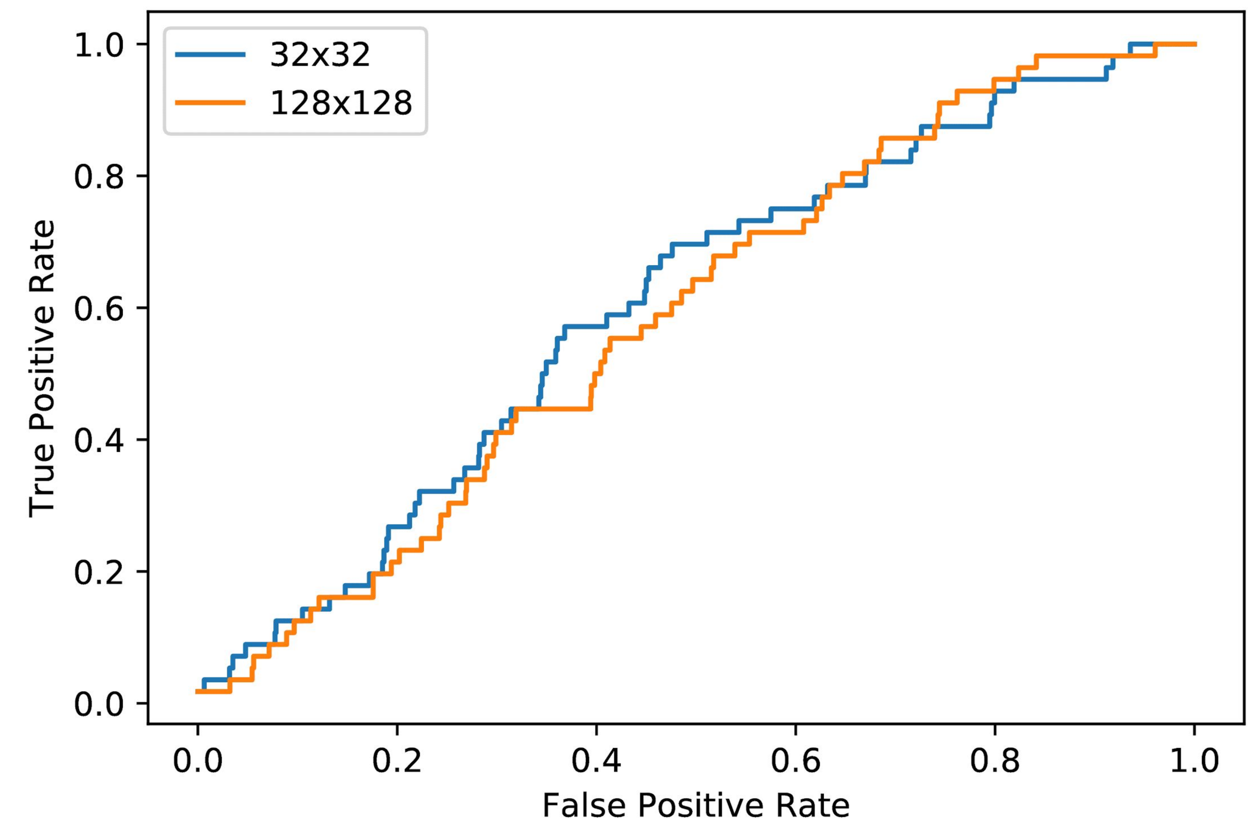

Our baseline treats each pixel as a feature in a logistic regression model. Using L2 regularization of 1500 and 3000, we achieve an AUC of 0.6037 and 0.5861 for 32x32 and 128x128 images, respectively.

Model: DenseNet-121[2]

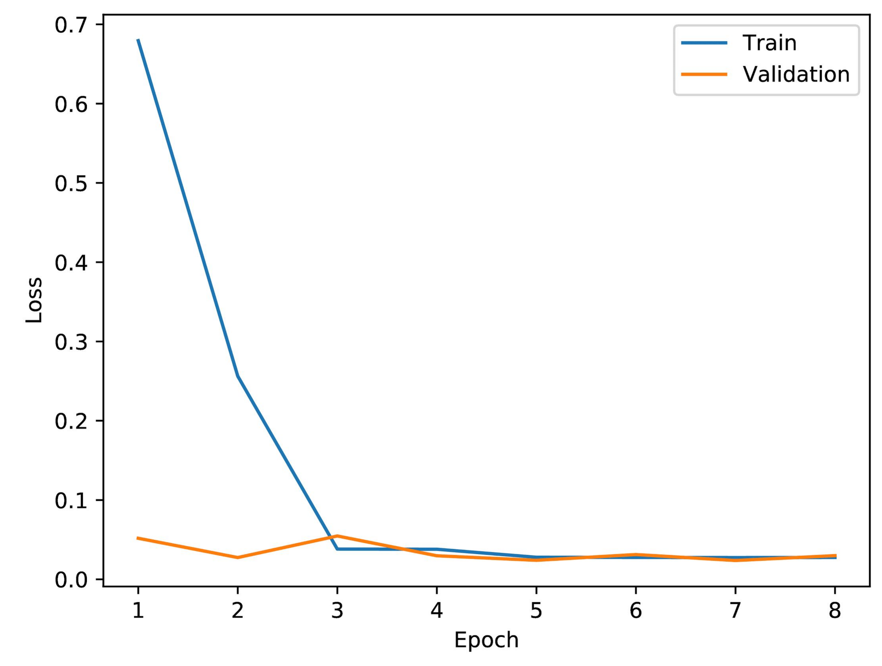

- Because this is such a computationally expensive task, we set up a GPU via Google Cloud to run our model
- We begin with a learning rate of 1 and anneal it as the training loss plateaus
- We utilize weighted cross-entropy loss:

$$L(y,\hat{y}) = -\frac{1}{B}\sum_{i=1}^{B} w^{(i)}\left[y^{(i)}log\hat{y}^{(i)} + (1-y^{(i)})log(1-\hat{y}^{(i)})\right]$$

Discussion

- Logistic regression does not capture complexity of dataset
- Tuning the hyperparameters for the CNN is difficult
- Imbalanced training set adds difficulty, meaning that accuracy alone isn't a good metric
- Challenging problem for expert radiologists, so we expected low performance metrics

Future Work

- Run the CNN for more epochs and evaluate accuracy, precision, and recall
- Explore other architectures such as residual networks
- Obtain the output from the final fully-connected layer and run t-SNE to cluster the output

References

[1] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[2] G. Huang, Z. Liu, L. V. D. Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.