Deep Convolutional Neural Networks for Lung Cancer Detection

Niranjan Balachandar, Peter Lu, Albert Chon

Background & Problem Statement

- Over 200,000 people are diagnosed with lung cancer every year.
- Lung cancer is the leading cause of death among all cancers.

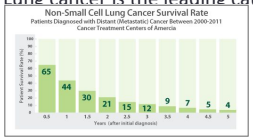

- Early detection is important for improved survival but early stage cancer nodules are small and hard to detect
- Our task is a binary classification problem to detect the presence of early stage lung cancer from patient CT scans.
- Dataset: Kaggle Data Science Bowl 2017 (70% labels are 0)
- Evaluation: Accuracy, Sensitivity, Specificity, AUC

Preprocessing (Segmentation)

- CT scan slices and corresponding histogram of radiodensities (Hounsfield Units) are shown below:

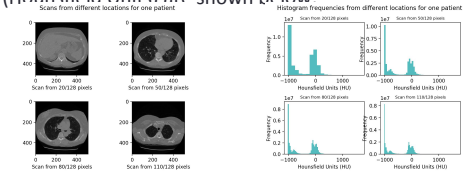

- We perform segmentation by thresholding
  - <--1000 HU: air (masked out)
  - >400 HU: Bone segment (masked out, left image)

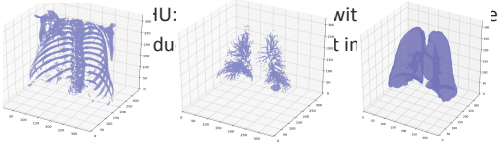

Baseline Models

- Linear Classifier and 2D CNN (architecture described at the right)
- Raw Pixels and HOG as feature representations

1. Convolution layer with 128 filters of size (3,3), strides (1, 1) and zero-padding
2. Batch normalization layer
3. 2D max normalization layer with pool size (2, 2) and strides of length 2
4. ReLu activation layer
5. Convolution layer with 128 filters of size (3,3), strides (1, 1) and zero-padding
6. Batch normalization layer
7. 2D max normalization layer with pool size (2, 2) and strides of length 2
8. ReLu activation layer
9. Fully connected layer of size 256
10. ReLu activation layer
11. Fully connected readout layer of size 2

Model	Accuracy Train / Val	Sensitivity Train / Val	Specificity Train / Val
Linear+pixels	0.895 / 0.663	0.956 / 0.520	0.869 / 0.710
Linear+HOG	0.957 / 0.584	0.989 / 0.400	0.944 / 0.644
2D-CNN+pixels	0.526 / 0.436	0.989 / 0.720	0.332 / 0.342
2D-CNN+HOG	0.707 / 0.752	0.089 / 0.080	0.967 / 0.974

Advanced Models

SUNet:

~19M parameters, maxpool after each conv layer

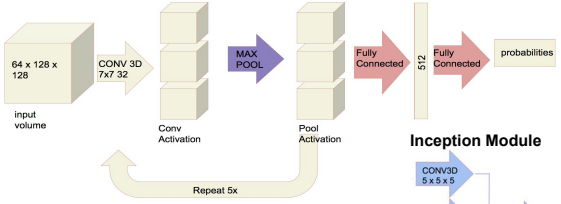

Modified GoogLeNet:

~ 23M parameters

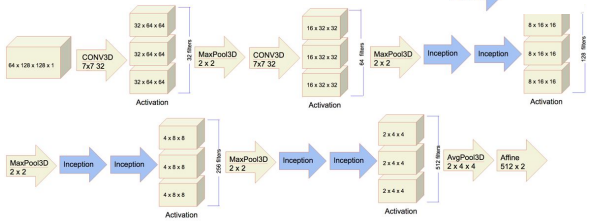

Advanced Models (cont.)

- Both Models:
  - Activation: Leaky ReLU
  - Loss: Softmax Cross Entropy
  - Gradient Descent Algorithm: Adam
  - Regularization: L2, dropout (P_{drop} = 0.2) after each conv2D and before final readout layer (P_{drop} = 0.5)

Preliminary Results

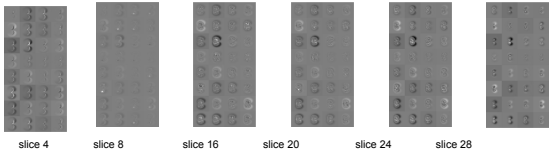

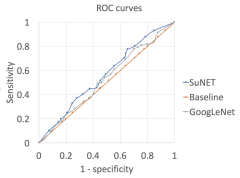

Evaluation: Area under ROC
SUNet AUC: 0.563
GoogleNet AUC: 0.528
Results slightly better than random guessing, but hold your horses! We are still training...

Conclusions/Future Work

- Current performance not great but models have not finished training.
- Use a more advanced segmentation algorithm (Watershed)
- Increase Regularization or Model Complexity
- Analyze saliency maps for better tumor detection
- Use ensemble models to achieve better performance