Segmentation of Breast Cancer Tumors using Deep Learning

Eric Matsumoto, Jason Lin

Motivation

Breast cancer is one of the leading causes of death for women in the United States. Many current treatments for breast cancer target certain hormone receptors in the breast cells or the HER2 protein, all of which can stop the spread of malignant breast cancer cells. However, if the cancer does not have hormone receptors and it also doesn't respond to the HER2 treatments, then it's considered a 'triple negative' cancer, meaning no known treatments will work on it. In this project, we attempt to develop methods to segment these tumors in patients with triple negative breast cancer.

U-Net Architecture

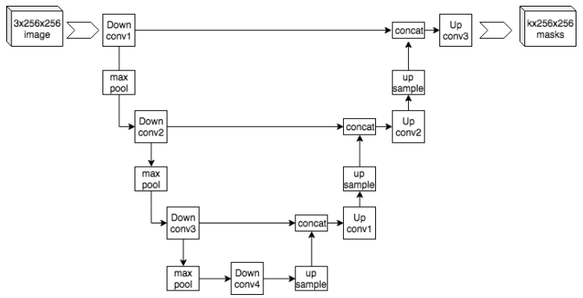

Dataset

Our dataset consists of 222 patient MRIs. This dataset was provided by our faculty sponsor, Haruka Itakura, at the Stanford School of Medicine.

The data is provided to us in the format of Nifti images, which are 3D images commonly used in the medical field. The data came in the shape of 512 x 512 x depth, where depth varied widely depending on the sample. To simplify, we sliced these into images of size 512 x 512. We also reduced the image size to 256 x 256 to enable faster training for subsequent experiments.

The dataset was very imbalanced since most pixels did not contain a tumor and the tumors were only represented in a subset of slices and a small area of each slice.

Approach

We use a U-Net architecture for segmentation. U-Net is a convolutional neural network aimed at segmenting biomedical imagery. It's a fully convolutional network, inspired by the fully convolutional network architecture proposed by Long and Shelhamer.

We trained the model using the Adam optimization algorithm for 5 epochs. We experimented with three different loss functions: binary cross entropy loss, Dice coefficient, and Jaccard distance. Binary cross entropy loss was the default loss of U-Net, and we chose to additionally use Dice and Jaccard because they are known to do well with imbalanced classes in image segmentation

We also created more "balanced" datasets by selecting slices that had tumor in them in an effort to correct our class imbalance. We ran the model on both the full and balanced datasets.

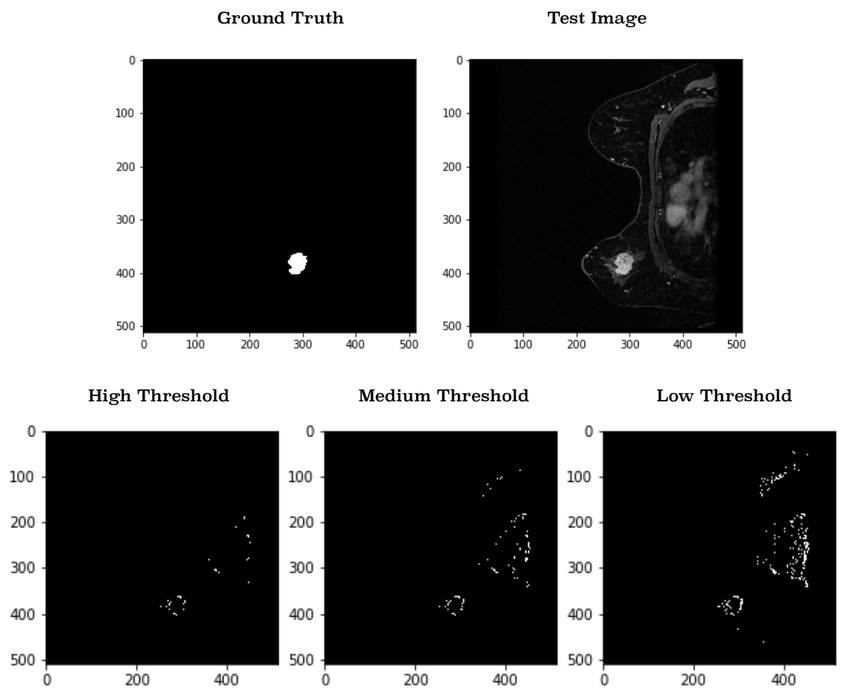

Results

Validation Performance	Accuracy	Precision	Recall	F1
Cross Entropy Full	0.997	0.005	0.078	0.009
Dice Full	0.976	0.000	0.000	0.000
Jaccard Full	0.976	0.000	0.000	0.000
Cross Entropy Balanced	0.18	0.12	0.99	0.21
Dice Balanced	0.43	0.13	1.000	0.22
Jaccard Balanced	0.18	0.13	1.000	0.22
Cross Entropy Full Lower Learning Rate	0.741	0.074	0.678	0.125

Conclusion

In conclusion, the segmentation of breast cancer tumors proved to be a difficult problem to tackle. A large part of this difficulty stems from the inherent class imbalance within the training data. Our initial efforts showed high accuracy but low recall, and our efforts to fix these issues created results that returned high recall but low accuracy, since it was over-predicting pixels as positive.

Additionally, we see that there is potential in decreasing the learning rate of the model. This suggests that our original learning rate caused the model to oscillate around or diverge from the optimal answer.

Future Work

We were constrained by the shared resources of Stanford Medicine's computing platform since the confidential nature of our data mandated keeping data on that platform. Given more time and resources we would hope to perform more experiments trying to balance our dataset or tuning hyperparameters to boost performance. We also hope to train on a model by Andriy Myronenko from Nvidia, which won first place in the 2018 BraTS challenge for brain tumor segmentation. We prepared the model and data but did not time to train.