Histological Image Segmentation: Cancer vs Stroma

Adam Abdulhamid

Overview

- Deep Learning, especially in the field of vision, has huge implications in the medical imaging domain.
- Goal: learn a model that outputs a pixel-wise classification of either cancer cell, stroma, or neither.
- The last 5 years have shown an incredible increase in convolutional networks ability on all vision tasks. I wanted to apply CNNs to a dataset hand generated by several doctors in the Stanford Medical School.

Problem Statement

- My model begins with the first 4 layers of VGG16 and adds on several more convolutional, transpose convolutional, and batch normalization layers in attempt to fine tune to this specific task.
- Evaluation of our model's performance is based on softmax cross entropy loss at each pixel, and is averaged over all pixels to give one loss value.

Data

- Consists of ~200 histological images of real breast cancer tumors, and the corresponding high quality labels generated by doctors at Stanford Medical School.
- Although small, each image is 1128x720 which provides roughly ~800k classifications per image.
- Augmented data set with flipped and mirrored images.

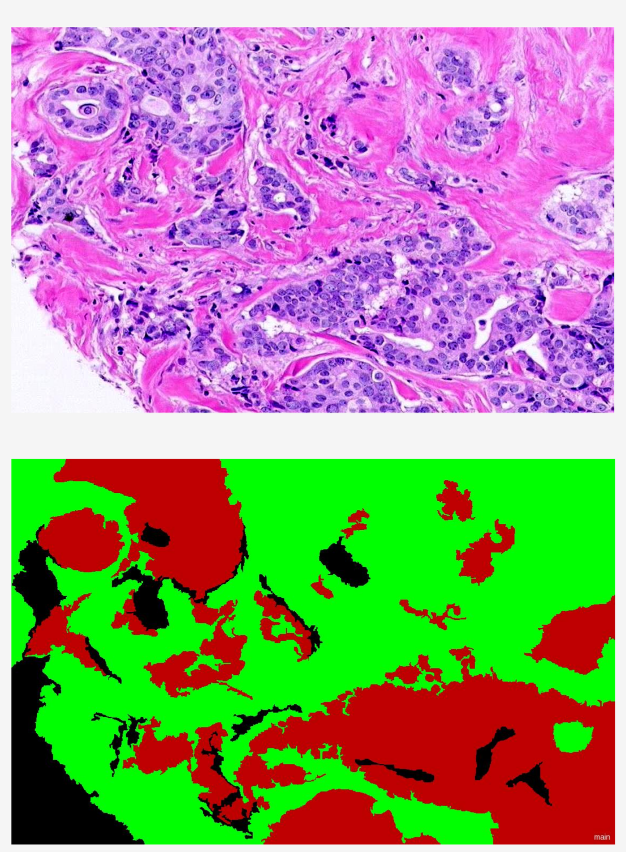

Network Architecture

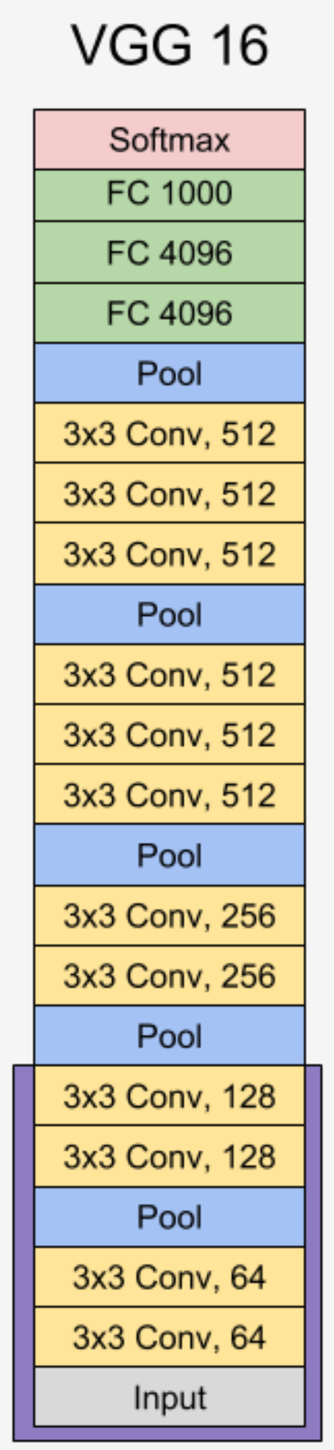

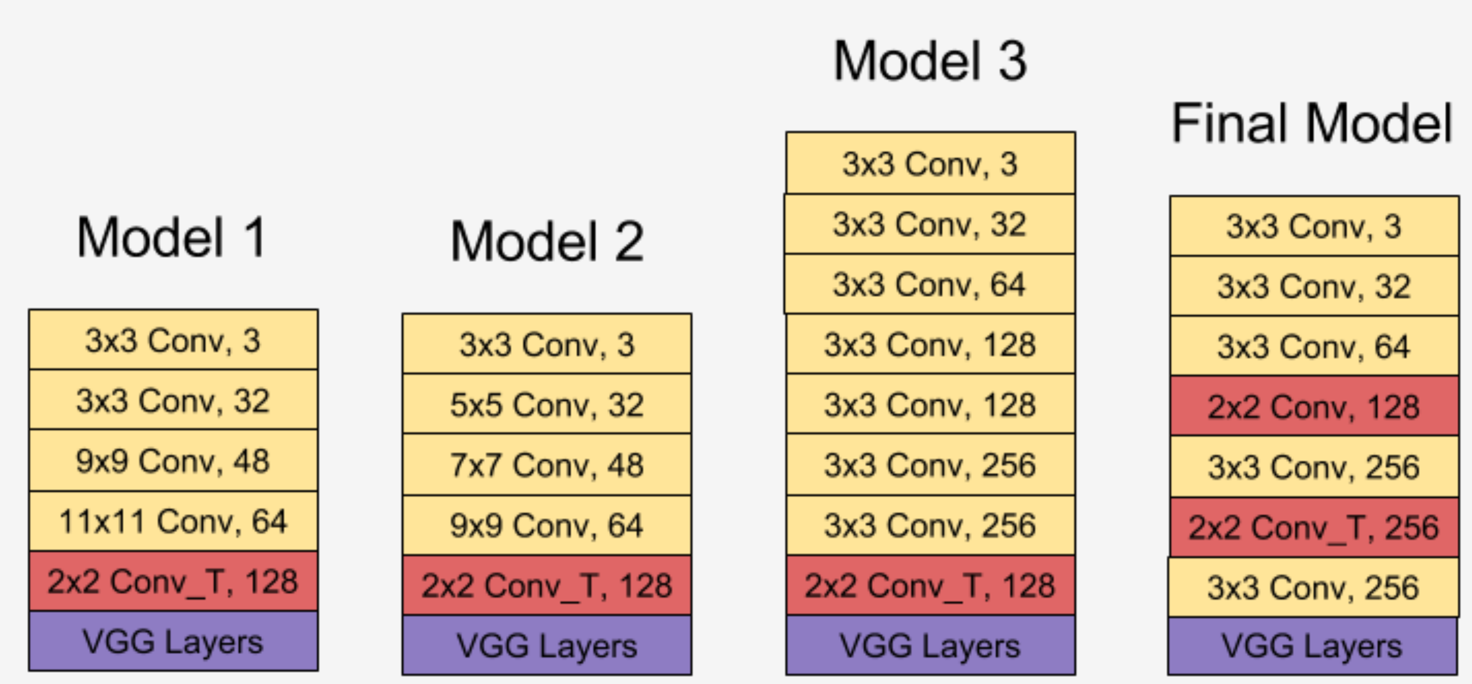

- All Conv layers use ReLu nonlinearity
- 2x2 Conv_T layer increases spatial resolution by 2x
- Model 3 introduces batch norm before every non-linearity
- Final model added 3 more layers taken from VGG
- All images end with same dimensions as inputs

Model Performance

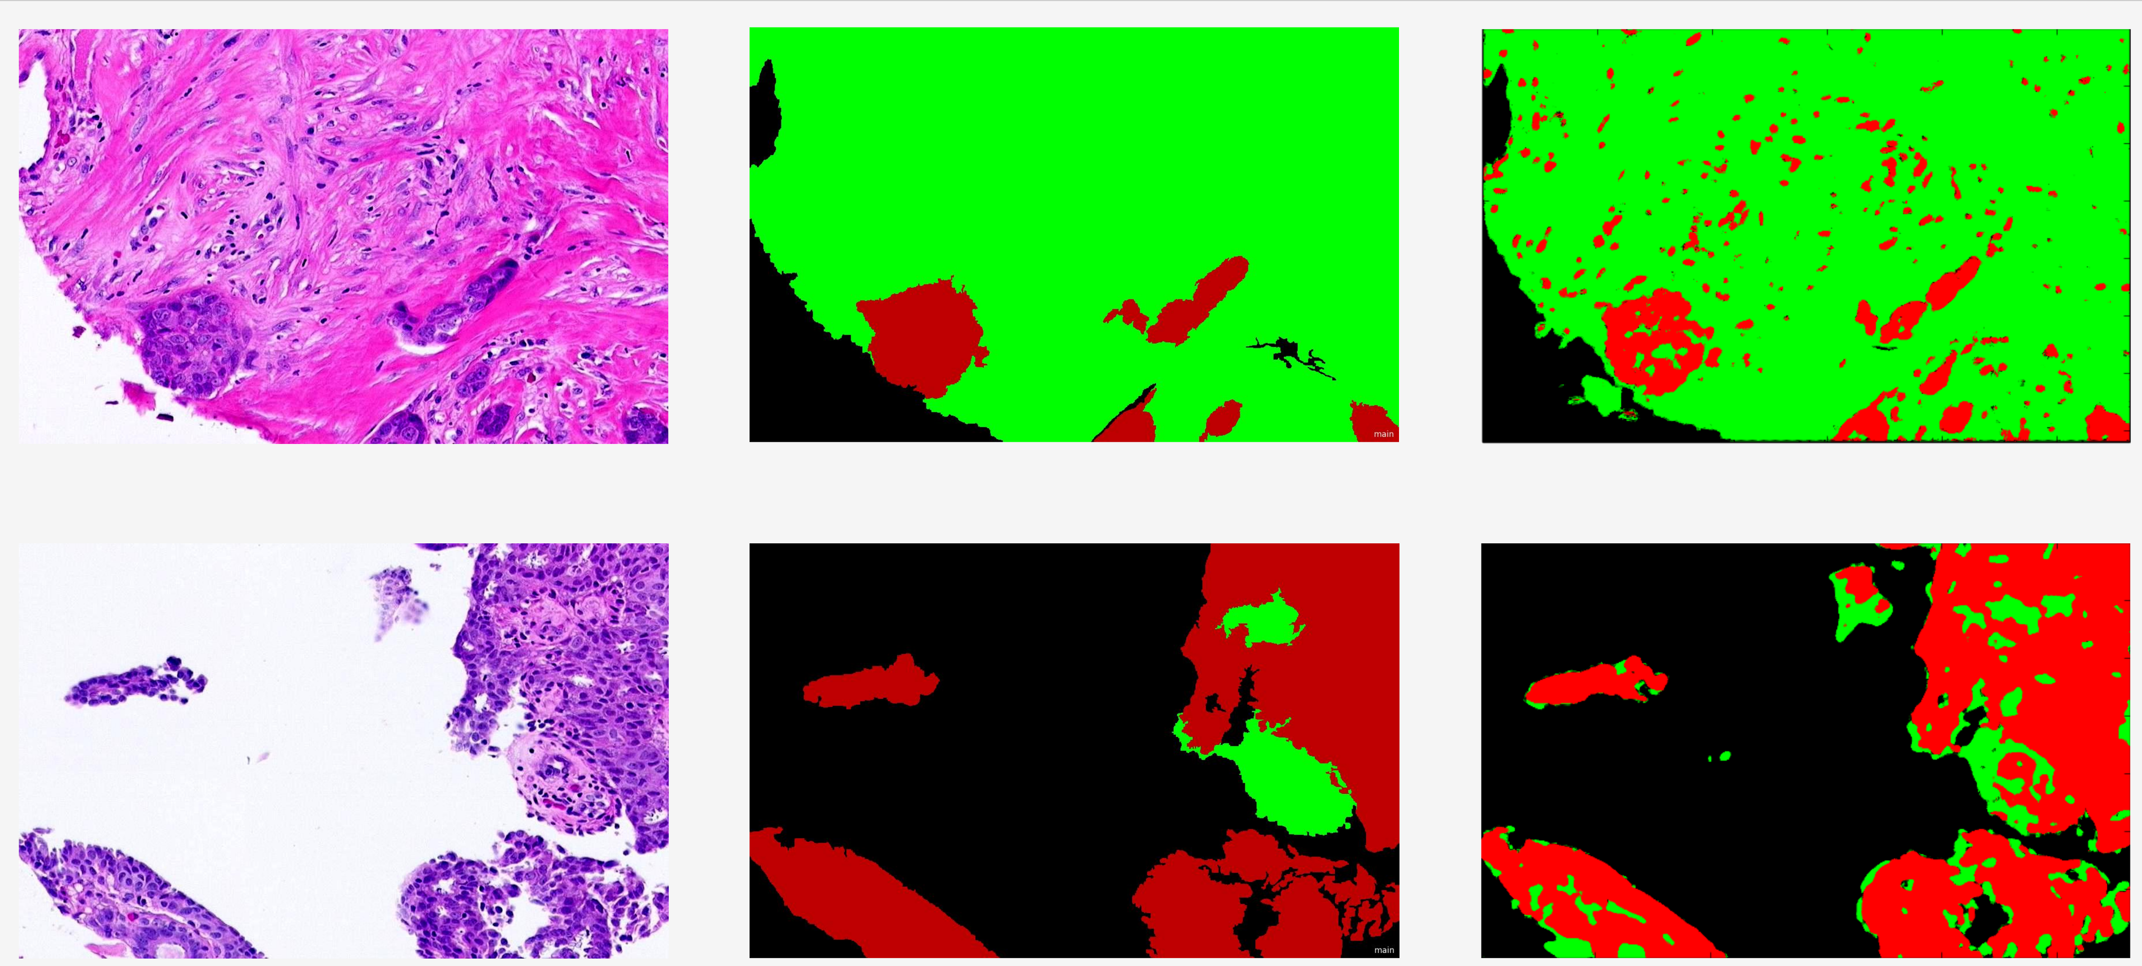

Experimental Findings

Histological Image Segmentation

Validation Loss: 0.697
Training Set Loss: 0.563
Test Set Loss: 0.685

Model	Validation Loss
Random	1.098
Model 1	0.849
Model 2	0.774
Model 3	0.748
Tuned Final Model	0.697

Conclusion

- Overall, good results were achieved, but with more noise than I would have liked.
- Clearly there is good signal here, which implies visual similarities across many different breast cancer tumors.
- Better performance was achieved when adding more complicated layers like batch norm and transpose convolution.

Future Work

- One big issue is lack of data. With a few hundred more examples there would be more signal.
- Experiment more with taking further layers of VGG and upsampling/downsampling.
- Do a more comprehensive hyper-parameter search given more time and computational resources.
- Take this to the doctors to see what they think!

References

Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. *CoRR* abs/1605.06211 *CVPR* 2016.

Hyeonwoo Noh, Seunghoon Hong, Bohyung Han. Learning Deconvolution Network for Semantic Segmentation. *CoRR* abs/1505.04366 *CVPR* 2016.

Alex Krizhevsky, Ilya Sutskever, Geoffrey Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *NIPS* 2016.

Karen Simonyan, Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556, *ICLR* 2017.

Stanford University

June 2017