Automatic Brain Tumor Segmentation

Michael Mernagh and Mihir Pendse

Background / Introduction

- Gliomas are the most common form of malignant brain tumor
- Treatment usually requires image-guided surgery or radiation therapy
- Gliomas are heterogeneous and contain subregions corresponding to different stages of the disease
- Accurate treatment requires a labeled image of the patient's brain depicting these stages
- Labeling done by highly trained radiologist → time-consuming, expensive

Problem Statement

- We will consider the MRI imaging modality and use four different types of contrast (T1, T2, T1c, and FLAIR) all taken on the same patient in the same position
- We wish to classify every voxel in the image as either (a) healthy tissue, (b) necrosis or nonenhancing tumor, (c) edema, or (d) enhancing tumor.
- We wish to use deep learning so the prediction can be done automatically and quickly without intervention from a radiologist.

Datasets

- The dataset, taken from the 2017 Brain Tumor Segmentation Challenge, consists of 3D image sets (four MRI contrasts and the label) for each of 210 patients with high grade gliomas (HGG)
- Image acquisition and manual labeling was performed across 19 institutions
- All images were coregistered, interpolated to 1mm isotropic resolution and skull stripped
- To address bias field variations across scans, N4ITK bias field correction was performed
- All images were normalized to have zero mean and unit variance
- Image Size = 240x240x155

Example of dataset for one subject

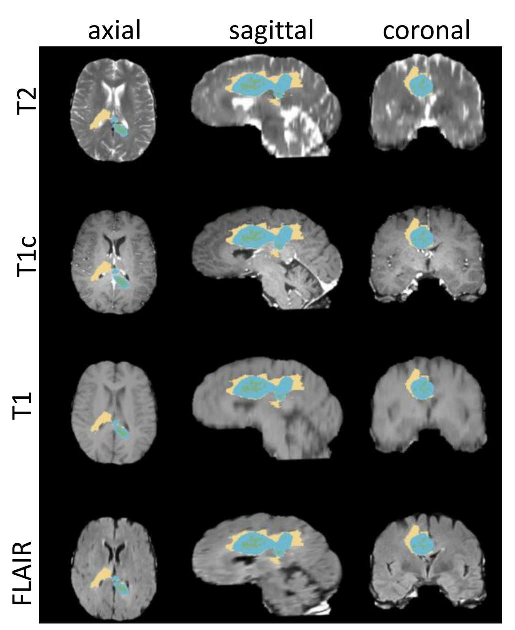

edema
enhancing tumor
necrosis

Methods / Algorithms / Models

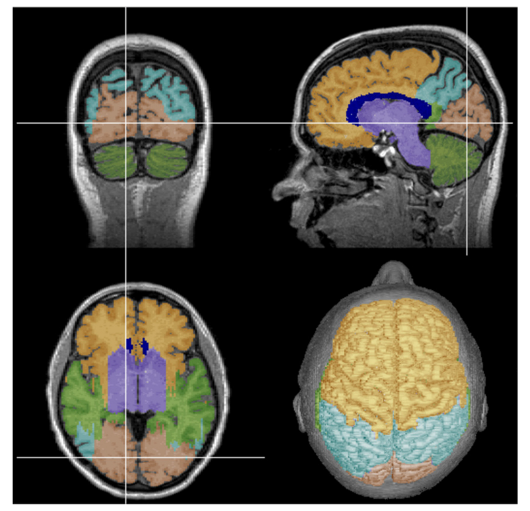

Triplanar views used as 2D input

The input to the model is a 3D volume, which is simultaneously split into 3 2D planes, each of which is passed through 6 convolutional layers.

The output of the three convolutional nets are joined together and 2 more convolutional layers follow, completed with a global averaging layer, and the softmax function.

Spatial batch normalization precedes each conv layer, and ReLU is used for activation. Dropout layers are used throughout.

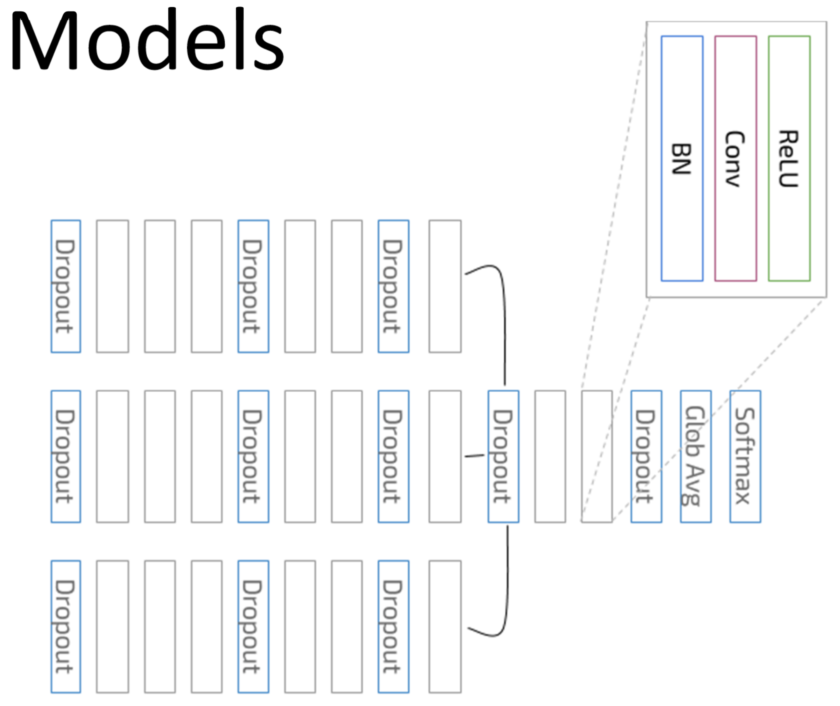

Architecture

Experimental Evaluation and Findings

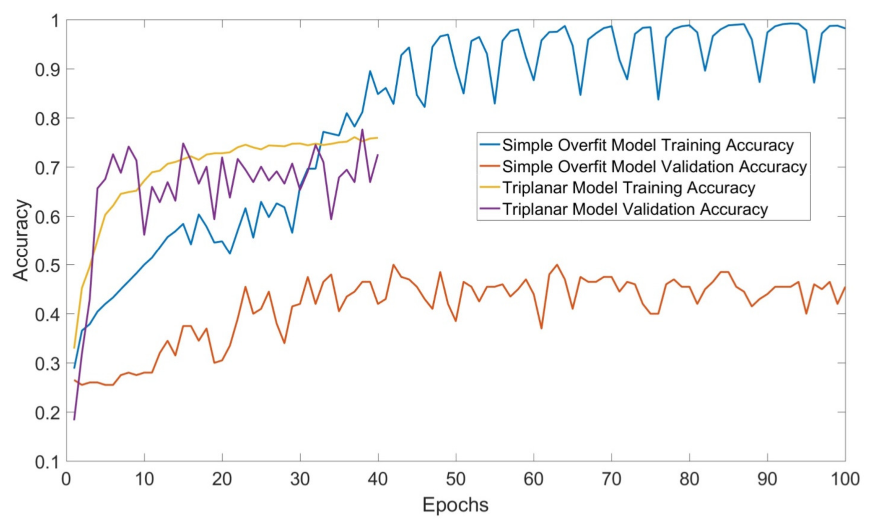

2x2x2 patch is classified from 90x90x90 triplanar image

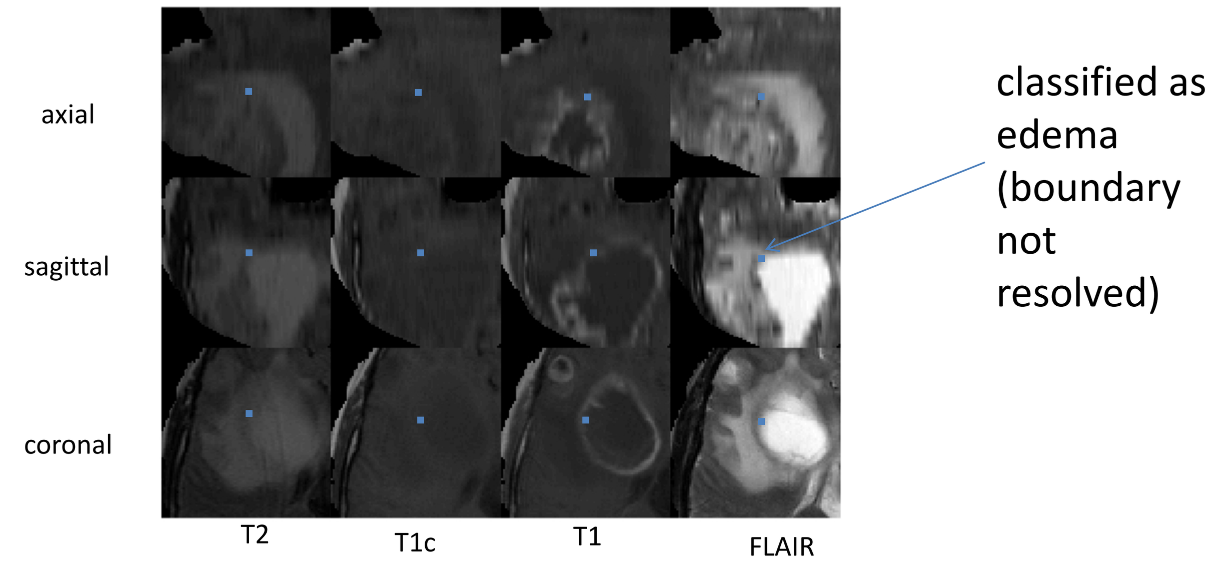

Conclusion / Future Directions

- The distribution of the patches in the training and validation sets was biased towards tumor classes as opposed to healthy tissue so there could be a nonnegligible number of false positives.
- Since boundaries between tissue classes are especially important, performance at boundaries could be improved by intentionally selecting training patches to be near boundaries