Deep Learning for Brain Tumor Segmentation

Charles Tsao and Neel Rao (chtsao and neelrao)@stanford.edu

Predicting

In this project, we wanted to build a tool to segment MRI brain scans into four parts: non-tumor, whole tumor, tumor core, and enhancing tumor core. We started by taking a pre-existing model, which attempted to solve the problem by chaining together three cascaded neural networks, each doing one step of the segmentation. We ran two experiments on this model. First, we added an extra layer to each neural network. Secondly, we experimented with the model performance by using a weighted cross-entropy loss function instead of the dice-coefficient loss function that was originally used. The extra layer and training with a different loss function did not improve on the original model

Data

We used a dataset from the 2017 Tumor Segmentation (BraTS) Challenges, held at the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI). The dataset consisted of 285 five-layered MRI scans taken in sagittal, axial, and coronal views, of brains that contained either low-grade gliomas (LGG) or high-grade gliomas (HGG).

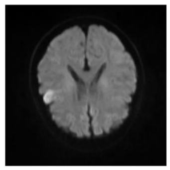

Features

The input data consisted solely of raw MRI brain scans. We did not extract any features from the scans, opting instead to let the model pick what features it paid attention to, minimizing potential for human error.

References

[1] Guotai Wang, Wenqi Li, Sebastien Ourselin, Tom Vercauteren. "Automatic Brain Tumor Segmentation using Cascaded Anisotropic Convolutional Neural Networks." In Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. Pages 179-190. Springer, 2018. https://arxiv.org/abs/1709.00382

[2] Eli Gibson, Wenqi Li, Carole Sudre, Lucas Fidon, Dzhoshkun I. Shakir, Guotai Wang, Zach Eaton-Rosen, Robert Gray, Tom Doel, Yipeng Hu, Tom Whyntie, Parashkev Nachev, Marc Modat, Dean C. Barratt, Sébastien Ourselin, M. Jorge Cardoso^, Tom Vercauteren^. "NiftyNet: a deep-learning platform for medical imaging." Computer Methods and Programs in Biomedicine, 158 (2018): 113-122. https://arxiv.org/pdf/1709.03485

Models

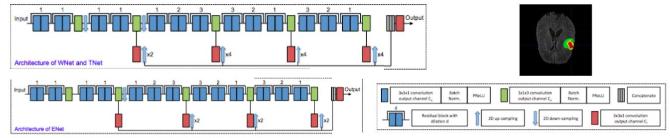

Results

Hausdorff Distance	Whole Tumor	Tumor Core	Tumor Core Enhancing
Original	3.78315	5.72536	7.25928
Added Layer	8.03155	22.60405	12.03764
New Loss Function	5.05962	5.93903	10.41551

Discussion

After training and testing all three models, we discovered that adding an extra layer to each neural network and using weighted cross-entropy as a loss function did not improve upon original performance. We believe that adding another layer may have made the model overfit to the training set, as it had higher training set accuracy, but was worse on the test set. The fact that a model trained using dice-coefficient loss as opposed to weighted cross entropy performed better on a testing set is not altogether unsurprising, because despite it being more difficult to backpropagate, dice loss performs well on class imbalanced datasets. We did find that the weighted cross entropy backpropagation had an easier time converging, however, as we would expect.

Future

From here, we would move on to the second step of the MICCAI BraTS challenge, which is to use the data provided to train a model to predict patient survival based on the MRI brain scan. We would utilize a fourth cascaded model added to the last segmentation. We currently don't have enough data to make a useful model for this task, but as more becomes available in successive years, we see this as a useful area of exploration.