Deep Learning Segmentation of Strokes in ATLAS

Roelant Kalthof, Andrew Lewis, Brian Triana
CS230: Deep Learning

Background and Motivation

Stroke is the leading cause of neurological disability and is estimated to be the leading cause of lost healthy life-years by 2020. Currently, brain images are manually reviewed to identify strokes in patients presenting with concerning symptoms. Applying deep learning for segmentation of strokes provides two potential sources of value.

1) Reliable segmentation of strokes can contribute towards generation of high quality data for clinical research purposes. Accurate identification of anatomical regions affected by stroke is key for creating accurate cohorts of patients to assess outcomes and the effect of potential therapies.
2) Rapid segmentation can be directly applied to the clinical setting, where rapid diagnosis and treatment is critical for improving patient outcomes. Approximately 2 million neurons are permanently lost per minute in which a stroke is untreated. In other words, time is brain. Rapid and accurate segmentation of stroke through deep learning can facilitate the clinical workflow pathway and help improve patient outcomes after stroke.

Dataset

We are utilizing the Anatomical tracings of Lesions After Stroke (ATLAS) dataset (release 1.1) for both our training and testing. This is an open-source dataset consisting of 304 T1-weighted MRIs with manually segmented lesions with additional metadata. Each MRI is represented as a stack of 2D 232x196 image slices, containing 188 axial images. Each patient also has 188 corresponding lesion masks with 232x196 binary labels for each lesion.

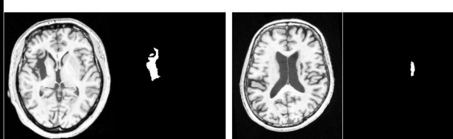

Figure 1. Example Axial Slices and Corresponding Lesion Masks

Objectives

- Develop a deep learning model to segment images from the ATLAS database
- Explore data splits by both slice and by patient
- Optimize the model using hyper parameter exploration

Model Architecture

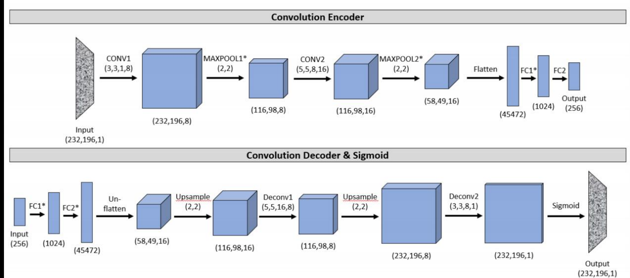

Figure 2. Diagram of Convolutional Neural Network

Training Results

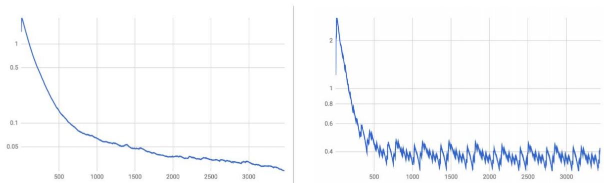

(i) (ii)

Figure 3. Training loss by iterations for split by slice (i) and split by patient (ii)
The loss for a split by patient remained highly variable and did not decrease steadily with increased epochs. We believe that this is a limitation of the sample size when split by patient.

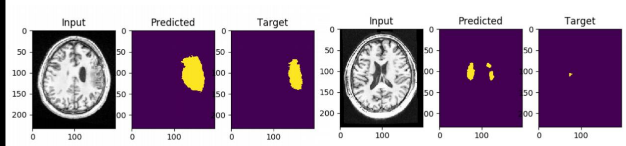

Figure 4. Example Prediction Results
Segmentation results for the model using split by slices are shown above. In general, the model seemed to overpredict the size of the lesion. Performance for the split by patient was very poor and did not accurately predict regions.

Training Results

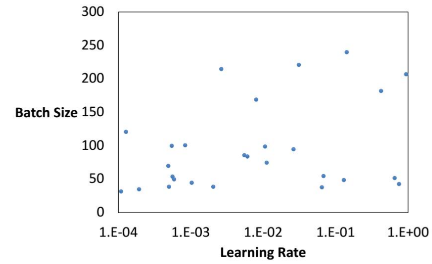

Figure 5. Randomized batch size and learning rates used for hyperparameter optimization

Batch Size	Learning Rate	Dropout	Training Dice Coefficient	Dev Dice Coefficient
0.002046	39	0.282	0.034	0.040
0.000487	70	0.046	0.436	**0.429**
0.000558	54	0.015	**0.442**	0.415
0.939068	207	0.212	0.027	0.026

Table 1. Select results of hyperparameter optimization with best results bolded

Architecture Variations

- We expanded the model with two additional FC layers and found the following results:
  Training Dice Coefficient: 0.351
  Dev Dice Coefficient: 0.332
- Although these results were worse than our initial model, this deeper model could likely be further optimized to match or improve performance of the existing model

Conclusions

- We have established feasibility of basic segmentation of stroke lesions in T1-weighted MRI images
- Additional hyperparameter optimization can be performed
- Additional architectures can be used to optimize the network, including a deeper model

Acknowledgements

Special thanks to David Eng for guidance on the project and the CS230 course faculty for their support.