

Post-stroke Lesion Detection using ATLAS Dataset

Andrew Zhang, Chenyao Yu, Helen Jiang

Abstract

Hundreds of thousands of people in the United States suffer from a stroke every year. A vast majority of stroke survivors have long-term problems that constantly affect their physical, emotional, and cognitive well-being, highlighting the importance of proper recovery. Accurate identification of stroke lesions from the MRI slices of the brain of a stroke victim can be very useful for research into more effective recovery. Deep learning has recently helped with improving automated lesion identification. In this paper, we try out the U-Net architecture on the new ATLAS dataset in the domain of post-stroke lesion detection.

Introduction

- Goal: automatic identification of lesions from MRI slices of the brains of stroke victims
- Current gold standard: manual segmentation
- Our method: 2D UNET with metadata, augmentation and batchnorm

Dataset

- ATLAS (Anatomical Tracings of Lesions After Stroke) Dataset
 - 229 MRI scans from different patients
 - Metalabels for each scan which include number of lesions, type of stroke and primary stroke location

Method

1. 2D U-Net model

- A network that combines a contracting path that learns higher level features with an expansive path that allows for the network to output a high-resolution segmentation map
- In these successive layers, the pooling operators are replaced by upsampling operators (fig. 1), which increases the resolution of the output




Figure 1. 2D U-Net architecture

2. Metadata features: context through number of strokes in each side of the brain
3. Batchnorm for accelerating convergence and improving performance
4. Augmentation: Random flips and distortions

Evaluation

We measure performance with the dice coefficient, the most common metric used to evaluate segmentations of volumetric imaging data.

$DICE = \frac{2TP}{2TP + FP + FN}$, where TP, FP, FN represent the true positive, false positive, and false negative pixel counts of the lesion mask predicted by the algorithm compared to the ground-truth lesion mask.

Results

	Baseline	w/ Metadata	w/ Batchnorm
Conv-Deconv Network	0.11	0.126	0.142
2D UNet	0.095	Still running	0.105

Analysis

The baseline conv-deconv model does not seem to perform very well on its own. The addition of metadata and batchnorm provide a small but noticeable increase.

The 2D UNet seems to underperform the conv-deconv model, but this is likely because the model is sensitive to hyperparameters that were carefully tuned by the original paper. We would also expect it to benefit a lot from data augmentation, which we will try to do for our final report.