Multi-modality 3D brain tissue segmentation with a fully convolutional neural network

Ryan Sanford and D. Louis Collins

Abstract

Segmentation and quantitative assessment of brain tissue from MRI can help reveal important clues regarding the mechanisms that may underlie a neurological disease. Despite numerous efforts dedicated to brain tissue segmentation, implementing a fully automatic method that provides accurate and reliable segmentations remains a challenge. To address this issue, we propose to implement a fully convolutional neural network that can automatically segment brain tissue. Our method was inspired by an existing implementation, V-net, and trained using T_1, T_1-IR and FLAIR data from the MRBrainS dataset.

Methods

Pre-processing

We applied minimal pre-processing to the MRBrainS dataset since the all scans were already bias corrected with N4ITK[1] and the T_1 and T_1-IR scans were aligned with the corresponding FLAIR scan. We simply sharpened all scans by subtracting a Gaussian (σ=5) smoothed image from each modality and then histogram matching was used to account for large intensity variations across the datasets[2]. The intensities were normalized with zero mean and unit variance before feeding them into the network.

Network architecture

The network architecture was based on the V-net [3], which was developed to segment volumetric medical images. We chose this architecture because it consists of contracting and expanding paths that are useful localizing high-resolution features in the brain. There are residual functions at various stages of the network that ensure the convergence. The V-net was modified for brain tissue classification by training on all available multiple image modalities (T_1, T_1-IR and FLAIR). The final layer of the network is a linear mapping of 32 features to four segmentation labels (background, cerebrospinal fluid, gray matter and white matter) using a 1x1x1 convolution followed by a softmax activation function.

Implementation details

Our method was implemented in Python using Lasagne[4] and Theano[5]. We trained our network on 3D patches (size=64x64x64) from the T_1 (thick slice), T_1-IR and FLAIR (size=240x240x48) scans with an adaptive learning rate (ADADELTA) using a batch size of 8 on data from the MRBrainS 2013 and 2018 dataset. The network was trained on a NVIDIA TITAN X GPU with 12 GB of memory. On-the-fly data augmentation was implemented to enhance the training samples, where randomly selected data was randomly rotated 90˚ or 180˚ or flipped. All data augmentations were applied along the axial slice. It took about one day to train the network and less than 2 minutes to process each test volume.

Post-processing

Given that the final output consists of four segmentation labels, most notably a background label, it is possible that tissues inside the brain will be labeled as background. To account for this, we extracted the whole brain using BEaST [6] and replaced all voxels labelled as background within the brain mask with the most appropriate label (white matter, gray matter or cerebrospinal fluid) based on the softmax probabilities outputted from the network.

References

1. Tustison NJ, Avants BB, Cook PA, et al. N4ITK: Improved N3 Bias Correction. IEEE Transactions on Medical Imaging **2010**; 29(6): 1310-20.
2. Nyúl LG, Udupa JK, Zhang X. New variants of a method of MRI scale standardization. IEEE Transactions on Medical Imaging **2000**; 19(2): 143-50.
3. Milletari F, Navab N, Ahmadi SA. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV), 2016:565-71.
4. Dieleman S, Schlüter J, Raffel C, et al. Lasagne: First release. **2015**.
5. Team TD. Theano: A Python framework for fast computation of mathematical expressions. arXiv e-prints **2016**; abs/1605.02688.
6. Eskildsen SF, Coupe P, Fonov V, et al. BEaST: Brain Extraction based on nonlocal Segmentation Technique. Neuroimage **2012**; 59(3): 2362–73.