Ensembling 3D U-Nets For Ischemic Stroke Lesion Segmentation

Andrew Beers[1], Ken Chang[1], James Brown[1], Emmett Sartor[2], Elizabeth Gerstner[3], Bruce Rosen[1], Jayashree Kalpathy-Cramer[1]

[1] Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital and Harvard Medical School, Boston, MA, USA
[2] Massachusetts General Hospital, Boston, MA, USA
[3] Athinoula A. Martinos Center for Biomedical Imaging, Department of Neuro-Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, MA, USA

Abstract. The 2017 ISLES competition challenges participants to develop a fully automatic ischemic stroke lesion segmentation pipeline based on coregistered patient diffusion maps (DWI, ADC) and perfusion maps (CBV, CBF, MTT, TTP, TMAX). Participants must predict a manually-drawn and coregistered stroke lesion label delineated on a T2 or FLAIR image acquired after lesion stabilization. We propose a novel deep learning architecture based on the 3D Convolutional U-Net, an architecture that has found success both in ISLES 2016 and a wide array of other tissue segmentation applications. A typical U-Net segmentation architecture operates by convolving and downsampling input data stepwise into a low-resolution representation, and then upsampling and deconvolving that representation into to a categorical labelmap. The downsampling arm of the U-Net is also concatenated at points to the upsampling arm, resulting in a densely-connected architecture. We improve upon previous implementations of the 3D U-Net both by increasing the number of layers and convolutional filters, and by adding multiple independent down-sampling arms to the network. The motivation for this chimeric structure is to increase accuracy by concatenating several unique and not necessarily correlated downsampled representations, thereby increasing the potential amount of relevant imaging biomarkers. We apply this architecture on a stacked representation of six of the seven given image maps (ADC, CBV, CBF, MTT, TTP, TMAX), and post-process the output by removing small islands and applying repeated segmentation erosions and dilations. The entire segmentation is computed using Python, Keras, and Tensorflow, and will be made publically available on Github.

Keywords: 3D CNN, U-Net, Ensembling, Deep Learning, Patches, Stroke Lesion, ISLES, Infarction