A Deep-Learning Based Approach for Ischemic Stroke Lesion Outcome Prediction

Ramandeep Randhawa, Ankit Modi, Parag Jain, and Prashant Warier

rrandhaw@usc.edu, University of Southern California,
{ankit.modi, prashant.warier}@fractalanalytics.com, Fractal Analytics,
paragjain@topiciq.ai, TopicIQ Inc.

The ISLES 2016 challenge aims to address two important aspects of Ischemic stroke lesion treatment prediction. The first aspect relates to segmenting the brain MRI to identify the areas with lesions and the second aspect relates to predicting the actual clinical outcome in terms of the patient's degree of disability. The input data consists of acute MRI scans and additional clinical such as TICI scores, Time Since Stroke, and Time to Treatment.

To address this challenge we take a deep-learning based approach. In particular, we first focus on the segmentation task and use an automatic segmentation model that consists of a Deep Neural Network (DNN). The DNN takes as input the MRI images and outputs the segmented image, automatically learning the latent underlying features during the training process. The DNN architectures we consider utilize many convolutional layers with small kernels, e.g., 3x3. This approach requires fewer parameters to estimate, and allows one to learn and generalize from the somewhat limited amount of data that is provided.

One of the architectures we are currently utilizing is based on the U-Net [1], which is an all-convolutional network. It acts as an auto-encoder, that first "encodes" the input image by applying combinations of convolutional and pooling operations. This is followed by the "decoding" step that up-scales the encoded images, while performing convolutions. The all-convolutional architecture of the U-Net allows it to handle input images of different dimensions as in the challenge dataset. In our experiments, we found that this architecture yielded excellent performance on the previous ISLES 2015 dataset. Although the modalities in the 2016 challenge are different, our initial training experiments have yielded promising segmentation results.

Our next steps involve addressing the regression challenge. There is limited amount of labeled data for this task. Our approach will be to include these outcomes as part of the segmentation training directly. This will allow the DNN to learn latent features that can directly help with the classification task.

References

1. O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.