Loss Ensembles for Intracranial Aneurysm Segmentation:

An Embarrassingly Simple Method

Jun Ma*[1], Department of Mathematics, Nanjing University of Science and Technology
Xingle An, China Electronics Cloud Brain (Tianjin) Technology CO., LTD

1. Brief description of the method

We develop a 3D U-Net [1] based method for Intracranial Aneurysm Segmentation. To deal with the high imbalance between foreground and background, we test two different loss functions: Dice loss + Cross entropy vs Dice loss + TopK loss.

The employed architecture is a naïve 3D U-Net which consists of one down-sampling path and one up-sampling path. Each path includes 5 convolution blocks and each block comprises of 3*3*3 convolution layer, instance normalization layer and leaky rectified linear unit. Long skip connections are also used in the same resolution between down-sampling and up-sampling paths. The implementation is based on the well-known nnU-Net [2].

Only preprocessed TOF images are used to train the U-Net. The pre-processing includes foreground (non-zero regions) cropping and Z-Score normalization. All the 113 training cases are used for training. Data Augmentation includes random rotation, scaling, mirroring, and gamma transformation. The optimizer is SGD with an initial learning rate of 0.01. To speed up the inference time, we disable the default test time augmentation during inference.

We apply five-fold cross validation for the two different loss functions. We train all the models on NVIDIA TITAN V100 GPU. Each fold costs about 5 days. Following is the results in terms of Dice scores. It can be found that none of the loss achieves the best Dice scores among all five folds. Thus, we select the best-fold models for ensembles.

Table 1. Quantitative segmentation results (Dice) of different loss functions.

Fold	Dice loss + Cross entropy	Dice loss + TopK loss
0	0.4370	**0.4921**
1	**0.5476**	0.4888
2	**0.5108**	0.4926
3	**0.6173**	0.5998
4	0.4404	**0.5240**

[1] junma@njust.edu.cn

The code and trained model are publicly available at https://github.com/JunMa11/ADAM2020.

2. Inference time

30s/case on NVIDIA TITAN V100.

[1] Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016, October). 3D U-Net: learning dense volumetric segmentation from sparse annotation. In International conference on medical image computing and computer-assisted intervention (pp. 424-432). Springer, Cham.
[2] Isensee, F., Petersen, J., Kohl, S. A., Jäger, P. F., & Maier-Hein, K. H. (2020). Automated Design of Deep Learning Methods for Biomedical Image Segmentation. arXiv preprint arXiv:1904.08128.