Fully Convolutional Neural Network for 3D Stroke Lesion Segmentation

Miguel Monteiro[1][*] and Arlindo L. Oliveira[2]

[1] INESC-ID, Lisboa, Portugal, `mab.mtr@gmail.com`
[2] INESC-ID / IST, University of Lisbon, Lisboa, Portugal, `aml@inesc-id.pt`

The segmentation of stroke lesions is an important step for the assessment of stroke outcome. This problem can be particularly difficult if the size of the lesion is very small when compared with the size of healthy tissue. Our solution consists of what has been denoted as a U-shaped [1] or V-shaped [2] fully convolutional neural network. This type of architecture has been shown to perform well for other medical image segmentation problems. To address the fact that the lesions can be very small sized, we propose a new loss function to optimize our network. This loss function consists of the sum between the cross entropy and a "soft" dice coefficient, which is the dice coefficient calculated before binarizing the outputs. We found that this loss function leads to much faster convergence than the simple cross entropy loss or just the dice coefficient as proposed in [2]. For post processing we used a dense Conditional Random Field [3] which has been shown to improve the results of segmentation by taking advantage of the spatial information contained in the classifier's output and the original image. Since our architecture was computationally intensive our implementation was developed using TensorFlow and Google Cloud ML engine in order to take advantage of parallelized training across multiple machines with GPUs.

Acknowledgements

References

1. O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9351, pp. 234–241, 2015.
2. F. Milletari, N. Navab, and S. A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," *Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016*, pp. 565–571, 2016.
3. P. Krähenbühl, V. Koltun, P. Krahenbuhl, P. Krähenbühl, and V. Koltun, "Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials," *Nips'11*, no. 4, pp. 1–9, 2011.

---

[*] I am the corresponding author of the abstract and in the name of all co-authors I declare that MICCAI has the right to distribute the submitted material to MICCAI members and workshop, challenge and MICCAI conference attendees.