3D Medical Image Segmentation based on 3D Convolutional Neural Networks

Alejandra Márquez Herrera

September 20th 2018

Abstract

A neural network is a mathematical model that is able to perform a task automatically or semi-automatically after learning the human knowledge that we provide it. Moreover, a Convolutional Neural Network (CNN) is an specific type of neural network that is able to efficiently learn tasks related to the area of image analysis (among other areas). One example of these tasks is image segmentation, which aims to find regions or separable objects within an image. A more specific type of segmentation called semantic segmentation, makes sure that each region has a semantic meaning by giving it a label or class. Since neural networks can automate the task of semantic segmentation of images, they have been very useful in the medical area, as they can perform segmentation of organs or abnormalities (tumors) in medical images, which otherwise would be a difficult and tedious task for humans. There has been recent work using CNNs for performing semantic segmentation of 2D images [6, 8], and volumetric medical images such as the 3D U-Net, V-Net, among others [3, 7, 9].

Loss functions [10], are a fundamental part of these networks since they are in charge of measuring how well a model fits the data, resulting in a number that indicates the degree of error of our network, so that we can manage to minimize it.

Different types of loss functions have been used for semantic segmentation, such as the well known Cross-entropy loss and its different variations for improving segmentation results [3, 2, 1], the Malis loss [11], the Tversky index [5], among others.

The purpose of this work is to propose a modification in the loss function formulation of a 3D CNN based method for the semantic binary segmentation of volumetric medical images, using as a base model the architecture of the V-Net [7] and applying the idea of using similarity coefficients as loss functions for segmentation. A general view of the pipeline is shown in Figure 1.

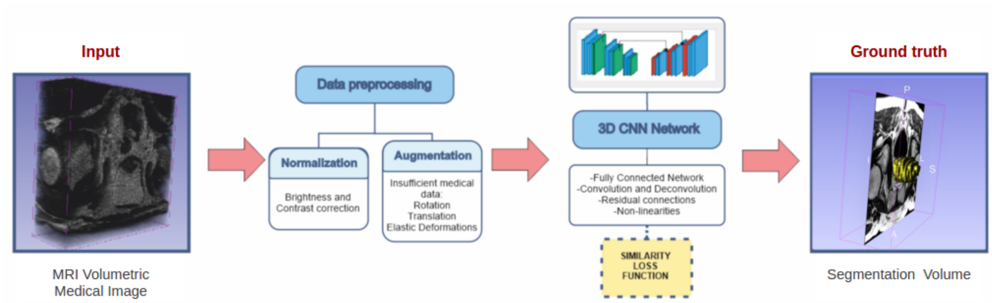

Figure 1: 3D medical image segmentation training pipeline

Even though we could think of modifying and adapting the classical Cross-entropy as an error function for a semantic segmentation problem, it might not always be efficient for performing pixel-wise classification in a segmentation context [11] .

That is why some other loss function alternatives for semantic segmentation emerge for improving the quality of the segmentation, taking advantage of the spatial context information of the pixels within an image, something that a Cross-entropy loss function does not do, for example. One of this alternatives is the idea of using similarity coefficients as actual loss functions. These coefficients, as its name suggests, are used for comparing the similarity between two objects and they have been adapted as loss functions in segmentation problems as a metric for comparing how similar a segmentation is to its ground truth.

One of the works that indroduced the idea of using a similarity coefficient as a loss function was introduced in [7], as the authors used a similarity coefficient called the Dice Similarity Coefficient and adapted it as a pixel-wise objective function. Some other work [12, 4], made some extensions of the loss function proposed in [7], and introduced modifications for the binary and multi-class cases. The authors proved that their loss functions performed better than the classical Cross-entropy for segmentation by proposing different ways of measuring the spatial overlap between the prediction and segmentation ground truth and then calculating the loss.

Based on these previous approaches, we propose the use of a metric called 'Kulckynsky coefficient' (Equation 1) which is a similarity coefficient based on operations using the True Positives (TP) and False Negatives (FN), as a loss function for guiding the semantic segmentation learning process of volumetrical medical images. We show results on how the choice of the loss function can affect the final quality of the segmentation.

$$\frac{1}{2} \times \left(\frac{\Theta_{TP}}{\theta_{TP} + \theta_{FP}} + \frac{\theta_{TP}}{\theta_{TP} + \theta_{FN}} \right) \tag{1}$$

For our experiments we used the PROMISE12[1] dataset, and we show comparisons in the quality of the final segmentation results obtained by training the network with some pre-existing loss functions such as the Cross-entropy loss, the Dice loss and the Wasserstein loss, and compare them against our proposed loss function. These comparisons are made in terms of the Dice Score.

[1] https://promise12.grand-challenge.org/

References

[1] Adrià Casamitjana, Santi Puch, Asier Aduriz, Elisa Sayrol, and Verónica Vilaplana. 3d convolutional networks for brain tumor segmentation. *Proceedings of the MICCAI Challenge on Multimodal Brain Tumor Image Segmentation (BRATS)*, pages 65–68, 2016.

[2] Patrick Ferdinand Christ, Florian Ettlinger, Felix Grün, Mohamed Ezzeldin A Elshaera, Jana Lipkova, Sebastian Schlecht, Freba Ahmaddy, Sunil Tatavarty, Marc Bickel, Patrick Bilic, et al. Automatic liver and tumor segmentation of ct and mri volumes using cascaded fully convolutional neural networks. *arXiv Preprint arXiv:1702.05970*, 2017.

[3] Ożgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016.

[4] Lucas Fidon, Wenqi Li, Luis C Garcia-Peraza-Herrera, Jinendra Ekanayake, Neil Kitchen, Sébastien Ourselin, and Tom Vercauteren. Generalised wasserstein dice score for imbalanced multi-class segmentation using holistic convolutional networks. *arXiv preprint arXiv:1707.00478*, 2017.

[5] Seyed Raein Hashemi, Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, Sanjay P. Prabhu, Simon K. Warfield, and Ali Gholipour. Tversky as a loss function for highly unbalanced image segmentation using 3d fully convolutional deep networks. *CoRR*, abs/1803.11078, 2018.

[6] Baris Kayalibay, Grady Jensen, and Patrick van der Smagt. CNN-based Segmentation of Medical Imaging Data. *arXiv preprint arXiv:1701.03056*, 2017.

[7] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 565–571. IEEE, 2016.

[8] Pim Moeskops, Jelmer M Wolterink, Bas HM van der Velden, Kenneth GA Gilhuijs, Tim Leiner, Max A Viergever, and Ivana Išgum. Deep learning for multi-task medical image segmentation in multiple modalities. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 478–486. Springer, 2016.

[9] Holger R Roth, Hirohisa Oda, Yuichiro Hayashi, Masahiro Oda, Natsuki Shimizu, Michitaka Fujiwara, Kazunari Misawa, and Kensaku Mori. Hierarchical 3D Fully Convolutional Networks for Multi-Organ Segmentation. *arXiv preprint arXiv:1704.06382*, 2017.

[10] Donald F Specht. Probabilistic neural networks. *Neural networks*, 3(1):109–118, 1990.

[11] Fabian Tschopp, Julien NP Martel, Srinivas C Turaga, Matthew Cook, and Jan Funke. Efficient convolutional neural networks for pixelwise classification on heterogeneous hardware systems. In *Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on*, pages 1225–1228. IEEE, 2016.

[12] Xin Yang, Lequan Yu, Shengli Li, Xu Wang, Na Wang, Jing Qin, Dong Ni, and Pheng-Ann Heng. Towards automatic semantic segmentation in volumetric ultrasound, 09 2017.