Inteneural at Aneurysm Detection And segMentation Challenge

Urszula Walińska 1, Michał Klimont M.D. 1,2, Marek Kraft Ph.D. 1,3, Dominik Pieczyński 1,3, Michał Mikołajczak 1, Mikołaj A. Pawlak M.D. Ph.D. 1,2

1. Inteneural Networks, Poznan, Poland; 2. Poznan University of Medical Sciences; 3. Poznan University of Technology

walinska@inteneural.com, mpawlak@ump.edu.pl

1. Introduction

ADAM challenge is a medical image analysis challenge as part of MICCAI 2020 (`http://adam.isi.uu.nl/`). The purpose of the challenge is to create a method of automatically detecting and segmenting unruptured intracranial aneurysms identified using Time of Flight Magnetic Resonance Angiographs (TOF-MRA). Our team constructed a solution consisting of three 2D U-Nets with pretrained EfficientNet [1] as backbone, which were subsequently fine-tuned for each axis (axial, coronal and sagittal). Finally, predictions were aggregated and further improved in post-processing. Additionally, method needed to be encapsulated in a Docker image. This poster briefly describes our solution.

2. Data preprocessing

In the preprocessing pipeline we first reorient each NIFTI image to Right-Anterior-Superior orientation and later resize the volume to match the voxel size of 0.4 mm in all dimensions. We compute skull-stripping masks to exclude all signal outside the brain. We use min-max normalization to obtain the values between 0 and 1. Subsequently, we equalize the shape of volume to be equal 512x512x512. Finally, we save only slices which contain some part of aneurysm as png files (in each axis) which we will use for training the model. We have combined slices of raw TOF scan and blood vessel segmentation as input (2 channels) and as gold standard we have aneurysm segmentation.

2.1. Skull-stripping and blood vessel segmentation

For the skull-stripping task a previously trained model was used. It was pretrained on a much bigger dataset and now was fine-tuned on 13 manually labeled examples from ADAM challenge dataset.
In order to enrich information used by the network to learn we developed automatic blood vessel segmentation tool. First, TOF images are preprocessed in the same way as described before. However, right before skull stripping, a new volume is derived by applying a threshold calculated as 99.5-th percentile of the original image and binarizing the response. Values above this threshold are now equal 1 and below are equal 0.
Subsequently, after performing skull stripping on a original volume a Jerman filter [2] is applied with subsequent parameter values: $sigmas\_range = (0.95, 1.3, 1.65, 2)$ – scales of image; $\tau = 0.5$ – Jerman threshold.
Finally, the result of applying a threshold (99.5-th percentile) and the result of applying Jerman filter are combined by using simply element-wise multiplication. In addition also binary dilation operation is performed (with default parameters) and a skull stripping mask is applied. Moreover, filtering out small volumes is performed namely, we keep only the biggest 4 ones unless they are smaller than a certain threshold.
The outcome of skull-stripping and blood vessel segmentation is presented below.

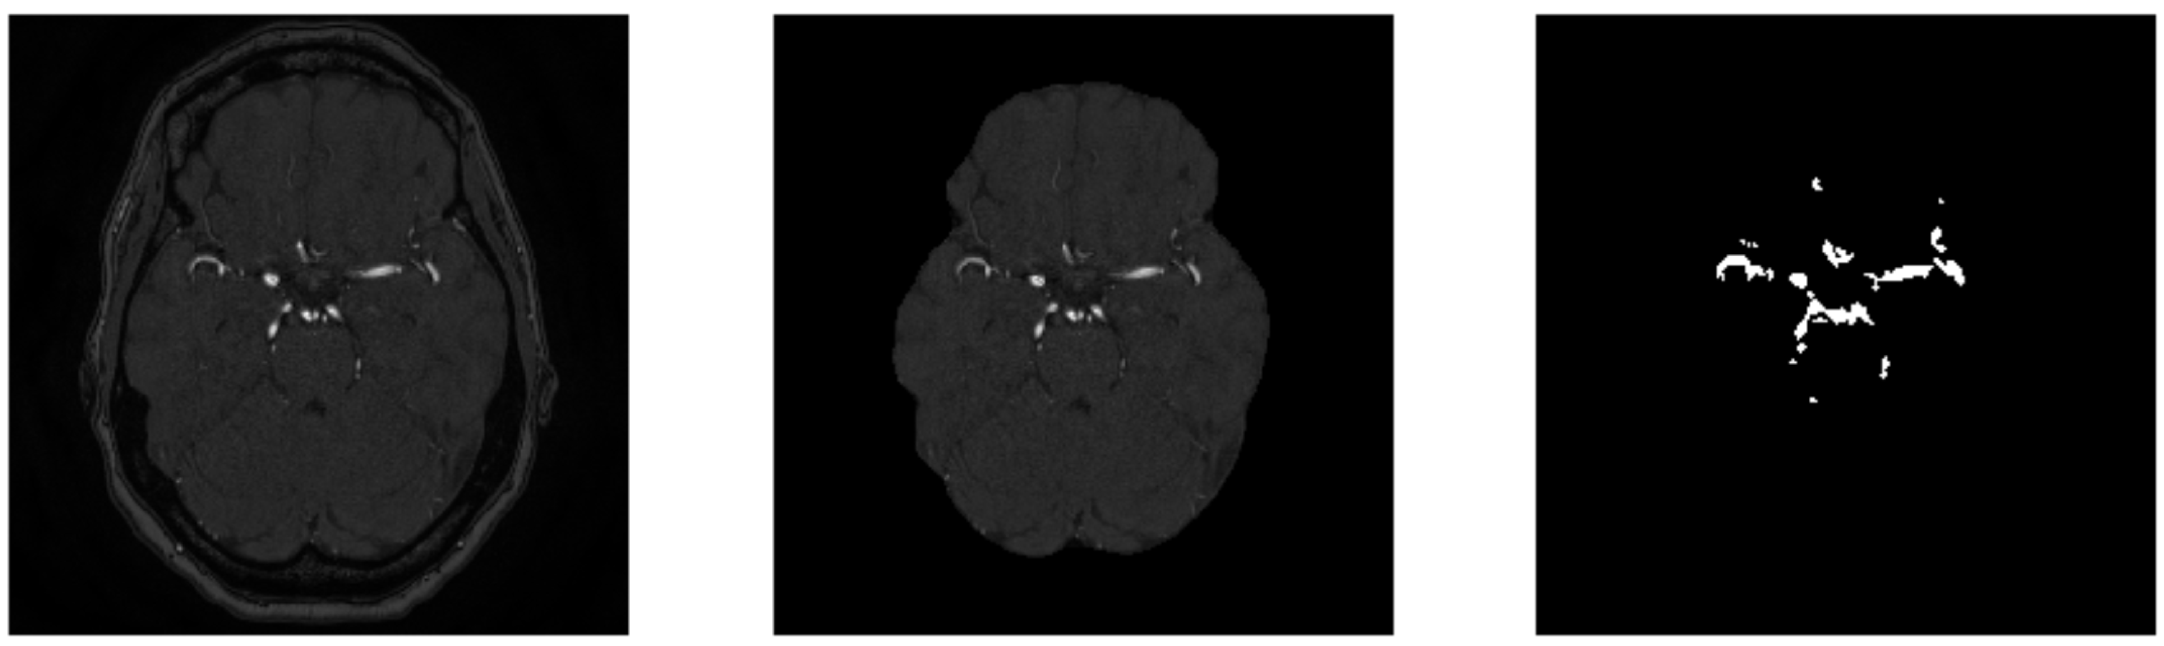

3. NN model

We use `segmentation_models` library for creating 3 already pre-trained NN models for aneurysm segmentation (one for each axis). Models' backbone is `efficientnetb1` which was trained on *ImageNet* benchmark.
In order to address a highly unbalanced segmentation problem we use generalized dice loss [3] and boundary loss [4] combined in the following way: $\alpha * GD\_loss + (1 - \alpha) * B\_loss$
The alpha coefficient was initially set to 1, and then decreased by 0.01 after each epoch (experiments showed this method yields the best results). Optimizer used for training the networks was Adam with default parameters. We split the dataset into training, validation and test in the proportion: 8:1:1. During the training we monitor loss on the validation set and keep track of the best model according to this value, in order to end up with the best model overall. Joint prediction is performed by averaging 3 individual models' predictions.

4. Post-processing and visualization

Finally, to further reduce false positive count we apply simple filtering, based on the knowledge of maximum aneurysm count in a single image and minimum volume of an aneurysm.
The outcome of described method is presented below. Original slice is on the left and the slice with aneurysm marked in red is on the right.

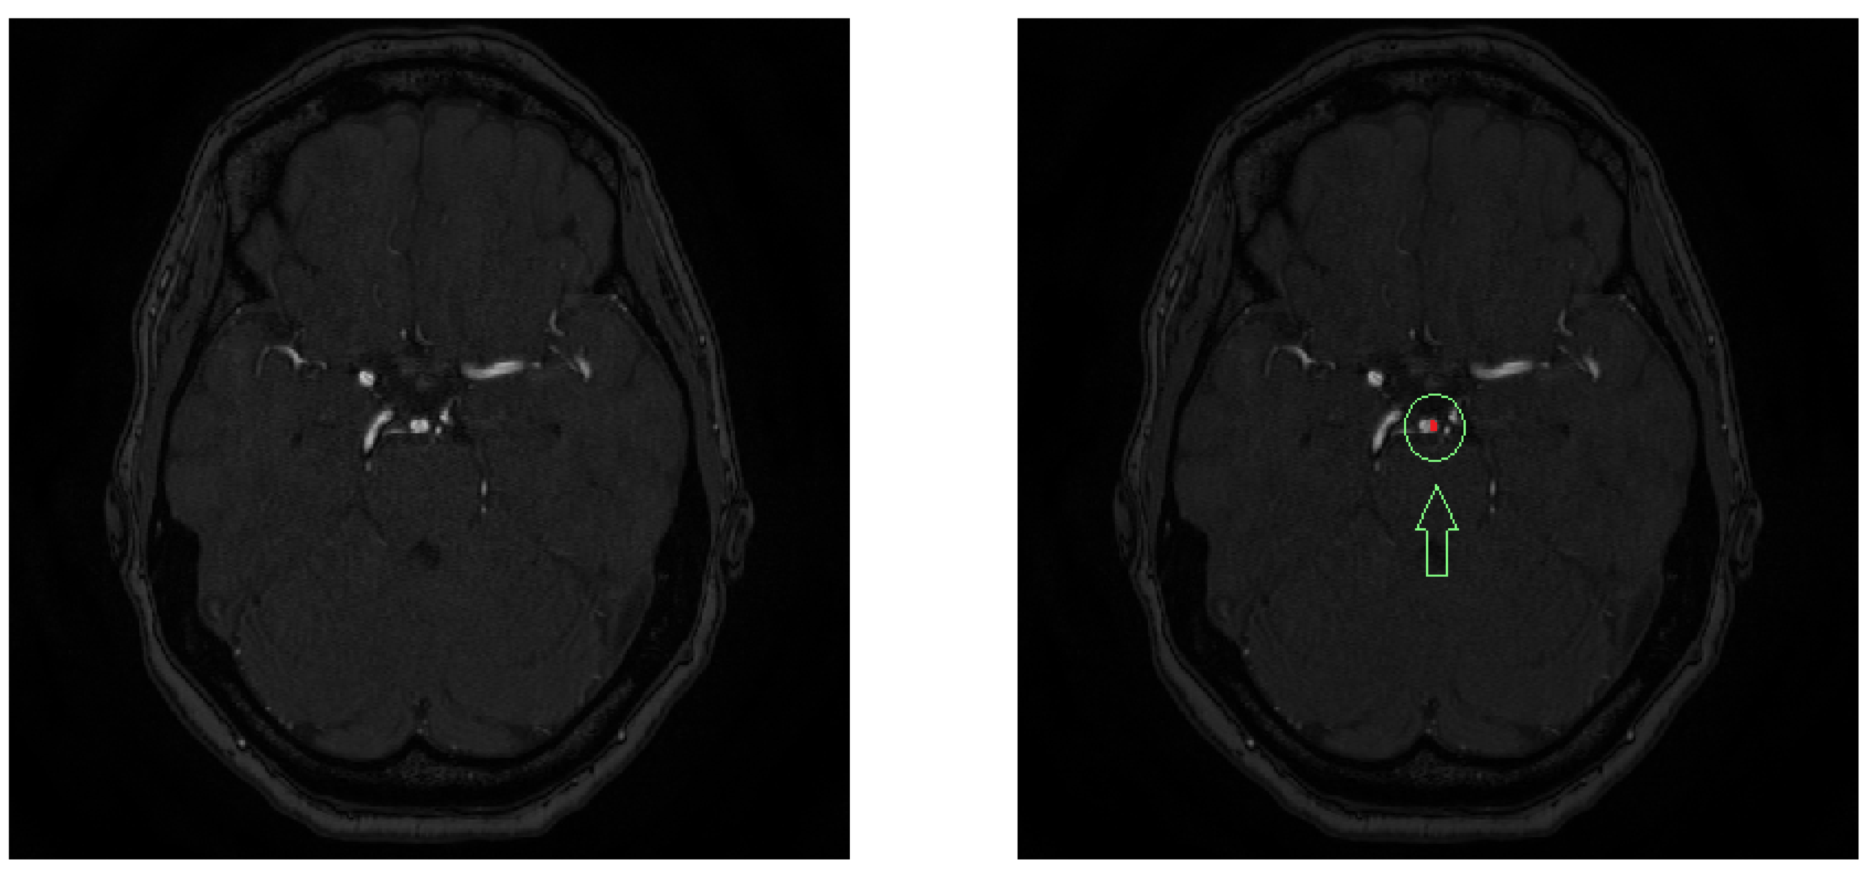

5. Results

Team: inteneural	Task 1 Rank:	0.17	Task 1 Place:	6 th
	Task 2 Rank:	0.39	Task 2 Place:	4 th

Task 1	False Positives	Sensitivity
Average	0.88	0.49
Rank	0.04	0.3

(lower rank is better)

Task 2	Dice Coefficient	Modified Hausdorff Distance (mm)	Volumetric Similarity
Average	0.17	23.98	0.36
Rank	0.61	0.27	0.29

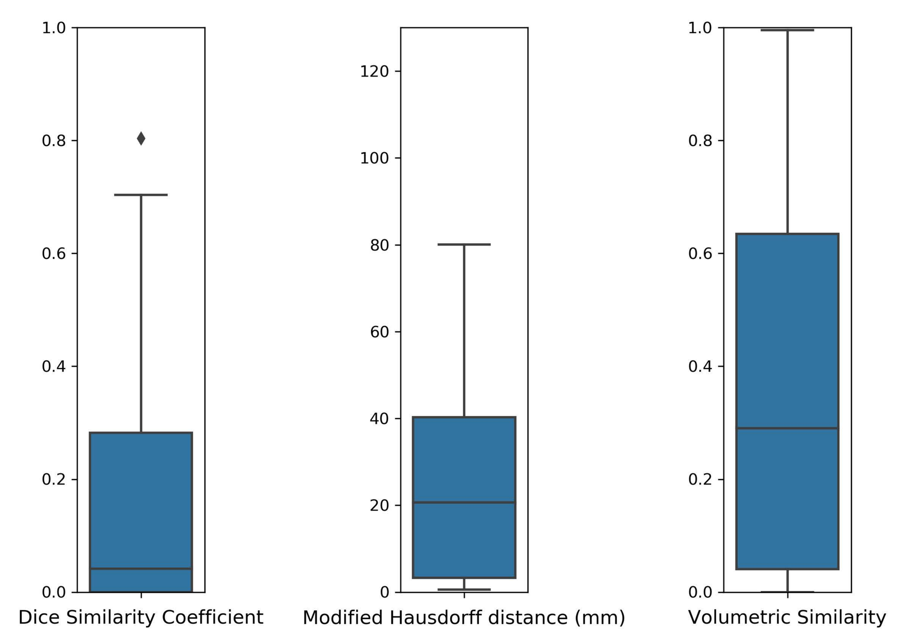

6. References

[1] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. 2019.

[2] Tim Jerman, Franjo Pernuš, Boštjan Likar, and Ziga Spiclin. Beyond frangi: An improved multiscale vesselness filter. 9413, 03 2015.

[3] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sébastien Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. 2017.

[4] Hoel Kervadec, Jihene Bouchtiba, Christian Desrosiers, Eric Granger, Jose Dolz, and Ismail Ben Ayed. Boundary loss for highly unbalanced segmentation. 2019.