Uncertainty-based Segmentation of Myocardial Infarction Areas on Cardiac MR images

Thomas Crozier*[1,2,3], Alexis Faure*[1,2,3], Daniel Bos[3,4], Marleen de Bruijne[1,3,5], and Robin Camarasa[1,3]

[1]Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands
[2]Ecole des Mines de Saint-Etienne, Saint-Etienne, France
[3]Department of Radiology and Nuclear Medicine, Erasmus MC, Rotterdam, The Netherlands
[4]Department of Epidemiology, Erasmus MC, Rotterdam, The Netherlands
[5]Department of Computer Science, University of Copenhagen, Copenhagen, Denmark
r.camarasa@erasmusmc.nl

Abstract. Every segmentation task is uncertain due to image resolution, artefacts, annotation protocol etc. Propagating those uncertainties in a segmentation pipeline can improve the segmentation. This article aims to assess if segmentation can benefit from uncertainty of an auxillary unsupervised task - the reconstruction of the input image. The method was applied to segmentation of myocardial infarction areas on cardiac magnetic resonance images.

1 Dataset

The presented methods were developed on the hundred Delayed Enhancement Magnetic Resonance (DEMR) exams of the segmentation dataset of the EMIDEC challenge [3]. From this data only the DEMR exams and the voxel-wise annotations of the five classes (background, cavity, normal myocardium, myocardial infarction and no-reflow) were used to tune the methods.

2 Methods

2.1 Method overview

Two methods, described in Fig. 1, are proposed for this challenge :

- **Baseline segmentation :** First a region of interest (ROI) U-net selects the ROI, a tile of shape $[n_x, n_y, n_z]$ centered on the four non-background classes. Then, a segmentation U-net segments the five classes of interest from the ROI. The selection of the ROI focuses the segmentation U-net on the non-background classes, which are the classes of interest.

*Equally contributed

– **Uncertainty-based segmentation :** The ROI is determined as in the baseline segmentation, using the ROI Unet. Then a probabilistic Auto-Encoder generates an uncertainty map of the reconstruction of the ROI. Finally, an uncertainty-based U-net segments the ROI from the uncertainty map and the ROI. This method is inspired from the work of Metha et al. [4], which demonstrates that cascading uncertainty in inference tasks can lead to improvement of downstream tasks.

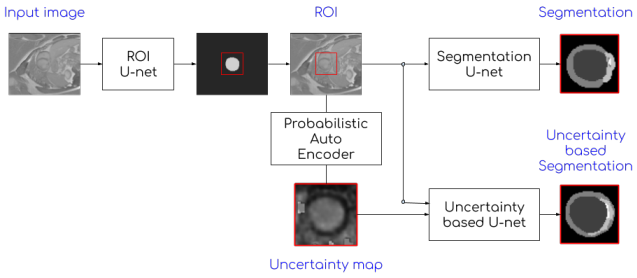

Fig. 1: Method overview

2.2 U-net

The different U-nets [6], described in this article, have n_{in} input channels, n_{out} output channels and depends on the number of features n_f as in Fig. 2. Note that due to the low resolution of the DEMR images in the z-axis compared to the other directions, the max-pooling and up-sampling are 2D instead of 3D [1].

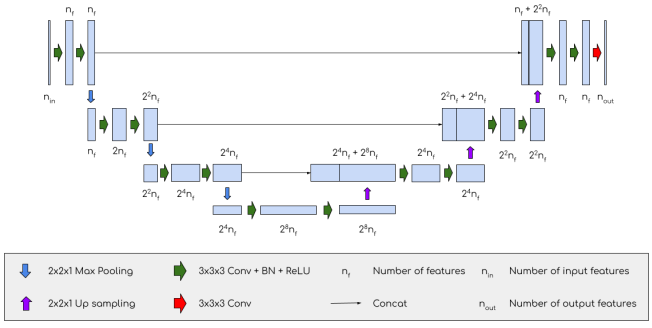

Fig. 2: U-Net structure

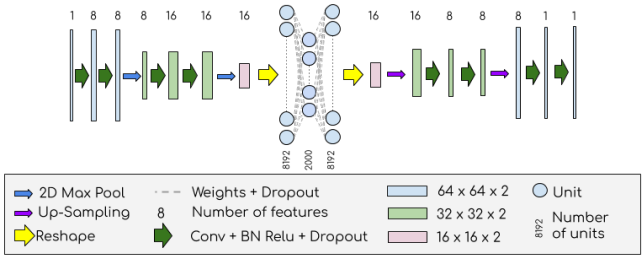

Fig. 3: Probabilistic Auto-Encoder structure

2.3 Probabilistic Auto-Encoder

The probabilistic Auto-Encoder, described in Fig. 3, uses Monte-Carlo dropout with a dropout rate p_d to compute the uncertainty map corresponding to ROI reconstruction [2]. Note that, similarly to the U-net, the up-sampling and max-pooling are 2D. To obtain several estimates of the reconstruction at test time, we sample $T = 25$ sets of parameters $(\theta_1, ..., \theta_T)$. From those T set of parameters, T outputs, $(f^{\theta_1}(x^t), ..., f^{\theta_T}(x^t))$, represent a sample of the region of interest reconstruction distribution $q(\hat{x}^t|x^t)$. From this sample, one can derive the variance of the output probabilities at voxel level, which is referred later on as the uncertainty map, as

$$\mathrm{Var}(q(\hat{x}_j^t|x^t)) = \frac{1}{T}\sum_{t=1}^{T} f_j^{\theta_t}(x^t)^2 - \mathbb{E}((q(\hat{x}_j^t|x^t))^2 \tag{1}$$

where f^θ is the auto-encoder with parameters θ, $q(\hat{x}_j^t|x)$ is the reconstructed distribution of the voxel j.

2.4 Data-augmentation

On the fly, three types of data-augmentation are performed in the following order: rotations, elastic deformations and flips. The rotation applied along the z-axis depends on an angle θ randomly sampled from $\mathcal{U}(\theta_{min}, \theta_{max})$. The random elastic deformation depends on a displacement grid of shape $[g_x, g_y, g_z]$ generated via the module elasticdeform[‡]. The values of the component of each displacement vector v_x, v_y and v_z are respectively sampled from the distribution $\mathcal{U}(-d_x, d_x)$, $\mathcal{U}(-d_y, d_y)$, and $\mathcal{U}(-d_z, d_z)$. Finally random flips along x and y axis are applied with a probability p. After data-augmentation, the resulting image intensities are linearly rescaled such that the minimum and the maximum are set to 0 and 1 respectively.

[‡]https://github.com/gvtulder/elasticdeform

Table 1: Implementation details

Parameters	ROI U-Net	Auto-Encoder	Baseline U-Net	Uncertainty-based U-Net
n_{in}	1	1	1	2
n_{out}	1	1	5	5
n_f	8	8	16	16
batch size	5	5	5	5
loss	Dice [5]	$\mathcal{L}_1$	Averaged dice [5]	Averaged dice [5]
optimizer	Adadelta [7]	Adadelta [7]	Adadelta [7]	Adadelta [7]
$[n_x, n_y, n_z]$	[104, 176, 8]	[64, 64, 2]	[64, 64, 8]	[64, 64, 8]
$[g_x, g_y, g_z]$			[20, 20, 3]	
$[d_x, d_y, d_z]$			[1, 1, 0.1]	
$[\theta_{min}, \theta_{max}]$			[-10, 10]	
p	0.5	0.5	0.5	0.5
p_d		0.1		
p_c			0.5	0.5

2.5 ROI U-Net

This U-net is trained to segment a binary mask of the Myocardial areas. Then the center of mass of the four largest connected components of the prediction determine the predicted center of the ROI. All input DEMR images were resized to the network input size $[n_x, n_y, n_z]$.

2.6 Segmentation U-net

This U-net is trained on tiles of shape $[n_x, n_y, n_z]$ containing the segmentation. At train time, a proportion p_c of the training sample had the ground truth segmentation centered.

2.7 Uncertainty map based U-net

This network is comparable to the segmentation U-net but takes a tile of shape $[n_x, n_y, n_z]$ of the input image and of the uncertainty map as input.

2.8 Parameters

The parameters and the implementation details of the networks are described in Tab. 1. Note that baseline U-net uses elastic deformations where uncertainty-based U-net did not.

2.9 Ensemble

The training dataset is randomly split in four subsets. Then, for all networks, four models are trained, each using three of the subsets as training set and the remaining subset as validation set. Segmentations on the test images are then obtained as the voxel-wise mean of the outputs of these four models.

Table 2: Performance of the baseline and the uncertainty-based method. (The best observed results per metrics are in bold)

	Baseline method	Uncertainty-based method
Myocardium		
Dice index (%)	**75.74**	68.86
Haussdorf distance (mm)	**25.44**	35.15
Volume difference (mm^3)	**17108.13**	19646.13
Infarction		
Dice index (%)	**30.79**	21.63
Volume difference (mm^3)	**4868.56**	7476.56
Volume difference ratio according to volume of myocardium (%)	**3.64**	6.09
No-reflow		
Dice index (%)	**60.52**	60.00
Volume difference (mm^3)	**867.86**	944.92
Volume difference ratio according to volume of myocardium (%)	**0.52**	0.57

3 Results

Challenge organizers evaluated both the uncertainty-based and the baseline method on an independent test set of fifty patients. The evaluated metrics vary from class to class: Dice index, Haussdorf distance and volume difference assessed the quality of the segmentation of myocardium; Dice index, volume difference and volume difference ratio according to volume of myocardium assessed the quality of the segmentation of infarction and no-reflow areas (table 2). Both the uncertainty-based and the baseline method segment the myocardium better than the no-reflow areas and the no-reflow areas better than infarction areas. For every metric, the baseline method out-performed the uncertainty-based method.

4 Discussion

We developed and compared two methods to segment myocardial infarction areas using deep learning. At the time of writing this article, the results cannot be compared to other methods as the results of the EMIDEC challenge are not released yet.

Two elements explain partly the poor performances of the uncertainty-based method compared to the baseline method. First, the comparison is not entirely fair as the implementation of the data augmentation of the uncertainty-based method did not include elastic deformations (Table 1). Second, fine-tuning the probabilistic auto-encoder properly would have required more time. A more in depth analysis would be required to draw definitive conclusion on the comparison of both methods.

5 Acknowledgments

This work was partly funded by Netherlands Organisation for Scientific Research (NWO) VICI project VI.C.182.042.

References

1. Camarasa, R., Bos, D., Hendrikse, J., Nederkoorn, P., Kooi, E., van der Lugt, A., de Bruijne, M.: Quantitative Comparison of Monte-Carlo Dropout Uncertainty Measures for Multi-Class Segmentation (2020), uncertainty for Safe Utilization of Machine Learning in Medical Imaging (UNSURE) workshop of MICCAI conference
2. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: international conference on machine learning. pp. 1050–1059 (2016)
3. Lalande, A., Chen, Z., Decourselle, T., Qayyum, A., Pommier, T., Lorgis, L., de la Rosa, E., Cochet, A., Cottin, Y., Ginhac, D., Salomon, M., Raphal Couturier Meriaudeau, F.: Emidec: A Database Usable for the Automatic Evaluation of Myocardial Infarction from Delayed-Enhancement Cardiac MRI. Data journal (2020)
4. Mehta, R., Christinck, T., Nair, T., Lemaitre, P., Arnold, D., Arbel, T.: Propagating Uncertainty Across Cascaded Medical Imaging Tasks for Improved Deep Learning Inference. In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures, pp. 23–32. Springer (2019)
5. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). pp. 565–571. IEEE (2016)
6. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
7. Zeiler, M.D.: Adadelta: an adaptive learning rate method. arXiv preprint arXiv:1212.5701 (2012)