Misclassification Loss for Segmentation of the Aortic Vessel Tree

Abbas Khan[1,2], Muhammad Asad[3], Alexander Zolotarev[2,4], Caroline Roney[2,4], Anthony Mathur[6], Martin Benning[2,5], and Gregory Slabaugh*[1,2]

[1] School of Electronic Engineering and Computer Science, Queen Mary University of London, United Kingdom (UK) g.slabaugh@qmul.ac.uk
[2] Queen Mary's Digital Environment Research Institute, UK
[3] School of Biomedical Engineering and Imaging Sciences, King's College London, UK
[4] School of Engineering and Materials Science, Queen Mary University of London, UK
[5] School of Mathematical Sciences, Queen Mary University of London, UK
[6] Centre for Cardiovascular Medicine and Devices, William Harvey Research Institute, Queen Mary University of London, London, UK

Abstract. Common pixel-based loss functions for image segmentation struggle with the fine-scale structures often found in the aortic vessel tree. In this paper, we propose a Misclassification Loss (MC loss) function, which can effectively suppress false positives and rescue the false negatives. A differentiable eXclusive OR (XOR) operation is implemented to identify these false predictions, which are then minimized through a cross-entropy loss. The proposed MC loss helps the network achieve better performance by focusing on these difficult regions. On the Segmentation of the Aorta SEG.A. 2023 challenge, our method achieves a Dice score of 0.93 and a Hausdorff Distance (HD) of 3.50 mm on a 5-fold split of 56 training subjects. We participated in the SEG.A. 2023 challenge, and the proposed method ranks among the top-six approaches in the validation phase-1. The pre-trained models, source code, and implementation will be made public.

Keywords: Segmentation · Aortic vessel trees · False Prediction Loss function for segmentation

1 Introduction

The heart supplies blood to all body parts through the aorta, which is an important anatomy for cardiovascular disease diagnosis. Aortic disorders are considered a significant cause of death for heart patients [15]. The initial screening of the aorta and all its related branches, collectively known as the aortic vessel tree (AVT), is performed using computed tomography (CT). To monitor the diseases associated with the aorta, the AVT structure is identified, and changes are analyzed over time [13]. However, manual reconstruction of the AVT is a

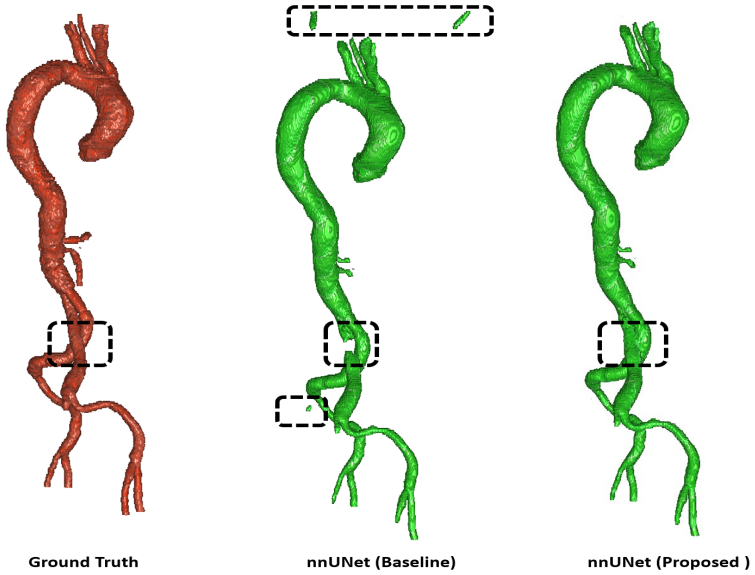

Fig. 1. Visual comparison of our proposed approach (a) Ground Truth (b) Segmentation from nnUNet trained on a Dice cross-entropy loss, and (c) Segmentation from our proposed approach. Please zoom in for details.

cumbersome and time-consuming task for clinicians. To overcome these difficulties, the Segmentation of the Aorta SEG.A. 2023 challenge [7] was proposed for bringing automated AI-based aorta segmentation methods into clinical practice [18].

Related work In this section, we briefly describe key approaches that utilize deep learning to segment the aorta. A cascaded two-stage convolutional neural network CNN-based approach is proposed by [14] to segment both the aorta and the true lumen. The CNNs extract features from the whole aortic lumen and true lumen, while the false lumen is reconstructed via a post-processing step. Cao et al. [3] proposed a serial multi-task model to segment the whole aorta, true lumen and false lumen. Both networks are 3D and use down-sampled images of size 128 $\times$ 128 $\times$ 256 to reduce computational complexity. Another multi-stage network is proposed by [5], where the first CNN coarsely segments the aorta, followed by three single-view CNNs to fine-segment the aorta along the axial, sagittal, and coronal planes. Finally, multi-view integration is performed to merge the predictions.

Segmentation of fine-scale structures is a challenging problem and is being addressed by many researchers. For example, clDice [22] finds the skeleton over

[7] https://multicenteraorta.grand-challenge.org/

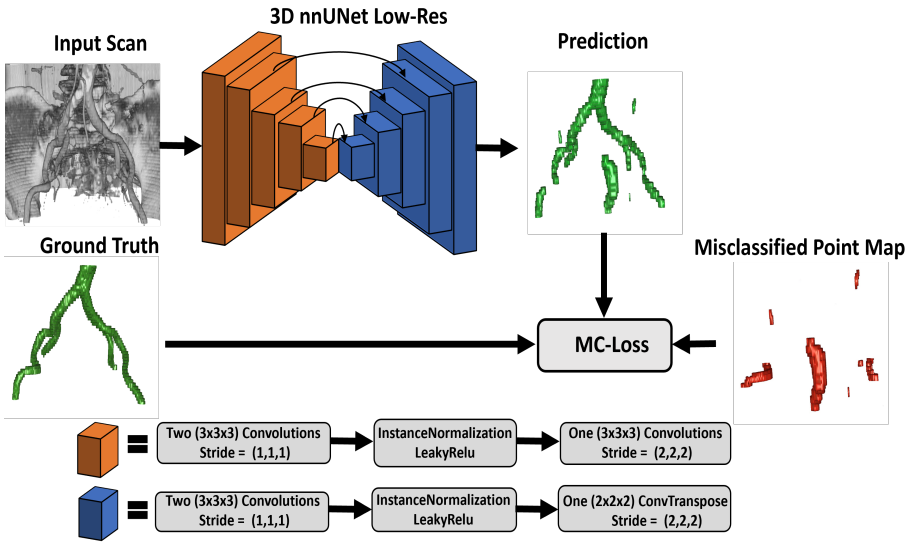

Fig. 2. Overview of our method. The proposed Misclassification (MC) Loss identifies the false positives and false negatives (misclassified points map). The MC Loss is defined based on this misclassified points map. Please zoom in for details.

the likelihood map. TopoNet [11] is based on persistent homology [4]; the identified erroneous predicted points from the likelihood map are re-weighted in the training loss. This method generated many irrelevant points and became very expensive in terms of training. DMT loss [9] is another skeleton-based loss and uses discrete Morse theory to find skeletons and related patches. Agata et al. [16] trained both UNet[21] and pre-trained VVG-16[23] iteratively to produce progressive refinements. The homotopy warping loss [10] uses distance transforms to find incorrect predictions and then incorporates them into the training to further refine the segmentation.

In this paper, we propose a plug-and-play loss function, which uses a simple yet effective approach to find falsely predicted voxels (false positives and false negatives). A differentiable XOR-gate is implemented between predicted foreground voxels and ground truth to find all points where the network predicts wrong voxels, i.e. false positives and false negatives collectively producing a *misclassified points map (MPM)*. As a second step, we map these points to the output logits and then fine-tune the segmentation network on misclassified voxels. The MC loss is meant to be used in conjunction with another loss (in our case, the Dice cross-entropy loss) that encourages true predictions (i.e. true positives and true negatives). It is therefore important to balance the MC loss and the original loss. When properly balanced, the MC loss can effectively remove the false positives and retrieve false negatives, as shown in Figure 1 compared to the original loss.

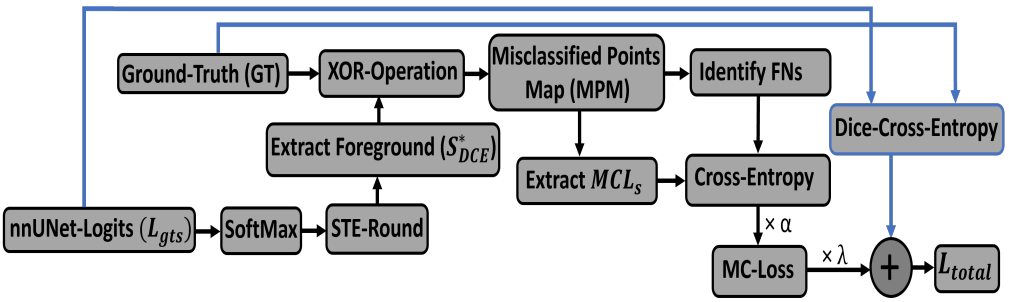

Fig. 3. Flowchart of proposed MC Loss. The blue color arrows represent the flow of Dice cross-entropy loss which is the first term of Equation 1, and the black color arrows show the different steps involved in the formulation of the proposed MC Loss (second term of Equation 1).

2 Methods

The proposed MC loss can be used with any segmentation network; Figure 2 illustrates the overall idea of incorporating it in an nnUNet network. The explanation of computing the proposed MC-loss is shown in Figure 3.

By identifying the misclassified voxels (false positives and false negatives) from the predictions of a trained segmentation network, we can fine-tune the same network by penalizing the misclassified voxels. This fine-tuning stage will reduce the model's errors without impairing the true positives and true negatives.

As a first step, we trained a nnUNet [12] with a combined Dice cross-entropy loss (L_{DCE}) to generate the initial segmentation map. As a second stage of training, we use the weights of this trained nnUNet and penalize it with the sum of L_{DCE} and proposed misclassification loss (L_{MC}), i.e.

$$L_{total} = L_{DCE} + \lambda \cdot L_{MC} \tag{1}$$

where λ is a positive balancing factor influencing the effect of two losses. A very high λ will force the network solely to focus on difficult points and misclassify the easy voxels already segmented in the first stage. A very low value of λ will not affect the second-stage training, and the network will retain its previous weights. Here, we argue that the value of λ varies from dataset to dataset and for the SEG.A. data we used λ = 5e-06. A detailed ablation study for choosing this balancing factor is provided in Section 3.1.

2.1 nnUNet Architecture

We use a 3D low-res nnUNet [12] network with 128 × 128 × 128 input dimensions, as shown in Figure 2. The maximum number of features is set to 320, and

two convolutional layers with a filter size of 3×3×3 were used at each stage of the encoder and decoder. Instance normalization and leaky Relu follow each convolution layer. Skip connections are employed at each encoder stage for better back-propagation of gradients. The features are downsampled with a convolution stride of 2 at the encoder side, and transposed convolutions are used at each decoder stage for upsampling and to match the output dimension with the input. The final segmentation map is generated using a $1 \times 1 \times 1$ convolutional filter followed by softmax activation.

2.2 Misclassification loss

This section explains how we formulated the proposed misclassification loss (MC loss); we also provide Figure 3 to illustrate the flow of different operations involved. The initial segmentation mask (S_{DCE}) is obtained from the logits (L_{gts}) of the trained nnUNet using softmax, followed by a round operation as follows:

$$S_{DCE} = \lfloor SoftMax(L_{gts}) \rceil \tag{2}$$

where $\lfloor \rceil$ represents the round operation. As the derivative for the round operation is not defined, we utilize a straight-through estimator (STE) to enable gradient forwarding through the round operation in the backward pass [2]. In this way, we can utilize the round operation within the chain rule and still be able to backpropagate from the proposed loss function. The ablation study in Section 3.1 explains the motivation for using the differentiable way of extracting the segmentation map from the logits.

The initial segmentation S_{DCE} contains the foreground and background segmentation mask. To obtain the misclassified points map (MPM) we only need the foreground segmentation mask denoted by S^*_{DCE}. A differentiable eXclusive OR (XOR) operation is performed between the foreground of the segmentation mask (S^*_{DCE}) and the ground-truth (GT) to obtain the misclassified points map.

$$MPM = S^*_{DCE} \otimes GT = S^*_{DCE} + GT - 2 \times S^*_{DCE} \cdot GT \tag{3}$$

where $\otimes$ represents the XOR operation.
The MPM contains all the voxels where the network was unable to predict the correct label for either foreground or background. This map includes false positives (FPs) and false negatives (FNs); we can further process the MPM and extract each of them as follows.

$$FNs = MPM \cdot GT \tag{4}$$

$$FPs = MPM \cdot (1 - GT) \tag{5}$$

At this stage, we have identified the errors made by the network; based on this, we can generate a new input and target on the fly to focus only on the misclassified logits (MCLs) as below:

$$MCLs = MPM \cdot L_{gts} \tag{6}$$

A cross-entropy loss is computed between misclassified logits (MCLs) and the false negatives (FNs). The FNs already contain the locations of FPs, with their corresponding values as zeros. So, the loss between MCLs and FNs tries to learn all FPs voxels in MCLs as background and all FNS voxels in MCLs as foreground. i.e.

$$L_{CE}(MCLs, FNs) = -(FNs \log(MCLs) + (1 - FNs) \log(1 - MCLs)) \tag{7}$$

Finally, the MC loss is calculated by multiplying $L_{CE}(MCLs, FNs)$ with the number of non-zero voxels in the misclassified points map denoted as α, which is the total number of false positives and false negatives.

$$L_{MC} = L_{CE}(MCLs, FNs) \cdot \alpha \tag{8}$$

3 Experiments

Dataset Description The dataset is provided by the SEG.A. MICCAI-2023 challenge, which intends to develop automatic AVT segmentation algorithms. The dataset comprises 56 computed tomography angiography (CTA) scans [20]. The scans are collected from three different locations and various hospitals, which include KiTS[7][8], RIDER [1], and Dongyang Hospital. The dataset comes with an in-plane spatial resolution of 512×512 and 512×666 and with varying numbers of Axial slices depending on the location and hospital. The dataset includes both healthy and patients with aortic dissections (AD), and abdominal aortic aneurysms (AAA). A semi-automatic method was used for AVT annotation, where a manual local threshold is first applied followed by an automatic GrowCut region growing algorithm [20]. 3D Slicer[8] was used for performing the above steps to acquire AVT segmentation ground truth. A more detailed overview of the dataset is provided in [20].

Training Details The proposed framework was implemented using PyTorch (1.11.0), and all the experiments were conducted using an NVIDIA A100 GPU with 40GB RAM [25]. As a pre-processing step, the CTA scans are resampled to a resolution of $1.88 \times 1.88 \times 3.67\ mm^3$. Each scan was normalized using z-score normalization. Different data augmentation strategies were utilized: including rotation, shifting, scaling, elastic deformation, Gaussian noise, Gaussian blur, and random bias field. We randomly shuffled the 56 samples and evaluated all models with 5-fold cross-validation in the validation phase, as suggested by the organizers. We also validated the proposed algorithm on the challenge validation set.

[8] https://www.slicer.org/

Table 1. Differentiable vs non-differentiable quantitative results for SEG.A. challenge data using 5-fold cross-validation split. The best results are in **bold**.

Differentiable XOR-Operation	**N**on-Differentiable XOR-Operation	**D**ifferentiable Segmentation Extraction	**N**on-Differentiable Segmentation Extraction	Training time (Hours) ↓	Dice Score ↑
✗	✓	✗	✓	75.7	0.927
✗	✓	✓	✗	74.8	0.928
✓	✗	✗	✓	74.9	0.931
✓	✗	✓	✗	**73.6**	**0.934**

3.1 Ablation Studies

We conducted three ablation studies to explore the different techniques contributing to the MC loss function. The experiments for the first ablation study (differentiable vs. non-differentiable components) were performed using 5-fold cross-validation, while the remaining two studies were conducted using only the first fold. For all the tables in the paper, the upward arrow (↑) indicates that higher is better while the downward arrow (↓) shows lower is better.

Differentiable vs. Non-differentiable components The proposed MC loss has two important components that can be implemented using either differentiable or non-differentiable techniques. (1) extraction of the segmentation map from predicted logits, and (2) the XOR operation to generate the misclassified points map. The argmax operation is utilized to get the indices of the maximum value from the logits. However, this operation is not differentiable, which limits its use within loss function design. In order to get a differentiable conversion of logits to segmentation, we replace argmax with a chain of operations that convert logits to segmentation in the forward pass and still have gradients for the backward pass. We achieve this by first applying softmax, followed by round operation. As the round operation is also non-differentiable, we utilize a straight-through estimator (STE) for its gradient propagation in the backward pass [2].

Logic gates play an important role in designing efficient neural networks. Conventionally, logic gates are non-differentiable and do not allow training with gradient descent. Inspired by [19] which proposes differentiable logic gate ideas, our MC loss employs a differentiable XOR operation to enable the back-propagation of gradients.

Table 1 showcases the outcome of this ablation study in terms of training time, and Dice score. All the experiments are performed using a single NVIDIA A100 GPU with 40GB RAM [25]. The training was performed for 1000 epochs, and the training time mentioned is for all five folds. We found that using differentiable components for the MC loss can reduce the training time $\sim$ 2 hours with improved Dice Score. Replacing the argmax operation with our proposed differentiable segmentation extraction method (given by Equation 2) reduces the training time by $\sim$ 1 hour.

Table 2. Choice of loss function to define the MC Loss between MCLs and FNs .

Loss Function	Dice Score↑
Mean square error	0.914
Dice Loss (MCLs, FNs) with mask = None	0.916
Dice Loss (L_{gts},GT) with mask = MPM	0.920
Cross-Entropy Loss	**0.924**

Choice of loss function for MC loss After identifying the misclassified logits, we can use a variety of loss functions to define the proposed MC loss. We experimented with four different losses here, mean square error (MSE), Dice loss, and Dice loss by using MPM as a mask and cross-entropy loss. Table 2 lists the quantitative results for the choice of the loss function to implement MC loss using data from the first fold. Empirically, we found that cross-entropy loss outperforms the others for the SEG.A. 2023 challenge dataset. However, we argue that it might depend on the dataset as well, so for different datasets, similar experiments can be repeated in future work.

Choice of the balancing factor λ As given by Equation 1, the proposed MC loss is trained in combination with L_{DCE}. Also, as indicated by Equation 8 the MC loss is multiplied by the number of non-zero voxels in the misclassified points map denoted by α, so the scale of the MC Loss and L_{DCE} are not the same. In order to bring them to the same scale we used the balancing factor λ (introduced in section 2). λ also controls the contributions of each term in Equation 1. We empirically found that using a large λ reduces the overall performance, as the network overly focusses on the misclassified points map, and the True Positives

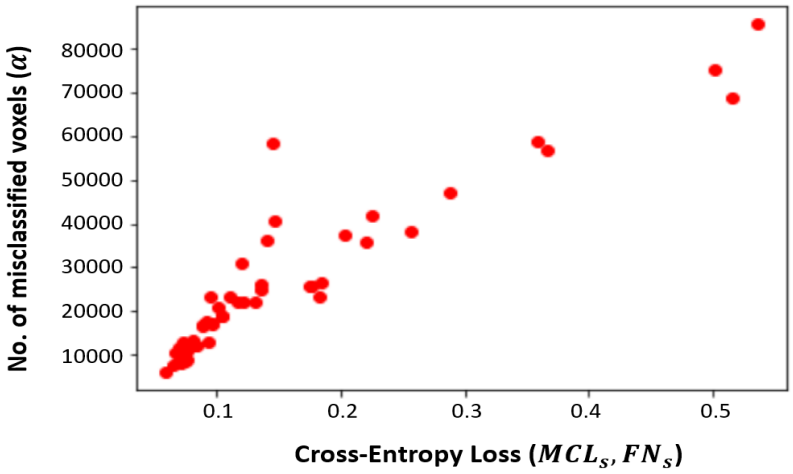

Fig. 4. Variation of cross-entropy loss against the number of non-zero voxels in the misclassified points map α.

Table 3. Ablation study for choosing the balancing factor λ to weight the effect of MC loss and L_{DCE}.

Value of balancing factor λ	Dice Score ↑	True Positives ↑	False Positives ↓	False Negatives ↓
0×10^{-5}	0.911	441,108	29,901	55,971
1×10^{-5}	0.918	449,123	26,189	53,359
2×10^{-5}	0.919	448,956	25,780	52,935
5×10^{-5}	**0.924**	**449,978**	**23,150**	50,280
8×10^{-5}	0.922	448989	25,761	**50,140**
1×10^{-4}	0.912	448958	28,254	57,989

(TPs), identified in the first training stage are affected, as shown in Table 3. We can also observe from the table that, using a very low value has no effect on the previously obtained performance. Figure 4 shows how the MC loss varies for network predictions having a different number of misclassified voxels. The graph is plotted for the validation samples of the first fold.

3.2 Results and Discussions

The experimental results are compared with the nnUNet baseline (trained with Dice cross entropy loss and without the proposed loss) and a UNet architecture. The results are evaluated using the Dice score and Hausdorff Distance (HD) [17]. Table 4 lists the quantitative comparison on the random 5-fold cross-validation split. The proposed method outperforms the nnUNet baseline and vanilla UNet both in terms of Dice score and HD. Although the improvement in terms of Dice score and HD is 1%, it highlights the importance of the proposed MC loss to retrieve the overall structure of the AVT, as shown in Figure 1. It can remove FPs that are not part of the main fine-scale structure and can recover the FNs as well. We found that these types of connections are very important in clinical applications where blood flow is estimated and requires flow distribution analysis in the entire aorta [6], [24]. All these tasks require segmenting the Aorta as a preliminary step in order to generate a tetrahedral mesh.

Additional visual results are provided in Figure 5 showcasing axial, coronal, and sagittal views of a computed tomography angiography (CTA) scan. Different colors are assigned to differentiate between predictions (red), ground truth (green), and overlap of predictions and ground truth (yellow). These correspond to true positives (yellow), false positives (red), and false negative (green) predictions.

The last column of Figure 5 demonstrates that the proposed MC loss can recover the FNs missed by the first stage training (row-1 and row-3), also it is evident that it can remove the FPs from row-2 of Figure 5.

4 Limitations

This section presents limitations of the proposed method, including cases where the MC loss does not yield an improvement in segmenting AVT. In some cases,

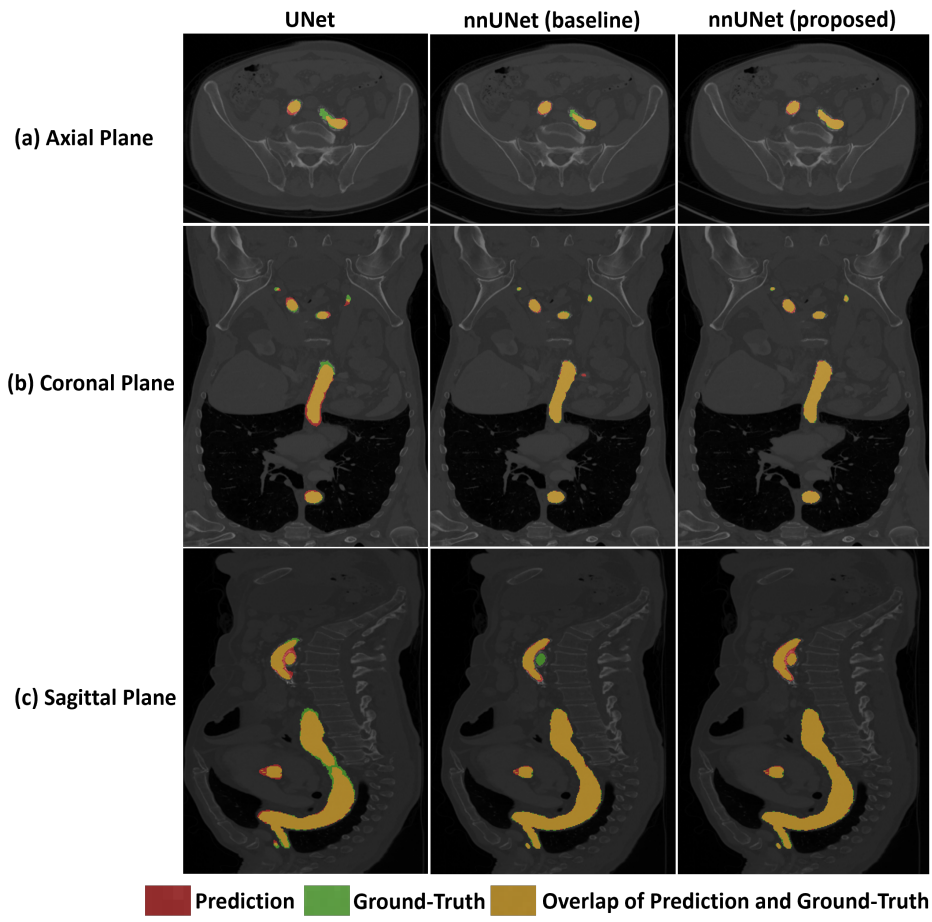

Fig. 5. Qualitative results of AVT segmentation along different views, (a) axial, (b) coronal, and (c) sagittal plane.

Table 4. Quantitative results for SEG.A. challenge data using 5-fold cross-validation split. The best results are in **bold**.

Segmentation Network	Loss function	Dice Score↑	Hausdorff Distance (mm) ↓
UNet	Dice-cross-Entropy	0.886	7.62
nnUNet	Dice-cross-Entropy	0.921	4.33
nnUNet	Dice-cross-Entropy + MC Loss	**0.934**	**3.55**

the proposed method is unable to fully remove all the FPs, as shown in Figure 6. Nevertheless, it still results in a partial reduction in FPs (as indicated by the red bounding box in Figure 6) and contributes to overall improved accuracy, as

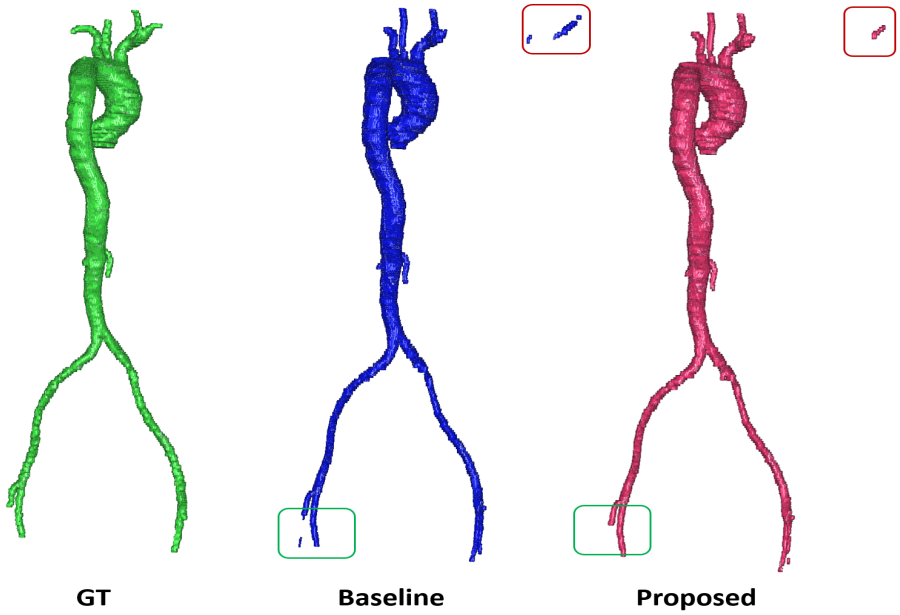

Fig. 6. Bad cases where the proposed MC Loss cannot completely remove the FPs (indicated by Red color Bounding Box), and for some voxels, it can help to get rid of all FPs (indicated by Green color Bounding Box) .

evident by the corresponding Dice Score and HD. We also note that in some regions, it can help to eliminate all FPs (highlighted by the green color box in

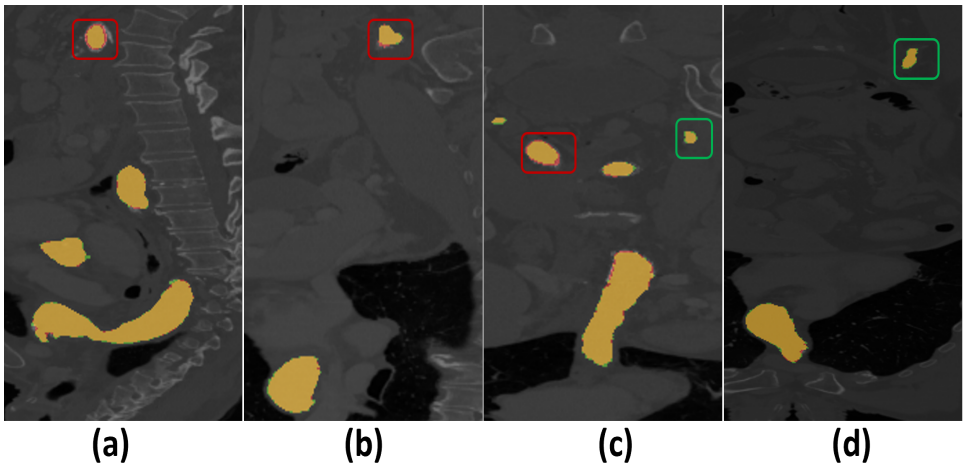

Fig. 7. Limitations of the proposed method for the cases where misclassified voxels are connected to TPs. A red color bounding box is drawn for FPs, and green color for FNs.

Figure 6).

Figure 7 showcases another limitation of the proposed approach, where it is unable to remove the FPs that are connected to the TPs (along the boundary of the segmented regions). In addition, it cannot successfully reduce the FNs in cases if they are found along the boundary. Here, we argue that the precise segmentation along the boundary regions might require additional enhancement modules to increase the learning capacity of MC Loss. This can be a potential future work, where we will extend the MC Loss to give more weight to the misclassified voxels along the boundary regions.

5 Conclusion

We propose a misclassification loss (MC loss) function by formulating a simple yet effective way of identifying the misclassified points from a pre-trained segmentation network. The proposed MC loss localizes the misclassified logits using a differentiable eXclusive-OR (XOR) operation. We employ MC loss for fine-tuning a network trained with existing loss functions. This fine-tuning using MC loss enables a reduction in the number of False Positives (FPs) and False Negatives (FNs) while increasing the True Positives (TPs). The efficacy of the proposed method is demonstrated by conducting experiments on the SEG.A. 2023 aorta segmentation challenge dataset. The method is ranked in the top six positions of the challenge leaderboard and achieved a Dice score of 0.926 and HD(mm) of 2.10 in the validation phase-1 of the challenge.

Acknowledgements This work acknowledges the support of the National Institute for Health Research Barts Biomedical Research Centre (NIHR203330). This research work is also funded by the mini-Centre for Doctoral Training (CDT) award through the Faculty of Science and Engineering, Queen Mary University of London, United Kingdom. The authors also thank mini-CDT partners, including NVIDIA Corporation, Circle Cardiovascular Imaging, and Conavi Medical. The proposed research work utilised Queen Mary's Andrena HPC facility, supported by QMUL Research-IT.

References

1. B. Zhao, M.K.: Data From RIDER Lung CT. The Cancer Imaging Archive. (2015)
2. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013)
3. Cao, L. et al.: Fully automatic segmentation of type B aortic dissection from CTA images enabled by deep learning. European journal of radiology **121**, 108713 (2019)
4. Edelsbrunner, Letscher, Zomorodian, Topological persistence and simplification. Discrete & Computational Geometry **28**, 511–533 (2002)

5. Fantazzini, A. et al.: 3D automatic segmentation of aortic computed tomography angiography combining multi-view 2D convolutional neural networks. Cardiovascular engineering and technology **11**, 576–586 (2020)
6. Garcia, J. et al.: Distribution of blood flow velocity in the normal aorta: effect of age and gender. Journal of Magnetic Resonance Imaging **47**(2), 487–498 (2018)
7. Heller, N. et al.: The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. arXiv preprint arXiv:1904.00445 (2019)
8. Heller, N. et al.: The state of the art in kidney and kidney tumor segmentation in contrast-enhanced CT imaging: Results of the KiTS19 challenge. Medical image analysis **67**, 101821 (2021)
9. Hu, X. et al.: Topology-Aware Segmentation Using Discrete Morse Theory. ICLR (2021)
10. Hu, X.: Structure-Aware Image Segmentation with Homotopy Warping. Advances in Neural Information Processing Systems **35**, 24046–24059 (2022)
11. Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. Advances in neural information processing systems **32** (2019)
12. Isensee, F. et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
13. Jin, Y. et al.: Ai-based aortic vessel tree segmentation for cardiovascular diseases treatment: status quo. arXiv preprint arXiv:2108.02998 (2021)
14. Li, Z. et al.: Lumen segmentation of aortic dissection with cascaded convolutional network. In: Statistical Atlases and Computational Models of the Heart. Atrial Segmentation and LV Quantification Challenges: 9th International Workshop, STACOM 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers 9, pp. 122–130 (2019)
15. Members, W.G. et al.: 2010 ACCF/AHA/AATS/ACR/ASA/SCA/SCAI/SIR/STS/SVM guidelines for the diagnosis and management of patients with thoracic aortic disease: a report of the American college of cardiology foundation/American heart association task force on practice guidelines, American association for thoracic surgery, American college of radiology, American stroke association, society of cardiovascular anesthesiologists, society for cardiovascular angiography and interventions, society of interventional radiology, society of thoracic surgeons, and society for vascular medicine. Circulation **121**(13), e266–e369 (2010)
16. Mosinska, A., Marquez-Neila, P., Koziński, M., Fua, P.: Beyond the pixel-wise loss for topology-aware delineation. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3136–3145 (2018)
17. Müller, D., Soto-Rey, I., Kramer, F.: Towards a guideline for evaluation metrics in medical image segmentation. BMC Research Notes **15**(1), 1–8 (2022)
18. Pepe, A. et al.: Detection, segmentation, simulation and visualization of aortic dissections: A review. Medical image analysis **65**, 101773 (2020)
19. Petersen, F., Borgelt, C., Kuehne, H., Deussen, O.: Deep differentiable logic gate networks. Advances in Neural Information Processing Systems **35**, 2006–2018 (2022)
20. Radl, L. et al.: AVT: Multicenter aortic vessel tree CTA dataset collection with ground truth segmentation masks. Data in brief **40**, 107801 (2022)
21. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, pp. 234–241 (2015)

22. Shit, S. et al.: clDice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16560–16569 (2021)
23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. ICLR (2015)
24. Sotelo, J. et al.: Fully Three-Dimensional Hemodynamic Characterization of Altered Blood Flow in Bicuspid Aortic Valve Patients With Respect to Aortic Dilatation: A Finite Element Approach. Frontiers in Cardiovascular Medicine **9**, 885338 (2022)
25. This research utilised Queen Mary's Andrena HPC facility, supported by QMUL Research-IT.