

Topologically faithful multi-class segmentation in medical images

Alexander H. Berger^{*1,4[0009–0004–8843–7684]}, Laurin Lux^{1,3[0009–0003–7359–6212]}, Nico Stucki^{1,2,3[0009–0005–9736–9895]}, Vincent Bürgin^{1[0000–0002–8563–4189]}, Suprosanna Shit^{1,6[0000–0003–4435–7207]}, Anna Banaszak¹, Daniel Rueckert^{1,2,3,4,5[0000–0002–5683–5889]}, Ulrich Bauer^{1,2,3[0000–0002–9683–0724]}, and Johannes C. Paetzold^{5[0000–0002–4844–6955]}

¹ School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

² Munich Data Science Institute, Technical University of Munich, Munich, Germany

³ Munich Center for Machine Learning, Munich, Germany

⁴ School of Medicine and Health, Klinikum rechts der Isar, Technical University of Munich, Munich, Germany

⁵ Department of Computing, Imperial College London, London, UK

⁶ Department of Quantitative Biomedicine, University of Zurich, Switzerland

Abstract. Topological accuracy in medical image segmentation is a highly important property for downstream applications such as network analysis and flow modeling in vessels or cell counting. Recently, significant methodological advancements have brought well-founded concepts from algebraic topology to binary segmentation. However, these approaches have been underexplored in multi-class segmentation scenarios, where topological errors are common. We propose a general loss function for topologically faithful multi-class segmentation extending the recent Betti matching concept, which is based on induced matchings of persistence barcodes. We project the N -class segmentation problem to N single-class segmentation tasks, which allows us to use 1-parameter persistent homology, making training of neural networks computationally feasible. We validate our method on a comprehensive set of four medical datasets with highly variant topological characteristics. Our loss formulation significantly enhances topological correctness in cardiac, cell, artery-vein, and Circle of Willis segmentation.⁷

Keywords: Topology · Multi-class Segmentation · Betti matching

1 Introduction

Topological correctness is crucial for many downstream tasks, such as blood flow modeling or cell counting. Modern segmentation networks achieve high pixel-wise accuracy but often cannot preserve important topological features [1, 10]. Recent

^{*} Corresponding author. a.berger@tum.de

⁷ The code is available at github.com/AlexanderHBerger/multiclass-BettiMatching




Fig. 1: Top: image, ground truth, and exemplary segmentation of the cardiac dataset. Bottom: ground truth (left) next to the segmentation (right), pairwise for every class. With our multi-class Betti-matching formulation, we match connected components (dim-0) and cycles (dim-1) in each individual class. Matched features in dim-0 and dim-1 are colored in a checkerboard pattern and have colored feature cycles, respectively. All classes’ matched and unmatched features guide our loss function for multi-class topology-preserving segmentation.

works addressed this issue via a variety of methods such as postprocessing [3, 14], injecting topological priors during training [3–5, 7], or topology-aware loss functions [8–10, 20, 22, 24]. These methods come with varying degrees of generalization capabilities and theoretical guarantees. For example, some persistent homology-based losses cannot guarantee the spatial matching of topological features [4, 10]. Recently, the *Betti matching* loss was proposed as a general loss function to train binary segmentation networks with strict topological guarantees and correct spatial alignment of topological features [24]. Despite these successes, the realm of topology-preserving multi-class segmentation has been under-explored. In this work, we fill this gap by proposing a generalizable multi-class segmentation loss extending the Betti matching concept to multi-class settings. Based on popular methods, specifically, cDice [22], and HuTopo [10], we establish other topology-aware multi-class baselines. Furthermore, we introduce a weighting term based on topological structures that can improve the topological correctness depending on the dataset characteristics. In experiments on four topologically variant datasets, we demonstrate our method’s utility and outperform all baselines.

Related works A plethora of methods have emphasized the importance of topologically correct segmentation in medical imaging. These methods can be

assigned to three categories. Firstly, post-processing frameworks that aim to fix topological errors of preliminary segmentations [14, 19, 26]. These approaches do not integrate with end-to-end training of neural networks. The second category uses fixed topological priors during training or post-processing [3]. Clough et al. [4] follow an approach where they defined the expected topology using Betti numbers on which they computed a loss term. Gupta et al. [7] apply physiological priors to segment various medical structures. Other studies use similar approaches for coronary artery segmentation [28, 29]. However, these methods do not generalize to other tasks and are only applicable when every sample for the segmentation task has the same topological features. The last category contains topologically faithful loss functions in which a topological loss is computed without using task-specific knowledge. These methods are often based on skeletonization for tubular structures [22, 20, 16]. In a foundational work, Hu et al. [10] proposed using a loss function for image segmentation, which minimizes the squared distance of matched points in persistence diagrams of dimension 1. However, this method does not guarantee that matched structures are spatially related in any sense, which has a significant negative impact on training [24]. A recent work overcame the limitations of these methods and showed that induced matchings can achieve a spatially correct matching between barcodes in a binary segmentation setting with formal guarantees [24]. Other approaches applied homotopy warping to identify critical pixels and measure topological differences between images [9, 11] or utilized discrete Morse theory to compare critical topological structures [8]. Such loss functions are beneficial because they do not require prior knowledge, they can be generalized across tasks and can be incorporated into end-to-end training. However, research regarding such general, homology-based loss functions in multi-class settings has been limited.

2 Method

This work presents a loss formulation that captures topological errors in multi-class segmentation masks. Our loss can be used to train arbitrary segmentation networks in an end-to-end fashion. We utilize persistent homology to capture topological features at multiple scales and generalize homology-based loss functions to multi-class segmentations.

Overview of persistent homology for image segmentation. Let $x \in \mathbb{R}^{W \times H}$ be a grayscale image, $y \in \{0, 1\}^{N \times W \times H}$ its segmentation in N mutually exclusive classes, and $\hat{y} \in [0, 1]^{N \times W \times H}$ a predicted likelihood map for each class with $\sum_{c=1}^N \hat{y}_{c,i,j} = 1$. In a binary segmentation setting, the predicted likelihood map $\hat{y} \in [0, 1]^{W \times H}$ can be seen as a function $\hat{y}: K^{W \times H} \rightarrow [0, 1]$ on a *cubical grid complex* $K^{W \times H}$ (see [24]) and its sublevel filtration $K_t = \hat{y}^{-1}((-\infty, t])$ as a model for its topology. Since the sublevel filtration depends only on one parameter $t \in \mathbb{R}$, the persistent homology of this filtration forms a 1-parameter persistence module, and its barcode decomposition can be used as a descrip-

tor for the topology of the prediction. For a thorough theoretical description of persistent homology for binary image segmentation, we refer the reader to [24].

Single-class loss functions for multi-class segmentations. In a multi-class segmentation task, the predicted likelihood map $\hat{y} \in [0, 1]^{N \times W \times H} \rightarrow [0, 1]$ corresponds to a function $\hat{y}: K^{W \times H} \rightarrow [0, 1]^N$. Since its sublevel filtration $K_{(t_n)} = \{x \in K^{W \times H} \mid \hat{y}(x) \leq (t_1, \dots, t_N)\}$ depends on N parameters $t_1, \dots, t_N \in \mathbb{R}$, its persistent homology forms an N -parameter persistence module. The complexity of computing descriptors for N -parameter persistent homology makes it infeasible to train segmentation networks with such a setup.

In order to circumvent this problem, consider a multi-class segmentation network $f: \mathbb{R}^{W \times H} \rightarrow [0, 1]^{N \times W \times H}$ and denote by

$$p_c: [0, 1]^{N \times W \times H} \rightarrow [0, 1]^{W \times H}, (\hat{y}_{n,i,j}) \mapsto (\hat{y}_{c,i,j}) \quad (1)$$

the restriction to the c -th channel. Then $f_c := p_c \circ f$ can be seen as the solution to a single-class segmentation task predicting the likelihood of a pixel being within class c . This way, we can break down the multi-class segmentation task with N classes into N single-class segmentation tasks, allowing us to apply single-class segmentation loss functions. Thus, any single-class segmentation loss $\mathcal{L}: [0, 1]^{W \times H} \times \{0, 1\}^{W \times H} \rightarrow \mathbb{R}$ can be extended to a multi-class segmentation loss by summing the individual losses of each channel:

$$\mathcal{L}: [0, 1]^{N \times W \times H} \times \{0, 1\}^{N \times W \times H} \rightarrow \mathbb{R}, (\hat{y}, y) \mapsto \sum_{c=1}^N \mathcal{L}(p_c(\hat{y}), p_c(y)) \quad (2)$$

Considering each channel c individually and computing the barcode of its sublevel filtration corresponds to slicing the N -parameter persistence module along the line

$$\{t \in \mathbb{R}^N \mid t_c = t, t_n = 1 - t \text{ for } n \neq c\} \quad (3)$$

to a 1-parameter module and computing its corresponding barcode decomposition, which scales the computational complexity linearly with the number of classes N compared to a binary setting. Therefore, the total computational complexity is dominated by the barcode computation that scales cubically with the number of pixels [24]. We utilize the efficient implementation for barcode computation presented in [23] and provide runtime information in the supplement. Notably, the presented formulation naturally extends to other topology-preserving losses, which allows us to extend and implement other binary topology-preserving losses to our setting and compare their generalizability, which we do in the experiment section. Furthermore, they extend to topological metrics, especially the Betti matching error [24].

Weighting. The Betti matching loss l_{BM} is based on the induced matching of persistence barcodes $\mu(\mathbf{L}, \mathbf{G})$ (as described in [24]), which consists of matched and unmatched features between ground truth and prediction (as depicted in Fig.

1). Therefore, l_{BM} can be decomposed into one loss component for matched l_{BM}^m and one for unmatched l_{BM}^u features, respectively. l_{BM}^m reinforces topological features that are correctly predicted by the network, and l_{BM}^u penalizes topological features that are predicted by the network but do not exist in the ground truth. In [24], these components are implicitly weighted equally. However, in a multi-class setting, with varying numbers of classes and topological features, one loss component can dominate the other, which may deteriorate performance (see Sec. 4). Hence, we propose a weighting of these components with the parameters γ^m and γ^u . Lastly, we combine our final multi-class topological loss with a standard pixel-based loss function (e.g., Dice) to the final loss function:

$$l_{\text{total}} = \alpha \cdot (\gamma^m \cdot l_{BM}^m + \gamma^u \cdot l_{BM}^u) + l_{\text{Dice}} \quad (4)$$

3 Data and Experimentation

Datasets. We evaluate our method on four different, public multi-class segmentation tasks. While all data is from different modalities, we choose tasks where topological correctness is a major objective (especially for downstream applications) and has proven to be challenging with traditional methods. For training/validation/test splits and further details, please see the Supplement.

First, in the Automated Cardiac Diagnosis Challenge (ACDC) challenge dataset [2], the task is cardiac segmentation, i.e., segmenting the myocardium and the left and right ventricles in 2D short-axis slices of Cardiac Magnetic Resonance (CMR) scans. Here, pathological samples often lead to failure modes that exhibit topological errors [3]. We use each slice, including the apex and base, to get data points with varying topological features. Second, we use a publicly available cellular and subcellular electron microscopy (EM) dataset (Platelet) with six classes (cell, mitochondrion, canalicular channel, alpha granule, dense granule, and dense granule core) [6]. Third, we test our method on artery-vein classification in OCTA images with the OCTA-500 dataset [15]. Last, we test our model on Circle of Willis (CoW) segmentation in Magnetic Resonance Angiography (MRA) scans with 15 vessel components as classes. The CoW has hypoplastic and absent components across different subjects, making correct segmentation challenging [27]. We project the MRA scan and the label to a 2D image and segmentation mask.

Baselines. We employ three baselines against which we evaluate the performance of our loss formulation. All baselines enable end-to-end training and are independent of the used network. First, the generalized Dice loss [25], which is still one of the most frequently used loss terms in biomedical image segmentation due to its inherent class balancing properties [17]. Thus, we choose the Dice loss as a standalone baseline. Second, we combine the Dice loss with the (multi-class extended) centerline Dice (clDice) loss [22]. clDice is especially effective for segmenting tubular structures. Lastly, we develop an additional baseline where we combine the Dice loss with a multi-class extension of the HuTopo loss [10] via the multiclass generalization introduced in Section 2.

Metrics We evaluate segmentation performance with a comprehensive set of topology and pixel-wise scores. Specifically: the *Betti matching error*, *Betti number errors*, *Dice*, and *clDice*. The Betti matching error (BM) is the most accurate indicator of faithful topological segmentations [24]. All metrics are averaged across all foreground classes without weighting.

Training and model selection We train a U-Net architecture [21] with residual units from scratch. We use the Adam optimizer [12], a fixed learning rate, and sigmoid scheduling for α (see Supplement). Note that our loss formulation is independent of the underlying network architecture. We perform 5-fold cross-validation and evaluate on an independent test set. We perform a random hyperparameter search with 50 runs on each of the splits and select the model that has the highest performance S on the validation set with S being defined as a balanced performance metric of pixel-wise accuracy and topological faithfulness:

$$S = \text{Dice} + \left(1 - \min\left(1, \frac{l_{BM}}{\beta_0 + \beta_1}\right)\right) \quad (5)$$

where β_i is the Betti number in dimension i . We report the mean performance and standard deviation on the independent test set across the five data splits. Please refer to the supplement for hyperparameters. We use the paired t-test between our model and each baseline to evaluate statistically significant performance (p -value < 0.05) improvements.

4 Results

Our results comprehensively demonstrate that our proposed multi-class segmentation loss outperforms all baselines across all datasets in Betti matching errors. Further, we outperform all baselines in Betti number errors 0 and 1 as well as clDice in the Platelet, TopCoW, and ACDC datasets; see Table 1. Even in the non-topology-aware Dice score, our method performs on par with the baselines across most datasets. Overall, we find that the (multi-class extended) HuTopo baseline [10] does not generalize well to the multi-class setting. We attribute this to an amplification of the possible incorrect matchings of persistence features as described in Stucki et al. [24]; i.e., the matched features do not correspond spatially. Further, we find specific results for the individual datasets:

ACDC. This dataset exhibits a mostly regular structure where the three classes have one unique topology, i.e., the Myocardium forms one connected component and mostly one cycle. Achieving good volumetric performance is rather simple; all methods show similar, high performance in Dice and clDice; however, our method still significantly improves the Betti matching error and the Betti number errors. Our method is especially beneficial in the irregular (e.g., basal) slices. There, the structures do not have their typical topology, making approaches based on topological post-processing unsuitable.

Table 1: Quantitative results. We show the performance of our multiclass Betti matching loss against multiple other topology-aware losses that we adapted to the multiclass segmentation setting. Best performances indicated in **bold**, statistical significance underlined (p -value < 0.05).

Dataset	Loss	Dice $\uparrow$	clDice $\uparrow$	BM $\downarrow$	B0 $\downarrow$	B1 $\downarrow$
ACDC	Dice	.868 $\pm$.023	<u>.636$\pm$.004</u>	<u>0.130$\pm$0.014</u>	<u>0.094$\pm$0.012</u>	<u>0.032$\pm$0.006</u>
	ClDice	.871$\pm$.020	<u>.629$\pm$.009</u>	<u>0.380$\pm$0.226</u>	<u>0.115$\pm$0.014</u>	<u>0.261$\pm$0.217</u>
	HuTopo	.862 $\pm$.005	.634 $\pm$.004	<u>0.219$\pm$0.049</u>	<u>0.164$\pm$0.044</u>	<u>0.049$\pm$0.010</u>
	Ours	.862 $\pm$.008	.641$\pm$.005	0.064$\pm$0.007	0.038$\pm$0.005	0.022$\pm$0.004
Platelet	Dice	.685 $\pm$.015	<u>.548$\pm$.009</u>	<u>1.289$\pm$0.147</u>	<u>0.643$\pm$0.072</u>	<u>0.280$\pm$0.045</u>
	ClDice	.682 $\pm$.016	<u>.544$\pm$.011</u>	<u>4.682$\pm$2.800</u>	<u>1.053$\pm$0.293</u>	<u>3.238$\pm$2.474</u>
	HuTopo	<u>.635$\pm$.042</u>	<u>.517$\pm$.018</u>	<u>1.420$\pm$0.408</u>	<u>0.684$\pm$0.218</u>	<u>0.301$\pm$0.136</u>
	Ours	.696$\pm$.019	.564$\pm$.012	0.978$\pm$0.126	0.483$\pm$0.068	0.191$\pm$0.024
OCTA-500	Dice	.829$\pm$.006	.567$\pm$.004	<u>34.258$\pm$2.929</u>	11.825 $\pm$ 2.518	0.018 $\pm$ 0.019
	ClDice	.794 $\pm$.013	.560 $\pm$.005	<u>23.510$\pm$0.498</u>	6.942$\pm$0.367	0.008$\pm$0.011
	HuTopo	.787 $\pm$.028	.554 $\pm$.009	<u>28.285$\pm$3.871</u>	8.255 $\pm$ 1.877	0.035 $\pm$ 0.021
	Ours	.798 $\pm$.013	.556 $\pm$.007	17.950$\pm$0.567	9.670 $\pm$ 0.667	0.055 $\pm$ 0.021
TopCoW	Dice	.625 $\pm$.148	.624 $\pm$.145	<u>1.791$\pm$1.176</u>	<u>1.495$\pm$1.114</u>	0.042 $\pm$ 0.021
	ClDice	.685 $\pm$.102	.693 $\pm$.109	<u>1.043$\pm$0.803</u>	<u>0.835$\pm$0.713</u>	0.032 $\pm$ 0.017
	HuTopo	.670 $\pm$.138	.675 $\pm$.141	1.442 $\pm$ 1.668	1.207 $\pm$ 1.585	0.042 $\pm$ 0.025
	Ours	.717$\pm$.021	.725$\pm$.025	0.773$\pm$0.304	0.616$\pm$0.231	0.027$\pm$0.017

Platelet. In this dataset, the aim is to segment round objects where the topology is described by "inclusion," e.g., a mitochondrion is always inside a cell segment. Our method is superior to all baselines across all metrics. The HuTopo and clDice baselines perform especially poorly, which we attribute to the large image size (200 $\times$ 200 pixels), making HuTopo's spatial mismatch more prevalent, and the lack of tubular structures, negatively impacting clDice performance. Qualitatively, we observe that the baselines often merge individual structures.

OCTA-500. Here, the Dice loss results in the best Dice score, which is in contrast to the other experiments. We hypothesize that the Dice loss is more robust against the dataset's annotation scarcity [18, 13]. However, our loss significantly decreases the BM error compared to all baselines, including the Dice loss. The qualitative examples show that all baselines suffer from a large number of discontinuous vessel segments, which can be critical for downstream applications.

TopCoW. Our method improves performance across all metrics. This dataset is challenging because of its 13 distinct classes, which partially overlap in the 2D projection. In this setting, all topology-aware losses perform more robustly than the Dice loss. We conclude that in such challenging settings, topological methods are especially useful. Furthermore, our method shows by far the lowest standard deviation, further emphasizing its robustness compared to other methods.




Fig. 2: Qualitative results on ACDC (a), Platelet (b), OCTA-500 (c), and Top-CoW (d) dataset. Our method improves topological correctness in all multi-class segmentation tasks. We indicate some topological errors with black arrows.

Ablation on weighting. In an ablation, we study the effect of our proposed weighting concept, see Section 2. We find that topological weighting can strongly affect the model’s performance regarding the Betti matching error while being robust in terms of pixel-wise accuracy, see Fig. 4. Here, we study the OCTA-500 dataset and show that a low weight for the loss of matched topological features l_{BM}^m (and thereby a relatively higher weight for the loss of unmatched topological features l_{BM}^u) drastically improves topological performance while having a negligible effect on Dice. We attribute this observation to the predominance of topological features in dimension 0 in the hierarchically structured vasculatures in OCTA images. Note that other datasets do not follow this trend indicating that the weighting parameter must be specifically tuned for the used dataset.

5 Conclusion

Our study introduces a novel loss for multi-class image segmentation, extending the Betti matching concept to preserve topology. By decomposing the N -class problem, we circumvent the use of multi-parameter persistent homology, facilitating neural network training. Our empirical findings are three-fold. First, we




Fig. 3: Betti Matching error (left) and Dice score (right) with varying γ^m

find that our proposed generalization can successfully extend topology-aware losses to multiclass problems. Second, we provide strong empirical evidence that our loss does not impede pixel-wise accuracy and significantly improves topological accuracy. Finally, we find that our loss function consistently outperforms other topology-preserving baselines across several datasets, rendering it suitable for topologically critical applications.

Acknowledgments. A. Berger is supported by the Stiftung der Deutschen Wirtschaft.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. BenTaieb, A., Hamarneh, G.: Topology aware fully convolutional networks for histology gland segmentation. In: International conference on medical image computing and computer-assisted intervention. pp. 460–468. Springer (2016)
2. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE transactions on medical imaging **37**(11), 2514–2525 (2018)
3. Byrne, N., Clough, J.R., Valverde, I., Montana, G., King, A.P.: A persistent homology-based topological loss for cnn-based multiclass segmentation of cmr. IEEE transactions on medical imaging **42**(1), 3–14 (2022)
4. Clough, J., Byrne, N., Oksuz, I., Zimmer, V.A., Schnabel, J.A., King, A.: A topological loss function for deep-learning based image segmentation using persistent homology. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
5. Clough, J.R., Oksuz, I., Byrne, N., Schnabel, J.A., King, A.P.: Explicit topological priors for deep-learning based image segmentation using persistent homology. In: International Conference on Information Processing in Medical Imaging. pp. 16–28. Springer (2019)
6. Guay, M.D., Emam, Z.A., Anderson, A.B., Aronova, M.A., Pokrovskaya, I.D., Storie, B., Leapman, R.D.: Dense cellular segmentation for em using 2d–3d neural network ensembles. Scientific reports **11**(1), 2561 (2021)

7. Gupta, S., Hu, X., Kaan, J., Jin, M., Mpoy, M., Chung, K., Singh, G., Saltz, M., Kurc, T., Saltz, J., et al.: Learning topological interactions for multi-class medical image segmentation. In: *European Conference on Computer Vision*. pp. 701–718. Springer (2022)
8. Hu, X., Wang, Y., Li, F., Samaras, D., Chen, C.: Topology-aware segmentation using discrete morse theory. In: *International Conference on Learning Representations (ICLR)* (2021)
9. Hu, X., Chen, C.: Image segmentation with homotopy warping. *arXiv preprint arXiv:2112.07812* (2021)
10. Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. *Advances in neural information processing systems* **32** (2019)
11. Jain, V., Bollmann, B., Richardson, M., Berger, D.R., Helmstaedter, M.N., Briggman, K.L., Denk, W., Bowden, J.B., Mendenhall, J.M., Abraham, W.C., et al.: Boundary learning by optimization with topological constraints. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 2488–2495. IEEE (2010)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
13. Kreitner, L., Paetzold, J.C., Rauch, N., Chen, C., Hagag, A.M., Fayed, A.E., Sivaprasad, S., Rausch, S., Weichsel, J., Menze, B.H., et al.: Synthetic optical coherence tomography angiographs for detailed retinal vessel segmentation without human annotations. *IEEE Transactions on Medical Imaging* (2024)
14. Li, L., Ma, Q., Ouyang, C., Li, Z., Meng, Q., Zhang, W., Qiao, M., Kyriakopoulou, V., Hajnal, J.V., Rueckert, D., et al.: Robust segmentation via topology violation detection and feature synthesis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 67–77. Springer (2023)
15. Li, M., Huang, K., Xu, Q., Yang, J., Zhang, Y., Ji, Z., Xie, K., Yuan, S., Liu, Q., Chen, Q.: Octa-500: a retinal dataset for optical coherence tomography angiography study. *Medical Image Analysis* **93**, 103092 (2024)
16. Lin, Z., Wei, D., Gupta, A., Liu, X., Sun, D., Pfister, H.: Structure-preserving instance segmentation via skeleton-aware distance transform. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 529–539. Springer (2023)
17. Ma, J., Chen, J., Ng, M., Huang, R., Li, Y., Li, C., Yang, X., Martel, A.L.: Loss odyssey in medical image segmentation. *Medical Image Analysis* **71**, 102035 (2021)
18. Menten, M.J., Kreitner, L., Paetzold, J.C., Hagag, A.M., Bassily, S.M., Sivaprasad, S., Rueckert, D., Fayed, A.E.: Synthetic data facilitates deep-learning-based segmentation of oct angiography images without human annotations. *Investigative Ophthalmology & Visual Science* **64**(8), 5450–5450 (2023)
19. Mosinska, A., et al.: Beyond the pixel-wise loss for topology-aware delineation. In: *CVPR*. pp. 3136–3145 (2018)
20. Qiu, Y., Li, Z., Wang, Y., Dong, P., Wu, D., Yang, X., Hong, Q., Shen, D.: Corseg-rec: a topology-preserving scheme for extracting fully-connected coronary arteries from ct angiography. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 670–680. Springer (2023)
21. Ronneberger, O., et al.: U-net: Convolutional networks for biomedical image segmentation. In: *MICCAI*. pp. 234–241. Springer (2015)
22. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylyka, A., Plum, J.P., Bauer, U., Menze, B.H.: 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. Stucki, N., Bürgin, V., Paetzold, J.C., Bauer, U.: Efficient betti matching enables topology-aware 3d segmentation via persistent homology (2024), <https://arxiv.org/abs/2407.04683>
24. Stucki, N., Paetzold, J.C., Shit, S., Menze, B., Bauer, U.: Topologically faithful image segmentation via induced matching of persistence barcodes. In: International Conference on Machine Learning. pp. 32698–32727. PMLR (2023)
25. Sudre, C.H., et al.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: MICCAI Workshop, pp. 240–248. Springer (2017)
26. Wang, X., Jiang, X.: Post-processing for retinal vessel detection. In: Tenth International Conference on Digital Image Processing (ICDIP 2018). vol. 10806, pp. 1442–1446. SPIE (2018)
27. Yang, K., Musio, F., Ma, Y., Juchler, N., Paetzold, J.C., Al-Maskari, R., Höher, L., Li, H.B., Hamamci, I.E., Sekuboyina, A., et al.: Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for cta and mra. arXiv preprint arXiv:2312.17670 (2023)
28. Zhang, X., Zhang, J., Ma, L., Xue, P., Hu, Y., Wu, D., Zhan, Y., Feng, J., Shen, D.: Progressive deep segmentation of coronary artery via hierarchical topology learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 391–400. Springer (2022)
29. Zhang, Z., Zhao, Z., Wang, D., Zhao, S., Liu, Y., Liu, J., Wang, L.: Topology-preserving automatic labeling of coronary arteries via anatomy-aware connection classifier. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 759–769. Springer (2023)