

Topology Optimization in Medical Image Segmentation with Fast χ Euler Characteristic

Liu Li, Qiang Ma, Cheng Ouyang, Johannes C. Paetzold, Daniel Rueckert, *Fellow, IEEE*,
Bernhard Kainz, *Senior Member, IEEE*

Abstract—Deep learning-based medical image segmentation techniques have shown promising results when evaluated based on conventional metrics such as the Dice score or Intersection-over-Union. However, these fully automatic methods often fail to meet clinically acceptable accuracy, especially when topological constraints should be observed, *e.g.*, continuous boundaries or closed surfaces. In medical image segmentation, the correctness of a segmentation in terms of the required topological genus sometimes is even more important than the pixel-wise accuracy. Existing topology-aware approaches commonly estimate and constrain the topological structure via the concept of persistent homology (PH). However, these methods are difficult to implement for high dimensional data due to their polynomial computational complexity. To overcome this problem, we propose a novel and fast approach for topology-aware segmentation based on the Euler Characteristic (χ). First, we propose a fast formulation for χ computation in both 2D and 3D. The scalar χ error between the prediction and ground-truth serves as the topological evaluation metric. Then we estimate the spatial topology correctness of any segmentation network via a so-called topological violation map, *i.e.*, a detailed map that highlights regions with χ errors. Finally, the segmentation results from the arbitrary network are refined based on the topological violation maps by a topology-aware correction network. Our experiments are conducted on both 2D and 3D datasets and show that our method can significantly improve topological correctness while preserving pixel-wise segmentation accuracy.

Index Terms—Euler characteristic, segmentation, topological violation detection, topology-aware correction

I. INTRODUCTION

Performance of automated medical image segmentation methods is commonly evaluated by pixel-wise metrics, *e.g.*,

L. Li was supported by a Lee Family Scholarship. Q. Ma was supported by a President's PhD Scholarship from Imperial College London. Support was also received by the ERC projects Medical Image Analysis with Normative Machine Learning (MIA-NORMAL) 101083647, Deep Learning for Medical Imaging: Learning Clinically Useful Information from Images (Deep4MI) 884622, and The Developing Human Connectome Project (dHCP) data collection 319456. B.K. acknowledges the high-tech agenda Bavaria and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU, b143dc and b180dc). NHR funding is provided by federal and Bavarian state authorities and partially funded by the German Research Foundation (DFG) - 440719683.

L. Li, Q. Ma, C. Ouyang, J. Paetzold, D. Rueckert, and B. Kainz are with the Department of Computing, Imperial College London, UK (e-mail: liu.li20@imperial.ac.uk). C. Ouyang is also with Department of Engineering Science, University of Oxford, UK and was with Institute of Clinical Sciences, Imperial College London. J. Paetzold is also with Weill Cornell Medicine, New York, USA, and Cornell Tech, New York, USA. D. Rueckert is the Chair for AI in Medicine and Healthcare, TUM University Hospital, Technical University of Munich, Germany. B. Kainz is also with SBEIS, King's College London, UK, and the Department AIBE of Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany.




Fig. 1. From left to right: image from the CREMI dataset [7], ground-truth map for the neuron boundary, nnU-Net [3] segmentation, segmentation error map and our topological violation map. Despite the high Dice accuracy of 88.97% and 88.85% for these two examples, the topological errors are very different, *e.g.*, the ring structure is broken in the upper example. In the segmentation error maps, we find that some errors will change the topological structure (highlight in purple) while some errors will not (highlight in blue). Most of these topology-related errors are essentially caused by imaging artifacts, *e.g.*, blurry boundaries, or dark patches. Based on these observations, we design a topology-aware segmentation method that can automatically detect these topology-related errors based on the Euler Characteristic χ . Resulting topological violation maps are shown in the last column.

the Dice score or the Intersection-over-Union [1]. Taking the average accuracy of each pixel/voxel into account, these evaluation methods are well suited to describe the overall segmentation performance. However, errors in some regions may be more important than others, *e.g.*, certain errors may lead to the wrong topology of the segmented object, highlighted in purple in Fig. 1 in the 4th column. In some cases, topological correctness is more important than pixel-wise classification correctness, for example, when segmentation is conducted as a fundamental step for surface reconstruction [2]. Although the state-of-the-art segmentation methods such as nnU-Net [3] and fine-tuned Segment Anything Models (SAM) [4] have proven to achieve high pixel-wise accuracy in many tasks, *e.g.*, brain tumor segmentation [3], prostate analysis [5], and breast evaluation [6], the topological correctness of the segmentation are usually not addressed.

As illustrated in Fig. 1, this paper is motivated by the understanding that only a part of real-world segmentation errors are related to topology. Accordingly, we propose a novel approach where a segmentation network learns to focus more on areas likely to contain topological errors, rather than uniformly addressing all pixel-wise errors equally in an image.

Current topology-aware segmentation techniques employ

persistent homology (PH) to identify critical points that are topologically significant. These techniques represent the segmentation probability map using a cubical complex. PH monitors the changes in topology across various filtrations [8, 9, 10, 11]. Topological correctness is encouraged by manipulating the probability of these critical points through an additional loss term.

Although integration of these methods is straight forward due to their differential properties, the main challenge arises from the polynomial computational complexity of PH, which limits their applicability to high-dimensional data such as 3D medical images. To address this challenge, we propose a fast method based on the Euler characteristic (χ) to detect topologically critical regions that likely contain topological errors.

Specifically, we first develop a DL-compatible pipeline for the bit-quads-based 2D χ calculation [12, 13, 14], and further extend this 2D formulation to 3D by introducing bit-octets in Section III-B. To identify topologically critical regions, we first calculate χ for both prediction and ground-truth (GT) locally to get two χ -maps. Analysing the L_1 difference between prediction and GT χ -maps, which we denote as χ error map, we can further reveal more refined details about topological violations by back-propagating the χ error with respect to the segmentation, resulting in a *topological violation map* that highlights topologically critical regions U .

Given this topological violation map, we propose an additional topology-aware feature synthesis (TFS) network to refine segmentations. The motivation for this design emerges from the observation that the topological error regions U for segmentation Y are correlated to imaging artifacts in the input images I , as illustrated in Fig. 1. Encouraging a segmentation network to directly learn from U through an additional loss term would be futile since the underlying image features $I_{ij}, (i, j) \in U$ are erroneous.

Therefore, instead of introducing additional loss terms, our approach relies on a generative TFS network to directly learn a topological prior, $p(Y)$, from regions with topologically correct representations, denoted as $Y_{ij}, (i, j) \in U^c$, where $U^c \cap U = \emptyset$.

The primary contributions of this paper are twofold. First, building upon the concept of DL-compatible χ computation as a tool to inform the network about topologically invariant properties [15], we discuss its applicability in high-dimensional medical data, such as 3D CT and MRI volumes. We define and categorize bit-octets into different groups based on their properties, and explore their combinations to enable fast 3D χ computation. Second, we conduct experiments on five datasets featuring diverse foreground structures, demonstrating the enhanced segmentation performance of our approach in terms of topological correctness.

This paper extends some of our initial ideas [15]. Firstly, we conduct a comprehensive review of related literature that incorporates priors into segmentation networks, categorizing these priors as either implicit shape priors or explicit topological constraints. Secondly, we broaden the scope of bit-quads-based Euler Characteristic (χ) computation, transitioning from 2D to 3D, by investigating the integration of bit-octets. Thirdly,

we evaluate on five different datasets, covering both 2D and 3D domains. Note that our approach can be applied to any existing segmentation pipeline as a plug-and-play module. Our code is publicly available¹.

II. RELATED WORK.

A. Implicit shape priors

Medical image segmentation aims to provide pixel-level delineation of anatomical structures, facilitating clinical interpretation and downstream analysis. U-Net [16] and its variants, such as Swin-Unet [17] and nnUNet [3], have achieved strong performance across various medical imaging tasks. More recently, inspired by the Segment Anything Model (SAM) [18] with self-attention mechanisms, MedSAM [4], MedSAM 2 [19], and ESP-MedSAM [20] have been developed for medical images, leveraging prompt-based strategies to refine segmentation predictions interactively.

Segmentation networks, however, are typically trained using binary cross-entropy or soft Dice loss functions. These metrics primarily assess pixel-wise overlap between the predicted and ground truth (GT) segmentations. However, in medical image segmentation, prioritizing pixel-wise accuracy can sometimes compromise the structural integrity of predictions, particularly in tasks such as downstream surface reconstruction.

While these methods are effective in certain aspects, they do not explicitly model topological constraints. To address this, shape priors have been explored, for example, by utilizing shape templates [21, 22] with diffeomorphic deformation. These shape priors are reinforced by constraining the deformation field to be smooth, which assumes a similarity between the shape prior and the target. However, a limitation of these methods is their reliance on predictions being extremely close to the shape priors, a condition that is often unattainable in practical scenarios.

Other approaches explicitly incorporate shape information into the training process [23]. Additionally, implicit shape awareness, based on boundary and centerline information that can model connectivity, has been explored in recent studies [24, 25, 26]. Despite these advancements, these methods do not inherently derive from topological principles and, as such, do not explicitly model topological information.

B. Explicit topological constraints

Topology in segmentation has been explored through the use of Persistent Homology (PH) [8, 9, 10, 27, 28, 29, 11, 30]. PH identifies topological features by tracking the emergence and dissolution of connected components, holes, and voids as filtration changes. In practice, PH-based methods detect critical points associated with topological changes and manipulate their values to align the predicted topology with the ground truth (GT). This approach can be made differentiable by modelling the probability of critical points in the segmentation. Recently, PH loss has also been applied to multi-category segmentation where it models the combinations of different labels to enforce topological constraints [31, 30].

¹<https://github.com/smilell/Topology-aware-Segmentation-using-Euler-Characteristic>




Fig. 2. Construction of a cubical complex from a 2D binary image. Yellow and black blocks represent foreground and background segmentations, respectively.

Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8	
0 0	1 0	0 1	0 0	0 0	1 1	1 0	0 0	Q_A
0 0	0 0	0 0	1 0	0 1	0 0	1 0	1 1	
Q_9	Q_{10}	Q_{11}	Q_{12}	Q_{13}	Q_{14}	Q_{15}	Q_{16}	Q_B
0 1	1 0	0 1	1 1	1 1	1 0	0 1	1 1	Q_C
0 1	0 1	1 0	1 0	0 1	1 1	1 1	1 1	

Fig. 3. Bit-quads for 2D χ computation.

Beyond segmentation tasks, PH has also been applied in image generation to guide the synthesis of images with topologically correct structures [32, 33].

Despite optimised computational packages for PH [34, 35], its polynomial computational complexity remains a major limitation, especially for large-scale or higher dimensional datasets [36]. Further, PH-based loss optimization in each iteration targets only a limited number of critical points. Newer approximation methods, such as those based on Discrete Morse Theory [37] and Homotopy Warping [36], have reduced computational complexity to $O(n \log n)$ and $O(n)$, respectively.

Nevertheless, these approaches primarily focus the network’s attention on regions with topological discrepancies when compared to the GT, still relying on a pixel-wise topology loss term instead of learning a global topological representation. Moreover, in medical image segmentation, topological errors often result from imaging artifacts, such as blurred boundaries. Training the network on these compromised inputs may not yield improved performance. To address this gap, we introduce a χ -based method that directly learns from the label map space, avoiding the lack of image space representation in those imaging artifact regions.

III. METHOD

A. Preliminaries

Algebraic topology studies the properties of topological spaces [38, 39]. This field aims to identify and characterize topological invariants, such as the number of holes, the connectivity, or the homotopy groups of a space [40, 41, 42]. In algebraic topology, high-dimensional data, such as point clouds, can be modeled using a simplicial complex $\mathcal{K}$, which is a set composed of points, line segments, triangles, tetrahedra, and higher-dimensional equivalents.

Image segmentation for 2D or 3D gray-scale images is defined on grids $I \in \mathbb{R}^{h \times w \times l}$. This can be naturally modeled by cubical complexes $\mathcal{K}$ to reduce the complexity of arbitrary possible connectivity between different vertices [42]. The basic elements of a 3-dimensional cubical complex $\mathcal{K}$ are points, line segments, squares, and cubes. For example, in a 3D binary image, each voxel is modeled by a cube, and every face of this cube, *i.e.*, 8 points, 8 line segments, and 4 squares, are also in this cubical complex $\mathcal{K}$, as shown in Fig. 2.

Betti Numbers. Betti numbers β_k are topological invariants that describe the property of geometric objects that remain invariant under homeomorphisms, *i.e.*, transformations include stretching, bending, and deforming without tearing or gluing [38]. Specifically, Betti numbers describe how many “holes” exist in a topology space X in dimension k . Mathematically, Betti numbers are defined as the ranks of the k -dimensional homology groups $H_k(X)$. Specifically, the k -dimensional homology group is a collection of k -dimensional non-boundary cycles, which can be intuitively understood as “holes” [42]. All non-boundary cycles can be grouped into homology classes based on whether they can be freely transformed into one another without altering their topological structure. The rank of $H_k(X)$ captures how many independent “holes” exist in this homology group. In dimension 0, 1 and 2, Betti numbers β_k represent the number of connected components, holes, and voids, respectively. Any k -dimensional cubical complex can only have $(k - 1)$ -dimensional non-zero Betti numbers, *i.e.*, $(k - 1)$ -dimensional holes. Betti numbers are also commonly used to evaluate the topological correctness. The Betti error e is defined by the averaged L_1 distance between two Betti numbers in all dimension k as

$$e_k = \|\beta_k^1 - \beta_k^2\|_1 \text{ and } e = \frac{1}{K} \sum_{k=0}^{K-1} e_k. \quad (1)$$

Euler Characteristic. $\chi \in \mathbb{Z}$ is a topological invariant that can be used to distinguish between different shapes of objects; it is commonly used to help understand and describe their structure [43]. χ is defined as the alternating sum of the number of k -dimensional components N_k , *e.g.*, the number of vertices, edges, faces, etc. χ also equals to the alternating sum of Betti numbers β_k :

$$\chi(\mathcal{K}) = \sum_{k=0}^n (-1)^k N_k = \sum_{k=0}^{n-1} (-1)^k \beta_k(\mathcal{K}). \quad (2)$$

Specifically, for 3-dimensional images $\chi(\mathcal{K}) = N_0 - N_1 + N_2 - N_3 = \beta_0 - \beta_1 + \beta_2$, where N_0 , N_1 , N_2 and N_3 are the number of vertices, edges, faces and squares.

Bit-quad-based method for χ computation. Conventional approaches such as the Gray algorithm have been developed to calculate χ in 2D binary images [12, 13]. Instead of directly counting the number of vertices N_0 , edges N_1 , and faces N_2 , the Gray algorithm proves that χ in 2D binary images equals to a linear combination of the occurrences of different 2×2 patterns $Q_i, i \in \{1, 2, \dots, 2^4\}$, called bit-quads, as shown in Fig. 3. Note that the combination of bit-quads can have more than one solution and not all the $2^4 = 16$ bit-quads will contribute to χ . For example in the Gray method, χ can be




Fig. 4. Bit-octets for 3D χ computation. $|O|$ represents the cardinality of the set O .




Fig. 5. Overview of our method. Given arbitrary pre-trained segmentation networks f and its predictions $\hat{Y}$, we generate the topology-aware masked segmentation predictions $\tilde{Y}$ (green) from the topological violation detection (TVD) block (blue), for training a topology-aware feature synthesis (TFS) network g_{θ} . During inference, the raw predictions from the test set $\hat{Y}_{\text{test}}$ are corrected by the trained TFS network g_{θ^*} . In TVD block, the χ -maps are computed as shown in the yellow region.

computed from the occurrences of only 10 bit-quads $\#Q_i$ as a linear combination f :

$$\chi = f(\#Q_i) = \frac{1}{4}(\#Q_A - \#Q_B - 2\#Q_C), \quad (3)$$

where

$$Q_A = \{Q_2, Q_3, Q_4, Q_5\} \quad (4)$$

$$Q_B = \{Q_{12}, Q_{13}, Q_{14}, Q_{15}\} \quad (5)$$

$$Q_C = \{Q_{10}, Q_{11}\}. \quad (6)$$

B. Extension of χ computation to 3D

Inspired by the Gray method for computing χ in 2D via a linear combination of the occurrences of bit-quads, we extend this concept to 3D by exploring whether a similar linear combination of the occurrences of 3D bit-octets can represent χ in 3D. We define 3D bit-octets as O_j , with their respective occurrences denoted as $\#O_j$.

The problem then reduces to determining whether there exists a coefficient vector $\omega = [\omega_1, \omega_2, \dots, \omega_N]$ such that for any 3D binary segmentation $\mathbf{Y}_i$, the following equation satisfies:

$$\chi_i = \sum_{j=1}^N \#O_{i,j} \omega_j, \quad (7)$$

where N is the number of bit-octets, χ_i is the Euler characteristic of segmentation $\mathbf{Y}_i$, and $\#O_{i,j}$ is the occurrence of the j^{th} bit-octet in segmentation $\mathbf{Y}_i$. Binary $2 \times 2 \times 2$ bit-octets have $N = 2^8$ different combinations. However, since χ is a topological invariant that it remains unchanged under affine transformations including rotation and flipping. Consequently, symmetric bit-octets must share the same coefficients ω_j . By categorizing bit-octets, the number of independent ω_j is reduced to $N = 22$, denoted as $\#O = [\#O_1, \#O_2, \dots, \#O_{22}]$. The representative bit-octets of each group are illustrated in Fig. 4.

Empirical Verification and Optimization. To verify the existence of a valid coefficient set $\omega = [\omega_1, \omega_2, \dots, \omega_{22}]$, we generate M pairs of $\{\chi_i, \#O_i\}$ from M 3D binary volumes, ensuring $M \gg N$. We solve the overdetermined linear system to obtain a feasible combination of coefficients that satisfies all M samples:

$$\chi = \frac{1}{8}(\#O_A + 2\#O_B + 3\#O_C + 4\#O_D - \#O_E - 2\#O_F - 3\#O_G - 5\#O_H), \quad (8)$$

where $O_A = \{O_2, O_{16}, O_{21}\}$, $O_B = \{O_{19}, O_{20}\}$, $O_C = \{O_{17}\}$, $O_D = \{O_{14}\}$, $O_E = \{O_6, O_8, O_{15}\}$, $O_F = \{O_4, O_{10}, O_{11}\}$, $O_G = \{O_7\}$, and $O_H = \{O_5\}$. As a result, we found a linear combination of bit-octets to calculate 3D χ that satisfies all samples in our dataset. This derived formula enables efficient computation of χ in 3D binary segmentations.

C. Realization

We illustrate our method in Fig. 5, which consists of three components: (1) a segmentation network, (2) a *topological violation detection* (TVD) block and (3) a *topology-aware feature synthesis* (TFS) network.

In (1), we utilize a U-Net to predict a segmentation probability map $\hat{\mathbf{Y}} \in [0, 1]^{h \times w \times d}$, with the supervision of GT segmentation $\mathbf{Y} \in \{0, 1\}^{h \times w \times d}$. The predicted segmentation map may contain violations of a given topological prior. (2) is the main contribution of our approach, consisting of two sub-nets: a χ -Net and a *Visualization Net*. The χ -Net takes the predicted $\hat{\mathbf{Y}}$ and GT segmentation $\mathbf{Y}$ as inputs and predicts their corresponding χ -maps $\hat{\mathbf{X}} \in \mathbb{Z}^{h \times w \times d}$ and $\mathbf{X} \in \mathbb{Z}^{h \times w \times d}$. We then measure the χ error $e \in \mathbb{R}$ as L_1 distance with $e = \|\mathbf{X} - \hat{\mathbf{X}}\|_1$. The Visualization Net takes e as input and produces a topology violation map $\mathbf{V} \in \mathbb{R}^{h \times w \times d}$, which highlights the regions with topological errors. Finally, in (3), we design a TFS network, which learns to fill the missing information in the erroneous regions. This subnetwork takes the predicted segmentation $\hat{\mathbf{Y}}$ and violation map $\mathbf{V}$ as input and generates a topology-preserving segmentation, which is the final prediction.

Topology-violation detection (TVD). TVD consists of two parts: a χ -Net and a Visualization Net (Fig. 5 bottom). The Gray algorithm to calculate χ cannot be directly integrated into the gradient-based optimization process as is not differentiable. To overcome this problem, we propose a DL-compatible χ -Net as a CNN-based method that leverages the Gray algorithm to calculate χ . The χ -Net serves as a function that maps the segmentation space to the χ space $g : \mathbf{Y} \rightarrow \chi$. It is worth mentioning that in order to preserve spatial information, for each 2D segmentation $\mathbf{Y}$, χ -Net locally produces Euler numbers χ_{ij} with the input of a segmentation patch $\mathbf{P}_{ij}^\Delta = \mathbf{Y}_{i:i+\Delta, j:j+\Delta}$, where Δ is the patch size. We can therefore obtain a χ -map $\mathbf{X}$ with each of its elements χ_{ij} . Similar process is also applied for 3D segmentation.

χ -Net consists of three parts: 1) fixed kernel CNN layers, 2) an average pooling layer and 3) a linear transformation f . For 2D segmentations, we first utilize three CNN layers with fixed kernels to localize the bit-quads in the segmentation following the Gray algorithm. The values of the kernels are the same as the bit-quads $\mathbf{K}_A \in \{-1, 1\}^{2 \times 2 \times 4}$, $\mathbf{K}_B \in \{-1, 1\}^{2 \times 2 \times 4}$, and $\mathbf{K}_C \in \{-1, 1\}^{2 \times 2 \times 2}$, as shown in Fig. 5. Note that we first binarize the prediction probability map $\hat{\mathbf{Y}}$ and further normalize it to $\{-1, 1\}$. Therefore, if and only if the prediction has the same pattern as the bit-quads, it will be activated to 4 after convolution. Subsequently, we apply an average pooling layer to obtain the local number of bit-quads $\#Q_{l,ij}$ at coordinate (i, j) . The process can be summarized as

$$\#Q_{l,ij} = \Delta^2 \cdot \text{AvgPool}(\mathbb{I}(\mathbf{P}_{ij}^\Delta * \mathbf{K}_l = 4)), \quad (9)$$

where $l \in \{A, B, C\}$ is the index of 2D kernel groups, $*$ represents the convolutional operation, and $\mathbb{I}(\cdot)$ is an indicator function that equals 1 if and only if the input is true. Note that the patch size of average pooling is the same as the patch size of the segmentation Δ . Finally, following Eq. 3, a linear transformation is used to calculate the χ -map $\mathbf{X}$ with each of its element as $\chi_{ij} = \frac{1}{4}(\#Q_{A,ij} - \#Q_{B,ij} - 2\#Q_{C,ij})$. During training, we separately take both $\mathbf{Y}$ and $\hat{\mathbf{Y}}$ as the input of χ -Net and obtain their corresponding χ -maps $\mathbf{X}$ and $\hat{\mathbf{X}}$. Similar process is also applied to 3D bit-octets.

In the second part of TVD, we measure the χ error as L_1 distance with $e = \|\mathbf{X} - \hat{\mathbf{X}}\|_1$. To obtain a high-resolution

indication map that highlights the regions of the input map contributing to the χ error, we calculate the gradient of e with respect to the segmentation maps. This gradient $\mathbf{V} = \partial e / \partial \mathbf{Y}$, referred to as the topology violation map, provides a detailed visualization of the χ error and is the output of TVD.

Topology-aware feature synthesis (TFS). Here we aim to improve the topological correctness of segmentation by utilizing the detected topology violation map. The motivation is that topological errors are often caused by poor feature representation in the input image, such as blurry boundary regions. These errors are difficult to correct when training directly from the image space. To address this, we propose a TFS network that explicitly learns how to repair the topological structures from the segmentation itself.

During training, the topology violation map is first binarized by a threshold t to highlight regions with topology violation:

$$\mathbf{V}^t = \mathbb{I}(\mathbf{V} \geq t), \quad (10)$$

where a lower threshold results in more pixels being classified as topologically incorrect. Then, we construct the masked segmentation map $\tilde{\mathbf{Y}}$ (with its element $\tilde{Y}_{i,j}$) by replacing pixels identified in $\mathbf{V}^t$ with Gaussian noise, while keeping the unmasked pixels unchanged. The noise $\epsilon \sim \mathcal{N}(0,1)$ is normalized to the range $(0,1)$ using a Sigmoid function:

$$\tilde{Y}_{i,j} = \begin{cases} \frac{1}{1 + e^{-\epsilon}}, & V_{i,j}^t = 1, \\ \hat{Y}_{i,j}, & V_{i,j}^t = 0. \end{cases} \quad (11)$$

The resulting masked segmentation map $\tilde{\mathbf{Y}}$ is then fed into the TFS network as input, which is trained to restore the masked regions under GT supervision with correct topology.

During inference, we feed the feature synthesis network with real predictions $\hat{\mathbf{Y}}$. We show evidence for the effectiveness of this design in Section IV.

IV. EVALUATION

A. Datasets

We conducted experiments on the following datasets:

- 1) Circuit Reconstruction from Electron Microscopy Images (CREMI) [7] is a dataset with electron microscopy images for neuron boundary segmentation. It consists of 3 subsets, each of which consists of 125 1250×1250 gray-scale images and corresponding label maps. We randomly selected 90 samples from each subset (270 samples in total) as the training set, 10 samples from each subset (30 samples in total) as the validation set, and use the remaining 25 samples (75 samples in total) as test set. We further divided each 1250×1250 image into 25 256×256 patches with an overlap of 8 pixels, in order to fit the GPU memory and enlarge the size of the training set. Thus, training, validation, and test sets consist of 6,750, 750 and 1,875 samples, respectively.
- 2) ISBI13 [44] consists of 100 electron microscopy images with neurites segmentation ground-truth. The resolution of each slice is 1024×1024 . To enlarge the dataset, we crop each slice into 16 256×256 patches, and split it

into 1200, 100, and 300 for training, validation and test, respectively.

- 3) Fundus Image Vessel Segmentation (FIVES) dataset [45] is a dataset of fundus images for retinal vessel segmentation. The dataset consists of 800 samples, and the origin resolution for each sample is 2048×2048 . To enlarge the dataset and keep the basic topological structure of vessels, we crop each slice into 4 1024×1024 patches. In total we have 2000, 400 and 800 slices for training, validation and test, respectively.
- 4) The Developing Human Connectome Project (dHCP²) dataset has 242 fetal brain T2 Magnetic Resonance Imaging (MRI) scans with gestational ages ranging from 20.6 to 38.2 weeks. All MR images were motion corrected and reconstructed to 0.8 mm isotropic resolution for the fetal head region of interest (ROI) [46, 47]. The images are affinely aligned to the MNI-152 coordinate space and clipped to the size of $144 \times 192 \times 192$. We randomly split the data into 145 samples for training, 73 for testing, and 24 for validation. We conduct cortical gray matter segmentation on this dataset that the label is first generated by DrawEM method [48] and then refined manually to improve the segmentation accuracy.
- 5) Topology-Aware Anatomical Segmentation of the Circle of Willis (TopCoW) dataset is the first 3D dataset with voxel-level annotations for the vessels of the Circle of Willis [49]. It contains magnetic resonance angiography (MRA) data with a magnetic field strength of 3 Tesla or 1.5 Tesla. We randomly split the 90 publicly available labeled volumes into 55 for training, 5 for validation, and 30 for testing. The training and evaluation are done on the ROI regions provided.

These datasets are used to evaluate different topology challenges, where neuron boundaries in CREMI [7] and ISBI13 [44] have a random and diverse number of holes β_1 and connected components β_0 , while the cortex in the dHCP dataset is expected to have fixed β_1 and β_0 , and the vessel structure in FIVES [45] and TopCoW [49] should remain connected tubular structures during segmentation.

B. Implementation

The architecture of the TFS network uses skip connection like a U-Net [16], with 3 blocks of the encoder layers and 3 blocks of the decoder layers. Each encoder block applies a downsampling operation, progressively reducing the spatial resolution to $1/8$ of the original input size, with the bottleneck feature containing 64 channels. ReLU activation is used in all convolutional layers except the final layer, which applies a sigmoid activation for segmentation. The TFS network is trained by the cross-entropy loss.

The training process of our pipeline is two steps, i.e., the segmentation network is first trained to obtain coarse segmentations, then the TFS network. Note that the parameters in TVD are fixed. We apply the Adam optimizer [50] with a weight decay of $1 \times e^{-4}$ to train the TFS network. The initial

²www.developingconnectome.org

TABLE I

SEGMENTATION PERFORMANCE OF OUR METHOD ON CREMI, ISBI13, FIVES, dHCP, AND TOPCoW DATASETS WITH TWO ABLATION EXPERIMENTS (OURS WITHOUT TVD AND TFS). e , e_0 AND e_1 REPRESENT BETTI ERROR, BETTI ERROR 0 AND BETTI ERROR 1. ASD IS AVERAGE SURFACE DISTANCE. MEAN VALUE AND STANDARD DEVIATION ARE REPORTED IN EACH TERM. STATISTICAL SIGNIFICANCE LEVELS ARE REPORTED BY P-VALUE. BEST VALUES ARE HIGHLIGHTED IN **BOLD**.

Data	Method	Dice $\uparrow$	e $\downarrow$	e_0 $\downarrow$	e_1 $\downarrow$	ASD $\downarrow$
CREMI (2D)	cl-Dice loss [25]	83.48 $\pm$ 7.31 ($p<0.001$)	10.403 $\pm$ 5.814 ($p<0.001$)	7.445 $\pm$ 5.224 ($p<0.001$)	2.958 $\pm$ 2.975 ($p<0.001$)	1.612 $\pm$ 4.448 ($p<0.001$)
	Boundary loss [24]	84.00 $\pm$ 7.20 ($p<0.001$)	9.877 $\pm$ 5.506 ($p<0.001$)	6.938 $\pm$ 4.840 ($p<0.001$)	2.939 $\pm$ 3.025 ($p<0.001$)	1.612 $\pm$ 4.351 ($p<0.001$)
	PH loss [8]	84.07 $\pm$ 7.30 ($p<0.001$)	9.103 $\pm$ 5.175 ($p<0.001$)	6.213 $\pm$ 4.667 ($p<0.001$)	2.889 $\pm$ 2.986 ($p<0.001$)	1.690 $\pm$ 4.607 ($p<0.001$)
	Warp loss [36]	83.03 $\pm$ 7.32 ($p<0.001$)	11.781 $\pm$ 6.572 ($p<0.001$)	8.877 $\pm$ 6.066 ($p<0.001$)	2.905 $\pm$ 2.996 ($p<0.001$)	1.634 $\pm$ 4.453 ($p<0.001$)
	Swin-Unet [17]	84.13 $\pm$ 5.24 ($p<0.001$)	7.127 $\pm$ 4.770 ($p<0.001$)	4.878 $\pm$ 4.044 ($p<0.001$)	2.249 $\pm$ 2.454 ($p<0.001$)	0.756 $\pm$ 2.562 ($p<0.001$)
	MedSAM [4]	82.24 $\pm$ 8.92 ($p<0.001$)	8.794 $\pm$ 25.827 ($p<0.001$)	6.483 $\pm$ 22.048 ($p<0.001$)	2.311 $\pm$ 4.716 ($p<0.001$)	1.934 $\pm$ 6.045 ($p<0.001$)
	Ours w/o TVD+TFS	83.92 $\pm$ 7.11 ($p<0.001$)	10.787 $\pm$ 5.985 ($p<0.001$)	7.926 $\pm$ 5.528 ($p<0.001$)	2.861 $\pm$ 2.913 ($p<0.001$)	1.571 $\pm$ 4.418 ($p<0.001$)
	Ours w/o TVD	85.19 $\pm$ 6.65 ($p<0.001$)	5.191 $\pm$ 3.615 ($p<0.001$)	3.009 $\pm$ 2.860 ($p<0.001$)	2.182 $\pm$ 2.417 ($p<0.001$)	1.298 $\pm$ 3.988 ($p<0.001$)
	Ours	85.24 $\pm$ 6.64	4.941 $\pm$ 3.418	2.799 $\pm$ 2.666	2.143 $\pm$ 2.404	1.307 $\pm$ 4.028
ISBI13 (2D)	cl-Dice loss [25]	80.66 $\pm$ 5.71 ($p<0.001$)	3.343 $\pm$ 2.359 ($p<0.001$)	1.813 $\pm$ 1.829 ($p<0.001$)	1.530 $\pm$ 1.372 ($p<0.001$)	0.751 $\pm$ 0.764 ($p<0.001$)
	Boundary loss [24]	80.18 $\pm$ 6.84 ($p<0.001$)	4.060 $\pm$ 2.766 ($p<0.001$)	2.567 $\pm$ 2.342 ($p<0.001$)	1.493 $\pm$ 1.340 ($p<0.001$)	0.852 $\pm$ 1.098 ($p<0.001$)
	PH loss [8]	80.26 $\pm$ 6.74 ($p<0.001$)	3.973 $\pm$ 2.752 ($p<0.001$)	2.403 $\pm$ 2.260 ($p<0.001$)	1.570 $\pm$ 1.390 ($p<0.001$)	0.783 $\pm$ 0.954 ($p<0.001$)
	Warp loss [36]	80.32 $\pm$ 6.69 ($p<0.001$)	5.420 $\pm$ 3.408 ($p<0.001$)	3.593 $\pm$ 3.151 ($p<0.001$)	1.827 $\pm$ 1.656 ($p<0.001$)	0.759 $\pm$ 0.888 ($p<0.001$)
	Swin-Unet [17]	80.18 $\pm$ 6.41 ($p<0.001$)	5.643 $\pm$ 3.784 ($p<0.001$)	2.797 $\pm$ 2.410 ($p<0.001$)	2.847 $\pm$ 2.577 ($p<0.001$)	0.514 $\pm$ 0.336 ($p<0.001$)
	MedSAM [4]	80.40 $\pm$ 5.64 ($p<0.001$)	5.190 $\pm$ 4.084 ($p<0.001$)	2.777 $\pm$ 2.857 ($p<0.001$)	2.413 $\pm$ 2.723 ($p<0.001$)	0.667 $\pm$ 0.618 ($p<0.001$)
	Ours w/o TVD+TFS	80.34 $\pm$ 6.75 ($p<0.001$)	5.947 $\pm$ 3.889 ($p<0.001$)	3.777 $\pm$ 3.200 ($p<0.001$)	2.170 $\pm$ 1.936 ($p<0.001$)	0.802 $\pm$ 0.990 ($p<0.001$)
	Ours w/o TVD	80.51 $\pm$ 6.66 ($p<0.001$)	2.957 $\pm$ 2.396 ($p<0.001$)	1.667 $\pm$ 1.828 ($p<0.001$)	1.290 $\pm$ 1.278 ($p=0.004$)	0.629 $\pm$ 0.578 ($p<0.001$)
	Ours	81.02 $\pm$ 6.30	2.330 $\pm$ 1.742	1.057 $\pm$ 1.312	1.273 $\pm$ 1.174	0.629 $\pm$ 0.579
FIVES (2D)	cl-Dice loss [25]	80.93 $\pm$ 16.31 ($p<0.001$)	47.966 $\pm$ 23.010 ($p<0.001$)	42.529 $\pm$ 21.807 ($p<0.001$)	5.438 $\pm$ 6.176 ($p<0.001$)	7.011 $\pm$ 18.422 ($p<0.001$)
	Boundary loss [24]	81.21 $\pm$ 16.80 ($p<0.001$)	44.921 $\pm$ 24.321 ($p<0.001$)	39.678 $\pm$ 22.819 ($p<0.001$)	5.244 $\pm$ 6.162 ($p=0.001$)	6.322 $\pm$ 17.955 ($p<0.001$)
	PH loss [8]	81.15 $\pm$ 16.88 ($p<0.001$)	39.345 $\pm$ 21.500 ($p<0.001$)	34.046 $\pm$ 20.063 ($p<0.001$)	5.299 $\pm$ 6.167 ($p=0.017$)	6.376 $\pm$ 18.164 ($p<0.001$)
	Warp loss [36]	81.31 $\pm$ 16.77 ($p<0.001$)	56.395 $\pm$ 33.951 ($p<0.001$)	50.474 $\pm$ 30.113 ($p<0.001$)	5.921 $\pm$ 7.074 ($p<0.001$)	6.158 $\pm$ 17.898 ($p<0.001$)
	Swin-Unet [17]	79.54 $\pm$ 12.88 ($p<0.001$)	40.124 $\pm$ 17.453 ($p<0.001$)	34.794 $\pm$ 15.467 ($p<0.001$)	5.330 $\pm$ 6.506 ($p<0.001$)	6.148 $\pm$ 19.665 ($p<0.001$)
	MedSAM [4]	80.53 $\pm$ 14.13 ($p<0.001$)	33.808 $\pm$ 13.623 ($p<0.001$)	28.793 $\pm$ 12.245 ($p<0.001$)	5.015 $\pm$ 5.901 ($p<0.001$)	6.773 $\pm$ 25.092 ($p<0.001$)
	Ours w/o TVD+TFS	81.45 $\pm$ 16.57 ($p<0.001$)	54.273 $\pm$ 31.352 ($p<0.001$)	48.733 $\pm$ 28.544 ($p<0.001$)	5.540 $\pm$ 6.526 ($p<0.001$)	6.193 $\pm$ 17.489 ($p<0.001$)
	Ours w/o TVD	83.36 $\pm$ 16.80 ($p<0.001$)	28.440 $\pm$ 16.649 ($p<0.001$)	23.402 $\pm$ 14.872 ($p<0.001$)	5.037 $\pm$ 6.032 ($p<0.001$)	5.165 $\pm$ 16.875 ($p<0.001$)
	Ours	84.07 $\pm$ 15.53	26.562 $\pm$ 15.255	21.399 $\pm$ 13.609	5.164 $\pm$ 6.183	5.039 $\pm$ 16.080
dHCP (2D)	cl-Dice loss [25]	88.26 $\pm$ 5.94 ($p<0.001$)	2.483 $\pm$ 2.573 ($p<0.001$)	1.012 $\pm$ 1.389 ($p<0.001$)	1.471 $\pm$ 1.914 ($p<0.001$)	0.230 $\pm$ 1.049 ($p<0.001$)
	Boundary loss [24]	88.52 $\pm$ 5.64 ($p<0.001$)	2.470 $\pm$ 2.574 ($p<0.001$)	0.999 $\pm$ 1.353 ($p<0.001$)	1.471 $\pm$ 1.940 ($p<0.001$)	0.210 $\pm$ 0.831 ($p<0.001$)
	PH loss [8]	88.17 $\pm$ 5.89 ($p<0.001$)	2.546 $\pm$ 2.676 ($p<0.001$)	1.075 $\pm$ 1.506 ($p<0.001$)	1.471 $\pm$ 1.936 ($p<0.001$)	0.232 $\pm$ 1.134 ($p<0.001$)
	Warp loss [36]	88.11 $\pm$ 5.99 ($p<0.001$)	2.495 $\pm$ 2.550 ($p<0.001$)	1.047 $\pm$ 1.418 ($p<0.001$)	1.448 $\pm$ 1.870 ($p<0.001$)	0.230 $\pm$ 0.927 ($p<0.001$)
	Swin-Unet [17]	87.26 $\pm$ 6.43 ($p<0.001$)	3.852 $\pm$ 3.693 ($p<0.001$)	2.280 $\pm$ 2.783 ($p<0.001$)	1.572 $\pm$ 1.965 ($p<0.001$)	0.315 $\pm$ 1.180 ($p<0.001$)
	MedSAM [4]	88.43 $\pm$ 5.66 ($p=0.001$)	2.221 $\pm$ 2.355 ($p<0.001$)	0.879 $\pm$ 1.293 ($p<0.001$)	1.342 $\pm$ 1.750 ($p<0.001$)	0.210 $\pm$ 0.858 ($p<0.001$)
	Ours w/o TVD+TFS	88.14 $\pm$ 5.87 ($p<0.001$)	2.542 $\pm$ 2.619 ($p<0.001$)	1.091 $\pm$ 1.500 ($p<0.001$)	1.451 $\pm$ 1.905 ($p<0.001$)	0.227 $\pm$ 0.834 ($p<0.001$)
	Ours w/o TVD	88.45 $\pm$ 6.14 ($p<0.001$)	2.220 $\pm$ 2.368 ($p<0.001$)	0.824 $\pm$ 1.140 ($p<0.001$)	1.395 $\pm$ 1.892 ($p<0.001$)	0.228 $\pm$ 0.922 ($p<0.001$)
	Ours	88.56 $\pm$ 6.26	2.032 $\pm$ 2.182	0.737 $\pm$ 1.050	1.295 $\pm$ 1.778	0.218 $\pm$ 0.845
TopCoW (3D)	cl-Dice loss [25]	94.48 $\pm$ 2.74 ($p=0.012$)	2.500 $\pm$ 1.544 ($p=0.028$)	1.433 $\pm$ 1.055 ($p=0.018$)	1.067 $\pm$ 1.436 ($p<0.001$)	0.113 $\pm$ 0.110 ($p<0.001$)
	Boundary loss [24]	94.55 $\pm$ 2.62 ($p=0.016$)	2.567 $\pm$ 1.564 ($p=0.009$)	1.600 $\pm$ 1.200 ($p=0.002$)	0.967 $\pm$ 1.402 ($p<0.001$)	0.110 $\pm$ 0.005 ($p<0.001$)
	PH loss [8]	94.62 $\pm$ 2.41 ($p=0.454$)	2.600 $\pm$ 1.705 ($p=0.018$)	1.600 $\pm$ 1.405 ($p=0.001$)	1.000 $\pm$ 1.461 ($p=0.005$)	0.099 $\pm$ 0.079 ($p=0.045$)
	Warp loss [36]	93.99 $\pm$ 3.23 ($p<0.001$)	2.900 $\pm$ 2.071 ($p<0.001$)	1.867 $\pm$ 1.454 ($p<0.001$)	1.033 $\pm$ 1.494 ($p=0.529$)	0.130 $\pm$ 0.141 ($p<0.001$)
	Swin-Unet [17]	92.62 $\pm$ 2.96 ($p<0.001$)	6.667 $\pm$ 3.458 ($p<0.001$)	5.133 $\pm$ 2.837 ($p<0.001$)	1.533 $\pm$ 1.454 ($p<0.001$)	0.181 $\pm$ 0.130 ($p<0.001$)
	MedSAM [4]	91.14 $\pm$ 4.20 ($p<0.001$)	2.933 $\pm$ 2.816 ($p=0.017$)	1.900 $\pm$ 1.640 ($p=0.011$)	1.033 $\pm$ 1.643 ($p<0.001$)	0.113 $\pm$ 0.062 ($p<0.001$)
	Ours w/o TVD+TFS	94.28 $\pm$ 3.11 ($p=0.016$)	2.800 $\pm$ 2.072 ($p<0.001$)	1.767 $\pm$ 1.453 ($p<0.001$)	1.033 $\pm$ 1.779 ($p<0.001$)	0.106 $\pm$ 0.073 ($p<0.001$)
	Ours w/o TVD	94.61 $\pm$ 2.54 ($p<0.001$)	2.700 $\pm$ 1.828 ($p<0.001$)	1.567 $\pm$ 1.283 ($p<0.001$)	1.133 $\pm$ 1.707 ($p<0.001$)	0.106 $\pm$ 0.099 ($p<0.001$)
	Ours	94.63 $\pm$ 2.52	2.267 $\pm$ 1.731	1.300 $\pm$ 1.130	0.967 $\pm$ 1.683	0.103 $\pm$ 0.094

learning rate is set to $1 \times e^{-4}$ and is reduced by a factor of 2 every 50 epochs. The model is trained for 200 epochs until converged.

We choose U-Net as the segmentation backbone for our method and for cl-Dice loss [25], boundary loss [24], warp loss [36] and PH loss [8] for comparison. However, due to the modular nature of our method it can also be incorporated into other segmentation frameworks, as discussed in Section IV-F Ablation study.

We use PyTorch 1.13.1 and calculate the Betti number with the GUDHI package [35]. The training time is evaluated on an NVIDIA RTX 3080 GPU with a batch size of 20.

C. Evaluation metrics

Segmentation performance is evaluated by Dice score and averaged surface distance (ASD), and the performance of topology is evaluated by Betti errors, which is defined as: $e_i = |\beta_i^{\text{pred}} - \beta_i^{\text{gt}}|$, where $i \in \{0, 1\}$ indicates the dimension. We also report the mean Betti error as $e = e_0 + e_1$.

D. Quantitative evaluation

We compare the segmentation performance of our approach against six baseline methods. Among them, four are designed to preserve shape and topology: clDice loss [25], Boundary




Fig. 6. Generated topology violation maps and qualitative comparison of segmentation results. We compare our segmentation results with four baseline methods (cl-dice loss, Boundary loss, PH loss, and Warp loss) and two ablations in the last three columns. Our method can correct disconnected structures better than existing methods.

loss [24], Warp loss [36], and PH loss [8]. Additionally, we compare our method with two segmentation pipelines: Swin-Unet [17] and MedSAM [4]. For pixel-wise accuracy, our method achieves the best Dice score on all five datasets, as demonstrated in Tab. I. Despite the comparable ASD scores, we achieve significant improvements in terms of Betti errors. Compared to all baseline methods, our method achieves better topological accuracy in terms of Betti error. For example, the average Betti error of our method is 4.941 on the CREMI dataset, representing an improvement of 3.853 (43.8%) over MedSAM and 2.186 (30.7%) over Swin-Unet in mean Betti error. We observe similar improvements for the ISBI13, FIVES, dHCP, and TopCoW datasets. We conducted Wilcoxon signed-rank tests [51] to compare our method against other baselines

on all five datasets. As shown in Tab. I, our method is statistically better than the baseline methods with in terms of Betti error with all p-values below 0.05.

The averaged training times of our method is 0.0134 s/item on CREMI dataset, with 0.0168 s/item, 0.0151 s/item, 0.1033 s/item, and 0.9748 s/item for cl-Dice loss [25], boundary loss [24], warp loss [36] and PH loss [8], respectively. Compared to the PH-based loss [8], which is a similarly well-grounded concept in algebraic topology, our method is 72.7 times faster than the PH-based loss [8].

To assess the effect of end-to-end training, we conducted additional experiments on the ISBI13 [44] dataset, where the segmentation and TFS networks were trained jointly by updating the parameters of the segmentation network and then




Fig. 7. Generated topology violation maps and the segmentation examples. Left: ISBI13 dataset, Right: FIVES dataset.

updating the parameters of the TFS network in each iteration. As shown in Tab. II, this approach consistently reduced Betti error while maintaining the Dice similarity of TFS-refined segmentations. The improvement can be attributed to the TFS network encountering a wider range of topological errors during early training stages, enhancing its ability to refine segmentation outputs. These results suggest that end-to-end training can be further beneficial.

E. Qualitative evaluation

As highlighted in Fig. 6 and Fig. 7, our method can effectively eliminate the topological errors in both 2D and 3D datasets, outperforming all the other methods. For instance, in the first row of the CREMI dataset, none of the baseline methods could segment the tiny neuron structure when the boundary of the cells is blurry. Similarly, in the dHCP dataset, all the baseline methods fail to segment the fetal cortex as a closed surface, whereas our method can successfully resolve this issue. Also in FIVES dataset, our method can capture the tiny vessel structure as an connected component while all other methods fails. In 3D TopoCoW dataset, Second, we show the topology violation maps from the TVD block in the third column in Fig. 6, which indicate the topology error regions between the GT (second column) and the prediction from our first segmentation network (eighth column). For example, in the evaluation on TopCoW dataset, we observe that the topology violation map can highlight the hole structure with topological error, therefore driving our method to correct these errors. Similar results can be observed on the example of other datasets.

F. Ablation study

We first evaluate the effectiveness of our TVD design. We train our pipeline without the TVD block. Instead, we

use the difference map between the prediction and GT as a substitute for the topology violation map. Qualitative results are shown in Tab. I between rows ‘Ours w/o TVD’ and ‘Ours’ across five datasets. Ours statistically outperforms the ablation setting of removing the TVD block in terms of Dice similarity and Betti errors, with p-value consistently lower than 0.001, demonstrating the effectiveness of our TVD design. We also summarize the qualitative results in Fig. 6 and Fig. 7. Feature synthesis with the difference map can correct some of the false negative errors, however, structures tend to remain incorrect for all datasets. In contrast, our approach successfully corrects these topological errors. Secondly, we further remove the TFS network. Quantitative results are provided in Tab. I. The topology performance without TVD+TFS is significantly inferior to our method in terms of Betti errors, which illustrates the effectiveness of our design.

Additionally, we conducted an ablation study to assess the impact of different decision thresholds on segmentation performance using the ISBI13 [44] dataset. As shown in Fig. 8, Dice similarity first increases and then decreases as the threshold increases, while Betti error follows the opposite trend. At a high threshold $t = 0.7$, only a small fraction of pixels are masked and further corrected by the TVD block. Conversely, a low threshold ($t = 0.1$) results in excessive masking, leading to insufficient input information for TVD to refine errors. The best performance is observed when t is between 0.2 and 0.5. Thus, we sample t from this range during training to enhance the diversity of synthetic data.

To assess the impact of patch size on topology refinement, we conducted an ablation study using patch sizes of 8, 16, 32, 64, 128, and 256 on the ISBI13 dataset. As shown in Fig. 8, Betti error initially decreases and then increases with larger patches, while Dice similarity declines monotonically for patch sizes above 64. This trend reflects a trade-off in local topological representation. Larger patches may cause topological




Fig. 8. Left: comparison of training time per sample (in seconds) on the CREMI dataset, shown on a logarithmic scale. Middle and right: ablation study of different patch sizes and decision thresholds in the TVD block on the ISBI13 dataset.

TABLE II

COMPARISON OF TRAINING STRATEGIES: END-TO-END TRAINING VS. SEQUENTIAL TRAINING FOR PRE-SEGMENTATION AND POST-TOPOLOGY CORRECTION NETWORKS

	Training type	Dice $\uparrow$	$e \downarrow$	$e_0 \downarrow$	$e_1 \downarrow$	ASD $\downarrow$
Pre-segmentation	Two-stage	80.34 ± 6.75	5.947 ± 3.889	3.777 ± 3.200	2.170 ± 1.936	0.802 ± 0.990
	End-to-end	79.00 ± 7.18	4.750 ± 2.904	3.010 ± 2.413	1.740 ± 1.602	0.910 ± 1.194
Post-TFS refined	Two-stage	81.01 ± 6.30	2.330 ± 1.742	1.057 ± 1.312	1.273 ± 1.174	0.629 ± 0.579
	End-to-end	80.98 ± 6.41	2.147 ± 1.492	0.943 ± 1.010	1.203 ± 1.150	0.660 ± 0.752

TABLE III

COMPARISON OF SEGMENTATION BACKBONES: INCORPORATING OUR TVD AND TFS NETWORKS INTO UNET [16] AND nnU-NET [3].

Data	Method	Dice $\uparrow$	$e \downarrow$	$e_0 \downarrow$	$e_1 \downarrow$	ASD $\downarrow$
ISBI13	U-Net	80.34 ± 6.75	5.947 ± 3.889	3.777 ± 3.200	2.170 ± 1.936	0.802 ± 0.990
	U-Net + Ours	81.02 ± 6.30	2.330 ± 1.742	1.057 ± 1.312	1.273 ± 1.174	0.629 ± 0.579
	nnU-Net	84.02 ± 5.55	1.307 ± 1.125	0.470 ± 0.665	0.837 ± 0.900	0.366 ± 0.473
	nnU-Net + Ours	84.23 ± 5.49	1.243 ± 1.091	0.457 ± 0.684	0.787 ± 0.865	0.365 ± 0.490
TopCoW	U-Net	94.28 ± 3.11	2.800 ± 2.072	1.767 ± 1.453	1.033 ± 1.779	0.106 ± 0.073
	U-Net + Ours	94.63 ± 2.52	2.267 ± 1.731	1.300 ± 1.130	0.967 ± 1.683	0.103 ± 0.094
	nnU-Net	95.15 ± 3.02	1.700 ± 1.696	0.700 ± 0.781	1.000 ± 1.693	0.074 ± 0.072
	nnU-Net + Ours	94.81 ± 2.66	1.433 ± 1.230	0.733 ± 0.727	0.700 ± 1.159	0.074 ± 0.066

errors in different regions to cancel out, while overly small patches introduce local Euler characteristic variations due to boundary mismatches between predictions and ground truth, despite the global topology remaining unchanged. We thereby selected a patch size of 32 as the optimal hyperparameter in our experiments.

To evaluate the generalizability of our method, we conduct experiments with an alternative segmentation backbone by replacing UNet [16] with nnU-Net [3]. We trained the TFS network using the predictions from nnU-Net. As shown in Tab. III, our method consistently improves the segmentation performance in terms of Betti error on both ISBI13 [44] and TopCoW [49] datasets, while maintaining Dice similarity. These results demonstrate the robustness of our design regardless of the segmentation backbone.

G. Discussion

We illustrate the effectiveness of our approach across five distinct datasets featuring varied topological structures. Our results uniformly show that our method not only addresses topological inaccuracies but also surpasses all comparative baseline methods in both quantitative and qualitative assessments. Additionally, we use the numerical Betti error for topological evaluation and employ a topology violation map

derived from χ to provide a visual representation of topological accuracy. This visual tool supports our TFS network in mitigating topological errors. An ablation study further substantiates the robustness of our method's design.

However, our approach retains a vulnerability to noise, as clearer image features and stronger image gradients yield better results. Typically, all topology-aware methods are also prone to errors from labelling noise; even minor mislabeling within an object can significantly disrupt topology. Strategies to mitigate this include label quality control or verification by multiple experts, but these measures are costly. Our method can be adapted to anticipate specific topological structures, such as being homeomorphic to a sphere, where the expected χ and Betti numbers are predefined.

We can extend our technique to address multi-class segmentation tasks by calculating χ for each class individually, similar to PH methods. While the computational time for both χ and PH increases linearly with more categories, our approach is more efficient per class and this efficiency is further amplified in multi-class segmentation.

Addressing topology characterization within non-linear yet Lipschitz continuous frameworks, like deep learning-based image segmentation, poses significant challenges. Counting functions are typically neither continuous nor differentiable. We overcome this by implementing fixed forward transforma-

tions with bit-quads and bit-octets to determine Betti numbers and χ . Learning directly through backpropagation may be more elegant, but this results in non-local image-level error functions and relies on the computationally intensive continuous differentiation of filtration levels, e.g., in strategies based on PH. Thus, further exploration into differentiable and localised representations of χ could be a fruitful future research direction.

V. CONCLUSION

In this work, we propose a method aimed at improving and evaluating topology-aware medical image segmentation. We observe that most existing methods either do not consider topological information, or are limited by their high computational complexity. As a computationally efficient approach, our method can be easily integrated into existing segmentation methods to improve the topological structure.

We propose a novel χ -based method to include topology constraints in the segmentation network. Different from PH-based approaches, our method has a distinct advantage in computational efficiency while providing improved performance.

We also generate a topology violation map, which can aid interpretability, and use this map directly in a post-processing feature synthesis network. We believe this map is valuable and can be applied to other tasks. In image classification and disease diagnosis, topological structures could serve as discriminative features. Additionally, in generative tasks where topological priors are required, such as histopathology image generation [32, 33], topology violation map can be used to assess the topological correctness of generated images and regularize loss functions accordingly.

Future work will focus on clinical translation and downstream evaluation for surface modelling and, e.g., flow or growth simulation applications.

REFERENCES

- [1] L. Maier-Hein, B. Menze et al., "Metrics reloaded: Pitfalls and recommendations for image analysis validation," *arXiv.org*, no. 2206.01653, 2022.
- [2] A. M. Dale, B. Fischl, and M. I. Sereno, "Cortical surface-based analysis: I. segmentation and surface reconstruction," *Neuroimage*, vol. 9, no. 2, pp. 179–194, 1999.
- [3] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [4] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," *Nature Communications*, vol. 15, no. 1, p. 654, 2024.
- [5] R. Wang, S. S. Chow, R. B. Serafin, W. Xie, Q. B. Han, E. Baraznenok, L. Lan, K. W. Bishop, and J. T. Liu, "Direct three-dimensional segmentation of prostate glands with nnu-net," *Journal of Biomedical Optics*, vol. 29, no. 3, pp. 036 001–036 001, 2024.
- [6] L. Huo, X. Hu, Q. Xiao, Y. Gu, X. Chu, and L. Jiang, "Segmentation of whole breast and fibroglandular tissue using nnu-net in dynamic contrast enhanced mr images," *Magnetic Resonance Imaging*, vol. 82, pp. 31–41, 2021.
- [7] J. Funke, F. Tschopp, W. Grisaitis, A. Sheridan, C. Singh, S. Saalfeld, and S. C. Turaga, "Large scale image segmentation with structured loss based deep learning for connectome reconstruction," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1669–1680, 2018.
- [8] X. Hu, F. Li, D. Samaras, and C. Chen, "Topology-preserving deep image segmentation," *Advances in neural information processing systems*, vol. 32, 2019.
- [9] J. R. Clough, N. Byrne, I. Oksuz, V. A. Zimmer, J. A. Schnabel, and A. P. King, "A topological loss function for deep-learning based image segmentation using persistent homology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 8766–8778, 2020.
- [10] J. R. Clough, I. Oksuz, N. Byrne, J. A. Schnabel, and A. P. King, "Explicit topological priors for deep-learning based image segmentation using persistent homology," in *Information Processing in Medical Imaging: 26th International Conference, IPMI 2019, Hong Kong, China, June 2–7, 2019, Proceedings 26*. Springer, 2019, pp. 16–28.
- [11] N. Stucki, J. C. Paetzold, S. Shit, B. Menze, and U. Bauer, "Topologically faithful image segmentation via induced matching of persistence barcodes," in *International Conference on Machine Learning*. PMLR, 2023, pp. 32 698–32 727.
- [12] S. B. Gray, "Local properties of binary images in two dimensions," *IEEE Transactions on computers*, vol. 100, no. 5, pp. 551–561, 1971.
- [13] B. Yao, L. He, S. Kang, Y. Chao, and X. Zhao, "A novel bit-quad-based euler number computing algorithm," *Springerplus*, vol. 4, pp. 1–16, 2015.
- [14] F. Arce Vega, J. H. Sossa Azuela, W. Gómez Flores, and L. G. Lira Vargas, "Learning an artificial neural network to discover combinations of bit-quads to compute the euler characteristic of a 2-d binary image," *Computación y Sistemas*, vol. 26, no. 1, 2022.
- [15] L. Li, Q. Ma, C. Ouyang, Z. Li, Q. Meng, W. Zhang, M. Qiao, V. Kyriakopoulou, J. V. Hajnal, D. Rueckert et al., "Robust segmentation via topology violation detection and feature synthesis," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023, pp. 67–77.
- [16] O. Ronneberger, P. Fischer, and T. Brox, "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*. Springer, 2015, pp. 234–241.
- [17] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, and M. Wang, "Swin-unet: Unet-like pure transformer for medical image segmentation," in *European conference on computer vision*. Springer, 2022, pp. 205–218.
- [18] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo et al., "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.
- [19] J. Zhu, A. Hamdi, Y. Qi, Y. Jin, and J. Wu, "Medical sam 2: Segment medical images as video via segment anything model 2," 2024.
- [20] Q. Xu, J. Li, X. He, Z. Liu, Z. Chen, W. Duan, C. Li, M. M. He, F. B. Tesema, W. P. Cheah et al., "Esp-medsam: Efficient self-prompting sam for universal domain-generalized medical image segmentation," *arXiv preprint arXiv:2407.14153*, 2024.
- [21] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: a survey," *Medical image analysis*, vol. 1, no. 2, pp. 91–108, 1996.
- [22] M. C. H. Lee, K. Petersen, N. Pawlowski, B. Glocker, and M. Schaap, "Tetris: Template transformer networks for image segmentation with shape priors," *IEEE transactions on medical imaging*, vol. 38, no. 11, pp. 2596–2606, 2019.
- [23] O. Oktay, E. Ferrante, K. Kamnitsas, M. Heinrich, W. Bai, J. Caballero, S. A. Cook, A. De Marva, T. Dawes, D. P. O'Regan et al., "Anatomically constrained neural networks (acnns): application to cardiac image enhancement and segmentation," *IEEE transactions on medical imaging*, vol. 37, no. 2, pp. 384–395, 2017.
- [24] H. Kervade, J. Bouchtiba, C. Desrosiers, E. Granger, J. Dolz, and I. B. Ayed, "Boundary loss for highly unbalanced segmentation," in *International conference on medical imaging with deep learning*. PMLR, 2019, pp. 285–296.
- [25] S. Shit, J. C. Paetzold, A. Sekuboyina, I. Ezhov, A. Unger, A. Zhylka, J. P. Pluim, U. Bauer, and B. H. Menze, "cldice-a novel topology-preserving loss function for tubular structure segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 560–16 569.
- [26] Y. Kirchhoff, M. R. Rokuss, S. Roy, B. Kovacs, C. Ulrich, T. Wald, M. Zenk, P. Vollmuth, J. Kleesiek, F. Isensee et al., "Skeleton recall loss for connectivity conserving and resource efficient segmentation of thin tubular structures," in *European Conference on Computer Vision*. Springer, 2024, pp. 218–234.
- [27] L. Li, Q. Ma, Z. Li, C. Ouyang, W. Zhang, A. Price, V. Kyriakopoulou, L. C. Grande, A. Makropoulos, J. Hajnal et al., "Fetal cortex segmentation with topology and thickness loss constraints," in *Ethical and Philosophical Issues in Medical Imaging, Multimodal Learning and Fusion Across Scales for Clinical Decision Support, and Topological Data*

- Analysis for Biomedical Imaging: 1st International Workshop, EPIMI 2022, 12th International Workshop, ML-CDS 2022, 2nd International Workshop, TDA4BiomedicalImaging, Held in Conjunction with MICCAI 2022, Singapore, September 18–22, 2022, Proceedings.* Springer, 2022, pp. 123–133.
- [28] P. de Dumast, H. Kebiri, C. Atat, V. Dunet, M. Koob, and M. B. Cuadra, "Segmentation of the cortical plate in fetal brain mri with a topological loss," in *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis: 3rd International Workshop, UNSURE 2021, and 6th International Workshop, PIPPI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, October 1, 2021, Proceedings 3.* Springer, 2021, pp. 200–209.
- [29] N. Byrne, J. R. Clough, I. Valverde, G. Montana, and A. P. King, "A persistent homology-based topological loss for cnn-based multiclass segmentation of cmr," *IEEE Transactions on Medical Imaging*, vol. 42, no. 1, pp. 3–14, 2022.
- [30] A. H. Berger, N. Stucki, L. Lux, V. Buerger, S. Shit, A. Banaszak, D. Rueckert, U. Bauer, and J. C. Paetzold, "Topologically faithful multi-class segmentation in medical images," *arXiv preprint arXiv:2403.11001*, 2024.
- [31] A. Santhirasekaram, M. Winkler, A. Rockall, and B. Glocker, "Topology preserving compositionality for robust medical image segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 543–552.
- [32] S. Abousamra, R. Gupta, T. Kurc, D. Samaras, J. Saltz, and C. Chen, "Topology-guided multi-class cell context generation for digital pathology," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 3323–3333.
- [33] M. Xu, S. Gupta, X. Hu, C. Li, S. Abousamra, D. Samaras, P. Prasanna, and C. Chen, "Topocellgen: Generating histopathology cell topology with a diffusion model," *arXiv preprint arXiv:2412.06011*, 2024.
- [34] S. Kaji, T. Sudo, and K. Ahara, "Cubical ripser: Software for computing persistent homology of image and volume data," *arXiv:2005.12692*, 2020.
- [35] C. Maria, J.-D. Boissonnat, M. Glisse, and M. Yvinec, "The gudhi library: Simplicial complexes and persistent homology," in *Mathematical Software–ICMS 2014: 4th International Congress, Seoul, South Korea, August 5–9, 2014. Proceedings 4.* Springer, 2014, pp. 167–174.
- [36] X. Hu, "Structure-aware image segmentation with homotopy warping," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 046–24 059, 2022.
- [37] X. Hu, Y. Wang, L. Fuxin, D. Samaras, and C. Chen, "Topology-aware segmentation using discrete morse theory," *arXiv preprint arXiv:2103.09992*, 2021.
- [38] L. Wasserman, "Topological data analysis," *Annual Review of Statistics and Its Application*, vol. 5, pp. 501–532, 2018.
- [39] S. Eilenberg and N. Steenrod, *Foundations of algebraic topology*. Princeton University Press, 2015.
- [40] F. Chazal and B. Michel, "An introduction to topological data analysis: fundamental and practical aspects for data scientists," *Frontiers in artificial intelligence*, vol. 4, p. 108, 2021.
- [41] E. Munch, "A user's guide to topological data analysis," *Journal of Learning Analytics*, vol. 4, no. 2, pp. 47–61, 2017.
- [42] D. Freedman and C. Chen, "Algebraic topology for computer vision," *Computer Vision*, pp. 239–268, 2009.
- [43] E. Richardson and M. Werman, "Efficient classification using the euler characteristic," *Pattern Recognition Letters*, vol. 49, pp. 99–106, 2014.
- [44] I. Arganda-Carreras, S. Seung, A. Vishwanathan, and D. Berger, "Snemi 3d: 3d segmentation of neurites in em images," 2013.
- [45] K. Jin, X. Huang, J. Zhou, Y. Li, Y. Yan, Y. Sun, Q. Zhang, Y. Wang, and J. Ye, "Fives: A fundus image dataset for artificial intelligence based vessel segmentation," *Scientific Data*, vol. 9, no. 1, p. 475, 2022.
- [46] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel, "Reconstruction of fetal brain mri with intensity matching and complete outlier removal," *Medical image analysis*, vol. 16, no. 8, pp. 1550–1564, 2012.
- [47] B. Kainz, M. Steinberger, W. Wein, M. Kuklisova-Murgasova, C. Malamatoniou, K. Keraudren, T. Torsney-Weir, M. Rutherford, P. Aljabar, J. V. Hajnal *et al.*, "Fast volume reconstruction from motion corrupted stacks of 2d slices," *IEEE transactions on medical imaging*, vol. 34, no. 9, pp. 1901–1913, 2015.
- [48] A. Makropoulos, E. C. Robinson, A. Schuh, R. Wright, S. Fitzgibbon, J. Bozek, S. J. Counsell, J. Steinweg, K. Vecchiato, J. Passerat-Palmbach *et al.*, "The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction," *Neuroimage*, vol. 173, pp. 88–112, 2018.
- [49] K. Yang, F. Musio, Y. Ma, N. Juchler, J. C. Paetzold, R. Al-Maskari, L. Höher, H. B. Li, I. E. Hamamci, A. Sekuboyina *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.
- [50] D. P. Kingma, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [51] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in Statistics: Methodology and Distribution*. Springer, 1992, pp. 196–202.