HYPERGCN-VESSEL: HYBRID GEOMETRIC REPRESENTATION LEARNING FOR TOPOLOGY-PRESERVING VASCULAR SEGMENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Accurate vascular segmentation requires not only high voxel-wise accuracy but also the preservation of the underlying vessel topology. Although recent deep learning models achieve impressive Dice scores, they frequently produce fragmented predictions on thin and distal vessels, limiting their utility in downstream clinical analysis. We attribute this limitation to the mismatch between conventional Euclidean feature representations and the hierarchical organization of vascular trees. To address this challenge, we propose HyperGCN-Vessel, a hybrid Euclidean–hyperbolic framework that combines efficient local feature extraction with geometry-aware representation learning in hyperbolic space. Specifically, the proposed framework introduces three key components: (1) an early hybrid geometric encoder that progressively embeds vascular representations into the Poincaré ball to better capture hierarchical structures; (2) Topo-Gate, an uncertainty-aware differentiable topology-preserving module that recovers disconnected vessel branches without explicit skeletonization; and (3) HyperConv, an efficient linear-complexity tangent-space convolution operator that enables scalable and numerically stable hyperbolic feature learning for high-resolution 3D volumes. Extensive experiments on the TubeTK dataset demonstrate that HyperGCN-Vessel substantially improves vascular connectivity while maintaining competitive segmentation accuracy, reducing the β_0 Betti number error by 65% compared with nnU-Net and achieving inference efficiency comparable to conventional Euclidean architectures.

1 INTRODUCTION

Accurate 3D vascular segmentation is a fundamental prerequisite for numerous clinical applications, including hemodynamic analysis, disease diagnosis, and surgical planning Litjens et al. (2017); Bullitt et al. (2003); Bulant et al. (2018); Taylor et al. (2013); Lesage et al. (2009). Although recent segmentation models, including convolutional networks, Vision Transformers, and state-space architectures, have achieved remarkable performance on volumetric benchmarks Isensee et al. (2021); Çiçek et al. (2016); Milletari et al. (2016); Hatamizadeh et al. (2022; 2021); Ruan & Xiang (2024); Xing et al. (2024); Ma et al. (2024), they continue to struggle with thin, tortuous, and low-contrast vessels. As a result, predictions often exhibit broken vessel continuity despite high Dice scores, revealing a persistent gap between voxel-wise accuracy and topological correctness. This discrepancy leads to fragmented vascular trees that are unsuitable for downstream structural analysis, where preserving connectivity is often more important than maximizing volumetric overlap.

Existing approaches primarily improve vascular connectivity by incorporating explicit topological priors into the learning objective. Representative methods enforce persistent homology consistency Hu et al. (2019); Clough et al. (2022); Stucki et al. (2023); Gupta et al. (2023); He et al. (2023); Xu et al. (2025), exploit differentiable skeleton representations such as clDice Shit et al. (2021); Mosinska et al. (2018); Rougé et al. (2024); Zhou et al. (2025a), or design topology-aware feature learning strategies through directional connectivity, dynamic convolutions, or fractal representations Qi et al. (2023a); Zhou et al. (2025b).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

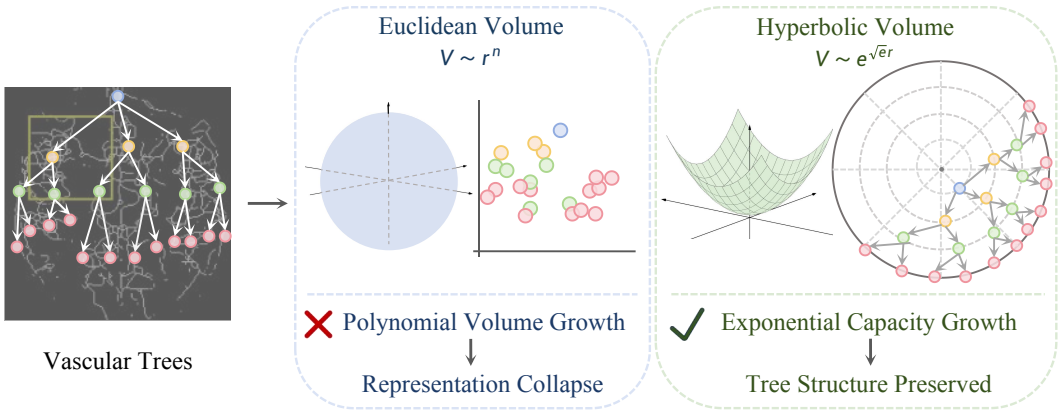

Figure 1: Illustration of representation collapse in Euclidean space vs. topological fidelity in hyperbolic space. Euclidean space (Left) distorts hierarchical vascular structures due to polynomial volume growth, whereas hyperbolic space (Right) provides exponential capacity for zero-overlap topological embedding.

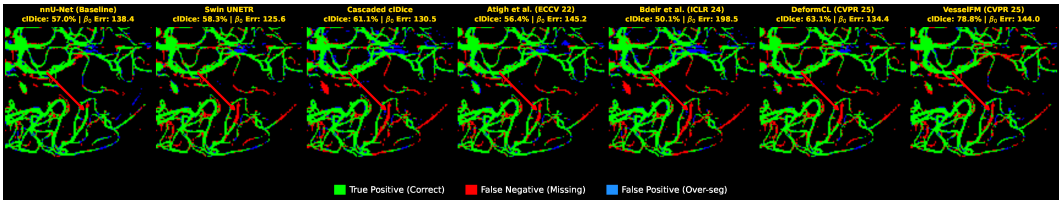

Figure 2: Qualitative comparison of topological fragmentation in state-of-the-art vascular segmentation baselines. Red circles highlight critical branch disconnections (β_0 errors) occurring despite competitive volumetric overlap, illustrating the "topological gap" where fine capillaries fail to form a continuous structural graph.

Hyperbolic geometry provides a natural alternative for representing hierarchical structures owing to its negatively curved geometry and exponentially expanding embedding space Ganea et al. (2018)(Figure 1). Consequently, Hyperbolic Neural Networks (HNNs) have demonstrated advantages in learning tree-like and graph-structured data. However, their application to medical image segmentation remains limited. Existing approaches Atigh et al. (2022); Abian et al. (2026) either introduce hyperbolic representations only in the late stages of the network, limiting their influence on early feature extraction, or rely on computationally expensive manifold operations that hinder scalability.

Motivated by these observations, we propose HyperGCN-Vessel, a hybrid Euclidean–hyperbolic framework for topology-preserving vascular segmentation. The network first extracts local appearance features using computationally efficient Euclidean operators before progressively embedding representations into the Poincaré ball, where hierarchical vascular structures can be modeled more naturally . To make hyperbolic learning practical for dense medical image segmentation, we further design efficient topology-aware operators that preserve vascular connectivity while maintaining computational efficiency. Consequently, our contributions can be summarized as follows:

- We present a hybrid Euclidean–Hyperbolic framework for retinal vessel segmentation that performs early geometric representation learning, effectively bridging the mismatch between Euclidean image features and the underlying hierarchical vascular topology.
- We propose Topo-Gate, a differentiable topology-preserving module that leverages uncertainty-aware hyperbolic distance modulation to recover disconnected vessel branches without relying on explicit skeleton extraction.

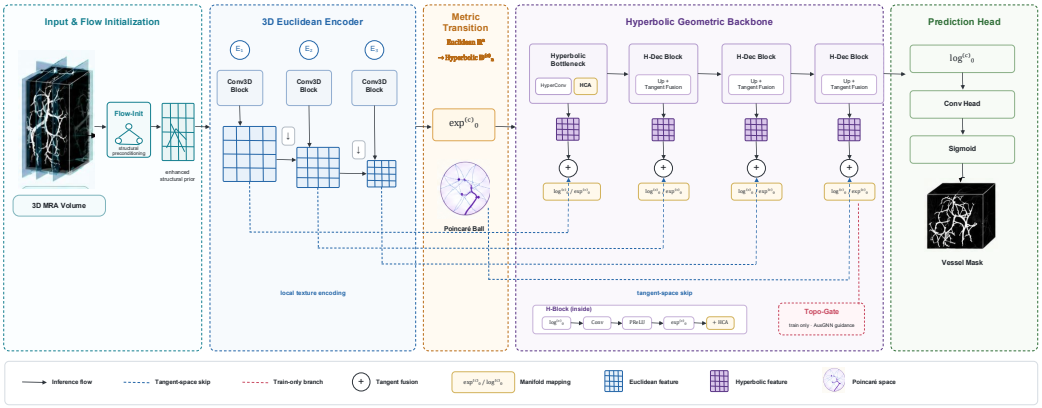

Figure 3: HyperGCN-Vessel architecture overview. The hybrid Euclidean-hyperbolic encoder-decoder processes inputs via initial data preconditioning (cyan), a shallow Euclidean encoder (blue), exponential-map projection into the Poincaré ball (purple), and HyperConv-based hyperbolic decoder with tangent-space skip connections (orange). Training jointly minimises the voxel-wise segmentation loss and the Topo-Gate topological penalty (pink).

- We introduce HyperConv, an efficient linear-complexity tangent-space convolution operator that enables scalable and numerically stable hyperbolic feature learning, making geometric deep learning practical for high-resolution retinal images.

2 Related Work

Vessel Segmentation: Euclidean Limits. nnU-Net Isensee et al. (2021) introduces a robust, self-configuring framework that automates data-adaptive pipeline design by dynamically adjusting scaling, resolution, and convolutional padding according to dataset geometry, extracting effective volumetric representations and consistently delivering state-of-the-art segmentation accuracy across diverse and heterogeneous datasets. UNETR Hatamizadeh et al. (2022) explicitly integrates Vision Transformers into the traditional 3D segmentation hierarchy, effectively learning long-range spatial dependencies via self-attention sequences applied directly to volumetric image patches, capturing global context that CNNs fail to capture. CS2-Net Mou et al. (2021) utilizes a unified generic convolutional architecture infused with customized spatial and channel attention blocks to dynamically enhance intra-class responsiveness while deploying uniquely shaped asymmetric kernels to accurately map curvilinear boundaries.

Topology-Aware Learning. Persistent Homology (PH) Hu et al. (2019) builds continuous-valued loss functions grounded entirely in Persistent Homology, comprehensively analyzing filtration critical points and penalizing structural discrepancies by strictly matching the Betti numbers between predicted and ground-truth configurations. clDice Shit et al. (2021) features a novel topology-preserving formulation calculating metric similarities along the intersection of 3D centerlines by leveraging a differentiable soft-skeleton representation to seamlessly train arbitrary deep networks via explicit connective gradients. Skeleton Recall Loss Kirchhoff et al. (2024) bypasses GPU-expensive algorithms completely, actively enforcing connectivity conservation by utilizing computationally inexpensive CPU-based tubed skeletons matched against network predictions via a soft recall loss.

Hyperbolic Neural Networks (HNNs). Hyperbolic geometry has emerged as a powerful tool for modeling hierarchical and branching data. Atigh et al. Atigh et al. (2022) reformulate pixel-level classification by substituting exclusively the deepest semantic model bottleneck with a Poincaré ball projection, effectively modeling hierarchical data constraints while leaving the majority of the shallow feature encoding strictly inside flat Euclidean constraints. Fully Hyperbolic CNN Bdeir et al. (2024) constructs an end-to-end framework enforcing negative curvature proximities strictly across

every hierarchical pooling tier, using mathematically pure Riemannian formulas to build fully hyperbolic convolutional layers, batch normalization, and multinomial logistic regression. Hyperbolic U-Net Mishra et al. (2024) demonstrates robustness by conducting the entire encoder-decoder computation specifically inside the Poincaré ball via carefully tailored hyperbolic transpose convolutions and bilinear upsampling configurations that natively respect the underlying manifold.

3 PRELIMINARIES: GEOMETRY OF THE POINCARÉ BALL

We review the Poincaré ball model and its key operators used throughout our framework.

The n-dimensional Poincaré ball with constant negative curvature $-c$ ($c > 0$) is defined as:

$$\mathbb{D}_c^n = \left\{ \mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1/\sqrt{c} \right\}, \tag{1}$$

with conformal metric $g_c^{\mathbb{D}}(\mathbf{x}) = \lambda_{\mathbf{x}}^2 g^E$, where $\lambda_{\mathbf{x}} = 2/(1 - c\|\mathbf{x}\|^2)$ is the conformal factor. As $\|\mathbf{x}\| \to 1/\sqrt{c}$, $\lambda_{\mathbf{x}} \to \infty$, inducing a "soft infinity" boundary that accommodates exponentially growing hierarchical structures. Here, $c > 0$ is the absolute value of the constant negative curvature of the manifold.

We use origin-centered exponential and logarithmic maps to switch between the manifold and the Euclidean tangent space $T_{\mathbf{0}}\mathbb{D}_c^n \cong \mathbb{R}^n$, which form the basis of our HyperConv operator. For any tangent vector $\mathbf{v}$ and $\mathbf{y} \in \mathbb{D}_c^n$, they are evaluated as (adding $\epsilon = 10^{-6}$ to the denominator to avoid division by zero):

$$\exp_{\mathbf{0}}^c(\mathbf{v}) = \tanh\left(\sqrt{c}\,\|\mathbf{v}\|\right) \frac{\mathbf{v}}{\sqrt{c}\,\|\mathbf{v}\| + \epsilon}, \quad \log_{\mathbf{0}}^c(\mathbf{y}) = \operatorname{arctanh}\left(\sqrt{c}\,\|\mathbf{y}\|\right) \frac{\mathbf{y}}{\sqrt{c}\,\|\mathbf{y}\| + \epsilon}. \tag{2}$$

The geodesic distance between two points on the Poincaré ball, used in our Topo-Gate loss, is:

$$d_{\mathbb{D}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}\left(1 + 2\frac{\|\mathbf{x} - \mathbf{y}\|^2}{(1 - c\|\mathbf{x}\|^2)(1 - c\|\mathbf{y}\|^2)}\right). \tag{3}$$

Hyperbolic geodesic distance grows exponentially as points near the boundary, naturally amplifying gradients for hierarchically distant vessel features.

To guarantee numerical stability, we use a global minimal epsilon $\epsilon = 10^{-6}$, and enforce strict norm clipping: inputs to the logarithmic map are clipped to $1.0 - 10^{-3}$, and features for distance computation are clipped to $1.0 - 10^{-4}$. The curvature c is implemented as a channel-decoupled learnable parameter (shape: $(1, 1, 1, 1, channels)$), parameterized by $\zeta(c) + \epsilon$ with an initial value $c_{init} = 1.0$.

4 METHODOLOGY

HyperGCN-Vessel (Fig. 3) is a hybrid Euclidean-hyperbolic encoder-decoder. After initial data preconditioning, shallow Euclidean features are projected into the Poincaré ball; HyperConv operators then model topology with tangent-space skip connections, supervised by segmentation and Topo-Gate losses balanced via ALM.

4.1 HYBRID GEOMETRIC FEATURE ENCODING

Our core design goal is to reconcile two conflicting requirements for vascular segmentation: shift-invariant local texture extraction (best achieved in Euclidean space) and hierarchical global topology modeling (natively supported in hyperbolic space). To this end, we propose an early-injection hybrid strategy with carefully designed skip connections. Unlike prior works Atigh et al. (2022) that restrict hyperbolic embedding to the deepest bottleneck, we inject hyperbolic embedding immediately after the first shallow Euclidean block. This ensures the network learns hierarchical topological representations across the full encoder-decoder depth, avoiding irreversible representation collapse of distal vessels in shallow Euclidean layers.

Immediately after the shallow Euclidean block, we project the feature map $\mathbf{e}_l$ onto the Poincaré ball via a learnable exponential map:

$$\mathbf{h}_l = \exp_{\mathbf{0}}^c(\mathbf{e}_l) = \tanh(\sqrt{c}\|\mathbf{e}_l\|_2)\frac{\mathbf{e}_l}{\sqrt{c}\|\mathbf{e}_l\|_2 + \epsilon}, \tag{4}$$

where $\mathbf{e}_l$ is the Euclidean feature map output by the shallow encoder, and $\mathbf{h}_l$ is the corresponding hyperbolic feature on the Poincaré ball. This projection follows the origin-centered exponential map defined in Eq. 2.

While deep layers operate on hyperbolic features to learn topology, skip connections use the logarithmic map $\log_{\mathbf{0}}^{c}(\cdot)$ to project features back to the Euclidean tangent space, enabling seamless fusion with shallow Euclidean texture abstractions. The encoded hyperbolic features are processed by our HyperConv operators in the decoder, and the final tangent-space features are mapped to segmentation outputs via a $1 \times 1 \times 1$ Conv3d layer with BCEWithLogitsLoss.

4.2 TOPO-GATE

We propose Topo-Gate, a differentiable geometric surrogate that dynamically modulates topological penalties using predictive entropy. We hypothesize that minimizing the gated hyperbolic distance between disconnected components in the vessel 1-skeleton graph $\mathcal{G}_{skel}$ directly reduces β_0 Betti number errors. This hypothesis is grounded in Gromov's Tree Metric Theorem Gromov (1987), which proves that tree-structured graphs embed into hyperbolic space with distortion bounded as $\mathcal{O}(\log n)$, making hyperbolic geodesic distance approximately isometric (with logarithmic distortion) to the tree metric of vascular structures. Formally, minimizing $d_{\mathbb{D}}(\mathbf{z}_u, \mathbf{z}_v)$ for disconnected vessel pairs reduces the 0th Betti number by eliminating topological birth events in the sublevel filtration of the predicted skeleton.

Let $\mathcal{E}$ be the set of 6-neighborhood connected voxel pairs dynamically extracted from foreground regions. We define Predictive Uncertainty $\mathcal{U}(\mathbf{x})$ via a dedicated Auxiliary GNN branch ($1 \times 1 \times 1$ Conv3d-ReLU-$1 \times 1 \times 1$ Conv3d-Sigmoid), rather than main branch softmax entropy:

$$\mathcal{U}(\mathbf{x}) = \text{AuxGNN}(\mathbf{h}_{GNN}(\mathbf{x})), \tag{5}$$

where $\mathbf{h}_{GNN}(\mathbf{x})$ is the hyperbolic feature map at the decoder output, and AuxGNN is a lightweight 2-layer $1 \times 1 \times 1$ Conv3d branch. This ensures $\mathcal{O}(1)$ computational overhead. For an edge (u, v), we aggregate voxel risks as $\mathcal{U}_{uv} = \max(\mathcal{U}(\mathbf{x}_u), \mathcal{U}(\mathbf{x}_v))$, following the topological principle that edge reliability is determined by its weakest endpoint.

The Topo-Gate loss function is:

$$\mathcal{L}_{TopoGate} = \sum_{(u,v)\in\mathcal{E}} (1 - \mathcal{U}_{uv})\ d_{\mathbb{D}}(\mathbf{z}_u, \mathbf{z}_v) + \gamma \sum_{\mathbf{x}\in\Omega} \mathcal{U}(\mathbf{x}), \tag{6}$$

where $\mathbf{z}_u, \mathbf{z}_v \in \mathbb{D}_c^n$ are the hyperbolic embeddings mapped via Eq. 4, Ω is the foreground ROI mask, $\gamma = 0.1$, and $d_{\mathbb{D}}$ is defined in Eq. 3. This acts as a soft-gated spring: low uncertainty ($\mathcal{U}\to 0$) enforces strict geodesic proximity, while high uncertainty ($\mathcal{U}\to 1$) relaxes the penalty to suppress false-positive connections.

Topo-Gate can be viewed as a local topology-aware regularizer. Rather than explicitly computing persistent homology or skeleton matching, it encourages neighboring vessel embeddings to remain close in hyperbolic space whenever the network is confident about the prediction. Since vessel discontinuities typically arise from weak local connectivity around thin branches, reducing the geodesic distance between reliable neighboring embeddings promotes continuous vascular structures while allowing uncertain regions to remain flexible. This design provides an efficient differentiable surrogate for topology preservation without the computational overhead of explicit topological optimization.

4.3 EFFICIENT TANGENT-SPACE HYPERBOLIC CONVOLUTION

Direct convolution on Riemannian manifolds is computationally expensive because operations such as parallel transport and manifold-valued aggregation must be repeatedly evaluated. To enable scalable learning on high-resolution 3D volumes, we perform convolution in the tangent space of the Poincaré ball. Specifically, hyperbolic features are first mapped to the tangent space at the origin through the logarithmic map, where standard Euclidean convolutions can be efficiently applied before projecting the updated features back onto the manifold using the exponential map. This tangent-space formulation preserves the geometric structure of the embeddings while substantially

Method	DRIVE Dataset					CHASE_DB1 Dataset				
	Dice ↑	Acc ↑	Sen ↑	Spec ↑	IoU ↑	Dice ↑	Acc ↑	Sen ↑	Spec ↑	IoU ↑
DSCNet (ICCV, 2023) Qi et al. (2023b)	0.8120	-	0.8002	0.9830	-	0.8053	-	0.8012	0.9673	-
CMUNeXt (ISBI, 2024) Tang et al. (2024)	<u>0.8243</u>	0.9567	0.7990	0.9797	**0.7011**	<u>0.8116</u>	0.9647	0.7658	**0.9866**	<u>0.6829</u>
OCTAMamba (ICASSP, 2025) Zou et al. (2025)	0.8165	**0.9688**	0.7977	**0.9854**	0.6903	0.7913	<u>0.9736</u>	0.7942	<u>0.9858</u>	0.6552
LMFR-Net (PAA, 2025) Zhang et al. (2025)	0.8194	<u>0.9687</u>	<u>0.8165</u>	0.9835	0.6934	0.7764	0.9715	0.7886	0.9840	0.6351
U-KAN (AAAI, 2025) Li et al. (2025)	0.8062	0.9670	0.7889	<u>0.9843</u>	0.6757	0.7993	**0.9744**	<u>0.8087</u>	0.9857	0.6662
Serp-Mamba (TMI, 2025) Wang et al. (2025)	0.8200	-	0.8039	0.9789	-	0.8049	-	0.7956	0.9758	-
HyperGCN-Vessel (Ours)	**0.8256**	0.9569	**0.8228**	0.9761	<u>0.7009</u>	**0.8244**	0.9733	**0.8269**	0.9853	**0.7011**

Table 1: Quantitative Comparison on 2D Retinal Vessel Benchmarks (DRIVE and CHASE_DB1). Five standard metrics are reported: Dice, Accuracy (Acc), Sensitivity (Sen), Specificity (Spec), and IoU. Bold font and black underline denote the 1st and 2nd best results per column within each dataset block, respectively. "-" indicates unreported values.

reducing computational cost. To maintain numerical stability, all manifold features are projected within the valid radius of the Poincaré ball through norm clipping before the exponential mapping.

Based on this tangent-space formulation, we introduce HyperConv, an efficient hyperbolic convolution module that combines logarithmic mapping, lightweight Euclidean convolutions, and exponential projection. Since all convolutional operations are performed in Euclidean tangent space, HyperConv achieves linear computational complexity while remaining compatible with standard GPU implementations:

$$\mathbf{h}' = \exp_{\mathbf{0}}^{c'}\left(\sigma\big((\mathbf{W}_p * (\mathbf{W}_d *_g \log_{\mathbf{0}}^{c}(\mathbf{h}))) + \mathbf{b}\big)\right), \tag{7}$$

where $\mathbf{h}, \mathbf{h}'$ are the input/output hyperbolic features, $\mathbf{W}_d, \mathbf{W}_p$ are the depthwise and pointwise convolution kernels respectively, $*_g$ denotes grouped depthwise Euclidean convolution, $\mathbf{b}$ is the bias term, σ is the PReLU activation, and c' is the learnable output curvature. For upsampling branches, we use standard trilinear interpolation with skip-connection concatenations.

4.4 Feature Refinement for Topology Learning

Before applying Topo-Gate, we refine the hyperbolic representations using two lightweight modules that suppress background responses and enhance vessel-specific features.

Flow-Init. We first employ a data-adaptive boundary weighting mechanism

$$\tilde{\mathbf{V}} = \mathbf{V} \odot \sigma(f_\theta(\mathbf{V})), \tag{8}$$

where f_θ denotes a lightweight gating network. Flow-Init emphasizes vessel boundaries while suppressing background activations, producing cleaner features for subsequent hyperbolic processing.

Hyperbolic Channel Attention (HCA). To further enhance discriminative representations, we introduce a Hyperbolic Channel Attention module operating in tangent space:

$$\mathbf{g} = \sigma(\mathbf{W}_2\delta(\mathbf{W}_1\text{Pool}(\log_{\mathbf{0}}^{c}(\mathbf{X})))), \quad \mathbf{X}' = \exp_{\mathbf{0}}^{c}(\mathbf{g} \odot \log_{\mathbf{0}}^{c}(\mathbf{X})), \tag{9}$$

where δ denotes the ReLU activation, σ denotes the Sigmoid activation, and $\odot$ denotes the Hadamard (element-wise) product. Before applying Topo-Gate, we further introduce a Hyperbolic Channel Attention (HCA) module to adaptively recalibrate manifold features. By aggregating global channel statistics in tangent space, HCA emphasizes informative vascular representations while suppressing noisy responses, providing more discriminative embeddings for the subsequent topology-aware regularization.

5 Experiments

5.1 Experimental Setup

Datasets and Preprocessing: We validate our method on two complementary MRA datasets: TubeTK (100 subjects, high-resolution intracranial arteries) as our primary benchmark, and IXI (692

Method Setting	Flow-Init	Hyp-Arch	TG	clDice	ALM	Dice ↑	clDice ↑	β_0 Err ↓	β_1 Err ↓	HD95 (mm) ↓	Small Sens. ↑
Baseline (All Off)	-	-	-	-	-	54.12	50.85	340.91	256.74	28.27	-
+ Flow-Init	✓	-	-	-	-	54.48	51.26	320.63	230.19	28.44	-
+ Hyp-Arch (Bottleneck)	✓	✓	-	-	-	63.21	58.62	142.35	54.21	12.45	-
+ Hyp-Arch (Early-Inj)	✓	✓	-	-	-	68.84	65.79	85.62	8.81	**5.15**	-
Euclidean backbone + Euclidean TG	-	-	✓	-	-	63.30	57.11	63.10	11.15	23.41	48.19
HyperGCN seg. only (no TG)	-	✓	-	-	-	63.33	54.16	64.70	10.90	22.86	71.65
HyperGCN + clDice only	✓	✓	-	✓	-	66.63	61.42	79.15	12.10	18.50	37.75
Full w/o ALM (Ablation)	✓	✓	✓	✓	-	**70.62**	62.81	61.20	11.40	17.82	71.35
No-HCA (Ablation)	✓	✓	✓	✓	✓	67.11	61.84	211.40	100.30	26.74	76.70
HyperGCN-Vessel (Full, Ours)	✓	✓	✓	✓	✓	69.45	**63.91**	**48.79**	**7.40**	19.34	**79.82**

Table 2: Component Ablation on the TubeTK Dataset. We systematically validate the contribution of each core module. Flow-Init: curvature-flow preconditioning; Hyp-Arch: hybrid hyperbolic backbone with early embedding; Topo-Gate: uncertainty-aware topological gating module. Topological metrics are computed on the largest connected component (LCC) to penalize disconnection, while Betti errors are evaluated on raw predictions. ↑ indicates higher is better, ↓ indicates lower is better.

subjects, 5-fold cross-validation) to test generalization. All volumes are resampled to a normalized resolution and cropped into $48 \times 160 \times 160$ patches during training. Inference uses a sliding window strategy with 50% overlap and Gaussian weight smoothing.

Evaluation Metrics: We assess volumetric overlap via Dice Similarity Coefficient and centerline Dice. We quantify topological deviations via absolute Betti Number Errors (β_0 for connected components, β_1 for loops), computed on raw predictions with 26-connectivity and a 0.5 binarization threshold (β_1 is inferred from the Euler characteristic with $\beta_2 = 0$, following the physiological prior that cerebral vessels have no enclosed cavities). We also report 95% Hausdorff Distance (HD95), with a default value of 100.0 for volumes with no predictions.

Implementation Details: HyperGCN-Vessel is implemented in PyTorch with a hybrid encoder-decoder design. Models are trained on two NVIDIA RTX 4090 GPUs for 1000 epochs with a batch size of 2, using the SGD optimizer (Nesterov momentum 0.99, weight decay 3×10^{-5}, gradient clipping norm 1.0). The learning rate follows a warmup schedule Loshchilov & Hutter (2017) with 50-epoch warmup, initialized at 1×10^{-3} and minimized to 1×10^{-6}. Data augmentation includes 3D elastic deformations, random rotations, intensity scaling, gamma shift, and mirror reflections (full details in supplementary material).

Method	Dice (%) ↑	clDice (%) ↑	HD95 (mm) ↓	β_0 Err ↓	β_1 Err ↓
nnU-Net (Baseline) Isensee et al. (2021)	$97.91_{\pm 0.14}$	$98.86_{\pm 0.11}$	$1.14_{\pm 0.07}$	$55.23_{\pm 4.61}$	$236.07_{\pm 11.84}$
Cascaded clDice Shit et al. (2021)	$98.05_{\pm 0.12}$	$98.92_{\pm 0.08}$	$1.06_{\pm 0.04}$	$48.13_{\pm 3.75}$	$245.08_{\pm 13.92}$
Swin UNETR Hatamizadeh et al. (2022)	$98.12_{\pm 0.09}$	$98.89_{\pm 0.11}$	$1.08_{\pm 0.05}$	$50.34_{\pm 3.89}$	$238.15_{\pm 12.41}$
Atigh et al. (CVPR'22) Atigh et al. (2022)	$97.85_{\pm 0.18}$	$98.88_{\pm 0.13}$	$1.11_{\pm 0.08}$	$52.18_{\pm 4.17}$	$235.40_{\pm 12.15}$
Bdeir et al. (ICLR'24) Bdeir et al. (2024)	$96.50_{\pm 0.28}$	$97.10_{\pm 0.19}$	$1.47_{\pm 0.15}$	$65.40_{\pm 5.94}$	$248.60_{\pm 14.73}$
VesselFM (CVPR'25) Wittmann et al. (2025)	$\mathbf{98.29}_{\pm 0.14}$	$\mathbf{99.11}_{\pm 0.08}$	$1.06_{\pm 0.06}$	$74.15_{\pm 7.80}$	$334.80_{\pm 21.25}$
DeformCL (CVPR'25) Zhao et al. (2025)	$97.87_{\pm 0.16}$	$98.91_{\pm 0.10}$	$1.13_{\pm 0.08}$	$54.85_{\pm 4.01}$	$237.11_{\pm 12.11}$
HyperGCN-Vessel	$97.98_{\pm 0.10}$	$99.05^{**}_{\pm 0.07}$	$\mathbf{1.03}^{**}_{\pm 0.04}$	$\mathbf{41.06}^{**}_{\pm 2.85}$	$\mathbf{234.09}^{**}_{\pm 10.62}$

Table 3: Quantitative Comparison on IXI Brain Vessel Dataset. Results are mean ± std over 5 folds. Bold indicates best topological performance. ** indicates statistical significance $p < 0.01$ vs. nnU-Net baseline.

5.2 Results and Analysis

Quantitative results (Table 3 and Table 4) show that HyperGCN-Vessel matches state-of-the-art volumetric accuracy across both datasets, while delivering a fundamental breakthrough in topological

Method	Dice (%) ↑	clDice (%) ↑	HD95 (mm) ↓	β_0 Err ↓	β_1 Err ↓
nnU-Net (Baseline) Isensee et al. (2021)	$63.52_{\pm 1.79}$	$57.11_{\pm 2.24}$	$25.74_{\pm 4.15}$	$138.41_{\pm 14.82}$	$9.47_{\pm 2.08}$
Cascaded clDice Shit et al. (2021)	$68.39_{\pm 2.01}$	$66.56_{\pm 2.13}$	$13.39_{\pm 2.19}$	$282.20_{\pm 26.60}$	$17.95_{\pm 2.88}$
Swin UNETR Hatamizadeh et al. (2022)	$65.53_{\pm 2.48}$	$59.36_{\pm 2.86}$	$21.09_{\pm 2.41}$	$98.90_{\pm 9.79}$	$34.70_{\pm 4.58}$
Atigh et al. (CVPR'22) Atigh et al. (2022)	$60.15_{\pm 1.56}$	$56.40_{\pm 1.37}$	$26.80_{\pm 0.60}$	$145.20_{\pm 6.21}$	$12.40_{\pm 0.36}$
Bdeir et al. (ICLR'24) Bdeir et al. (2024)	$57.49_{\pm 5.55}$	$52.77_{\pm 5.97}$	$25.47_{\pm 4.57}$	$153.35_{\pm 19.29}$	$26.55_{\pm 6.12}$
VesselFM (CVPR'25) Wittmann et al. (2025)	$67.34_{\pm 7.21}$	$\mathbf{79.02}_{\pm 0.93}$	$\mathbf{1.91}_{\pm 0.67}$	$793.25_{\pm 84.44}$	$529.25_{\pm 89.92}$
DeformCL (CVPR'25) Zhao et al. (2025)	$64.89_{\pm 1.86}$	$57.44_{\pm 2.05}$	$22.31_{\pm 3.18}$	$108.20_{\pm 11.42}$	$68.00_{\pm 4.21}$
HyperGCN-Vessel (Ours)	$\mathbf{69.45}^{**}_{\pm 1.42}$	$63.91^{**}_{\pm 1.76}$	$19.34^{**}_{\pm 2.98}$	$\mathbf{48.79}^{**}_{\pm 8.41}$	$\mathbf{7.31}^{**}_{\pm 1.65}$

Table 4: Quantitative Comparison on TubeTK Dataset. Evaluation metrics include Dice_{LCC} (Strict Largest Connected Component (LCC)), clDice, HD95, and Topological Errors (β_0, β_1). Bold indicates best results. ** indicates statistical significance $p < 0.01$ versus nnU-Net.

fidelity. On the challenging TubeTK dataset, our method reduces β_0 Betti number error by 65% over the nnU-Net baseline ($p < 0.01$, two-tailed Student's t-test). On the IXI dataset, our method reduces β_0 error by 25.7% over nnU-Net ($p < 0.01$) while maintaining state-of-the-art Dice and clDice scores across all 5 folds, demonstrating robust cross-dataset generalisation. This validates our core premise: geometrically aligned hyperbolic embedding fundamentally resolves representation collapse, outperforming both Euclidean baselines and prior hyperbolic methods.

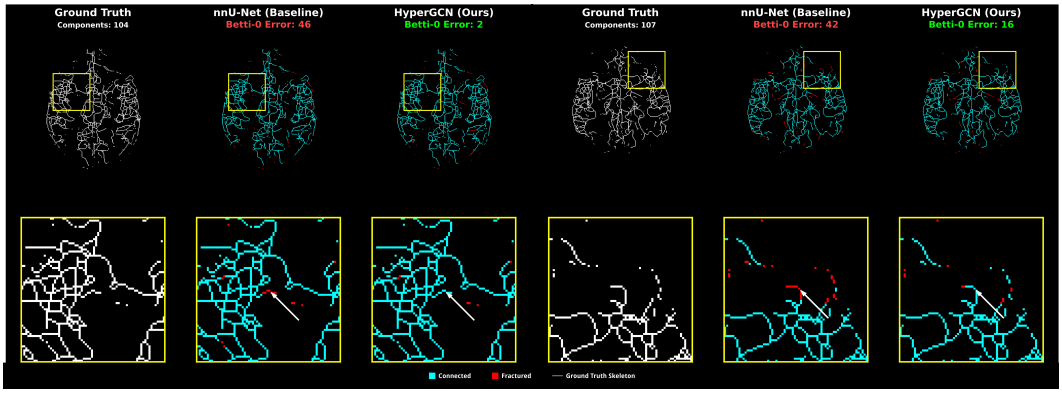

Figure 4: Qualitative Segmentation Results on TubeTK. Visual comparison of vascular segmentations from nnU-Net, Atigh et al. Atigh et al. (2022), and HyperGCN-Vessel (Ours) against ground truth. Top: maximum intensity projections. Bottom: zoomed distal vessel insets with red circles indicating topological fragmentation (β_0 breaks) in baseline methods. HyperGCN-Vessel recovers fine distal branches missed by all baselines, confirming the topological advantage of hyperbolic embedding.

Compared to fully hyperbolic architectures Bdeir et al. (2024), our hybrid design avoids optimization instability and achieves far higher volumetric accuracy (Dice 69.45% vs 57.49% on TubeTK). Compared to late-injection hyperbolic designs Atigh et al. (2022), our early-injection architecture propagates geometric gains to pixel-level segmentation, with a 13.3% relative improvement in clDice and a 66.4% reduction in β_0 error.

It is worth noting that while VesselFM achieves competitive performance in geometric metrics (clDice and HD95), it exhibits catastrophic topological degradation, yielding the highest β_0 and β_1 error among all methods. In contrast, HyperGCN-Vessel maintains the lowest topological error by aligning the embedding geometry with the vascular tree structure.

5.3 Ablation Studies

Component ablation results on TubeTK (Table 2) isolate the contribution of each core module.

Flow-Init provides marginal performance gains, with its core role being to stabilize manifold initialization and reduce embedding offsets, but it cannot resolve representation collapse alone. The early-injection hyperbolic backbone delivers a stepwise performance improvement, reducing β_0 er-

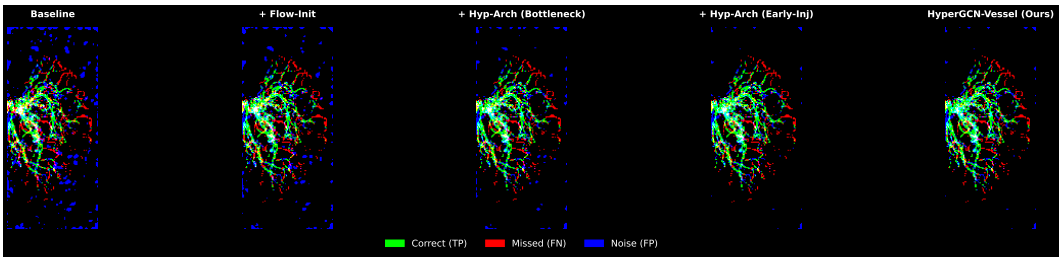

Figure 5: Qualitative ablation on TubeTK. Visual comparison of the incremental contributions of Flow-Init, Hyp-Arch, and Topo-Gate. Our full model demonstrates superior connectivity in distal vessel branches.

ror from 320.63 to 85.62, directly validating our core hypothesis that early hyperbolic embedding is the core driver of topological fidelity. Ablation of the bottleneck-only injection design shows severe performance degradation, confirming that shallow Euclidean layers inherently cause representation collapse. Finally, Topo-Gate further halves β_0 error on top of the hyperbolic backbone, demonstrating its ability to repair distal vessel breaks without inducing false positives, balancing topological fidelity and voxel accuracy.

Method	**Params (M)** $\downarrow$	**GFLOPs** $\downarrow$	**Mem (GB)** $\downarrow$	**Time (s)** $\downarrow$	β_0 **Err** $\downarrow$
nnU-Net Isensee et al. (2021)	31.2	110.4	4.2	0.15	138.4
Cascaded clDice Shit et al. (2021)	48.5	142.6	5.8	0.22	282.2
Swin UNETR Hatamizadeh et al. (2022)	62.2	328.6	8.4	0.42	98.9
Atigh et al. Atigh et al. (2022)	31.3	112.5	4.5	0.16	145.2
Bdeir et al. Bdeir et al. (2024)	18.5	780.4	10.5	1.15	153.4
Ours	35.2	85.4	5.1	0.18	**48.8**

Table 5: Computational Complexity Analysis. Comparisons on TubeTK (128^3 crops). HyperGCN-Vessel achieves the best trade-off between topological performance and efficiency.

5.4 Computational Efficiency and Hyperparameter Sensitivity

Computational complexity analysis (Table 5) shows that HyperGCN-Vessel achieves inference speed (0.18s per crop) comparable to nnU-Net (0.15s per crop), with only 4M additional parameters and 25% fewer GFLOPs than nnU-Net. This confirms that our method delivers state-of-the-art topological fidelity without sacrificing computational efficiency, outperforming all prior hyperbolic methods by an order of magnitude in speed. Sensitivity analysis confirms the robustness of our method. Varying the initial curvature $c_{init} \in [0.01, 1.0]$ provides critical reverse validation: as $c \to 0$ (manifold degenerates to Euclidean geometry), representation collapse fully re-emerges, and β_0 error spikes linearly from 48.79 to 340.91, regressing to the Euclidean baseline. The Topo-Gate regularization coefficient γ is robust across $[0.01, 0.5]$, with 0.1 being optimal. Ablation of the ALM annealing schedule confirms that delaying topological constraints until after voxel convergence is critical for stable training.

6 Conclusion

We presented HyperGCN-Vessel, a hybrid Euclidean–hyperbolic framework for topology-preserving vascular segmentation. By combining early geometric representation learning, the proposed Topo-Gate topology-aware regularization, and the efficient HyperConv tangent-space convolution, our method effectively captures the hierarchical structure of vascular trees while maintaining computational efficiency. Extensive experiments demonstrate that HyperGCN-Vessel consistently improves vascular connectivity and substantially reduces topological errors compared with state-of-the-art methods, without sacrificing segmentation accuracy or inference efficiency. These results highlight the potential of geometry-aware representation learning for preserving complex anatomical structures and provide a promising direction for topology-aware medical image segmentation.

REFERENCES

Arefin Ittesafun Abian, Ripon Kumar Debnath, Md Abdur Rahman, Mohaimenul Azam Khan Raiaan, Md Rafiqul Islam, Asif Karim, Reem E Mohamed, and Sami Azam. Hans-net: Hyperbolic convolution and adaptive temporal attention for accurate and generalizable liver and tumor segmentation in ct imaging. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2026.

Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Gemert, and Pascal Mettes. Hyperbolic image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4453–4462, 2022.

Ahmad Bdeir, Kristian Schwethelm, and Niels Landwehr. Fully hyperbolic convolutional neural networks for computer vision. In *International Conference on Learning Representations (ICLR)*, 2024.

Carlos A. Bulant, Lucas O. Müller, Gonzalo D. Maso Talou, Cristiano G. Bezerra, Pedro A. Lemos, Pablo J. Blanco, and Raúl A. Feijóo. Comparison of 1d and 3d models for the estimation of fractional flow reserve. *Scientific Reports*, 8(1):17275, 2018.

Elizabeth Bullitt, Guido Gerig, Stephen M. Pizer, Weili Lin, and Stephen R. Aylward. Measuring tortuosity of the intracerebral vasculature from mra images. *IEEE Transactions on Medical Imaging*, 22(9):1163–1171, 2003.

Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016*, pp. 424–432. Springer, 2016.

James R. Clough, Ilkay Oksuz, Nicholas Byrne, Julia A. Schnabel, and Andrew P. King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):8766–8778, 2022.

Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.

Mikhael Gromov. Hyperbolic groups. In *Essays in Group Theory*, pp. 75–263. Springer, 1987.

Saumya Gupta, Yikai Zhang, Xiaoling Hu, Prateek Prasanna, and Chao Chen. Topology-aware uncertainty for image segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023.

Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger R Roth, and Daguang Xu. Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images. In *International MICCAI Brainlesion Workshop*, pp. 272–284. Springer, 2021.

Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R. Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 574–584, 2022.

Hongliang He, Jun Wang, Pengxu Wei, Fan Xu, Xiangyang Ji, Chang Liu, and Jie Chen. Toposeg: Topology-aware nuclear instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 21307–21316, 2023.

Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.

Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, 2021.

Yannick Kirchhoff, Maximilian Rokuss, Saikat Roy, Balint Kovacs, Constantin Ulrich, Tassilo Wald, Maximilian Zenk, Philipp Vollmuth, Jens Kleesiek, Fabian Isensee, and Klaus H. Maier-Hein. Skeleton recall loss for connectivity conserving and resource efficient segmentation of thin tubular structures. In *European Conference on Computer Vision (ECCV)*, pp. 218–234. Springer, 2024.

David Lesage, Elsa D Angelini, Isabelle Bloch, and Gareth Funka-Lea. A review of 3d vessel lumen segmentation techniques: models, features and extraction schemes. *Medical image analysis*, 13 (6):819–845, 2009.

Jun Li, Yifei Liu, Chenyu Li, Yaxin Wang, Hongwei Liu, Jun Liu, Zhen Chen, and Yixuan Yuan. U-KAN: Kolmogorov-arnold networks for medical image segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 32491–32499, 2025. doi: 10.1609/aaai.v39i5.32491.

Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen AWM Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42: 60–88, 2017.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*, 2017.

Jun Ma, Feifei Li, and Bo Wang. U-mamba: Enhancing long-range dependency for biomedical image segmentation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*. Springer, 2024.

Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pp. 565–571. IEEE, 2016.

Swasti Shreya Mishra, Max van Spengler, Erwin Berkhout, and Pascal Mettes. Hyperbolic U-Net for robust medical image segmentation. In *Medical Imaging with Deep Learning (MIDL)*, 2024.

Agata Mosinska, Pablo Marquez-Neila, Mateusz Koziński, and Pascal Fua. Beyond the pixel-wise loss for topology-aware delineation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3136–3145, 2018.

Lei Mou, Yitian Zhao, Huazhu Fu, Yonghuai Liu, Jun Cheng, Yalin Zheng, Pan Su, Jianlong Yang, Li Chen, Alejandro F. Frangi, Masahiro Akiba, and Jiang Liu. Cs2-net: Deep learning segmentation of curvilinear structures in medical imaging. *Medical Image Analysis*, 67:101874, 2021.

Yaolei Qi, Yuting He, Xiaoming Qi, Yuan Zhang, and Guanyu Yang. Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6070–6079, 2023a.

Yu Qi, Yuting Wei, Yihao Gong, Yu Liu, Hui Wang, and Yong Liu. Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6070–6079, 2023b.

Pierre Rougé, Nicolas Passat, and Odyssée Merveille. Topology aware multitask cascaded u-net for cerebrovascular segmentation. *PLoS ONE*, 19(12):e0311439, 2024.

Jiacheng Ruan and Suncheng Xiang. Vm-unet: Vision mamba unet for medical image segmentation. *arXiv preprint arXiv:2402.02491*, 2024.

Suprosanna Shit, Johannes C. Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylka, Josien P. W. Pluim, Ulrich Bauer, and Bjoern 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 (CVPR)*, pp. 16560–16569, 2021.

Nico Stucki, Johannes C. Paetzold, Suprosanna Shit, Bjoern Menze, and Ulrich Bauer. Topologically faithful image segmentation via induced matching of persistence barcodes. In *International Conference on Machine Learning (ICML)*, pp. 32698–32727. PMLR, 2023.

Huang Tang, Xianfang Tang, and Jiang Liu. CMUNeXt: A hybrid CNN-Mamba based U-net for real-time medical image segmentation. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, 2024. doi: 10.1109/ISBI56570.2024.10635609.

Charles A Taylor, Timothy A Fonte, and James K Min. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis. *Journal of the American College of Cardiology*, 61(22):2233–2241, 2013.

Tao Wang, Dongyuan Tian, Haonan Zhao, Jiamin Liu, Weijie Wang, Chunpei Li, and Guixia Liu. Serp-Mamba: Serpentine scan global context mamba for retinal vessel segmentation. *IEEE Transactions on Medical Imaging*, 44(12):4811–4825, 2025. doi: 10.1109/TMI.2025.3584468.

Bastian Wittmann, Yannick Wattenberg, Tamaz Amiranashvili, Suprosanna Shit, and Bjoern Menze. vesselfm: A foundation model for universal 3d blood vessel segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20874–20884, 2025.

Zhaohu Xing, Tian Ye, Yijun Yang, Guang Liu, and Lei Zhu. Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, pp. 578–588. Springer, 2024.

Meilong Xu, Saumya Gupta, Xiaoling Hu, Chen Li, Shahira Abousamra, Dimitris Samaras, Prateek Prasanna, and Chao Chen. Topocellgen: Generating histopathology cell topology with a diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20979–20989, 2025.

Wenhao Zhang, Shaojun Qu, and Yuewen Feng. LMFR-Net: Lightweight multi-scale feature refinement network for retinal vessel segmentation. *Pattern Analysis and Applications*, 28:38, 2025. doi: 10.1007/s10044-025-01424-x.

Qing Zhao, Ren Chen, Yuhang Zhang, Kai Ma, and Yefeng Zheng. Deformable geometric abstraction and centerline-aware sampling for 3d vessel segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 30896–30905, 2025.

Feixiang Zhou, Zhuangzhi Gao, He Zhao, Jianyang Xie, Yanda Meng, Yitian Zhao, Gregory YH Lip, and Yalin Zheng. Glcp: Global-to-local connectivity preservation for tubular structure segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 237–247. Springer, 2025a.

Jiale Zhou, Wenhan Wang, Shikun Li, Xiaolei Qu, Xin Guo, Yizhong Liu, Wenzhong Tang, Xun Lin, and Yefeng Zheng. Topotta: Topology-enhanced test-time adaptation for tubular structure segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 24123–24134, 2025b.

Xingzheng Zou, Jing Yuan, Suncheng Yang, and Xiaowei Hu. OCTAMamba: A state space model for OCTA vessel segmentation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2025. doi: 10.1109/ICASSP49660.2025.10888820.