Cupido Loss for Structure-Aware Pulmonary Airway and Artery Segmentation

Chang Yuwen[1,✉][0009-0000-6864-6885], Litao Yang[1], Dwarikanath Mahapatra[2][0000-0001-9749-7858], Behzad Bozorgtabar[3,4][0000-0002-5759-4896], and Zongyuan Ge[1,][0000-0002-5880-8673]

[1] Monash University, Clayton, VIC, Australia
[2] Khalifa University, Abu Dhabi, UAE
[3] Aarhus University, Aarhus, Denmark
[4] École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
{chang.yuwen, litao.yang, zongyuan.ge}@monash.edu,
dwarikanath.mahapatra@ku.ac.ae, behzad@ece.au.dk

Abstract. Accurate segmentation of pulmonary tubular structures is essential for lung surgical planning and quantitative pulmonary analysis, yet remains challenging due to distal branch sensitivity, complex geometries, and low-intensity contrast. Although recent deep learning methods improve accuracy, they often fail to preserve structural integrity and rely on heavy architectures or extensive annotations. Designing loss functions that explicitly encode structural priors offers a promising alternative. We propose Cupido (Consistent Unified Prior for Intricate Directional Optimization), a plug-and-play loss for thin, complex tubular structures. Cupido integrates two constraints: a directional consistency term enforcing local angular smoothness, and a geometry-aware union term unifying skeleton- and distance-based priors to enhance global coherence. We evaluate Cupido with three backbones (nnUNet, UMambaBot, and UxLSTMEnc) on two public pulmonary segmentation benchmarks comprising three evaluation subsets, training all models under identical settings for 50 epochs to evaluate robustness across backbones and datasets. Experiments show consistent improvements across backbones and datasets. On the ATM'22 Batch2 dataset, Cupido achieves 91.59% Dice, 86.38% NSD, and 86.80% clDice with UxLSTMEnc, exceeding the baseline by +1.78%, +2.75%, and +3.26%, respectively. It improves fine-branch delineation and structural consistency. By encoding directional and geometric priors, Cupido enables robust, anatomically consistent predictions with minimal training cost for clinically relevant tubular anatomy analysis.

Keywords: Pulmonary tubular structures · Lung surgical planning · Structural integrity · Cupido

1 Introduction

Medical image segmentation [10,16] plays a fundamental role in computer-aided diagnosis and interventional planning. In particular, segmentation of tubular

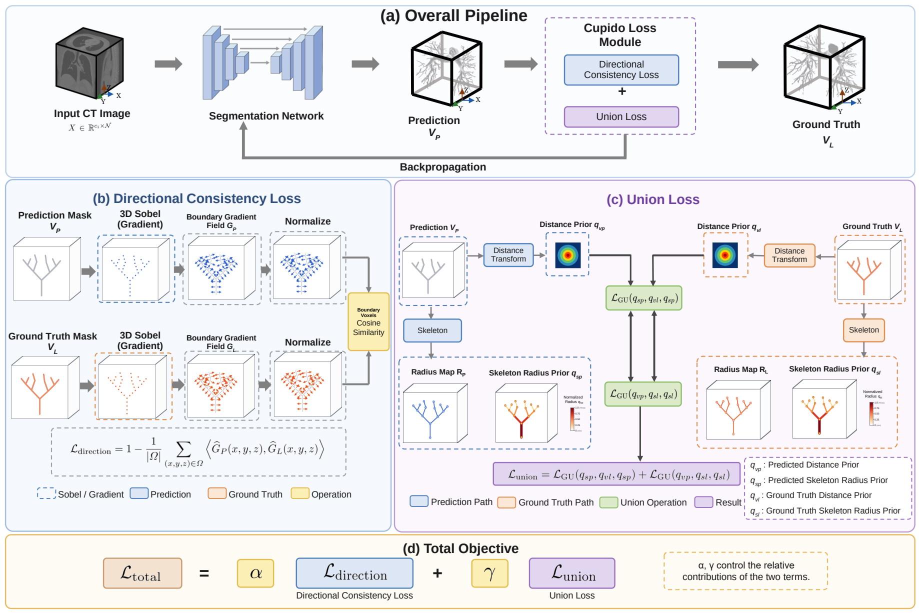

Fig. 1. Overview of the proposed Cupido Loss framework. (a) Overall optimization pipeline for tubular structure segmentation. (b) Directional Consistency Loss, which enforces boundary-normal alignment between prediction and ground truth. (c) Union Loss, which jointly optimizes distance and radius priors derived from skeleton representations. (d) Final objective function combining the two complementary loss terms.

structures [8] is essential for pulmonary analysis, where accurate delineation of bronchial and arterial trees supports assessment of airway obstruction and surgical planning [15]. Beyond visualization, reliable tubular segmentation facilitates downstream tasks such as quantitative disease measurement and computational modeling [4]. Consequently, clinically meaningful segmentation requires not only high voxel-wise accuracy but also structural consistency that respects anatomical structure and geometric continuity. Despite substantial progress driven by deep neural networks, tubular segmentation remains inherently challenging. Pulmonary tubular structures exhibit elongated and branching morphologies with significant diameter variation across hierarchical levels. Thin peripheral branches are particularly susceptible to fragmentation, while major trunks may dominate optimization due to severe class imbalance [9]. Preserving directional continuity across bifurcations and maintaining sensitivity to fine distal branches are persistent difficulties [5, 19]. In practical settings, imaging noise and limited spatial resolution further exacerbate these challenges, often leading to broken centerlines or structural inconsistencies.

Recent advances have incorporated structural priors beyond region overlap. Skeleton-based supervision such as clDice [19] enhances connectivity, while topology-driven methods impose global structural constraints [2]. Geometry-aware approaches based on dynamic structural modeling [14] and multi-scale regularization [3] refine contour alignment and tubular balance. However, existing strategies typically emphasize either topology preservation or geometric refinement, without jointly addressing directional continuity and adaptive balance across tubular scales. The formulation of the loss function plays a central role in resolving these limitations. Conventional objectives such as Dice and Tversky losses [13, 17] optimize region similarity but lack explicit structural priors. Recent developments such as General Union Loss [22] attempt to balance small and large structures through adaptive weighting, yet directional alignment remains insufficiently addressed. As a result, existing formulations struggle to simultaneously preserve geometric fidelity and maintain optimization stability in thin tubular regions.

To address these challenges, we propose **Cupido Loss**, a unified structure-aware optimization framework that jointly optimizes boundary-direction consistency and scale-aware geometric priors. The main contributions of this work are summarized as follows: (1) a directional consistency loss that aligns boundary-normal directions between predicted and reference structures, encouraging accurate boundary localization and preserving local geometric continuity; (2) a union-based geometric constraint that integrates distance and radius-aware skeleton priors to provide scale-adaptive geometric supervision and enhance structural coherence across tubular structures with varying diameters; (3) a unified optimization objective that combines these complementary components to provide explicit geometric supervision while balancing multi-scale structural representations; and (4) extensive experiments on multiple pulmonary tubular structure segmentation datasets demonstrating consistent improvements across different segmentation backbones.

2 Methodology

2.1 Preliminaries

The overall architecture is illustrated in Fig. 1. Let $X \in \mathbb{R}^{c_i \times \mathcal{N}}$ denote a 3D medical image volume, where $\mathcal{N} = w \times h \times d$ is the total number of voxels. The segmentation network produces a probabilistic prediction $Y \in \mathbb{R}^{c_o \times \mathcal{N}}$, which is thresholded to obtain a binary mask $V \in \mathbb{R}^{\mathcal{N}}$. A corresponding skeleton map $S \in \mathbb{R}^{\mathcal{N}}$ is extracted from V via a skeletonization operation. Predicted and reference masks and skeletons are denoted by (V_P, S_P) and (V_L, S_L), respectively.

The proposed Cupido Loss combines two complementary structure-aware objectives: a directional consistency term $\mathcal{L}_{\text{direction}}$ for boundary-normal alignment and a union term $\mathcal{L}_{\text{union}}$ for geometry-aware structural regularization. The overall objective is defined as:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{direction}} + \gamma \mathcal{L}_{\text{union}} \tag{1}$$

where α, γ control their relative contributions.

2.2 Directional Consistency Loss

Tubular anatomical structures are characterized by smooth and coherent boundary geometry. Conventional region-based losses primarily optimize volumetric overlap and often fail to explicitly constrain local boundary morphology, which may lead to irregular contours. To encourage geometric consistency between the predicted and reference structures, we introduce a directional consistency loss based on local boundary-normal alignment.

Specifically, we compute three-dimensional Sobel gradients on the predicted and reference binary masks:

$$G_P(x,y,z) = \left(\frac{\partial V_P}{\partial x}, \frac{\partial V_P}{\partial y}, \frac{\partial V_P}{\partial z}\right) \qquad G_L(x,y,z) = \left(\frac{\partial V_L}{\partial x}, \frac{\partial V_L}{\partial y}, \frac{\partial V_L}{\partial z}\right) \quad (2)$$

Since gradients are only informative near structural boundaries, directional supervision is restricted to the union of valid gradient locations from both prediction and reference masks:

$$\Omega = \{(x,y,z) \mid \|G_P(x,y,z)\| > 0 \text{ or } \|G_L(x,y,z)\| > 0\} \quad (3)$$

By taking the union of valid gradient locations from both prediction and reference masks, the proposed formulation captures discrepancies between predicted and reference boundaries, enabling supervision of omitted structures as well as spurious boundary regions.

The gradients are normalized into unit boundary-normal vectors:

$$\widehat{G}_P(x,y,z) = \frac{G_P(x,y,z)}{||G_P(x,y,z)|| + \epsilon} \qquad \widehat{G}_L(x,y,z) = \frac{G_L(x,y,z)}{||G_L(x,y,z)|| + \epsilon} \quad (4)$$

where ϵ is a small constant introduced for numerical stability and $\|\cdot\|$ denotes the Euclidean norm:

$$||G_P(x,y,z)|| = \sqrt{G_{Px}^2 + G_{Py}^2 + G_{Pz}^2} \qquad ||G_L(x,y,z)|| = \sqrt{G_{Lx}^2 + G_{Ly}^2 + G_{Lz}^2} \quad (5)$$

The directional consistency loss is then defined as:

$$\mathcal{L}_{\text{direction}} = 1 - \frac{1}{|\Omega|} \sum_{(x,y,z)\in\Omega} \left\langle \widehat{G}_P(x,y,z), \widehat{G}_L(x,y,z) \right\rangle \quad (6)$$

This formulation encourages alignment between predicted and reference boundary-normal directions while remaining invariant to gradient magnitude. By explicitly regularizing local boundary geometry, the proposed loss complements region-overlap objectives and promotes smoother and more coherent tubular structures, particularly around bifurcations and thin peripheral branches.

2.3 Union Loss

To improve geometric fidelity, we construct distance-based and radius-weighted skeleton priors under a predefined radius constraint $R_{\max}$, following the skeleton-weighting strategy of cbDice proposed by Shi et al. [18]. Let D_P and D_L denote the Euclidean distance transforms of the predicted and reference masks, respectively. Both distance maps are clipped at $R_{\max}$ and normalized as:

$$q_{vp} = \frac{\min(D_P, R_{\max})}{R_{\max}} \quad q_{vl} = \frac{\min(D_L, R_{\max})}{R_{\max}} \tag{7}$$

These normalized maps encode spatial proximity to the tubular boundary within a bounded radius.

Skeleton maps S_P and S_L are obtained via three-dimensional thinning. Corresponding radius maps R_P and R_L are clipped to the interval $[1, R_{\max}]$ and normalized by $R_{\max}$. The skeleton-based priors are then defined as:

$$q_{sp} = \frac{R_P}{R_{\max}} \quad q_{sl} = \frac{R_L}{R_{\max}} \tag{8}$$

The four priors serve complementary geometric roles. Specifically, q_{vp} and q_{vl} are distance-based tubular priors derived from the predicted and reference distance transforms, respectively, encoding the relative distance of each voxel to the tubular centerline. In contrast, q_{sp} and q_{sl} are skeleton-radius priors extracted from the predicted and reference centerlines, respectively, encoding local tubular scale information along the skeleton. Together, these priors provide complementary descriptions of tubular geometry from interior-region and centerline perspectives.

These priors incorporate tubular diameter information and emphasize structurally relevant regions along the centerline. The union objective aggregates two General Union terms:

$$\mathcal{L}_{\text{union}} = \mathcal{L}_{\text{GU}}(q_{sp}, q_{vl}, q_{sp}) + \mathcal{L}_{\text{GU}}(q_{vp}, q_{sl}, q_{sl}) \tag{9}$$

where

$$\mathcal{L}_{\text{GU}}(A, B, C) = 1 - \frac{\sum_{i=1}^{\mathcal{N}} C_i \left((A_i + \epsilon)^{\delta} B_i\right)}{\sum_{i=1}^{\mathcal{N}} C_i (\lambda A_i + \mu B_i)} \tag{10}$$

The first term, $L_{GU}(q_{sp}, q_{vl}, q_{sp})$, encourages consistency between the predicted skeleton representation and the reference distance prior, while the second term, $L_{GU}(q_{vp}, q_{sl}, q_{sl})$, aligns the predicted distance representation with the reference skeleton prior. By enforcing these two complementary alignments, the proposed formulation improves geometric consistency between tubular centerlines and surrounding regions.

In all experiments, we set $\lambda = 0.1$, $\mu = 0.9$, $\delta = 0.7$, and $\epsilon = 10^{-4}$. The weighting term C, the focal factor $(A + \epsilon)^{\delta}$, and the interaction between A and B jointly emphasize structurally important regions while alleviating class imbalance. By integrating distance and radius-aware skeleton priors within the General

Union formulation, the proposed objective enforces alignment between centerline and surrounding regions, preserves structural consistency across varying tubular diameters, and provides scale-adaptive geometric regularization. Overall, these properties promote coherent and anatomically plausible tubular segmentation.

3 Experiments

3.1 Datasets

Airway Tree Modeling Challenge 2022 (ATM'22 [21]): The dataset consists of chest CT scans with expert airway annotations. The image resolution is 512×512 pixels, with a slice thickness of 0.450–1.000 mm and 157–1125 slices. **Pulmonary Artery Segmentation Challenge (PARSE'22 [11])**: The dataset consists of chest CT scans with expert pulmonary artery annotations. The images have a resolution of 512×512 pixels, a slice thickness of 1 mm, and 228–408 slices.

Table 1. Experimental results on the PARSE'22 artery dataset and ATM'22 Batch1 and Batch2 subsets. Metrics include average Dice (%), NSD (mm), and clDice (%).

Dataset	Model	nnUNet [6]			UMambaBot [12]			UxLSTMEnc [1]		
		DSC (↑)	NSD (↑)	clDSC (↑)	DSC (↑)	NSD (↑)	clDSC (↑)	DSC (↑)	NSD (↑)	clDSC (↑)
Artery	Baseline	82.81±3.1	61.53±5.2	75.56±4.8	83.87±2.8	64.11±4.9	79.87±4.2	82.46±3.3	59.11±5.5	73.65±5.1
	CLDC [19]	84.80±2.2	67.73±3.8	80.93±3.1	85.89±2.0	70.05±3.5	83.27±2.8	84.63±2.4	67.48±3.9	81.13±3.3
	CBDC [18]	84.75±2.3	65.14±4.1	80.39±3.4	84.60±2.5	64.73±4.2	80.75±3.6	84.40±2.6	63.63±4.5	79.42±3.8
	BDoU [20]	82.21±3.5	59.54±5.8	74.78±5.2	83.00±3.3	61.98±5.8	77.49±4.7	82.88±3.2	60.42±5.6	75.16±5.3
	CLMDC [18]	83.60±2.9	61.09±4.9	76.61±4.5	84.86±2.4	64.65±4.3	80.16±3.7	84.10±2.7	62.87±4.8	78.07±4.2
	Skeleton [7]	72.75±5.8	47.29±8.1	50.69±9.4	77.75±4.5	53.66±6.8	58.05±7.2	77.74±4.6	53.64±6.9	57.10±7.5
	Cupido (Ours)	**85.11**±1.8	**68.70**±3.2	**82.64**±2.5	**86.20**±1.6	**70.12**±2.9	**85.18**±2.1	**85.29**±1.9	**68.85**±3.1	**82.43**±2.6
Batch1	Baseline	88.52±1.5	82.72±2.4	84.96±2.8	88.87±1.4	84.29±2.1	86.15±2.5	89.28±1.4	84.84±2.0	86.44±3.1
	CLDC [19]	89.08±1.3	85.15±1.9	**88.41**±1.8	89.72±1.2	85.54±1.8	88.30±2.0	89.80±1.3	86.15±1.7	87.85±1.9
	CBDC [18]	88.12±1.8	81.61±2.7	83.90±3.2	86.40±2.1	79.96±3.1	82.07±3.5	88.78±1.6	83.08±2.4	85.66±2.8
	BDoU [20]	87.97±1.9	83.29±2.2	84.69±2.9	86.45±2.2	81.27±2.9	79.10±4.1	89.40±1.4	84.45±2.1	85.88±2.6
	CLMDC [18]	88.37±1.6	81.21±2.8	83.51±3.3	85.54±2.4	80.60±3.2	81.23±3.6	89.11±1.5	83.13±2.5	84.73±2.9
	Skeleton [7]	87.22±2.1	81.14±2.9	80.90±3.8	87.07±2.1	79.94±3.3	77.56±4.5	87.99±1.8	81.87±2.8	80.41±3.7
	Cupido (Ours)	**89.45**±1.2	**85.39**±1.9	87.18±2.2	**90.52**±1.0	**85.83**±1.8	**89.12**±1.7	**91.51**±0.9	**86.62**±1.5	**89.21**±1.6
Batch2	Baseline	90.29±1.1	83.20±2.2	82.75±2.6	88.48±1.5	83.01±2.3	83.86±2.5	89.81±1.2	83.63±2.1	83.54±2.4
	CLDC [19]	90.24±1.2	**85.48**±1.8	86.20±2.2	91.10±0.9	85.32±1.7	86.09±1.9	90.67±1.1	85.55±1.8	**86.81**±1.5
	CBDC [18]	89.96±1.4	82.87±2.4	83.57±2.7	**91.97**±0.8	84.26±2.1	84.13±2.3	90.35±1.3	83.60±2.4	84.07±2.5
	BDoU [20]	90.43±1.1	80.84±2.9	77.22±3.5	90.31±1.2	81.57±2.6	78.46±3.2	89.88±1.2	81.95±2.7	81.78±3.0
	CLMDC [18]	91.27±0.9	83.11±2.3	82.90±2.5	89.92±1.3	83.21±2.4	81.71±3.0	90.67±1.1	83.97±2.2	85.12±2.3
	Skeleton [7]	85.79±2.5	74.40±3.8	70.92±4.2	88.18±1.8	78.23±3.1	76.48±3.6	88.55±1.5	77.91±3.2	75.60±3.8
	Cupido (Ours)	**90.46**±1.0	85.11±1.9	**86.45**±2.1	91.46±1.0	**86.54**±1.5	**88.07**±1.6	**91.59**±0.8	**86.38**±1.4	**86.80**±1.8

3.2 Implementation Details

For the bronchial segmentation task, we used the 300 publicly annotated training cases from the ATM'22 dataset, which are officially provided as Batch1 and Batch2. The remaining 200 cases belong to the test set without public annotations. Similarly, for the pulmonary artery segmentation task, we used the 100

Table 2. Ablation study on the Artery dataset investigating the impact of Directional Consistency Loss (D) and Union Loss (U).

Loss Components		**nnUNet [6]**			**UMambaBot [12]**			**UxLSTMEnc [1]**		
D	U	Dice (↑)	NSD (↑)	clDice (↑)	Dice (↑)	NSD (↑)	clDice (↑)	Dice (↑)	NSD (↑)	clDice (↑)
✓		84.78±2.4	68.22±4.2	81.71±3.8	85.62±2.1	69.34±3.9	84.51±3.3	85.28±2.3	**69.31**±3.6	83.24±3.5
	✓	85.00±2.3	68.45±4.1	82.12±3.7	85.71±2.0	69.69±3.7	84.00±3.5	85.00±2.4	68.12±4.5	82.01±4.1
✓	✓	**85.11**±1.8	**68.70**±3.2	**82.64**±2.5	**86.20**±1.6	**70.12**±2.9	**85.18**±2.1	**85.29**±1.9	68.85±3.1	82.43±2.6

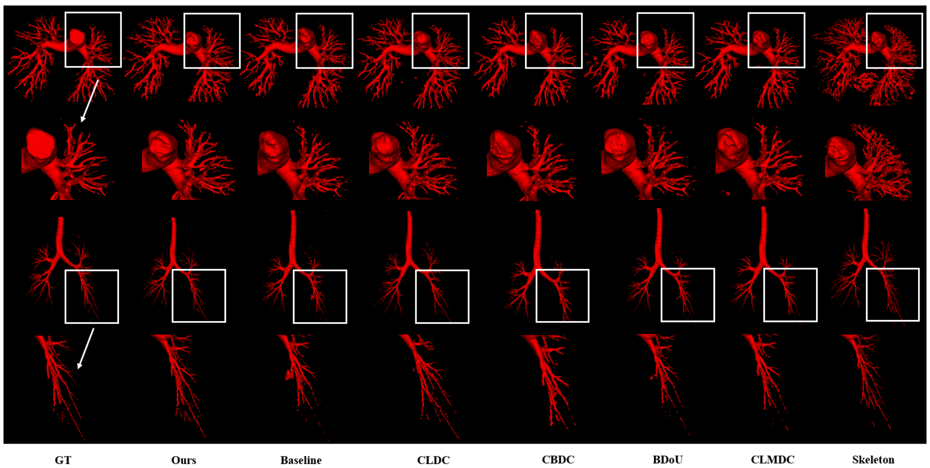

Fig. 2. Qualitative comparison on representative cases from the Artery and Bronchus datasets. Cupido preserves thin branches and bifurcation continuity more effectively than competing methods.

publicly annotated training cases from the PARSE dataset, while the remaining 103 cases belong to the test set without public annotations. Both datasets were split into training and testing sets with an 80:20 ratio. We utilized the default image preprocessing pipeline of nnU-Net [6] and incorporated deep supervision in all experiments. For fair comparison, all methods were trained under identical settings for 50 epochs following a fixed-budget evaluation protocol. This setting evaluates the ability of the proposed loss to improve segmentation performance without requiring additional training iterations. Other hyperparameters followed the default configurations of each dataset. For the bronchial task, the initial learning rate was set to 0.001. For the arterial task, the initial learning rate of UxLSTMEnc was set to 0.001, while those of the other two networks were set to 0.01. Tubular structure segmentation performance was evaluated using the Dice Similarity Coefficient (DSC), the Normalized Surface Distance (NSD) with a 0 mm tolerance, and the centerline Dice Similarity Coefficient (clDSC) as defined in [19].

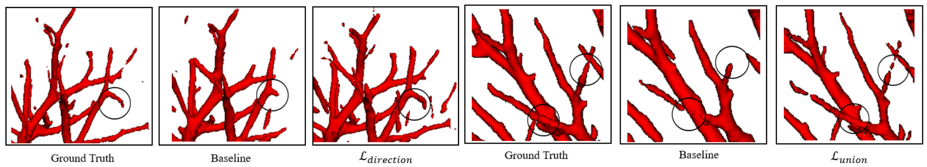

Fig. 3. A representative arterial case for ablation analysis demonstrating the effectiveness of **Cupido Loss** components in preserving coherent tubular structures.

3.3 Results

Performance Evaluation. We use nnUNet, UMambaBot, and UxLSTMEnc with default Dice and Cross-Entropy loss functions as baselines, adopting weight settings from [18] for experiments. For Cupido Loss, bronchial segmentation uses weights $\alpha = 0.5, \gamma = 0.1$ across all networks, while arterial segmentation uses $\alpha = 1.0, \gamma = 1.0$ for nnUNet and UxLSTMEnc, and $\alpha = 1.3, \gamma = 0.2$ for UMambaBot.

Quantitative Evaluation. As shown in Table 1, on the Artery dataset with nnUNet, Cupido increases Dice from 82.81 to 85.11 (+2.30) and clDice from 75.56 to 82.64 (+7.08), while NSD improves from 61.53 to 68.70 (+7.17). Similar trends are observed on Batch1, where Cupido improves UxLSTMEnc Dice from 89.28 to 91.51 and NSD from 84.84 to 86.62, indicating substantial gains in distal branch preservation. The ablation results in Table 2 further show that combining both loss components yields the best performance (e.g., nnUNet Dice 85.11 vs. 84.78 with directional loss only), confirming that each loss contributes complementary structural improvements.

Qualitative Visualization Analysis. Qualitative comparisons (Fig. 2) demonstrate that Cupido better preserves fine branch continuity while reducing false positives. Moreover, the ablation in Fig. 3 shows that our combined loss components effectively resolve structural failure modes, achieving optimal anatomical consistency.

4 Conclusion

We presented **Cupido Loss**, a unified structure-aware optimization framework for tubular segmentation that enhances boundary directional consistency and multi-scale geometric consistency. The proposed formulation integrates directional alignment and skeleton–distance union constraints to preserve structure and fine-branch morphology. Extensive experiments across two public benchmarks and three evaluation subsets demonstrate consistent improvements in both overlap and structure-sensitive metrics. Future work will explore distance-field-based directional supervision and its application to transformer and diffusion segmentation frameworks.

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

References

1. Chen, T., Ding, C., Zhu, L., Xu, T., Wang, Y., Ji, D., Zang, Y., Li, Z.: xlstm-unet can be an effective backbone for 2d 3d biomedical image segmentation better than its mamba counterparts. In: 2024 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI) (2024)
2. Clough, J.R., Byrne, N., Oksuz, I., Zimmer, V.A., Schnabel, J.A., King, A.P.: 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 (2020)
3. Cui, H., Yuwen, C., Jiang, L., Xia, Y., Zhang, Y.: Multiscale attention guided u-net architecture for cardiac segmentation in short-axis mri images. Computer Methods and Programs in Biomedicine **206**, 106142 (2021)
4. Glocker, B., Robinson, R., Castro, D.C., Dou, Q., Konukoglu, E.: Machine learning with multi-site imaging data: An empirical study on the impact of scanner effects. arXiv preprint arXiv:1910.04597 (2019)
5. Hu, X., Li, F., Samaras, D., Chen, C.: Topology-preserving deep image segmentation. Advances in neural information processing systems **32** (2019)
6. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
7. Kirchhoff, Y., Rokuss, M.R., Roy, S., Kovacs, B., Ulrich, C., Wald, T., Zenk, M., Vollmuth, P., Kleesiek, J., Isensee, F., et al.: Skeleton recall loss for connectivity conserving and resource efficient segmentation of thin tubular structures. In: European Conference on Computer Vision. pp. 218–234. Springer (2024)
8. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. Medical image analysis **13**(6), 819–845 (2009)
9. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
10. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A., van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. Medical image analysis **42**, 60–88 (2017)
11. Luo, G., Wang, K., Liu, J., Li, S., Liang, X., Li, X., Gan, S., Wang, W., Dong, S., Wang, W., et al.: Efficient automatic segmentation for multi-level pulmonary arteries: The parse challenge. arXiv preprint arXiv:2304.03708 (2023)
12. Ma, J., Li, F., Wang, B.: U-mamba: Enhancing long-range dependency for biomedical image segmentation. arXiv preprint arXiv:2401.04722 (2024)
13. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). pp. 565–571. Ieee (2016)
14. Qi, Y., He, Y., Qi, X., Zhang, Y., Yang, G.: Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6070–6079 (2023)

15. Qin, Y., Zheng, H., Gu, Y., Huang, X., Yang, J., Wang, L., Yao, F., Zhu, Y.M., Yang, G.Z.: Learning tubule-sensitive cnns for pulmonary airway and artery-vein segmentation in ct. IEEE transactions on medical imaging **40**(6), 1603–1617 (2021)
16. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015)
17. Salehi, S.S.M., Erdogmus, D., Gholipour, A.: Tversky loss function for image segmentation using 3d fully convolutional deep networks. In: International workshop on machine learning in medical imaging. pp. 379–387. Springer (2017)
18. Shi, P., Hu, J., Yang, Y., Gao, Z., Liu, W., Ma, T.: Centerline boundary dice loss for vascular segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 46–56. Springer (2024)
19. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P., Bauer, U., Menze, B.H.: cldice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16560–16569 (2021)
20. Sun, F., Luo, Z., Li, S.: Boundary difference over union loss for medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 292–301. Springer (2023)
21. Zhang, M., Wu, Y., Zhang, H., Qin, Y., Zheng, H., Tang, W., Arnold, C., Pei, C., Yu, P., Nan, Y., et al.: Multi-site, multi-domain airway tree modeling. Medical Image Analysis **90**, 102957 (2023)
22. Zheng, H., Qin, Y., Gu, Y., Xie, F., Yang, J., Sun, J., Yang, G.Z.: Alleviating classwise gradient imbalance for pulmonary airway segmentation. IEEE transactions on medical imaging **40**(9), 2452–2462 (2021)