Dynamic Collaborative Continual Test-Time Adaptation for 3D Vessel Segmentation

Xiang Li[1], Yuqi Fang[1,2], Dan Zhang[3], Jiong Zhang[4(✉)], and Caifeng Shan[1(✉)]

[1] School of Intelligence Science and Technology, Nanjing University, Suzhou, China
[2] National Institute of Healthcare Data Science at Nanjing University, Nanjing University, Nanjing, China
[3] School of Cyber Science and Engineering, Ningbo University of Technology, Ningbo, China
[4] School of Biomedical Engineering, Nanjing University, Suzhou, China
jiong.zhang@ieee.org; caifeng.shan@gmail.com

Abstract. Morphological characteristics of vessels serve as important biomarkers for numerous diseases, making accurate vessel segmentation fundamental to computer-aided diagnosis. However, cross-center domain shifts often lead to substantial segmentation performance degradation in real-world deployment. Continual Test-Time Adaptation (CTA) addresses this challenge by updating models online using streams of unlabeled test data, while mitigating error accumulation and catastrophic forgetting. Nevertheless, when applied to 3D vessel segmentation, CTA often introduces topology-breaking artifacts, such as vessel disconnections and spurious branches. To address these challenges, we propose DyCo-CTA, a dynamic collaborative continual test-time adaptation framework tailored for 3D vessel segmentation. Unlike conventional CTA methods with fixed teacher–student roles, we recognize that the teacher is not always more reliable than the student for complex 3D vessel structures, which may introduce bias and reduce adaptation quality. To mitigate this issue, we design a dynamic collaborative network that enables the two models to adaptively exchange teacher–student roles based on their relative reliability. In addition, we introduce a pseudo-break transformation to construct complementary teacher–student inputs that explicitly encourage structural integrity during adaptation. To further preserve global vessel connectivity, we propose a topological regularization loss that aligns topological critical points beyond pixel-level consistency. Extensive experiments on 3D cerebrovascular and cardiovascular segmentation benchmarks demonstrate that DyCo-CTA outperforms state-of-the-art methods, achieving improvements of 3.05% and 2.66%, respectively. Our code is available at: `https://github.com/lixiangcog/DyCo-CTA`.

Keywords: Continual Test-Time Adaptation · 3D Vessel Segmentation · Dynamic Collaborative Learning · Domain Shift · Cardio-Cerebrovasculature

1 Introduction

Morphological alterations in vascular networks act as key disease biomarkers, highlighting the importance of accurate vessel segmentation in computer-aided

diagnosis [1]. However, clinical deployment is often hindered by domain shift, as imaging data from different centers, devices, and patient populations may follow varying distributions [2,19,13]. Consequently, models trained in one institution suffer degradation when applied to unseen environments [30]. Moreover, due to privacy and data storage constraints that limit access to source data, models must be capable of adapting to new domains without retraining.

To address this issue, Continual Test-Time Adaptation (CTA) [20] has emerged as a practical paradigm: it updates a deployed model online using a non-stationary stream of unlabeled target-domain data, aiming to maintain performance under ongoing distribution shifts and potentially mitigate catastrophic forgetting. Most CTA methods perform self-supervised updates and can be categorized by their adaptation mechanisms into: 1) approaches that leverage batch normalization (BN) statistics to guide updates [11,26,22]; and 2) teacher–student modeling schemes [20,21,16,5,8,28]. BN-statistics-based adaptation is often ineffective for 3D volumetric data, as memory and computational constraints require small batch sizes that produce unstable statistics, yielding limited improvement or even performance degradation. In this context, an effective strategy for 3D vessel segmentation is to adopt a teacher-student framework [7] and perform online self-supervised updates with entropy-minimization objectives [18,24,14].

Although CTA has shown promise across a wide range of tasks, it still faces challenges in 3D vessel analysis. Because adaptation often relies on self-supervised signals, the model can reinforce its own errors. Vessels typically appear as fine tubular structures characterized by strong connectivity and topological constraints. Under continual adaptation, error reinforcement leads to topology-breaking artifacts, such as vessel disconnections and spurious branches [29]. These errors extend beyond pixel-level inaccuracies, disrupting global connectivity and further accelerating error accumulation [30]. In addition, static teacher-student schemes assume that the teacher consistently outperforms the student. However, this assumption often breaks down due to domain shift, image-quality variation, and anatomical complexity in 3D vessel segmentation. Unreliable teacher pseudo-labels can bias supervision, mislead online optimization, and amplify structural errors, resulting in more severe topological violations. [10].

To address these challenges, we propose DyCo-CTA, a dynamic collaborative continual test-time adaptation method tailored for 3D vessel segmentation. To our knowledge, this is the first work to explore CTA in the 3D vessel segmentation setting. Compared with the static teacher–student framework adopted by conventional CTA methods, DyCo-CTA incorporates dynamic collaboration, pseudo-break transformations, and topology-aware constraints, enabling online adaptation under continual domain shifts. The main contributions are as follows:

1) We propose DyCo-CTA, the first dynamic collaborative CTA framework tailored for 3D vessel segmentation, which alleviates the cognitive bias caused by the static teacher-student assumption. **2)** We further design a pseudo-break transformation to construct teacher-student inputs that explicitly focus on preserving structural integrity. **3)** We propose a topological regularization loss that aligns topological critical points beyond pixel-level consistency, reducing the ac-

cumulation of errors such as vessel disconnections and spurious branches. **4)** Experiments on 3D cerebrovascular and cardiovascular benchmarks show DyCo-CTA outperforms state-of-the-art methods by 3.05% and 2.66%, respectively.

2 Proposed method

As shown in Fig. 1, the proposed DyCo-CTA method is built upon three key components: a dynamic collaborative framework that mitigates the cognitive bias induced by the static teacher-student assumption (i.e., the teacher model is always superior to the student), the pseudo-break transformation, and a topological regularization constraint. The latter two components help preserve vascular structural integrity during continual updates and prevent error accumulation.

2.1 Dynamic Collaborative (DyCo) CTA

To mitigate the cognitive bias in the static teacher-student framework (Fig. 1 (a)), where the teacher is assumed to be consistently more reliable and may generate erroneous supervisory signals that are further amplified during online updates, we propose a dynamic collaborative network (Fig. 1 (b,c)). This network comprises two sub-networks, M_1 and M_2, which dynamically supervise each other and collaboratively improve performance. Both M_1 and M_2 are initialized with the weights of the source model.

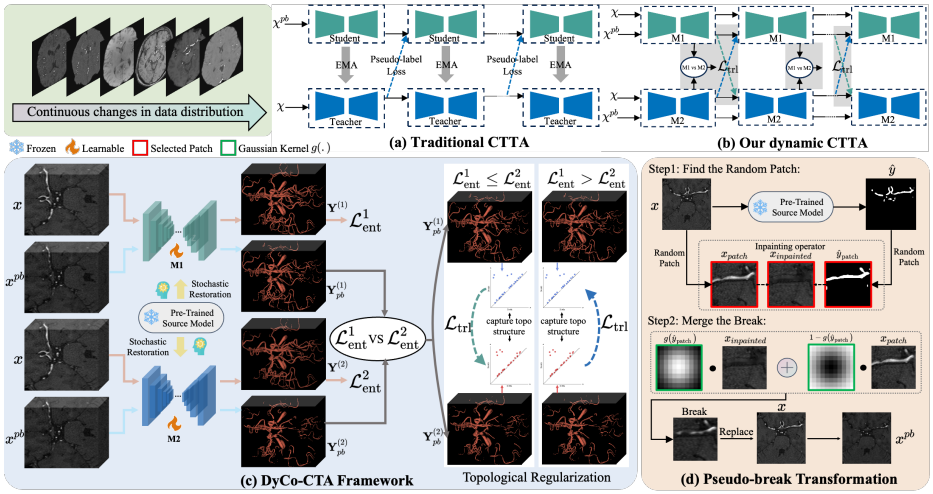

Fig. 1: Overview of DyCo-CTA. It comprises three components: (c) a dynamic collaborative CTA framework for 3D vessel segmentation; (d) the pseudo-break transformation; and the topological regularization constraint $\mathcal{L}_{\text{trl}}$.

As shown in Fig. 1 (b), at each adaptation step, we take the original volume x and its pseudo-break counterpart x^{pb} as inputs, where $x^{pb} = \mathcal{T}_{pb}(x)$, and $\mathcal{T}_{pb}$ denotes the pseudo-break transformation (Sec. 2.2). Both x and x^{pb} are fed into M_1 and M_2 to produce predictions:

$$\mathbf{Y}^{(i)}, \mathbf{Y}^{(i)}_{pb} = M_i(x),\ M_i(x^{pb}), \quad i = \{1, 2\}. \tag{1}$$

We compute prediction entropy on x: $\mathcal{L}^i_{ent} = \ell_{ent}\left(\mathbf{Y}^{(i)}\right), i \in \{1, 2\}$, where lower entropy indicates higher confidence, providing an effective criterion for assigning teacher-student roles. We define ℓ_{ent} as the volume-averaged voxel-wise entropy:

$$\ell_{ent}\left(\mathbf{Y}^{(i)}\right) = -\frac{1}{|\Omega|} \sum_{\mathbf{v}\in\Omega} \sum_{c=1}^{2} p_c^{(i)}(\mathbf{v}) \log p_c^{(i)}(\mathbf{v}), \tag{2}$$

where Ω denotes the voxel set, $c \in \{1, 2\}$ indexes the two classes, and $p_c^{(i)}(\mathbf{v})$ is the softmax probability at voxel $\mathbf{v}$; minimizing $\ell_{ent}\left(\mathbf{Y}^{(i)}\right)$ is also an optimization objective that encourages confident predictions [8].

We assign roles by comparing $\mathcal{L}^1_{ent}$ and $\mathcal{L}^2_{ent}$: if $\mathcal{L}^1_{ent} \leq \mathcal{L}^2_{ent}$, M_1 is the teacher and M_2 is the student; otherwise, the roles are swapped. Accordingly, the pseudo-label on x^{pb} and the prediction to be supervised are defined as:

$$\hat{\mathbf{Y}}_{pb} = \begin{cases} \mathbf{Y}^{(1)}_{pb}, & \mathcal{L}^1_{ent} \leq \mathcal{L}^2_{ent}, \\ \mathbf{Y}^{(2)}_{pb}, & \text{otherwise}, \end{cases} \qquad \mathbf{Y}^{o}_{pb} = \begin{cases} \mathbf{Y}^{(2)}_{pb}, & \mathcal{L}^1_{ent} \leq \mathcal{L}^2_{ent}, \\ \mathbf{Y}^{(1)}_{pb}, & \text{otherwise}. \end{cases} \tag{3}$$

Then, the loss $\mathcal{L}_{trl}$ (Eq. 11) is computed by comparing $\mathbf{Y}^{o}_{pb}$ with $\hat{\mathbf{Y}}_{pb}$. This loss is used to update the current student model via gradient descent. We use the teacher prediction on x^{pb} as a detached pseudo-label $\hat{\mathbf{Y}}_{pb}$ to supervise the student prediction $\mathbf{Y}^{o}_{pb}$, and update only the student network. In this way, our method enables dynamic teacher-student supervision between the two sub-networks. Through mutual supervision, M_1 and M_2 can dynamically refine their respective capabilities, leading to improved overall performance and a more fine-grained, deeper understanding of the complex characteristics of vessel data.

2.2 Pseudo Break Transformation (PBT)

As shown in Fig. 1 (d), to explicitly simulate topology-breaking perturbations and encourage restoration-oriented adaptation, we introduce the PBT. We draw an unlabeled target image x from the incoming stream and obtain a pseudo mask using the source model f_θ: $\hat{y} = \mathbb{I}(f_\theta(x) > \tau), \hat{y} \in \{0, 1\}^d$, where τ is a fixed threshold. We then remove the pseudo-labeled vessel structures via an inpainting operator [17] with a mask dilated three times to ensure smooth boundaries:

$$x_{inpainted} = \text{Inpaint}(x,\ \text{dilate}(\hat{y}, \text{times} = 3)). \tag{4}$$

We generate a pseudo-break sample by locally replacing selected structure regions in x with $x_{inpainted}$. Let $C = \{q_1, \ldots, q_m\}$ denote a set of pixel coordinates

selected from the current teacher prediction, where each $q \in C$ is treated as a patch center for local erasing. For any pixel coordinate u, we blend $x_{inpainted}$ into x within a local patch centered at each q using a Gaussian kernel:

$$g(u, q, \sigma) = \exp\big(-\|u - q\|_2^2 / 2\sigma^2\big) \in [0, 1], \tag{5}$$

where $\|u - q\|_2$ is the Euclidean distance between u and the center q, and σ controls the spatial decay of the blending weight. The patch operator $\varphi(q, s)$ denotes the set of coordinates within an $s \times s$ square patch centered at q, and the indicator $\mathbb{I}[u \in \varphi(q, s)]$ activates blending inside that patch. We aggregate contributions from all selected centers by defining a blending mask $M \in [0, 1]^d$:

$$M_u = \max_{q \in C} \mathbb{I}\big[u \in \varphi(q, s)\big] \cdot g(u, q, \sigma), \tag{6}$$

so that each pixel u takes the strongest local blending weight among candidate patches. Finally, the pseudo-break image is obtained by:

$$x^{pb} = x + M \odot (x_{inpainted} - x), \tag{7}$$

where $\odot$ denotes element-wise multiplication. These steps are visualized in Fig. 1 (d).

2.3 Topological Regularization Constraint ($\mathcal{L}_{\mathbf{trl}}$)

To reduce error accumulation such as vessel disconnections and spurious branches, we propose a topological regularization ($\mathcal{L}_{\text{trl}}$) in addition to pixel-level consistency (CE loss [15]). Inspired by persistence-diagram-based [23] unsupervised regularization, we develop a teacher-student objective for CTTA, where the teacher and student roles are assigned dynamically (Sec. 2.1). Let f_t and f_s be the teacher and student likelihood maps, with persistence diagrams $Dgm(f_t)$ and $Dgm(f_s)$. Direct matching may introduce spurious correspondences, so we split each diagram into *signal* and *noise* by persistence and apply constraints.

For any point $p \in Dgm(f)$, its persistence is defined as $per(p) = death(p) - birth(p)$, where $birth(p)$ and $death(p)$ [23] denote the filtration values at which the corresponding topological structure appears and disappears, respectively. Using a threshold φ, we split the diagram as:

$$\begin{aligned} Dgm(f)^{signal} &= \{p \in Dgm(f) \mid per(p) > \varphi\}, \\ Dgm(f)^{noise} &= \{p \in Dgm(f) \mid per(p) \leq \varphi\}. \end{aligned} \tag{8}$$

This decomposition is applied to both $Dgm(f_t)$ and $Dgm(f_s)$. We use a fixed φ in all experiments. After decomposition, we encourage the student to match the teacher on meaningful topological signals by computing an optimal matching on the signal diagrams. Specifically, we solve: $\gamma^* = \arg\min_{\gamma \in \Gamma} \sum_{p \in Dgm(f_s)^{signal}} \|p - \gamma(p)\|_2^2$, where Γ denotes the set of bijections between generalized persistence diagrams, and unmatched points are matched to the diagonal. $\gamma(p)$ is the matched point of p in the teacher diagram. The corresponding signal consistency loss is:

$$\mathcal{L}^U_{\text{topo-cons}} = \sum\nolimits_{p \in Dgm(f_s)^{signal}} \|p - \gamma^*(p)\|_2^2 . \tag{9}$$

While the signal term aligns stable structures, we suppress unstable artifacts by minimizing the total persistence of the student's noise diagram:

$$\mathcal{L}^{U}_{\text{topo-rem}} = \sum\nolimits_{p \in Dgm(f_s)^{noise}} \big[death(p) - birth(p)\big]^2. \tag{10}$$

The final noise-aware topological consistency loss $\mathcal{L}_{\text{trl}}$ becomes the sum of the two topology-aware loss terms:

$$\mathcal{L}_{\text{trl}} = \mathcal{L}^{U}_{\text{topo-cons}} + \mathcal{L}^{U}_{\text{topo-rem}}. \tag{11}$$

We adopt a persistence diagram construction so that gradients can backpropagate through $birth(\cdot)$ and $death(\cdot)$.

3 Experiments

3.1 Dataset and Evaluation

Cerebrovascular segmentation task. We evaluate test-time adaptation on six cerebral vessel segmentation datasets collected from different sources, denoted as domain A (IXI-HH[5]), B (IXI-Guys[5]), C (IXI-IOP[5]), D (COSTA-LocH1 [12]), E (ICBM[6]), and F (CereVessMRA-CN [4]). We treat one of the six domains as the labeled source domain and regard the remaining domains as unlabeled target domains, in order to assess the method's generalization and online adaptation ability under distribution shifts. For consistent preprocessing, all MRA volumes are resampled to a unified voxel spacing and reformatted to a consistent in-plane size of 512×512 (with variable depth along the slice dimension), and are normalized following standard practice in cerebral vessel segmentation [12].
Cardiovascular segmentation task. We validate DyCo-CTA on two CTA segmentation datasets from different sources: ImageCAS [25] and ASOCA [3]. All CTA volumes are resized to 512×512 (with variable depth along the slice dimension) and normalized using the same protocol as the cerebral vessel task. For the above tasks, we report DICE and HD95 [9] as the evaluation metrics.

3.2 Implementation Details

All experiments were conducted on a workstation equipped with an NVIDIA GeForce RTX 4090 (24GB) GPU, using a containerized environment based on PyTorch 2.4.0 and CUDA 12.4. For each task, we independently trained a source model on the source domain (using ResUNet3D [6] and setting the Adam learning rate to 0.01), and performed CTA evaluation on the remaining target domains; the final results were reported as the average performance across target domains. For online continual test-time adaptation, we used four update steps per test batch with a batch size of 1 and a learning rate of 1×10^{-4}. The weight Stochastic

[5] `http://brain-development.org/ixi-dataset/`
[6] `https://www.nitrc.org/projects/icbmmra`

Restoration mentioned in Fig. 1 follows [20]. In the pseudo-label transformation module, the mask binarization threshold τ was set to 0.5 and the patch size was set to 15×15. In topological regularization, the persistence threshold φ for decomposing persistence diagrams was set to 0.7.

3.3 Experimental Results

Cerebrovascular and cardiovascular segmentation results. As shown in Table 2 (cerebrovasculature), DyCo-CTA achieves the best performance across all target domains (A–F), obtaining the highest average Dice (0.7529) and the lowest average HD95 (13.67). It significantly outperforms the "Source Only" baseline (0.6799/18.55) and further surpasses the strongest prior teacher-student method. Table 3 (cardiovasculature) shows DyCo-CTA achieves the best mean Dice/HD95 (0.7628/10.19), outperforming Source Only and MoASE [27].

Long-term adaptation performance. To evaluate the performance of DyCo-CTA in the CTA setting [20], we conducted multiple-round experiments on the cerebrovascular segmentation task, and report the results of each round in Table 1. This experiment aims to verify the model's ability to cope with error accumulation and to suppress catastrophic forgetting over multiple adaptation rounds. DyCo-CTA achieves the highest average Dice score while exhibiting the smallest performance degradation.

Qualitative comparison. As shown in Fig. 2, DyCo-CTA yields predictions closest to the ground truth across cross-domain settings, with improved vessel continuity and fewer spurious branches than other methods.

Visualization of the PBT effect. As shown in Fig. 3, we visualize the effect of our PBT by showing patch-level comparisons across three views: the original image patch, the inpainted patch, and the resulting pseudo-break patch.

Table 1: Under the long-term CTA setting, the performance of DyCo-CTA and other methods on the cerebrovascular segmentation task is evaluated by DICE. "Perform. Degra." denotes the difference between Round 1 and the average over subsequent rounds. Round 1 results match Table 2 and are omitted for brevity.

Time	→																
Round	2							3							Avg. $DICE\uparrow$	Perform. Degra. ↓
Methods	A	B	C	D	E	F	Ave.	A	B	C	D	E	F	Ave.		
No Adapt	0.6843	0.6297	0.6931	0.6504	0.7013	0.7206	0.6799	0.6843	0.6297	0.6931	0.6504	0.7013	0.7206	0.6799	0.6799	0.0000
TENT	0.6534	0.5821	0.6017	0.4839	0.6243	0.6132	0.5931	0.5982	0.5234	0.5576	0.3984	0.5837	0.5428	0.5340	0.5636	0.0854
DLTTA	0.7018	0.6521	0.7302	0.6974	0.7236	0.7412	0.7077	0.6871	0.6432	0.7148	0.6803	0.7107	0.7324	0.6948	0.7013	0.0117
CoTTA	0.7146	0.6661	0.7634	0.7324	0.7407	0.7482	0.7276	0.7084	0.6591	0.7571	0.7264	0.7352	0.7416	0.7213	0.7245	0.0062
MoASE	0.7094	0.6603	0.7576	0.7291	0.7351	0.7523	0.7239	0.7118	0.6521	0.7527	0.7234	0.7292	0.7486	0.7196	0.7218	0.0083
DyCo-CTA (Ours)	**0.7327**	**0.7064**	**0.7812**	**0.7432**	**0.7587**	**0.7743**	**0.7494**	**0.7274**	**0.7031**	**0.7796**	**0.7408**	**0.7561**	**0.7684**	**0.7459**	**0.7477**	**0.0052**

3.4 Ablation Studies

In Table 4, to assess the contribution of each module to DyCo-CTA on the cerebrovascular segmentation task, we conduct ablation studies covering the dynamic collaborative framework, the pseudo-break transformation, and the topological

Table 2: Performance comparison of DyCo-CTA, the "Source Only" baseline, and other SOTA methods on the cerebrovascular segmentation task. The best and second-best results in each column are highlighted in **bold** and underlined.

Methods	Category	Domain A		Domain B		Domain C		Domain D		Domain E		Domain F		Average	
		DICE↑	*HD95*↓	*DICE*↑	*HD95*↓	*DICE*↑	*HD95*↓	*DICE*↑	*HD95*↓	*DICE*↑	*HD95*↓	*DICE*↑	*HD95*↓	*DICE*↑	*HD95*↓
Source Only (ResUNet3D)	–	0.6843	22.31	0.6297	21.74	0.6931	18.42	0.6504	20.27	0.7013	14.86	0.7206	13.72	0.6799	18.55
DUA (CVPR 2022) [11]	BN-based	0.6974	21.08	0.6472	20.03	0.7068	17.39	0.6735	19.04	0.7126	14.07	0.7199	12.49	0.6929	17.35
DomainAdaptor (CVPR 2023) [26]	BN-based	0.7028	20.43	0.6536	19.37	0.7197	16.48	0.6894	18.16	0.7231	14.55	0.7429	11.82	0.7053	16.80
SiCTA (MedIA 2026) [22]	BN-based	0.7115	19.74	0.6594	18.43	0.7338	15.06	0.7026	17.31	0.7342	13.06	0.7523	11.18	0.7156	15.80
TENT (ICLR 2021) [18]	Entropy Min.	0.6792	23.68	0.6193	24.73	0.6482	28.34	0.5831	35.41	0.6732	21.26	0.6904	20.03	0.6489	25.58
DLTTA (TMI 2022) [24]	Entropy Min.	0.7132	19.81	0.6124	18.72	0.7438	15.31	0.7129	17.06	0.7397	12.63	0.7568	11.02	0.7131	15.76
SAR (ICLR 2023) [14]	Entropy Min.	0.7214	19.36	0.6652	18.34	0.7496	14.83	0.7197	16.74	0.7368	12.94	0.7529	11.16	0.7243	15.56
CoTTA (CVPR 2022) [20]	Teacher-student	0.7201	19.42	0.6693	18.26	0.7672	14.35	0.7364	16.14	0.7415	12.79	0.7493	10.97	0.7306	15.32
UPL-TTA (IPMI 2023) [21]	Teacher-student	0.7168	19.73	0.6635	18.82	0.7526	14.77	0.7231	16.76	0.7394	13.04	0.7534	11.39	0.7248	15.75
ECoTTA (CVPR 2023) [16]	Teacher-student	0.7094	20.08	0.6591	19.03	0.7437	15.04	0.7136	17.04	0.7324	13.36	0.7468	11.76	0.7175	16.05
BECoTTA (ICML 2024) [5]	Teacher-student	0.7124	19.86	0.6627	18.74	0.7513	14.81	0.7226	16.83	0.7365	13.14	0.7541	11.47	0.7233	15.81
MoASE (AAAI 2026) [27]	Teacher-student	0.7118	19.49	0.6624	17.36	0.7628	13.44	0.7327	16.19	0.7473	14.68	0.7635	10.82	0.7301	15.33
DyCo-CTA (Ours)	Teacher-student	**0.7352**	**17.76**	**0.7094**	**16.38**	**0.7841**	**12.73**	**0.7473**	**14.42**	**0.7612**	**11.29**	**0.7804**	**9.46**	**0.7529**	**13.67**

Table 3: Comparison of different SOTA methods on the cardiovascular dataset.

Methods	Category	*ImageCAS* → *ASOCA*		*ASOCA* → *ImageCAS*		Avg.	
		DICE ↑	*HD95* ↓	*DICE* ↑	*HD95* ↓	*DICE* ↑	*HD95* ↓
Source Only (ResUNet3D)	–	0.7024	15.34	0.6731	18.23	0.6878	16.79
DUA (CVPR 2022)	BN-based	0.7143	14.29	0.6826	17.14	0.6985	15.72
DomainAdaptor (CVPR 2023)	BN-based	0.7238	13.56	0.6942	16.28	0.7090	14.92
SiCTA (MedIA 2026)	BN-based	0.7351	12.89	0.7064	15.42	0.7208	14.16
TENT (ICLR 2021)	Entropy Min.	0.6813	21.04	0.6428	28.35	0.6621	24.70
DLTTA (TMI 2022)	Entropy Min.	0.7416	11.93	0.7123	14.67	0.7270	13.30
SAR (ICLR 2023)	Entropy Min.	0.7482	11.43	0.7196	14.03	0.7339	12.73
CoTTA (CVPR 2022)	Teacher-student	0.7564	10.36	0.7284	13.25	0.7424	11.81
UPL-TTA (IPMI 2023)	Teacher-student	0.7493	11.27	0.7182	14.14	0.7338	12.71
ECoTTA (CVPR 2023)	Teacher-student	0.7381	12.42	0.7047	15.68	0.7214	14.05
BECoTTA (ICML 2024)	Teacher-student	0.7462	11.54	0.7163	14.34	0.7313	12.94
MoASE (AAAI 2026)	Teacher-student	0.7548	10.63	0.7312	12.88	0.7430	11.76
DyCo-CTA (Ours)	Teacher-student	**0.7742**	**9.14**	**0.7513**	**11.24**	**0.7628**	**10.19**

Table 4: Ablation study of the DyCo-CTA on the cerebrovascular datasets.

Components			Avg. *DICE* ↑	Avg. *HD95* ↓
DyCo	PBT	$\mathcal{L}_{trl}$		
			0.6799	18.55
✓			0.7342	15.38
✓	✓		0.7426	14.64
✓	✓	✓	**0.7529**	**13.67**

regularization term. When the pseudo-break transformation is disabled, we replace it with the conventional TTA data augmentation strategy [8]. When the topological regularization term is removed, we use the CE loss instead.

4 Conclusion

We proposed DyCo-CTA, a dynamic collaborative continual test-time adaptation framework for 3D vessel segmentation under domain shift. By dynamically swapping teacher–student roles, introducing a pseudo-break transform to emphasize structural integrity, and enforcing topological regularization beyond pixel-level consistency, DyCo-CTA effectively reduces topology-breaking artifacts such as disconnections and spurious branches during online adaptation. Experiments on 3D cerebrovascular and cardiovascular benchmarks demonstrate consistent improvements over prior CTA methods, validating the benefit of topology-aware dynamic collaboration for robust real-world deployment.

Acknowledgments. This work was supported by the National Science Foundation Program of China (62371442), in part by the Zhejiang Provincial Natural Science Foundation (LQ23F010002, LR24F010002), and in part by the Young Scientists Fund of the National Natural Science Foundation of China (13001024).

Disclosure of Interests. The authors declare no competing interests.

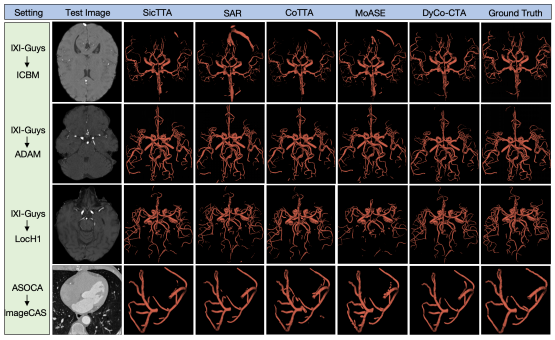

Fig. 2: Visualization of cerebrovascular and cardiovascular segmentation results for DyCo-CTA and other state-of-the-art methods.

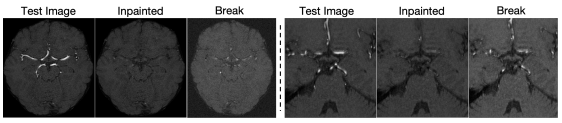

Fig. 3: Visualization of pseudo-break transformation. We show the pseudo-break patch (image-inpainted-break) to illustrate the effect of the transformation.

References

1. Elsayed, N., Osman, Y.B.M., Li, C., Liu, J., Si, W., Zhang, J., Wang, S.: Automating vessel segmentation in the heart and brain: A trend to develop multi-modality and label-efficient deep learning techniques. Biomedical Signal Processing and Control **109**, 108028 (2025)
2. Fang, Z., Lu, J., Liu, F., Zhang, G.: Semi-supervised heterogeneous domain adaptation: Theory and algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(1), 1087–1105 (2022)
3. Gharleghi, R., Adikari, D., Ellenberger, K., Ooi, S.Y., Ellis, C., Chen, C.M., Gao, R., He, Y., Hussain, R., Lee, C.Y., et al.: Automated segmentation of normal and diseased coronary arteries–the asoca challenge. Computerized Medical Imaging and Graphics **97**, 102049 (2022)
4. Guo, B., Chen, Y., Lin, J., Huang, B., Bai, X., Guo, C., Gao, B., Gong, Q., Bai, X.: Self-supervised learning for accurately modelling hierarchical evolutionary patterns of cerebrovasculature. Nature Communications **15**(1), 9235 (2024)
5. Lee, D., Yoon, J., Hwang, S.J.: Becotta: Input-dependent online blending of experts for continual test-time adaptation. In: International Conference on Machine Learning (2024)
6. Lee, K., Zung, J., Li, P., Jain, V., Seung, H.S.: Superhuman accuracy on the snemi3d connectomics challenge. arXiv:1706.00120 (2017)
7. Li, X., Fang, H., Liu, M., Xu, Y., Duan, L.: Diffusion-enhanced transformation consistency learning for retinal image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 221–231 (2024)

8. Li, X., Fang, H., Wang, C., Liu, M., Duan, L., Xu, Y.: Cache-driven spatial test-time adaptation for cross-modality medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 146–156 (2024)
9. Livne, M., Rieger, J., Aydin, O.U., Taha, A.A., Akay, E.M., Kossen, T., Sobesky, J., Kelleher, J.D., Hildebrand, K., Frey, D., et al.: A u-net deep learning framework for high performance vessel segmentation in patients with cerebrovascular disease. Frontiers in neuroscience **13**, 97 (2019)
10. Lyu, F., Du, K., Li, Y., Zhao, H., Hu, F., Zhang, Z., Liu, G., Wang, L.: Variational continual test-time adaptation. arXiv:2402.08182 (2024)
11. Mirza, M.J., Micorek, J., Possegger, H., Bischof, H.: The norm must go on: Dynamic unsupervised domain adaptation by normalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14765–14775 (2022)
12. Mou, L., Lin, J., Zhao, Y., Liu, Y., Ma, S., Zhang, J., Lv, W., Zhou, T., Liu, J., Frangi, A.F., Zhao, Y.: Costa: A multi-center tof-mra dataset and a style self-consistency network for cerebrovascular segmentation. IEEE Transactions on Medical Imaging **43**(12), 4442–4456 (2024)
13. Ning, M., Lu, D., Xie, Y., Chen, D., Wei, D., Zheng, Y., Tian, Y., Yan, S., Yuan, L.: Madav2: Advanced multi-anchor based active domain adaptation segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(11), 13553–13566 (2023)
14. Niu, S., Wu, J., Zhang, Y., Wen, Z., Chen, Y., Zhao, P., Tan, M.: Towards stable test-time adaptation in dynamic wild world. arXiv:2302.12400 (2023)
15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. arXiv:1505.04597 (2015)
16. Song, J., Lee, J., Kweon, I.S., Choi, S.: Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11920–11929 (2023)
17. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2149–2159 (2022)
18. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. arXiv:2006.10726 (2020)
19. Wang, J., Li, X., Cheng, Y.: Towards an extended efficientnet-based u-net framework for joint optic disc and cup segmentation in the fundus image. Biomedical Signal Processing and Control **85**, 104906 (2023)
20. Wang, Q., Fink, O., Van Gool, L., Dai, D.: Continual test-time domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7201–7211 (2022)
21. Wu, J., Gu, R., Lu, T., Zhang, S., Wang, G.: Upl-tta: Uncertainty-aware pseudo label guided fully test time adaptation for fetal brain segmentation. In: International Conference on Information Processing in Medical Imaging. pp. 237–249 (2023)
22. Wu, J., Liu, X., Wang, G., Zhang, S.: Sictta: Single image continual test time adaptation for medical image segmentation. Medical Image Analysis **108**, 103859 (2026)
23. Xu, M., Hu, X., Gupta, S., Abousamra, S., Chen, C.: Semi-supervised segmentation of histopathology images with noise-aware topological consistency. In: European Conference on Computer Vision. pp. 271–289 (2024)

24. Yang, H., Chen, C., Jiang, M., Liu, Q., Cao, J., Heng, P.A., Dou, Q.: Dltta: Dynamic learning rate for test-time adaptation on cross-domain medical images. IEEE Transactions on Medical Imaging **41**(12), 3575–3586 (2022)
25. Zeng, A., Wu, C., Lin, G., Xie, W., Hong, J., Huang, M., Zhuang, J., Bi, S., Pan, D., Ullah, N., et al.: Imagecas: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. Computerized Medical Imaging and Graphics **109**, 102287 (2023)
26. Zhang, J., Qi, L., Shi, Y., Gao, Y.: Domainadaptor: A novel approach to test-time adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18971–18981 (2023)
27. Zhang, R., Cheng, A., Luo, Y., Dai, G., Yang, H., Liu, J., Xu, R., Du, L., Du, Y., Jiang, Y., Zhang, S.: Decomposing the neurons: Activation sparsity via mixture of experts for continual test time adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence (2026)
28. Zhang, R., Cheng, A., Luo, Y., Dai, G., Yang, H., Liu, J., Xu, R., Du, L., Du, Y., Jiang, Y., et al.: Decomposing the neurons: Activation sparsity via mixture of experts for continual test time adaptation. arXiv:2405.16486 (2024)
29. Zhou, F., Gao, Z., Zhao, H., Xie, J., Meng, Y., Zhao, Y., Lip, G.Y., Zheng, Y.: Glcp: Global-to-local connectivity preservation for tubular structure segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 237–247 (2025)
30. Zhou, J., Wang, W., Li, S., Qu, X., Guo, X., Liu, Y., Tang, W., Lin, X., Zheng, Y.: Topotta: Topology-enhanced test-time adaptation for tubular structure segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 24123–24134 (2025)