

TG-OT: Topology-guided CCTA-IVUS registration via optimal transport matching

Rudolf L. M. van Herten^{1,2,3}, José P. Henriques⁴, R. Nils Planken⁵,
Joost Daemen⁶, Eline M. J. Hartman⁶, Jolanda J. Wentzel⁶,
Johannes C. Paetzold^{1,2}, and Ivana Išgum^{3,4,5}

¹ Department of Radiology, Weill Cornell Medicine, New York, USA
rlv4001@med.cornell.edu

² Cornell Tech, New York, USA

³ Department of Biomedical Engineering and Physics, Amsterdam UMC,
University of Amsterdam, Amsterdam, The Netherlands

⁴ Amsterdam Cardiovascular Sciences, Amsterdam, The Netherlands

⁵ Department of Radiology, Mayo Clinic, Rochester, USA

⁶ Department of Cardiology, Erasmus University Medical Center, Rotterdam,
The Netherlands

Abstract. Registering coronary CT angiography (CCTA) and intravascular ultrasound (IVUS) enables comprehensive coronary analysis that neither modality can provide alone, yet their fusion remains challenging due to differences in imaging geometry, resolution, and artifact profiles. Existing methods depend on pre-computed lumen or vessel wall segmentations that are unreliable under IVUS acoustic shadowing from calcifications, limiting their clinical applicability. We propose TG-OT, a fully automatic CCTA-IVUS registration framework that eliminates this dependency by integrating trained feature detectors directly into the registration pipeline. Lightweight CNNs are trained to predict calcifications, bifurcations, and lumen radii on the topological (θ, z) cylinder, encouraging topologically coherent detections without requiring explicit segmentation. Registration is formulated as an optimization over centerline warping parameters, driven by an unbalanced Sinkhorn optimal transport loss on the cylindrical geometry that provides spatially informative gradients even for spatially disjoint predictions, complemented by a lumen matching term. Evaluated on $N=47$ paired CCTA-IVUS cases in a 5-fold cross-validation setup, TG-OT achieves strong longitudinal ($\overline{\text{Dice}}_{\text{cti}}=0.99$), rotational ($\overline{S}_c=0.96$), and lumen alignment ($\overline{\text{Dice}}_L=0.69$) without manual interaction or prior segmentation, marking a meaningful step toward clinical integration of automatic CCTA-IVUS fusion.

Keywords: Multimodality image registration · Topology · Optimal transport · Coronary CT angiography · Intravascular ultrasound.

1 Introduction

Coronary CT angiography (CCTA) and intravascular ultrasound (IVUS) present complementary modalities for the analysis of coronary artery disease [15,28]:

CCTA captures the full 3D arterial geometry non-invasively, while IVUS offers high-resolution cross-sectional imaging of vessel wall morphology and plaque composition [7,20]. Registering these modalities enables comprehensive coronary analyses that neither modality can provide alone [13,19], yet multimodality fusion remains challenging due to fundamental differences in imaging geometry, resolution, and artifact profiles [11,30].

CCTA-IVUS registration requires extracting a coronary artery centerline from CCTA to generate multiplanar reformatted (MPR) images, which are then aligned to the IVUS pullback. Early methods relied on manual landmark selection and interactive rotation [10,18], while later semi-automatic approaches reduced user burden through seedpoint selection and segmentation [2,22]. More recently, Kadry et al. [14] and Li et al. [16] proposed differentiable and Transformer-based frameworks respectively, but both depend on pre-computed lumen or vessel wall segmentations—a condition particularly difficult to satisfy for IVUS, where acoustic shadowing from calcifications obscures vessel wall boundaries [1].

In this work, we propose **topology-guided image registration via optimal transport (TG-OT)**, a fully automatic CCTA-IVUS registration framework that integrates feature detection directly into the registration pipeline, eliminating the need for prior segmentation. Our contributions are:

- (i) We train lightweight CNNs to predict calcifications, bifurcations, and lumen radii on the topological (θ, z) cylinder of the artery, circumventing the ill-defined delineation of calcified regions under IVUS acoustic shadowing.
- (ii) We introduce an unbalanced Sinkhorn optimal transport (OT) loss [25] on cylindrical geometry for cross-modal feature matching, stabilizing voxel-wise losses with spatially aware transport costs that tolerate differences in predicted feature mass between modalities.
- (iii) We demonstrate that the integrated pipeline achieves competitive registration accuracy on $N=47$ paired CCTA-IVUS cases without requiring any prior segmentation.

2 Method

We present TG-OT, a flexible framework for automatic CCTA-IVUS registration. Our method comprises two main components. First, CNNs are trained to identify key anatomical features along with lumen radii from polar-transformed MPR or IVUS data, while additionally learning to detect guidewire artifacts in IVUS images. Second, once trained, the CNNs are integrated into a differentiable image registration module which optimizes a set of transformation parameters that optimally align an IVUS pullback to a corresponding automatically extracted coronary artery centerline in CCTA [12]. An overview of our registration method is presented in Fig. 1. Code is available on [GitHub](#).

2.1 Feature detection CNN

Aligning CCTA and IVUS requires identifying anatomical features visible across both modalities. We detect lumen radii, calcifications, and bifurcation locations




Fig. 1. Overview of our TG-OT framework for CCTA-IVUS image registration. A CCTA volume is sampled in polar coordinates along an extracted centerline, with warping governed by trainable parameters ϕ_T . Frozen CNNs f_ξ predict lumen radii (L), bifurcations (B), and calcifications (C) on the topological cylinder for both modalities, bypassing explicit segmentation. Registration is driven by $\mathcal{L}_{\text{feat}}$, combining unbalanced Sinkhorn optimal transport with a Dice loss to match sparse features across modalities with spatially informed gradients, and $\mathcal{L}_{\text{lumen}}$ denoting a lumen radii matching term.

from polar-transformed MPR and IVUS data, operating in the polar (r, θ, z) domain where classifying the *presence* of structures along each radial ray is well-defined, even when full segmentation is ill-posed under IVUS acoustic shadowing. We adapt FanCNN [29] and train separate networks f_ξ for MPR and IVUS, with the IVUS network additionally predicting guidewire artifacts. Each f_ξ is supervised with a combination of soft Dice and Betti matching loss [4,27] ($\alpha_{\text{BM}} = 0.05$), the latter penalizing topological errors in detected features, a useful property for reliable downstream OT matching. Lumen radii are simultaneously regressed using a standard L_1 loss. An overview is shown in Fig. 2.

2.2 Registration pipeline

Registration is formulated as a centerline warping problem: given a CCTA centerline and an IVUS pullback, we optimize transformation parameters that align the two through a combination of sparse anatomical feature matching, intensity-based similarity, and lumen area agreement. The pipeline proceeds in three stages: a transformation model defines the warping parameterization, rigid pre-alignment initializes global pose, and deformable optimization jointly refines all parameters, driven primarily by an unbalanced OT objective on the cylindrical geometry. We denote the CCTA-derived MPR as the warped moving image M_w and the IVUS pullback as the fixed image F throughout.




Fig. 2. Overview of the feature detection CNN f_ξ . IVUS (top) and MPR (bottom) data are transformed from Cartesian (u, v, z) to polar (r, θ, z) coordinates via $P \odot$. A 3D CNN consisting of N ResNet blocks processes the polar-transformed data, followed by average pooling along the radial (r) dimension to produce 2D (θ, z) predictions of lumen radii (L), bifurcation locations (B), and calcification presence (C). The resulting (θ, z) map represents the unwound topological cylinder $\mathbb{S}^1 \times \mathbb{R}$, with periodic boundary conditions at $\theta = 0 \equiv 2\pi$.

Transformation model. Given a CCTA centerline $\mathcal{C} \in \mathbb{R}^{n \times 3}$ and an IVUS pullback, we seek transformation parameters ϕ_T comprising global scaling s_z and translation t_z along the vessel axis, and per-frame local rotation θ_p and in-plane displacement $(t_{u,p}, t_{v,p})$. Following [14], local rotations are parameterized via a cumulative B-spline [9], naturally enforcing smoothness and capturing the gradual twist of the IVUS catheter during pullback. Each centerline point $\mathbf{x}_p \in \mathbb{R}^3$ is transformed by:

$$\mathbf{x}'_p = \mathbf{T}_g \mathbf{T}_{l,p} \mathbf{x}_p, \quad (1)$$

where $\mathbf{T}_g \in \mathbb{R}^{4 \times 4}$ embeds the global 1D affine parameters (s_z, t_z) into a 3D transformation, and $\mathbf{T}_{l,p} \in \text{SE}(2)$ encodes local in-plane rotation and displacement, embedded analogously with an identity component along z .

Rigid pre-alignment. Pre-alignment is performed in two stages. First, the global scaling factor s_z is initialized analytically as the ratio of the CCTA centerline length to the IVUS pullback length. Longitudinal alignment is then determined by grid search over t_z while minimizing a combined objective:

$$\mathcal{L}_{\text{init}} = \sum_{k \in \{B, C\}} \mathcal{L}_{\text{Dice}}(\ell_{M_w}^k(z), \ell_F^k(z)) + \mathcal{L}_{\text{lumen}}(A_{M_w}, A_F), \quad (2)$$

where $\ell^k(z) = \max_\theta f_\xi(P(\cdot))_k(\theta, z)$ aggregates the 2D classifier output for class k along the angular dimension to produce a 1D longitudinal signal per feature type, $\mathcal{L}_{\text{Dice}}$ denotes the soft Dice loss, and $\mathcal{L}_{\text{lumen}}$ matches cross-sectional lumen

areas (A) computed from predicted radii via a radial polygon approximation, using a combination of absolute error and normalized cross-correlation [3]. Subsequently, $\theta_0 \in [0, 2\pi)$ is found by grid search (step size = 5°), shifting classifier outputs along θ and evaluating $\mathcal{L}_{\text{Dice}}$ summed over $k \in \{B, C\}$.

Deformable optimization. Starting from the rigid initialization, all parameters ϕ_T are jointly refined. The objective combines sparse anatomical feature matching, lumen radius agreement, and parameter regularization:

$$\hat{\phi}_T = \arg \min_{\phi_T} [\mathcal{L}_{\text{feat}} + \alpha_{\text{lumen}} \mathcal{L}_{\text{lumen}} + \alpha_{\text{reg}} \mathcal{L}_{\text{reg}}]. \quad (3)$$

where $\mathcal{L}_{\text{feat}}$ is the feature matching term detailed below, $\mathcal{L}_{\text{lumen}}$ provides complementary lumen-based alignment, and $\mathcal{L}_{\text{reg}}$ enforces spatial smoothness of the local transformation parameters.

Feature matching via unbalanced optimal transport. A voxel-wise loss such as Dice penalizes feature mismatch uniformly regardless of spatial proximity, providing no gradient signal when predicted features do not overlap. We instead formulate feature matching as an optimal transport problem on the topological cylinder $\mathbb{S}^1 \times \mathbb{R}$, where the transport cost encodes geodesic distance and thus provides informative gradients even for spatially disjoint predictions.

Given classifier probability maps derived from polar-transformed $P(\cdot)$ data

$$\mu^k = f_\xi(P(M_w))_k, \quad \nu^k = f_\xi(P(F))_k, \quad k \in \{B, C\}, \quad (4)$$

for the warped moving image M_w and fixed image F respectively, the feature matching loss combines an OT term with a Dice overlap term:

$$\mathcal{L}_{\text{feat}} = \sum_{k \in \{B, C\}} [\mathcal{L}_{\text{Dice}}(\mu^k, \nu^k) + \alpha_{\text{OT}} \mathcal{L}_{\text{OT}}(\mu^k, \nu^k)]. \quad (5)$$

The Dice term encourages overlap between predicted feature mass per class, while the OT term captures the spatial cost of transporting misaligned features across the cylinder.

Transport cost. The cost matrix is defined via geodesic cylindrical distance:

$$C_{ij} = \sqrt{d_\theta(i, j)^2 + d_z(i, j)^2}, \quad (6)$$

where d_θ is the circular (wraparound) distance on $\mathbb{S}^1$, and d_z is the Euclidean distance along the vessel axis.

Unbalanced Sinkhorn formulation. The optimal transport plan is computed via the Sinkhorn algorithm [8] with entropic regularization ε and unbalanced marginal relaxation parameter $\tau \in (0, 1)$:

$$\mathcal{L}_{\text{OT}} = \langle \mathbf{C}, \mathbf{P}^* \rangle, \quad \text{where } \mathbf{P}^* = \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{v}), \quad (7)$$

Table 1. Cylindrical feature detection performance of f_ξ on IVUS and MPR across training loss configurations. Reported metrics are Dice score for lumen, bifurcation, calcification, and guidewire detection, Betti matching error (τ_{err}) for bifurcations and calcifications, and centerline Dice (clDice) for guidewire continuity. Results are reported as mean (std) over 5-fold cross-validation, with the best score indicated in **bold**.

Loss	Lumen		Bifurcations		Calcifications		Guidewire	
	Dice $\uparrow$		Dice $\uparrow$	$\tau_{\text{err}} \downarrow$	Dice $\uparrow$	$\tau_{\text{err}} \downarrow$	Dice $\uparrow$	clDice $\uparrow$
IVUS	$\mathcal{L}_{\text{BCE}}$	0.94 (0.02)	0.46 (0.17)	4.53 (3.33)	0.68 (0.25)	4.34 (3.79)	0.55 (0.17)	0.60 (0.19)
	$\mathcal{L}_{\text{Dice}}$	0.90 (0.04)	0.53 (0.14)	4.53 (2.38)	0.64 (0.26)	3.66 (3.60)	0.67 (0.08)	0.71 (0.11)
	$\mathcal{L}_{\text{BCE}+\text{Dice}}$	0.92 (0.02)	0.52 (0.16)	4.23 (2.71)	0.67 (0.24)	4.09 (3.83)	0.64 (0.14)	0.68 (0.16)
	$\mathcal{L}_{\text{BM}+\text{Dice}}$	0.92 (0.02)	0.56 (0.13)	4.15 (2.65)	0.67 (0.23)	3.89 (2.98)	0.69 (0.09)	0.72 (0.11)
MPR	$\mathcal{L}_{\text{BCE}}$	0.90 (0.04)	0.53 (0.17)	2.61 (1.87)	0.60 (0.27)	1.65 (2.34)	-	-
	$\mathcal{L}_{\text{Dice}}$	0.90 (0.04)	0.62 (0.13)	2.33 (1.71)	0.65 (0.25)	1.45 (1.87)	-	-
	$\mathcal{L}_{\text{BCE}+\text{Dice}}$	0.89 (0.05)	0.59 (0.16)	2.37 (1.80)	0.62 (0.26)	1.65 (2.14)	-	-
	$\mathcal{L}_{\text{BM}+\text{Dice}}$	0.91 (0.03)	0.61 (0.12)	2.33 (1.89)	0.63 (0.26)	1.39 (1.94)	-	-

with Gibbs kernel $K_{ij} = \exp(-C_{ij}/\varepsilon)$, marginals μ, ν given by the per-class feature maps μ^k, ν^k of Eq. (4), and scaling vectors $(\mathbf{u}, \mathbf{v})$ obtained by iterating:

$$\mathbf{u} \leftarrow \left(\frac{\mu}{\mathbf{K}\mathbf{v}} \right)^\tau, \quad \mathbf{v} \leftarrow \left(\frac{\nu}{\mathbf{K}^\top \mathbf{u}} \right)^\tau. \quad (8)$$

Here $\tau \in (0, 1)$ directly controls marginal relaxation; the corresponding KL divergence penalty weight on the marginal constraints is $\rho = \varepsilon\tau/(1-\tau)$ [25], with $\tau \rightarrow 1$ recovering balanced OT. We utilize the unbalanced formulation, as the CNN may predict features more prominently in one modality than the other due to image resolution and contrast discrepancies, so enforcing exact marginal constraints would penalize legitimate asymmetries in detection confidence.

Auxiliary losses. Lumen radius agreement is measured via normalized cross-correlation $\mathcal{L}_{\text{NCC}}$ between the predicted radii of the warped MPR and IVUS, providing an area-based alignment signal in pullback segments with sparse anatomical landmarks. Both $\mathcal{L}_{\text{NCC}}$ and image similarity terms are computed exclusively over regions without guidewire interference, as identified by the IVUS guidewire predictions of f_ξ . Additionally, the regularization term $\mathcal{L}_{\text{reg}}$ penalizes the squared gradient of local parameters $\{\theta_p, t_{u,p}, t_{v,p}\}$ to enforce spatial smoothness.

3 Experiments

Data. All experiments are conducted on a subset of the IMPACT study ($N=47$) from the Erasmus University Medical Center [13] for which both CCTA and IVUS data are available. For all IVUS pullbacks, a reference CCTA coronary artery centerline with tangent orientation was manually identified by matching

Table 2. Registration results across longitudinal initialization (z_{init}), rotational initialization (θ_{init}), and deformable refinement. Rows marked † are baselines representing standard registration approaches; remaining rows ablate components of the proposed method. The best prior-stage configuration is carried forward at each stage; best results are denoted in **bold**. Metrics are centerline overlap Dice_{ctl} , cosine similarity of cross-sectional normal vectors S_c , and lumen overlap Dice_L , reported as median (Q1, Q3) over 5-fold cross-validation. The top row (–) denotes the unregistered baseline.

	$\mathcal{L}_{\text{feat}}$	$\mathcal{L}_{\text{lumen}}$	$\mathcal{L}_{\text{image}}$	$\text{Dice}_{\text{ctl}} \uparrow$	$S_c \uparrow$	$\text{Dice}_L \uparrow$
–	–	–	–	0.67 (0.56, 0.81)	0.21 (–0.36, 0.55)	0.47 (0.40, 0.55)
z_{init}	$\text{Dice}^\dagger$	–	–	0.97 (0.87, 1.00)	0.16 (–0.54, 0.64)	0.57 (0.49, 0.62)
	$\text{Dice}^\dagger$	NCC	–	0.98 (0.93, 1.00)	0.16 (–0.56, 0.69)	0.58 (0.52, 0.62)
	Dice	NCC + L_1	–	0.99 (0.97, 1.00)	0.11 (–0.56, 0.65)	0.58 (0.54, 0.64)
θ_{init}	$\text{Dice}^\dagger$	–	–	0.99 (0.97, 1.00)	0.94 (0.75, 0.97)	0.63 (0.56, 0.68)
	OT	–	–	0.99 (0.97, 1.00)	0.94 (0.84, 0.98)	0.63 (0.56, 0.68)
	OT + Dice	–	–	0.99 (0.97, 1.00)	0.95 (0.87, 0.97)	0.63 (0.56, 0.68)
deform	$\text{Dice}^\dagger$	–	NMI	0.98 (0.96, 1.00)	0.95 (0.86, 0.97)	0.66 (0.59, 0.71)
	$\text{Dice}^\dagger$	NCC	–	0.99 (0.97, 1.00)	0.96 (0.90, 0.98)	0.68 (0.65, 0.74)
	$\text{Dice}^\dagger$	NCC	NMI	0.98 (0.96, 1.00)	0.95 (0.71, 0.97)	0.69 (0.64, 0.74)
	OT	NCC	–	0.99 (0.97, 1.00)	0.94 (0.80, 0.97)	0.68 (0.65, 0.74)
	OT + Dice	NCC	–	0.99 (0.97, 1.00)	0.96 (0.92, 0.98)	0.69 (0.64, 0.74)

large side branches visible in both modalities. Lumen segmentations were made available for both modalities, from which lumen radii were derived in the polar domain. Additionally, the presence of calcifications and bifurcations was manually annotated for each radial direction in the polar-transformed data, yielding binary labels on the (θ, z) cylinder for both MPR and IVUS data.

Evaluation. All experiments were evaluated in a 5-fold cross-validation setup, ensuring metrics are reported for patients unseen during classifier training. Since no directly comparable fully automatic CCTA-IVUS registration method exists, and direct comparison with semi-automatic approaches such as Kadry et al. [14] is precluded by segmentation label unavailability and modality mismatch, we evaluate against two natural baselines: the unregistered case, and a standard Dice feature matching loss in place of the proposed OT objective. We additionally ablate the lumen alignment term and evaluate NMI as an image similarity term $\mathcal{L}_{\text{image}}$ [21]. CNN training loss configurations ($\mathcal{L}_{\text{BCE}}$, $\mathcal{L}_{\text{Dice}}$, $\mathcal{L}_{\text{BCE}+\text{Dice}}$, $\mathcal{L}_{\text{BM}+\text{Dice}}$) are evaluated separately, reporting Dice and Betti matching error τ_{err} in dimension-0 [5,27] for bifurcation and calcification detection, and cIDice [26] for guidewire continuity. Registration accuracy is assessed by centerline overlap (Dice_{ctl}) following Schaap et al. [23], cosine similarity of cross-sectional normal vectors (S_c), and lumen overlap (Dice_L), reported as median (Q1, Q3).

We set $\alpha_{\text{OT}}=0.1$, $\alpha_{\text{lumen}}=1.0$, $\alpha_{\text{image}}=10^2$, $\alpha_{\text{reg}}=10^3$, $\varepsilon=0.1$, and $\tau=0.8$ for all registration experiments where the respective loss terms are included.




Fig. 3. Qualitative registration result for a representative case exhibiting catheter rotational drift. Cross-sectional normal vectors are color-coded by cosine similarity S_c . Columns show, from left to right: the extracted CCTA centerline (unregistered), the pre-deformable alignment, deformable optimization with Dice-only feature matching, deformable optimization with OT+Dice feature matching, and the reference centerline. The Dice-only model fails to recover rotational alignment throughout the pullback, as the rigid rotation initialized at θ_{init} captures only part of the catheter twist. The OT+Dice model successfully corrects this drift.

4 Results

Feature detection. Table 1 compares classification loss configurations. Lumen Dice is stable across configurations, as expected given the shared $\mathcal{L}_1$ regression objective. For bifurcations and calcifications, $\mathcal{L}_{\text{BM}+\text{Dice}}$ consistently yields low Betti matching errors across modalities while maintaining competitive Dice scores, indicating topologically coherent detections critical for downstream OT matching [17]. We therefore adopt $\mathcal{L}_{\text{BM}+\text{Dice}}$ for all registration experiments.

Registration. Table 2 reports registration accuracy across all conditions. Rotational initialization (θ_{init}) identifies the approximate global minimum of the rotational objective, placing subsequent deformable optimization in the correct basin of convergence. The Dice-only baseline achieves competitive performance when combined with NCC lumen matching (Dice_{ctl}: 0.99, S_c : 0.96, Dice_L: 0.68), while adding $\mathcal{L}_{\text{NMI}}$ does not improve results, consistent with limited contrast correspondence between MPR and IVUS. The proposed OT+Dice objective yields a statistically significant improvement in rotational alignment over the Dice-only baseline (Wilcoxon signed-rank $p < 0.01$), driven by more consistent performance across the cohort (S_c Q1: 0.92 vs. 0.90) and successful recovery of cases with severe catheter rotational drift where the Dice-only baseline fails (see Fig. 3). Optimization took 90 seconds per case on a single RTX 2080 Ti GPU.

5 Discussion and Conclusion

TG-OT achieves strong CCTA-IVUS registration performance without requiring any manual interaction or prior 3D image segmentation, outperforming Dice-only and NMI baselines in rotational accuracy—particularly for cases with severe catheter drift, where voxelwise losses lack the spatial gradients to recover alignment. Failure cases were observed in two categories: longitudinal initialization failed in $N=4$ cases due to incorrectly truncated automatic centerline extraction from CCTA, while rotational alignment degraded in $N=2$ cases with extreme landmark sparsity. Furthermore, because registration is driven by features detected on the resampled CCTA volume, its reliability is bounded by that of the feature detection step, for which foundation models may offer greater robustness at scale [6,24]. Nevertheless, all $N=47$ cases are included in the reported metrics. A broader challenge is the absence of automatic reference standards for CCTA-IVUS registration, making the constructed baselines the most meaningful available comparison. By removing the dependency on error-prone 3D intermediate segmentations, TG-OT addresses a key barrier to scalable multi-modal coronary analysis and brings fully automatic CCTA-IVUS fusion closer to clinical deployment.

Disclosure of Interests. Ivana Išgum has received research funding from Philips Healthcare, the Dutch Research Council, Pie Medical Imaging BV, Esaote Europe BV, Abbott Vascular Inc, Horizon Europe, and Innovative Health Initiative. José P. Henriques has received research funding from Health Holland, Infrarex Inc, B. Braun Medical Inc, and Nipro Medical Europe NV. All other authors declare no competing interests.

References

1. Arora, P., Singh, P., Girdhar, A., Vijayvergiya, R.: A state-of-the-art review on coronary artery border segmentation algorithms for intravascular ultrasound (IVUS) images. *Cardiovascular Engineering and Technology* **14**(2), 264–295 (2023)
2. Athanasiou, L., et al.: Three-dimensional reconstruction of coronary arteries and plaque morphology using CT angiography—comparison and registration with IVUS. *BMC Medical Imaging* **16**, 1–13 (2016)
3. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis* **12**(1), 26–41 (2008)
4. Berger, A.H., Lux, L., Stucki, N., Bürgin, V., Shit, S., Banaszak, A., Rueckert, D., Bauer, U., Paetzold, J.C.: Topologically faithful multi-class segmentation in medical images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 721–731. Springer (2024)
5. Berger, A.H., Lux, L., Weers, A., Menten, M.J., Rueckert, D., Paetzold, J.C.: Pitfalls of topology-aware image segmentation. In: *International Conference on Information Processing in Medical Imaging*. pp. 297–312. Springer (2025)

6. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025)
7. Cury, R.C., et al.: CAD-RADS™ 2.0–2022 coronary artery disease-reporting and data system: an expert consensus document of the society of cardiovascular computed tomography (SCCT), the American college of cardiology (ACC), the American college of radiology (ACR), and the North America society of cardiovascular imaging (NASCI). *Cardiovascular Imaging* **15**(11), 1974–2001 (2022)
8. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems* **26** (2013)
9. Eilers, P.H., Marx, B.D.: Flexible smoothing with B-splines and penalties. *Statistical Science* **11**(2), 89–121 (1996)
10. Van der Giessen, A.G., et al.: 3D fusion of intravascular ultrasound and coronary computed tomography for in-vivo wall shear stress analysis: a feasibility study. *The International Journal of Cardiovascular Imaging* **26**, 781–796 (2010)
11. Haddou, S.B., Van Herten, R.L.M., Bezzina, C.R., Planken, N., Daemen, J., Wentzel, J.J., Henriques, J.P., Isgum, I.: CyclePhase: Robust phase detection in cardiovascular imaging through cyclic motion estimation. In: *Medical Imaging with Deep Learning* (2026)
12. Hampe, N., et al.: Graph neural networks for automatic extraction and labeling of the coronary artery tree in CT angiography. *Journal of Medical Imaging* **11**(3), 034001–034001 (2024)
13. Hartman, E.M.J., et al.: Lipid-rich plaques detected by near-infrared spectroscopy are more frequently exposed to high shear stress. *Journal of Cardiovascular Translational Research* **14**, 416–425 (2021)
14. Kadry, K., et al.: Morphology-based non-rigid registration of coronary computed tomography and intravascular images through virtual catheter path optimization. *IEEE Transactions on Medical Imaging* **44**(2), 880 (2025)
15. Komilovich, E.B.: Coronary artery disease. *European Journal of Modern Medicine and Practice* **3**(12), 81–87 (2023)
16. Li, C., et al.: AutoFOX: An automated cross-modal 3D fusion framework of coronary X-ray angiography and OCT. *Medical Image Analysis* **101**, 103432 (2025)
17. Lux, L., Berger, A.H., Weers, A., Stucki, N., Rueckert, D., Bauer, U., Paetzold, J.: Topograph: An efficient graph-based framework for strictly topology preserving image segmentation. In: *International Conference on Learning Representations*. vol. 2025, pp. 80207–80231 (2025)
18. Marquering, H.A., et al.: Coronary CT angiography: IVUS image fusion for quantitative plaque and stenosis analyses. In: *Medical Imaging 2008: Visualization, Image-Guided Procedures, and Modeling*. vol. 6918, pp. 499–508. SPIE (2008)
19. Maurovich-Horvat, P., et al.: CT or invasive coronary angiography in stable chest pain. *New England Journal of Medicine* **386**(17), 1591–1602 (2022)
20. Mintz, G.S., Guagliumi, G.: Intravascular imaging in coronary artery disease. *The Lancet* **390**(10096), 793–809 (2017)
21. Pluim, J.P., Maintz, J.A., Viergever, M.A.: Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging* **22**(8), 986–1004 (2003)
22. Qian, Z., Vazquez, G., Voros, S.: Intermodal registration of CTA and IVUS-VH, and its application on CTA-based plaque composition analysis. In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. pp. 419–423. IEEE (2011)

23. Schaap, M., et al.: Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical Image Analysis* **13**(5), 701–714 (2009)
24. Scheinfeld, A., Zhang, H., Mu, S., Van Herten, R.L.M., Stoffl, L., Erturk, A., Wu, Z., Paetzold, J.C.: A multimodal 3D foundation model for light sheet fluorescence microscopy enables few-shot segmentation, classification, and deblurring. *arXiv preprint arXiv:2605.26026* (2026)
25. Séjourné, T., Peyré, G., Vialard, F.X.: Unbalanced optimal transport, from theory to numerics. *Handbook of Numerical Analysis* **24**, 407–471 (2023)
26. Shit, S., et al.: clDice - A novel topology-preserving loss function for tubular structure segmentation. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2021*. pp. 16555–16564. IEEE Computer Society (2021)
27. Stucki, N., Paetzold, J.C., Shit, S., Menze, B., Bauer, U.: Topologically faithful image segmentation via induced matching of persistence barcodes. In: *International Conference on Machine Learning*. pp. 32698–32727. PMLR (2023)
28. Van Herten, R.L.M., Lagogiannis, I., Leiner, T., Išgum, I.: The role of artificial intelligence in coronary CT angiography. *Netherlands Heart Journal* **32**(11), 417–425 (2024)
29. Van Herten, R.L.M., et al.: Automatic coronary artery plaque quantification and CAD-RADS prediction using mesh priors. *IEEE Transactions on Medical Imaging* **43**(4), 1272–1283 (2023)
30. Van Herten, R.L.M., Graf, R., Bitzer, F., Kirschke, J., Paetzold, J.C.: GeoReg: Direct biplanar DSA-to-CTA registration with geodesic consistency for acute ischemic stroke. In: *Medical Imaging with Deep Learning* (2026)