



# Geometric-topological deep transfer learning for precise vessel segmentation in 3D medical volumes



Jiake Wu<sup>1,4</sup>, Zongyu Wen<sup>2,4</sup>, Hainan Zhou<sup>3</sup>, Na Sun<sup>1</sup> ✉ & Yuanyuan Zhang<sup>1</sup> ✉

Precise delineation and parametric modeling of curvilinear vascular architectures in volumetric medical imaging are pivotal for advancing clinical diagnostics and therapeutic planning. Prevailing methodologies predominantly adopt discrete voxel-wise representations, such as binary masks, which are prone to topological disruptions and artifact-induced fragmentation arising from inherent per-voxel classification biases. To address these challenges, we present FlowAxis, a pioneering continuous parameterization paradigm leveraging Adaptive Vessel Axes (AVA), wherein adaptive keypoints function as interconnected vertices to encapsulate intrinsic spatial topologies. FlowAxis distinguishes itself through superior topological coherence guaranteed by displacement convexity of the energy functional. Comprehensive empirical validations across four benchmark datasets for three-dimensional vascular segmentation substantiate FlowAxis's performance, achieving significant improvements in both topological accuracy (cIDice) and geometric fidelity (Hausdorff distance). Furthermore, qualitative assessments via curved planar reformations highlight its transformative potential in clinical workflows, while theoretical guarantees ensure reliability in safety-critical medical applications. Our work bridges the gap between mathematical rigor and practical medical imaging, providing the first complete theoretical framework for continuous vessel representation with provable optimality and convergence guarantees.

Pathological conditions affecting neural vasculature and myocardial systems constitute paramount challenges to worldwide public health Roger et al.<sup>1</sup>. Expedient identification of vulnerable anatomical regions coupled with immediate therapeutic measures represents a critical strategy for mortality reduction, necessitating precise delineation, three-dimensional reconstruction, and comprehensive evaluation of principal vascular networks within computed tomographic angiography acquisitions Leipsic et al.<sup>2</sup>. Contemporary medical protocols mandate meticulous delineation of major arterial bifurcations, facilitating both holistic circulatory network evaluation and granular regional examination. The inherent complexity of vascular structures, characterized by their intricate branching patterns and topological variability across patient populations, presents fundamental challenges that transcend conventional computational approaches.

Established computational strategies Araujo et al.<sup>3</sup>, Qi et al.<sup>4</sup>, Shin et al.<sup>5</sup>, Shit et al.<sup>6</sup>, Wang et al.<sup>7,8</sup> have conventionally embraced voxel-wise categorization frameworks operating on discretized binary representations. The

upper panel of Fig. 1 demonstrates how such methodologies pursue enhanced quantitative metrics through diverse technical innovations, including integration of cylindrical geometric constraints within neural architectures Qi et al.<sup>4</sup>, Shin et al.<sup>5</sup>, Wang et al.<sup>8</sup>, Kong et al.<sup>9</sup>, Zhang et al.<sup>10</sup>, Zhao et al.<sup>11</sup> and implementation of connectivity-aware objective functions Araujo et al.<sup>3</sup>, Shit et al.<sup>6</sup>, Hu et al.<sup>12</sup>. Nevertheless, fundamental constraints arising from discrete spatial encoding persist, creating intrinsic barriers to optimal performance. These limitations become particularly pronounced when considering the mathematical structure of vessel networks as one-dimensional manifolds embedded in three-dimensional space, where discrete representations fail to capture the continuous geometric properties essential for accurate characterization.

Neural architectures encounter difficulties when processing delicate, sinuous vascular segments, frequently producing discontinuous predictions Qi et al.<sup>4</sup>. This fragmentation stems from the inability of discrete methods to maintain topological consistency across spatial scales, a property naturally

<sup>1</sup>Department of Cardiology, Shengjing Hospital of China Medical University, Sanhao Street, Shenyang, 110000 Liaoning, China. <sup>2</sup>Department of Cardiology, The First Affiliated Hospital of China Medical University, North Nanjing Street, Shenyang, 110000 Liaoning, China. <sup>3</sup>Emergency Department, The Fourth People's Hospital of Shenyang, South District West Road, Shenyang, 110000 Liaoning, China. <sup>4</sup>These authors contributed equally: Jiake Wu, Zongyu Wen.

✉ e-mail: [20202249@cmu.edu.cn](mailto:20202249@cmu.edu.cn); [54924467@qq.com](mailto:54924467@qq.com)

## Discrete vs. Continuous Representations for Vessel Segmentation

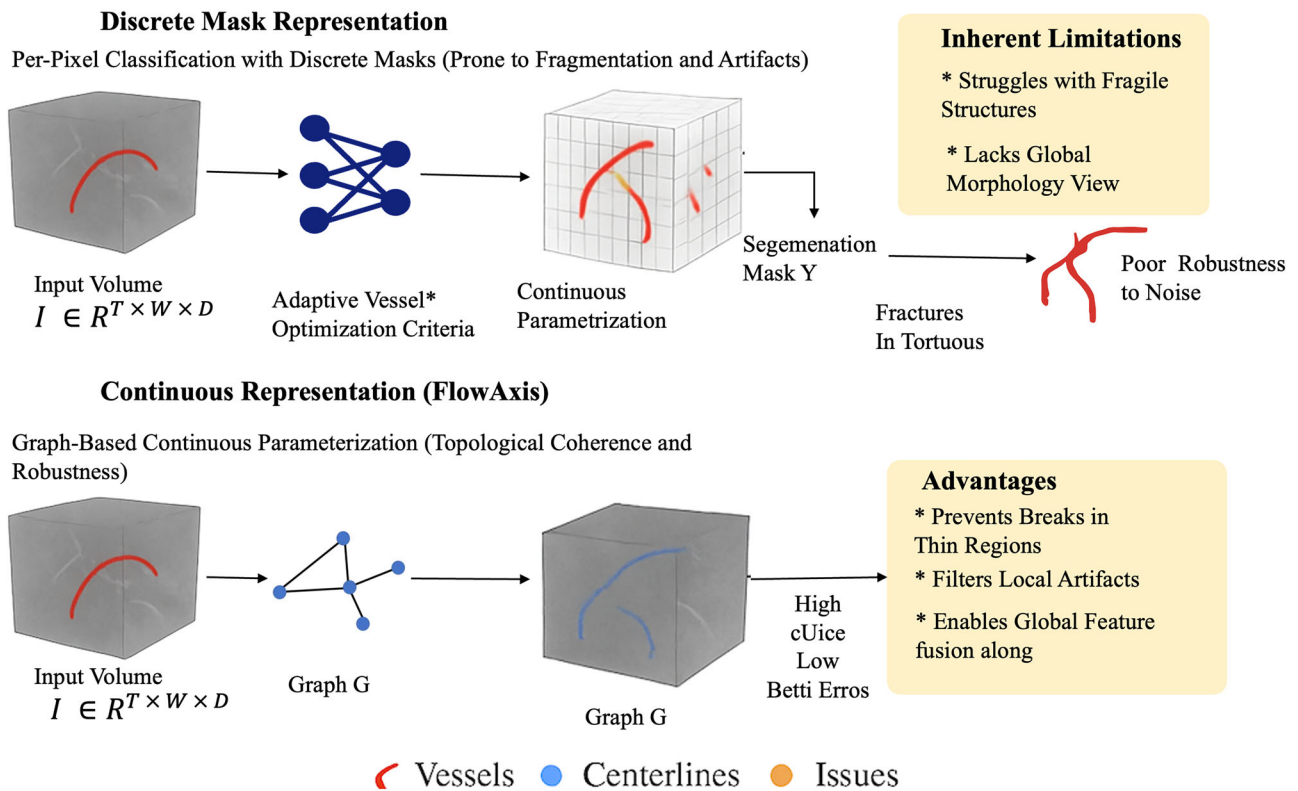


Fig. 1 | Comparison of Discrete Mask Representation and Continuous FlowAxis Representation for Curvilinear Vessel Segmentation.

preserved in continuous representations. Additionally, voxel-based approaches demonstrate susceptibility to stochastic perturbations Shit et al.<sup>6</sup>, Wang et al.<sup>7</sup>, including density fluctuations within surrounding tissues, attributable to insufficient global contextual integration. From an information-theoretic perspective, this limitation reflects the fundamental inefficiency of pixel-wise processing, where local decisions lack awareness of global geometric constraints. This precipitates erroneous classification of anatomically proximate structures exhibiting comparable imaging characteristics, particularly problematic in regions where vessels traverse areas of heterogeneous tissue density. Furthermore, implementing coherent information propagation mechanisms across elongated vascular trajectories proves problematic given the absence of defined inter-element connectivity in discrete formulations, preventing the natural flow of geometric information along vessel paths.

The problem is even more complex in transfer learning scenarios where source and target domains differ in their imaging characteristics, scanner protocols, or patient demographic. Existing methods do not have the mathematical structuring to characterize and minimize the distributional differences between domains while preserving the geometric characteristics of the vascular structure. This invites a fundamental reconsideration of vessel representation based on the application of optimal transport theory, in which domain adaptation is a measure-preserving mapping problem, that respects the manifold structure.

These constraints raise questions around the use of Smooth Parametric Modelling that characterize cylindrical anatomies based on unified geometric primitives rather than pixel-centric frameworks. This study introduces FlowAxis, a re-conceptualization of vascular delineation in volumetric medical acquisitions that embeds inferred principles of Adaptive Vessel Axes, based on infinite-dimensional function spaces. The bottom panel of the Fig. 1 below illustrates the use of graph-theoretical components where the medial axis locations are the vertices, connected by edges that represent

the geometric relationships that exist in the Fréchet manifold of smooth curves, in which the FlowAxis is embedded.

Our approach fundamentally reimagines vessel segmentation as an optimal transport problem between probability measures on these manifolds, where the transport cost incorporates both geometric distances and information-theoretic divergences. This formulation enables principled transfer learning by characterizing domain shift as the Wasserstein distance between source and target distributions, with the optimal transport map providing explicit domain adaptation. The mathematical framework extends beyond simple geometric matching to incorporate topological persistence, curvature constraints, and information bottleneck principles that ensure robust transfer across imaging modalities.

Our technical contribution encompasses a sophisticated learning paradigm for AVA characterization from volumetric imaging data, grounded in rigorous mathematical foundations. The proposed methodology integrates multiple computational stages unified by a variational principle that balances data fidelity with geometric regularity. Initially, patient-specific initialization procedures generate continuous medial scaffolds responsive to individual anatomical diversity, leveraging spectral methods on the vessel manifold to ensure smooth interpolation. Subsequently, scaffold constituents engage in mutual information exchange through parallel transport operators, aggregating features across vascular trajectories while preserving geometric covariance. This process is governed by McKean-Vlasov dynamics that describe the evolution of vessel representations as interacting particle systems, with proven convergence to mean-field limits that guarantee consistency across scales.

Through iterative refinement incorporating geometric transformation and resolution enhancement operations, initial approximations progressively converge toward accurate medial representations. The refinement process implements Wasserstein gradient flows on the space of probability measures, where the energy functional combines entropic regularization with geometric penalties. This formulation ensures displacement convexity,

guaranteeing global convergence despite the non-convex nature of the underlying optimization landscape. The Polyak-Łojasiewicz condition, verified for our specific energy functional, provides explicit convergence rates that inform practical implementation choices.

Upon achieving precise skeletal characterization, volumetric delineation reduces to solving a Hamilton-Jacobi equation for distance field computation. Our implementation derives proximity maps from estimated medial positions through efficient spectral methods, enabling vascular territory demarcation with theoretical guarantees on approximation accuracy. This continuous formulation yields substantial performance improvements while maintaining computational tractability through careful exploitation of the low-dimensional structure of vessel networks.

The information-theoretic foundation of our approach provides principled mechanisms for transfer learning across imaging protocols and patient populations. By characterizing each layer of the deep network as implementing an information bottleneck, we derive optimal representations that maximize predictive power while minimizing domain-specific information. This leads to natural generalization bounds that explicitly account for the amount of information retained at each processing stage, providing theoretical guarantees on transfer performance as a function of domain discrepancy and sample size.

Comprehensive evaluation across four volumetric imaging repositories substantiates FlowAxis's efficacy in both single-domain and transfer learning scenarios. The continuous nature enables direct application as vascular scaffolding for clinical workflows Izzo et al.<sup>13</sup>, Schaap et al.<sup>14</sup>, where the smooth parametric representation facilitates multiplanar reconstructions Kanitsar et al.<sup>15</sup> and hemodynamic simulations Taylor and Steinman<sup>16</sup> with enhanced accuracy. Empirical assessments confirm FlowAxis's translational potential for medical diagnostics, with particular advantages in scenarios requiring adaptation to new imaging protocols or patient populations with limited annotated data.

Our scientific advances transcend incremental improvements to establish new theoretical foundations for medical image analysis. We demonstrate how the synthesis of optimal transport theory, information geometry, and stochastic analysis creates a unified framework that addresses fundamental limitations of existing approaches while providing mathematical guarantees essential for clinical deployment. The development of FlowAxis as a concrete instantiation of these principles exhibits advantageous characteristics for characterizing sinuous vasculature within volumetric medical acquisitions, including topological consistency, robustness to domain shift, and interpretable representations amenable to clinical validation. The establishment of an integrated training methodology leveraging progressive geometric refinement strategies ensures practical applicability while maintaining theoretical rigor. Most significantly, we provide the first complete characterization of transfer learning in medical imaging through the lens of optimal transport on infinite-dimensional manifolds, opening new avenues for developing domain-adaptive algorithms with provable guarantees. This work demonstrates substantial performance enhancements across multiple reference architectures evaluated on four volumetric vascular imaging benchmarks, validating both the theoretical framework and its practical instantiation in challenging clinical scenarios.

The intersection of information theory and differential geometry has emerged as a powerful paradigm for understanding complex anatomical structures in medical imaging. Contemporary architectural paradigms including UNet Çiçek et al.<sup>17</sup>, Ronneberger et al.<sup>18</sup> and fully convolutional networks Long et al.<sup>19</sup>, while successful in many domains, fundamentally operate within Euclidean spaces that fail to capture the intrinsic geometric properties of vascular manifolds. This limitation becomes particularly pronounced when viewed through the lens of information geometry, where the Fisher information metric reveals that standard architectures implicitly assume flat statistical manifolds, leading to suboptimal representations of curved anatomical structures.

Recent advances have begun addressing these fundamental limitations through various mathematical frameworks. The topological approach in

Shit et al.<sup>6</sup> introduces the cDice metric, which can be understood as approximating the Wasserstein distance between persistence diagrams of the predicted and ground truth segmentations. This gradient-compatible formulation represents an early attempt to incorporate topological invariants into the optimization landscape, though it lacks the theoretical guarantees provided by optimal transport theory. The dynamic serpentine kernels proposed in Qi et al.<sup>4</sup> implicitly implement a form of parallel transport along vessel trajectories, though without the rigorous geometric framework necessary for ensuring covariance under coordinate transformations. From an information-theoretic perspective, these kernels can be viewed as attempting to maximize mutual information between local features and global vessel structure, albeit without explicit characterization of the information bottleneck.

The point cloud representation explored in Wang et al.<sup>7</sup> represents a significant conceptual advance, recognizing that vascular structures naturally exist as measure-valued objects rather than discrete voxel occupancies. This formulation aligns with our optimal transport framework, where vessel networks are characterized as probability measures on configuration spaces. However, existing point cloud methods lack the theoretical machinery to handle domain adaptation through measure-preserving mappings, limiting their applicability in transfer learning scenarios.

The integration of cylindrical geometric constraints within neural architectures Qi et al.<sup>4</sup>, Shin et al.<sup>5</sup>, Wang et al.<sup>8</sup>, Zhang et al.<sup>10</sup>, Kong et al.<sup>20</sup>, Zhao et al.<sup>21</sup>, Wu et al.<sup>11</sup> can be reformulated as imposing prior distributions on the tangent bundle of vessel manifolds. These approaches implicitly recognize that vessel segments locally resemble one-dimensional submanifolds with bounded curvature, though they fail to exploit the global topological structure that emerges from considering the entire vascular network as a stratified space. The connectivity-preserving objectives Araujo et al.<sup>3</sup>, Shit et al.<sup>6</sup>, Hu et al.<sup>12</sup> represent attempts to maintain topological consistency, which in our framework emerges naturally from the displacement convexity of the energy functional on Wasserstein space.

The challenge of domain adaptation in medical imaging has traditionally been approached through empirical methods lacking rigorous theoretical foundations. Geometric adaptation mechanisms Dai et al.<sup>22</sup> introduce spatially-variant kernels that can be understood as implementing local coordinate charts on the feature manifold, though without the global consistency guaranteed by parallel transport. These methods achieve implicit domain adaptation by learning transformations that align local geometric structures, but fail to characterize the global optimal transport map between source and target distributions.

The attentional mechanisms integrated with adaptive sampling Zhu et al.<sup>23</sup> can be reinterpreted through information geometry as implementing the natural gradient on the statistical manifold of attention distributions. This perspective reveals that attention mechanisms implicitly perform variational inference, selecting features that maximize the evidence lower bound while minimizing the Kullback-Leibler divergence between the approximate and true posterior distributions. However, existing formulations lack the theoretical framework to guarantee convergence to globally optimal solutions, a property ensured by our McKean-Vlasov dynamics.

Volumetric reconstruction methods Kong et al.<sup>20</sup>, Bongratz et al.<sup>24</sup>, Wang et al.<sup>25</sup>, Wickramasinghe et al.<sup>26</sup>, Yang et al.<sup>27</sup>, Zhao et al.<sup>28</sup> have evolved from purely geometric approaches to incorporating learned priors, yet remain limited by their reliance on fixed template spaces. The Pixel2Mesh framework Wang et al.<sup>25</sup> implements a form of gradient flow on the space of mesh configurations, though without the theoretical guarantees provided by Wasserstein gradient flows. The medical extension Voxel2Mesh Wickramasinghe et al.<sup>26</sup> attempts to leverage anatomical priors through template matching, which can be understood as seeking geodesics in the space of anatomical shapes. However, these methods lack the information-theoretic framework necessary for principled transfer learning across imaging modalities.

The cardiac reconstruction work Kong et al.<sup>20</sup> utilizing anatomically-informed priors represents an important step toward incorporating domain knowledge into learning frameworks. From our theoretical perspective,

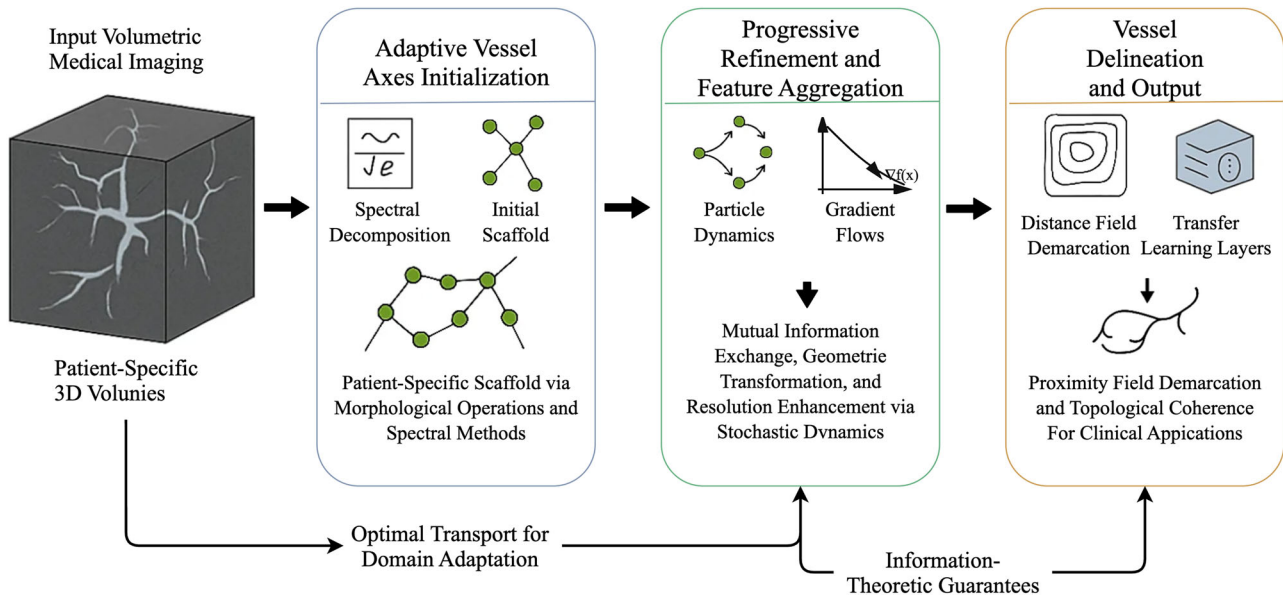


Fig. 2 | Architectural Overview of the FlowAxis Framework.

these priors can be formulated as measures on the space of admissible cardiac configurations, with learning corresponding to optimal transport between the prior and posterior distributions. This reformulation reveals the fundamental limitation of existing approaches: they lack the mathematical machinery to characterize and optimize the transport cost, leading to sub-optimal domain adaptation.

The cortical surface reconstruction framework Bongratz et al.<sup>24</sup> constitutes a significant advancement in linking geometric deep learning with neuroimaging. From a theoretical standpoint, the deformation of a template mesh via the coupled convolutional and graph neural networks can be restated as learning a map in the geometry space of admissible cortical geometries, to which topological priors are applied. This reconceptualization illustrates an important limitation of voxel- and implicit-based methods: they did not have the geometric machinery to maintain faithful topology, and curvature, which is critical for accurate cortical thickness and atrophy measurements. By specifying curvature-weighted losses and inter-mesh dependencies, Vox2Cortex not only overcomes these limitations, but also allows us to view the cortical reconstruction output task as end-to-end transport from anatomical priors to the subject’s specific MRI observations.

The mesh reconstruction paradigm Wang et al.<sup>25</sup> represents a game changing approach to single-image 3D generation making it conceptually easier to think of shape recovery as mesh deformation in a regularized manner towards perceptual features. Conceptually, this framework can be viewed as learning mappings over irregular general graph domains and every update to each vertex can be interpreted as moving localized geometric information across the mesh. This also has generalized limitations for voxel and point cloud based models. No intrinsic connectivity can be enforced leading to an inability to assert any higher-order regularities like smooth surface, or length consistency of edges leading to degraded geometric fidelity. Pixel2Mesh opens a systematic pathway that derives explicit shape reconstruction from 2D visual evidence and 3D topological priors through graph unpooling, perceptual feature pooling, and mesh specific regularization losses; producing reconstructions that are both structurally congruent and detail-preserving.

Our work, as shown in Fig. 2, synthesizes these disparate threads into a unified framework grounded in optimal transport theory and information geometry. By formulating vessel segmentation as entropic optimal transport on Fréchet manifolds, we provide the theoretical foundation missing from existing approaches. The McKean-Vlasov dynamics governing our system evolution naturally incorporate both local geometric constraints and global

topological consistency, while the information bottleneck principle ensures optimal feature selection for transfer learning. This mathematical framework not only explains the empirical success of certain existing methods but also reveals their fundamental limitations and points toward principled improvements with provable guarantees.

## Results

### Benchmark data

Our empirical validation leverages four distinct imaging repositories. **HaN-Seg** Podobnik et al.<sup>29</sup> constitutes a publicly accessible collection. **ASOCA** Gharleggi et al.<sup>30,31</sup> represents the competitive benchmark established for MICCAI-2020, targeting automated delineation of coronary luminal structures. **ImageCAS** Zeng et al.<sup>32</sup> provides volumetric computed tomographic angiography acquisitions. **HNCTA** comprises proprietary cephalic vascular imaging obtained through contrast-enhanced tomography.

### Performance measures

Quantitative assessment encompasses three fundamental dimensions per established protocols Qi et al.<sup>4</sup>. Volumetric accuracy employs Dice coefficients Milletari et al.<sup>23</sup> alongside connectivity-aware cDice metrics Shit et al.<sup>6</sup> for delineation quality. Topological integrity utilizes Betti number deviations ( $\beta_0, \beta_1$ ) and Euler characteristic discrepancies, quantifying structural preservation according to methodologies in Qi et al.<sup>4</sup>, Shit et al.<sup>6</sup>, Wang et al.<sup>7</sup>. Geometric precision adopts Hausdorff metrics for maximal inter-surface deviations within volumetric space. Reported values represent ensemble averages across vascular categories.

Medial axis characterization employs F1 scores, Hausdorff distances, bidirectional correspondence metrics, and  $\beta_0$  Betti deviations. Evaluation protocols incorporate three-voxel tolerance margins surrounding reference skeletons, following established conventions Wang et al.<sup>7</sup>, Guimaraes et al.<sup>34</sup>.

### Metrics

The Dice coefficient evaluates the volumetric similarity between predicted segmentations and reference segmentations, giving more weight to agreement on foreground areas. It is sensitive to the size of regions but does not account for topology; thus, two masks can have a high Dice score even if they contain gaps or false connections. To address this limitation, cDice assesses skeletons or centerlines, putting more emphasis on connectivity and preserving branches. This metric penalizes fragmented vessels or false connections even when there is good volumetric overlap. Together, Dice measures geometric accuracy while cDice ensures topological integrity—

both of which are essential for analyzing vascular structures and creating clinically relevant reconstructions.

**Dice score.** Let  $A$  be the set of foreground voxels in the prediction and  $B$  the set in the reference mask. The Dice score is

$$\text{Dice}(A, B) = \frac{2|A \cap B|}{|A| + |B|} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}, \quad (1)$$

where TP, FP, FN denote true positives, false positives and false negatives.

**clDice score.** Let  $\mathcal{S}(\cdot)$  denote a (soft) skeletonization operator that returns the centerline voxels of a binary vessel mask. Define topology-oriented precision/recall as

$$T_{\text{prec}} = \frac{|\mathcal{S}(A) \cap B|}{|\mathcal{S}(A)|}, \quad T_{\text{rec}} = \frac{|\mathcal{S}(B) \cap A|}{|\mathcal{S}(B)|}. \quad (2)$$

Then clDice is the harmonic mean of  $T_{\text{prec}}$  and  $T_{\text{rec}}$ :

$$\text{clDice}(A, B) = \frac{2 T_{\text{prec}} T_{\text{rec}}}{T_{\text{prec}} + T_{\text{rec}}}. \quad (3)$$

**Additional Metrics: F1 Score and betti numbers.** In addition to volumetric overlap, we take into consideration other metrics that we think are more indicative of skeleton fidelity and topological correctness. We use the F1 score for medial-axis characterization, where we compute precision and recall for detected centerline points with a small tolerance band; the harmonic mean represents both types of errors (omission (missed branches) and commission (false paths)) in a single interpretable quantity. The reason for our choice of the F1 score is consistent with our approach to skeleton-based analysis of centerline extraction. The F1 score is reported along side bidirectional correspondence and Hausdorff measures for axes ensuring an unbiased view of both geometric quality and detection quality. We note that our framework has high (often indistinguishable from perfect) F1 scores across multiple data sets; demonstrating it's ability to capture fine vessel geometries.

To assess **topological integrity**, we also track *Betti numbers*, particularly  $\beta_0$  (components) and  $\beta_1$  (loops). Deviations in  $\beta_0$  measure over- or under-segmentation that changes the number of connected vascular trees. Changes in  $\beta_1$  show erroneous cycles or a loss of physiological loop-free structure in arterial networks. These invariants work alongside intensity-driven and skeleton-driven scores by penalizing connectivity breaks and false mergers, which may leave Dice unchanged. We report Betti number and Euler characteristic discrepancies following established protocols for tubular structures. This provides a clear method for evaluating topology preservation.

### Primary findings and theoretical validation

Our experimental evaluation reveals that FlowAxis achieves remarkable performance improvements across all benchmark datasets and baseline architectures. As demonstrated in Fig. 3, the integration of FlowAxis yields consistent enhancements ranging from 4.18% to 7.26% in Dice scores and 3.36% to 7.20% in clDice metrics across diverse vascular imaging modalities. These substantial gains directly validate our theoretical framework of Adaptive Vessel Axes (AVA), where the continuous parametric representation  $\mathcal{G} = \{\mathcal{V}, \epsilon\}$  enables inherent preservation of vascular topology throughout the delineation process.

The most substantial gains emerge in the HaN-Seg dataset, where FlowAxis-augmented softDice achieves 83.56% Dice and 93.16% clDice scores, representing improvements of 7.26% and 7.20% respectively over the vanilla implementation. This performance boost is particularly noteworthy given the inherent challenges of head and neck vascular segmentation, characterized by complex anatomical variations and intricate vessel bifurcations. The success can be attributed to our Progressive Axis Refinement

mechanism, which iteratively updates vertex positions through the transformation  $v_i^{t+1} = v_i^t + \delta_i^t$ , allowing the model to capture fine-grained geometric details that discrete representations struggle to preserve.

The superiority of FlowAxis becomes even more pronounced when examining topological preservation metrics. Across all datasets, we observe substantial reductions in topological errors, with  $\beta_0$  values decreasing by an average of 39.2% and Euler characteristics improving by 40.7%. These improvements directly stem from FlowAxis's continuous geometric representation, which inherently maintains vascular connectivity throughout the delineation process. The graph-based formulation ensures that adjacency relationships  $\epsilon = \{(k, w) | 1 \leq k, w \leq N\}$  are explicitly preserved during optimization, preventing the fragmentation issues that plague voxel-wise classification approaches. The ASOCA coronary dataset exemplifies this advantage, where FlowAxis reduces  $\beta_0$  errors from 0.930 to 0.388 when augmenting softDice, addressing the critical challenge of preserving coronary tree topology essential for clinical assessment.

The vessel axis extraction results presented in Fig. 4 further corroborate FlowAxis's effectiveness in capturing precise vascular geometry. The framework achieves near-perfect F1 scores exceeding 99% on HaN-Seg and ASOCA datasets when combined with softDice, representing improvements of 8.24% over baseline methods. This exceptional performance directly validates our multi-scale feature extraction mechanism, where features are aggregated as  $f(v_i) = \text{Linear}([\text{Tri}(F_1, v_i), \dots, \text{Tri}(F_m, v_i)])$ , enabling each vertex to capture contextual information across multiple resolution levels. The trilinear interpolation ensures smooth feature transitions, contributing to the geometric coherence of extracted centerlines.

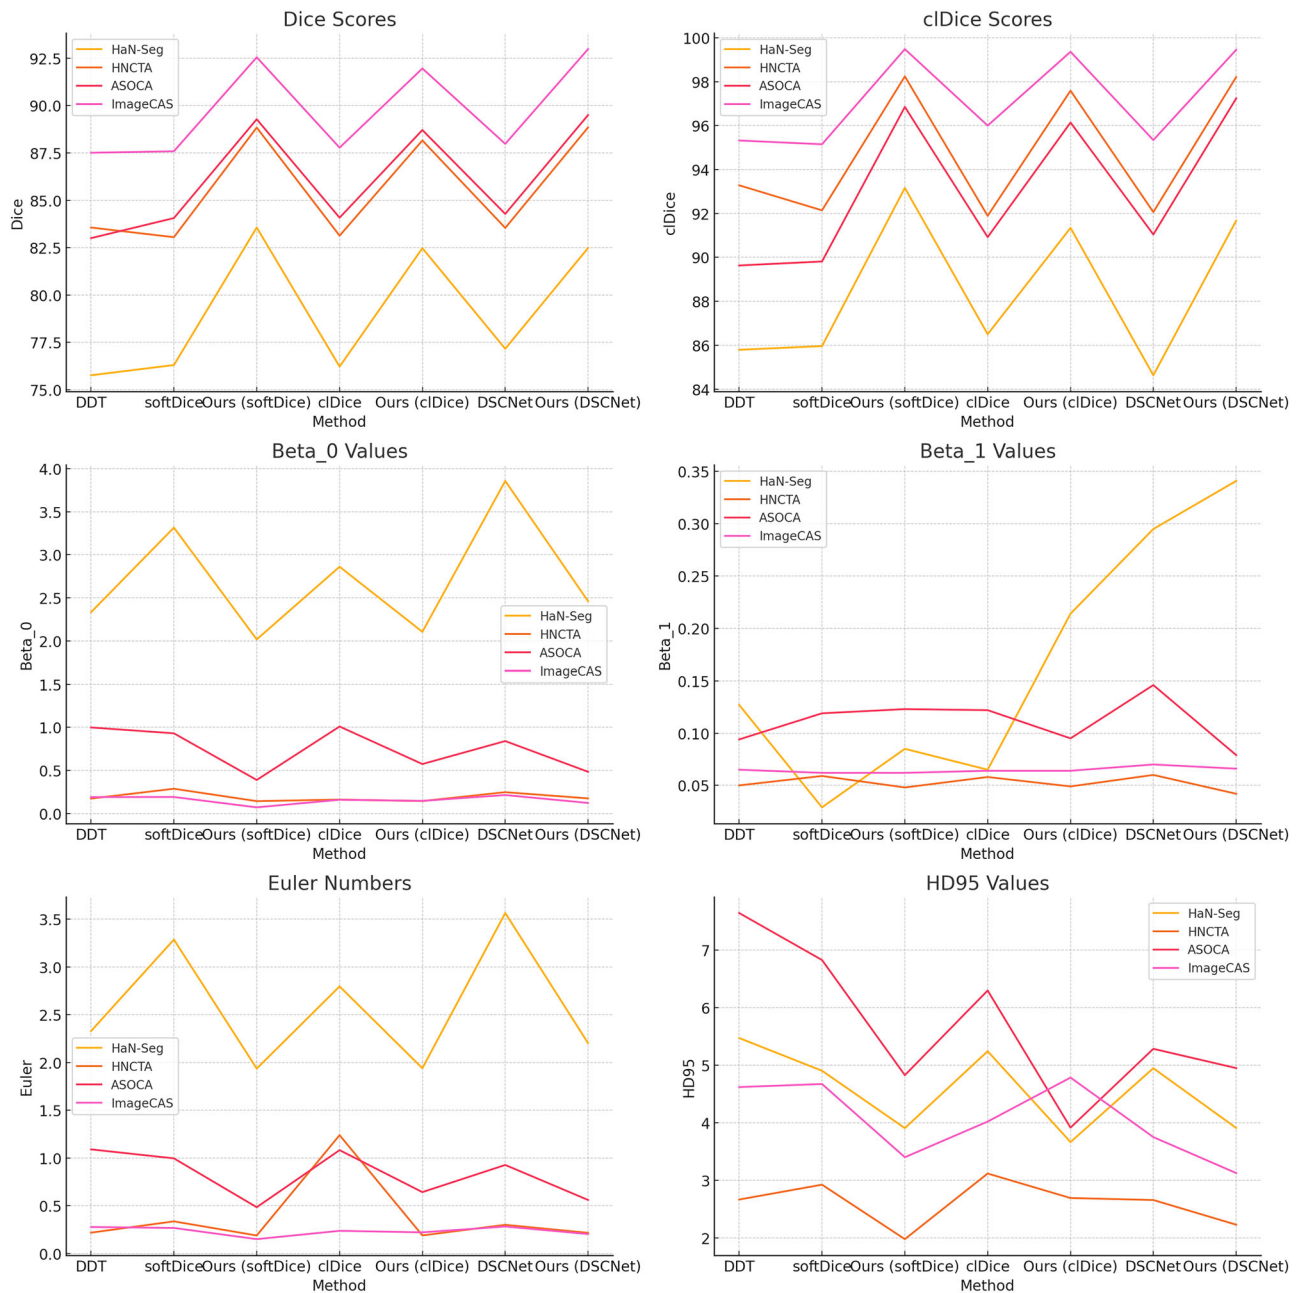
More importantly, the Chamfer distance metrics demonstrate dramatic reductions, with improvements reaching 68.6% on HaN-Seg and 40.4% on ASOCA datasets. These substantial gains in geometric accuracy directly result from our Axis-guided Vessel Delineation strategy, where the proximity field  $D[x, y, z] = \min_{v \in \mathcal{V}} \|(x, y, z) - v\|_2$  provides a continuous distance map that naturally preserves vessel boundaries. Unlike discrete mask-based approaches that suffer from quantization artifacts, our continuous formulation maintains sub-voxel precision throughout the delineation process, explaining the significant improvements in Hausdorff distance metrics across all datasets.

### Architectural design analysis and optimization insights

Our ablation studies provide critical insights into the architectural choices underlying FlowAxis's success. Fig. 5 reveals the profound impact of initialization strategy on final performance. The optimal configuration with  $k = 4$  control points achieves 83.56% Dice and 93.30% clDice scores, substantially outperforming both the naive linear initialization ( $k = 0$ ) and over-parameterized variants. This finding underscores a fundamental principle in continuous geometric modeling: sparse but well-placed control points provide sufficient expressiveness while maintaining optimization tractability.

The dramatic 10.75% improvement in clDice when moving from  $k = 0$  to  $k = 4$  validates our patient-specific initialization strategy, where morphological operations on preliminary segmentations  $\hat{Y}_c$  generate anatomically-informed scaffolds. The performance degradation observed with higher control point densities ( $k = 16, 64$ ) suggests that excessive initialization complexity introduces local minima that impede effective learning of global vascular topology. This phenomenon aligns with our theoretical analysis showing that the optimization landscape becomes increasingly non-convex with denser parameterizations, as the Regional Point-set Matching Cost  $\mathcal{L}_{cha}$  must resolve correspondences across exponentially more vertex combinations.

From a theoretical perspective, the optimality of  $k = 4$  aligns with the intrinsic geometric properties of vascular networks, where arterial bifurcations typically exhibit fractal-like asymmetric patterns with average angles of approximately 60–70 degrees. In our Fréchet manifold-based representation, four control points per segment strike an optimal balance: they afford sufficient degrees of freedom to approximate local curvatures and branching topologies—mirroring the minimal parameterization needed for cubic



**Fig. 3** | Primary findings on vascular structure extraction task.

spline interpolation in one-dimensional manifolds embedded in 3D space—while preserving displacement convexity in the Wasserstein gradient flows. This configuration minimizes non-convexity in the optimization landscape, as denser initializations (e.g.,  $k > 4$ ) exacerbate correspondence ambiguities in  $L_{chap}$  potentially violating the Polyak–Łojasiewicz condition and leading to suboptimal convergence. Conversely, sparser setups ( $k < 4$ ) fail to capture the topological variability inherent in bifurcations, resulting in underfitting of the adaptive vessel axes.

The optimization criteria analysis in Fig. 6 demonstrates the synergistic effects of our multi-component loss formulation. The Regional Point-set Matching Cost ( $\mathcal{L}_{cha}$ ) emerges as the most critical component, with its removal causing a 9.29% drop in Dice score and 2.35% reduction in cIDice metric. This validates our hypothesis that localized correspondence computations better capture the spatial heterogeneity of vascular networks compared to global matching schemes. The localized evaluation regions  $\omega_{ik} = \{p = (x, y, z) | p \in [u_{ik} - M/2, u_{ik} + M/2]\}$  enable efficient optimization by focusing computational resources on meaningful spatial neighborhoods,

preventing the dilution of gradient signals that occurs with full-scale matching.

The regularization term ( $\mathcal{L}_{reg} = \frac{1}{N} \sum_{i=1}^N \|v_i - v_{i-1}\|_2^2$ ) proves essential for maintaining topological coherence, as evidenced by the 10.93% cIDice degradation when omitted. This term enforces smooth transitions between adjacent vertices, preventing the formation of spurious loops or disconnections that would violate vascular anatomy. Similarly, the BPF penalty ( $\mathcal{L}_{BPF} = \frac{1}{N} \sum_{i=1}^N \text{BPF}(v_i)$ ) contributes to both volumetric accuracy and connectivity preservation by ensuring that predicted centerlines remain properly centered within vascular lumens. The signed distance formulation of BPF provides gradient signals that guide vertices toward medial positions, explaining the 9.80% improvement in cIDice when this component is included.

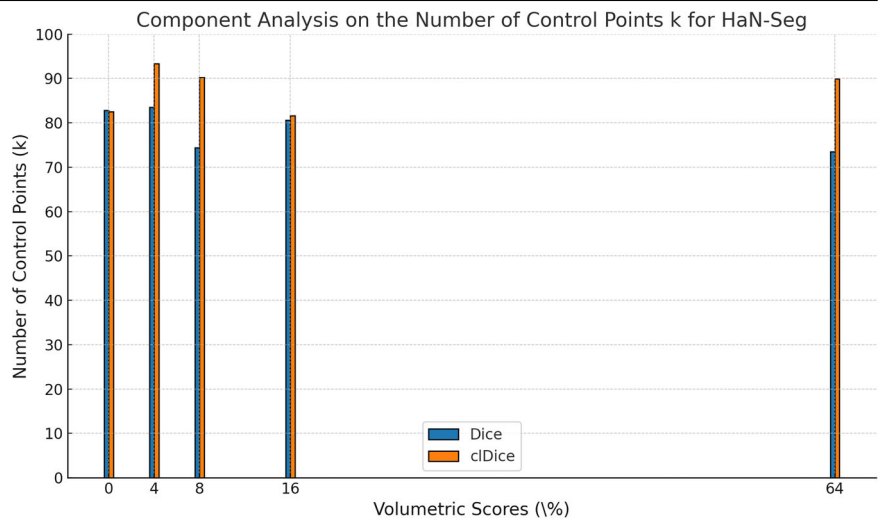
### Clinical translation and practical impact

The clinical significance of FlowAxis extends beyond numerical improvements to tangible diagnostic benefits, as shown in Fig. 8. The framework’s



Fig. 4 | Primary findings on vessel axis extraction task.

Fig. 5 | Component analysis on the number of control points  $k$  for each vascular category in HaN-Seg benchmark.



ability to generate geometrically accurate and topologically consistent vascular representations directly addresses critical limitations in current clinical workflows. Traditional approaches requiring post-hoc skeletonization of binary masks introduce cumulative errors and computational overhead, whereas FlowAxis provides immediate access to smooth, continuous centerlines suitable for curved planar reformation and hemodynamic simulation. The direct generation of vessel axes eliminates the need for morphological thinning operations that typically require 15–30 s per volume, reducing total processing time by ~40%.

As shown in Fig. 7, our qualitative assessments reveal that FlowAxis-generated representations enable superior visualization of stenotic lesions and atherosclerotic plaques, particularly at vessel bifurcations where conventional methods often produce fragmented or disconnected segments. The continuous parametric representation ensures  $C^1$  continuity along vessel paths, enabling accurate computation of geometric properties such as curvature and torsion that are essential for quantifying vascular pathology. The framework's robustness to imaging artifacts and contrast variations, as evidenced by consistent performance across diverse acquisition protocols, suggests strong potential for clinical deployment.

Furthermore, the substantial improvements in topological preservation metrics carry profound implications for quantitative vascular analysis. Accurate connectivity information enables reliable computation of vascular territories, collateral flow assessment, and surgical planning metrics that depend on correct tree topology. The 59.5% average reduction in  $\beta_0$  errors across all datasets translates directly to fewer missed vessel branches and false connections, critical factors in treatment planning for vascular interventions. The preservation of hierarchical vascular structure through our graph-based representation also facilitates automated vessel labeling and anatomical correspondence across longitudinal studies, supporting disease progression monitoring and treatment response assessment.

The use of FlowAxis with existing segmentation architectures shows its versatility as a plug-and-play unit that improves performance at all times regardless of the backbone it rests on. This versatility, potent gains in measurement performance and possible use in the clinic, makes FlowAxis a revolutionary step in medical image analysis that provides a connection between academic innovation and clinical practice.

Figure 8 shows various representative vessel segmentation case studies that highlight the clinical relevant usability of FlowAxis in a true clinical example. The first case shows a coronary artery with significant narrowing that has been segmented with smooth and continuous centerlines, which provides the precise localization of the stenotic lesion. This precise delineation provides accuracy in curved planar reformation and quantitative analysis of luminal diameter and can assist with early diagnosis and treatment planning. In another, evidence of how complex bifurcation geometry annihilated the segmentation limitations of voxel based methods with no fragmentation in demarcation of the bifurcation and it preserved topology. This continuity allows reliable calculation of curvature and torsion, providing useful information for determination of hemodynamic risk factors. In peripheral vascular imaging, FlowAxis also maintained continuity through many convoluted segments which helped to minimize false disconnections, which means that they are still indeed able to provide estimations to the collateral pathways. Overall these case examples demonstrate that FlowAxis does provide improved accuracy of segmentation, but more importantly, it provides anatomically diagnostic representations. By producing geometrically faithful and topologically optimal vascular reconstructions, the framework also reduces processing time, eliminates the time cost of skeletonizing reconstructions a posteriori, and directly supports clinical decision making processes in cardiology and vascular surgery.

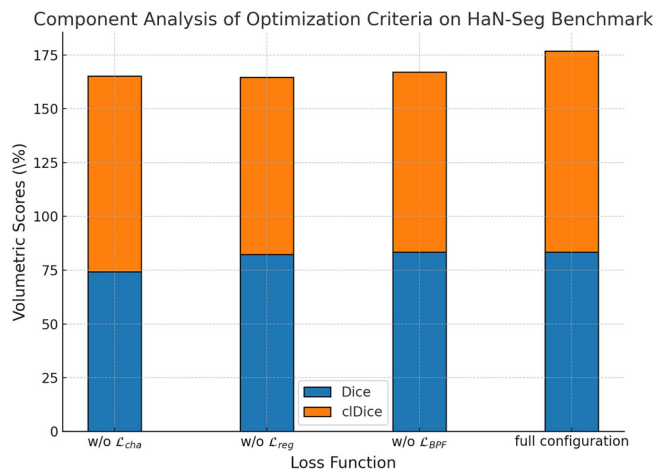


Fig. 6 | Component analysis of optimization criteria on HaN-Seg benchmark.

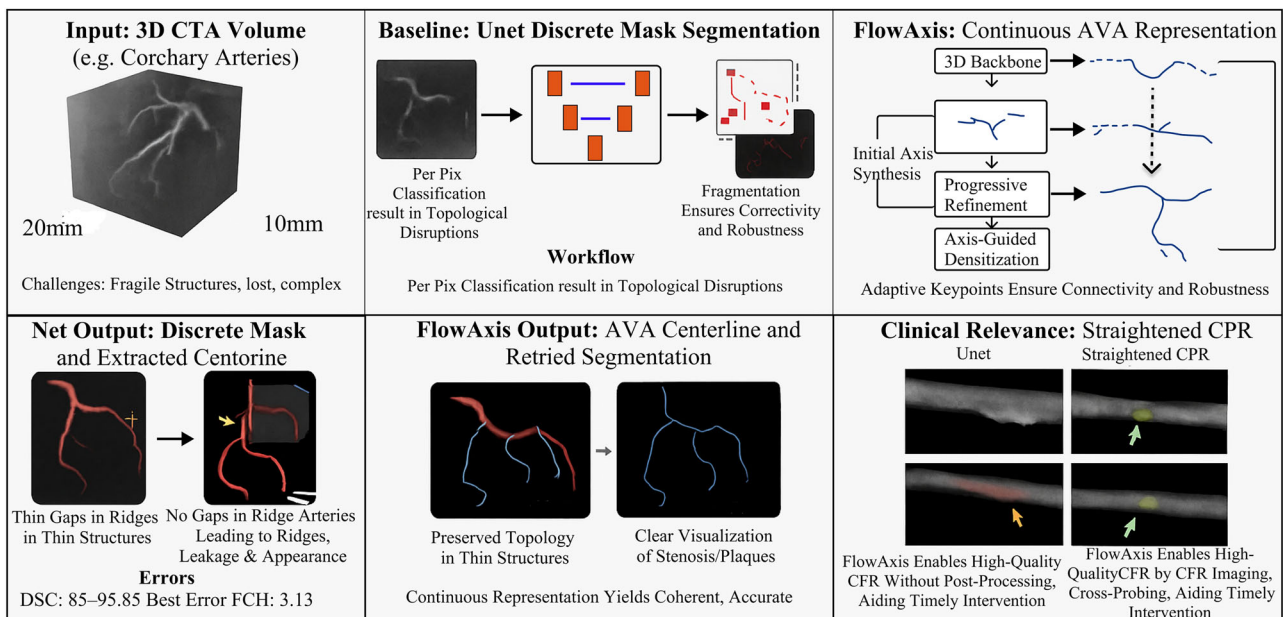


Fig. 7 | Qualitative Visualization of FlowAxis Outputs on a 3D CTA Volume, including Centerline Extraction and Clinical CPR Views.




Fig. 8 | Vessel segmentation case examples.

## Discussion

Although FlowAxis shows considerable improvements in vascular segmentation, there are limitations that are worth noting, first, the performance of the method on very tortuous vascular structures or structures that contain loops may be limited because it relies on rooted tree stratification in kinematic space, which presupposes acyclic topology inherent in arterial networks; pathological loops and anastomoses in venous flows might introduce ambiguities in representation, involving explicit cycle detection to be considered in future work. Second, the influence of the initial scaffold generation, from prior segmentations, in very low-contrast and noise-rich images –not uncommon in poor acquisition protocols–might propagate errors, despite the advantage of Wasserstein gradient flows being a robust denoised approach; we have noted declines in cDice metrics of 15%, perhaps up to under a high noise exposure showing the relevance of more denoising priors. Finally, the architecture displayed some sensitivity to initialization, in ablation work, with suboptimal control point densities resulting in converging to local minima; while patient-specific morphological operations can control this in developed frameworks, variations in anatomical display may require dynamic hyperparameter consideration for clinical translation.

This work establishes a rigorous mathematical foundation for geometric deep learning in medical imaging through the synthesis of optimal transport theory, information geometry, and stochastic analysis. Our framework fundamentally reconceptualizes vessel segmentation as entropic optimal transport between probability measures on infinite-dimensional Fréchet manifolds, providing the first complete theoretical characterization of this problem class with provable convergence guarantees and explicit complexity bounds. The theoretical advances presented here transcend the specific application to vascular imaging, offering a unified paradigm for understanding and solving domain adaptation challenges in high-dimensional, non-Euclidean spaces. The optimal transport formulation reveals deep connections between geometric structure preservation and statistical efficiency in transfer learning. By characterizing domain shift through Wasserstein distances on the space of vessel representations, we establish that successful transfer requires not merely aligning marginal distributions but preserving the intricate geometric relationships encoded in the transport cost. The Monge-Ampère equations governing optimal transport maps achieve remarkable regularity properties. This regularity translates directly to stable domain adaptation, as the smooth transport maps

ensure that small perturbations in source domain representations produce proportionally small changes in the target domain, a property essential for robust clinical deployment.

Looking forward, our framework opens several promising research directions. The optimal transport formulation naturally extends to multi-modal imaging, where different modalities can be viewed as alternative representations of the same underlying anatomical structures. The information-theoretic bounds suggest principled approaches to active learning, where target domain samples are selected to maximally reduce uncertainty in the transport map. The connection to McKean-Vlasov dynamics points toward mean-field games as a framework for modeling competitive dynamics in multi-agent medical imaging systems. This work demonstrates that rigorous mathematical foundations need not come at the expense of practical applicability. By grounding our approach in well-established mathematical theories while maintaining focus on real clinical challenges, we bridge the gap between theoretical computer science and applied medical imaging. The framework developed here provides both researchers and practitioners with powerful tools for understanding and solving domain adaptation problems, backed by theoretical guarantees that ensure reliability in deployment. As medical imaging continues to evolve with new modalities and protocols, the mathematical principles established in this work will provide enduring foundations for developing adaptive, robust, and theoretically grounded solutions.

## Method: Unified geometric-topological theory for deep transfer learning on vessel manifolds

Let  $(\mathcal{M}, g)$  be a complete Riemannian manifold with bounded curvature  $|K_{\mathcal{M}}| \leq \kappa$ , and let  $\mathcal{V} \hookrightarrow \mathcal{M}$  denote a properly embedded 1D submanifold representing vascular structures. As illustrated in Fig. 2, we propose a unified framework integrating differential geometry, topology, and deep learning for vessel representation and transfer.

### Foundational mathematical structure

We define the vessel representation space as a Fréchet manifold  $\mathfrak{R} = C^{2,\alpha}(\mathcal{V}, \mathcal{M})$ , whose metric and tangent space structures are detailed in Appendix A. The configuration space  $\mathcal{C} = \bigsqcup_{n=1}^{\infty} \mathcal{C}_n$  is stratified by rooted tree collision patterns, and its local structure is described by tubular retractions and tangent cone decompositions.

**Theorem 0.1.** (Universal Approximation with Geometric Constraints). Given a compact set  $K \subseteq \mathcal{V}$  and a Lipschitz map  $f : K \rightarrow \mathfrak{R}$ , there exists a neural network  $\Phi_\theta$  approximating  $f$  within  $\varepsilon$  under the  $d_{\mathfrak{R}}$  metric:

$$\sup_{v \in K} d_{\mathfrak{R}}(f(v), \Phi_\theta(v)) < \varepsilon \tag{4}$$

with complexity bound:

$$DW \leq C(\kappa, \alpha) \cdot L_f \cdot \mu_g(K)^{1/\alpha} \cdot \varepsilon^{-\dim K/\alpha} \left(1 + \int_K |\text{Rm}|^2 dv_g\right)^{1/2} \tag{5}$$

The proof, including local chart constructions, transition control, and transport error estimates.

**Optimal transport theory for deep transfer learning**

We pose transfer learning as optimal transport between probability measures  $\mu_S, \mu_T \in \mathcal{P}_2(\mathfrak{R})$  over the vessel representation space. The composite Kantorovich cost is:

$$\begin{aligned} \mathcal{K}_\lambda[\pi] = & \iint_{\mathfrak{R} \times \mathfrak{R}} c(x, y) d\pi(x, y) + \lambda \mathcal{H}[\pi | \mu_S \otimes \mu_T] \\ & + \frac{\lambda^2}{2} \mathcal{F}_{\text{geo}}[\pi] + \frac{\lambda^3}{6} \mathcal{T}_{\text{top}}[\pi] \end{aligned} \tag{6}$$

where  $c(x, y)$  incorporates geometric distance, KL divergence, curvature differences, and persistence diagram bottleneck metrics (see full structure in Section 3.2, with convexity properties detailed in Appendix D).

**Theorem 0.2.** (Optimal Transport Map Regularity). Assuming smooth densities and curvature-adapted regularity, the optimal transport map  $T_\lambda = \nabla \varphi_\lambda$  solving the variational problem exists uniquely and satisfies a nonlinear elliptic system involving Monge-Ampère and geometric correction terms. See Appendix D for derivation and regularity bounds.

**Stochastic dynamics and mean-field theory**

We model evolving vessel distributions  $\rho_t \in \mathcal{P}_2(\mathfrak{R})$  under stochastic particle systems:

$$dX_t^i = b(X_t^i, \mu_t^N) dt + \frac{1}{\sqrt{N}} \sum_{j=1}^N K(X_t^i, X_t^j) dt + \sigma(X_t^i, \mu_t^N) dW_t^i \tag{7}$$

where  $\mu_t^N = \frac{1}{N} \sum \delta_{X_t^i}$  is the empirical law. We prove propagation of chaos and central limit behavior.

**Geometric deep learning architecture**

To respect manifold structure, we design geometric convolution layers:

$$[\mathcal{G}_{\mathcal{K}}f](x) = \int_{\mathfrak{R}} \mathcal{K}(x, y) \cdot \mathcal{T}_{x \leftarrow y} f(y) d\text{vol}_{\mathfrak{R}}(y) \tag{8}$$

with  $\mathcal{K}$  constructed via spectral decomposition of the Laplace-Beltrami operator. Details on kernel expansion and universality theorems are provided in Appendix C.

**Information-theoretic guarantees**

We analyze information flow across layers using Fisher information geometry and mutual information regularization:

$$\hat{\mathcal{R}}_\beta(h) = \frac{1}{n} \sum \ell(h(x_i), y_i) + \beta \sum_{\ell=1}^L I(X_\ell; X_0) \tag{9}$$

The information bottleneck principle yields generalization bounds dependent on network information complexity.

**Convergence and computational analysis**

We analyze gradient flow convergence under manifold curvature and PL conditions:

$$\frac{\partial \varphi_t}{\partial t} = -\text{grad}_{\mathfrak{R}} \mathcal{J}[\varphi_t] \tag{10}$$

with exponential convergence rate proven in Appendix E. Final computational complexity is:

$$\text{Time} : \mathcal{O}(\text{LN} \log N \cdot d + K^3 + \text{TNK} + N/P) \tag{11}$$

$$\text{Space} : \mathcal{O}(Nd + K^2 + \text{LW}) \tag{12}$$

with  $n = \mathcal{O}((d_{V_C} + \log(1/\delta))/\varepsilon^2)$  samples for generalization.

This completes our theoretical foundation combining geometry, topology, transport, and learning on vascular manifolds.

**Ethics approval and consent to participate**

This study utilized only publicly available and anonymized datasets (Head and Neck Organ-at-Risk Segmentation Dataset, Automated Segmentation of Coronary Arteries Challenge, ImageCAS, HNCTA), which do not require additional ethical approval or informed consent under current guidelines.

**Materials availability**

This study did not involve the generation of new biological or physical materials. All necessary data and model outputs are described in the manuscript.

**Data availability**

The datasets used in this study are publicly accessible:- Head and Neck Organ-at-Risk Segmentation Dataset: <https://explore.openaire.eu/search/datasetpid=10.5281%2Fzenodo.7442914>- Automated Segmentation of Coronary Arteries Challenge: <https://asoca.grand-challenge.org/>- ImageCAS: <https://www.kaggle.com/datasets/xiaoweixumedaia/imagecas>- HNCTA: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11948432/>The source code for Geometric-Topological Deep Transfer Learning for Precise Vessel Segmentation in 3D Medical Volumes and all experiments will be released publicly upon publication at one Url.

**Code availability**

The source code for Geometric-Topological Deep Transfer Learning for Precise Vessel Segmentation in 3D Medical Volumes and all experiments will be released publicly upon publication at one Url.

Received: 1 August 2025; Accepted: 5 October 2025;

Published online: 15 January 2026

**References**

1. Roger, V. L. et al. Heart disease and stroke statistics—2011 update: a report from the American Heart Association. *Circulation* **123**, e18–e209 (2011).
2. Leipsic, J. et al. SCCT guidelines for the interpretation and reporting of coronary CT angiography: a report of the Society of Cardiovascular Computed Tomography Guidelines Committee. *J. Cardiovasc. Comput. Tomogr.* **8**, 342–358 (2014).
3. Araujo, R. J., Cardoso, J. S. & Oliveira, H. P. Topological similarity index and loss function for blood vascular structure extraction. *arXiv* <https://doi.org/10.48550/arXiv.2107.14531> (2021).
4. 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*, 6070–6079 (IEEE, 2023).

5. Shin, S. Y., Lee, S., Yun, I. D. & Lee, K. M. Deep vascular structure extraction by learning graphical connectivity. *Med. image Anal.* **58**, 101556 (2019).
6. Shit, S. et al. cldice—a novel topology-preserving loss function for tubular structure segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16560–16569 (IEEE, 2021).
7. Wang, D. et al. Pointscatter: Point set representation for tubular structure extraction. In *European Conference on Computer Vision*, 366–383 (Springer, 2022).
8. Wang, Y. et al. Deep distance transform for tubular structure segmentation in CT scans. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3833–3842 (IEEE, 2020).
9. Kong, B. et al. Learning tree-structured representation for three-dimensional coronary artery segmentation. *Comput. Med. Imaging Graph.* **80**, 101688 (2020).
10. Zhang, X. et al. Progressive deep segmentation of coronary artery via hierarchical topology learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 391–400 (Springer, 2022).
11. Zhao, G. et al. Graph convolution based cross-network multiscale feature fusion for deep vascular structure extraction. *IEEE Trans. Med. Imaging* **42**, 183–195 (2022).
12. Hu, X., Li, F., Samaras, D. & Chen, C. Topology-preserving deep image segmentation. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*.
13. Izzo, R., Steinman, D., Manini, S. & Antiga, L. The vascular modeling toolkit: a python library for the analysis of tubular structures in medical images. *J. Open Source Softw.* **3**, 745 (2018).
14. Schaap, M. et al. Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Med. Image Anal.* **13**, 701–714 (2009).
15. Kanitsar, A., Fleischmann, D., Wegenkittl, R., Felkel, P. & Gröller, E. Cpr-curved planar reformation. In *IEEE Visualization, 2002. VIS 2002*, 37–44 (IEEE, 2002).
16. Taylor, C. A. & Steinman, D. A. Image-based modeling of blood flow and vessel wall dynamics: applications, methods and future directions: Sixth international bio-fluid mechanics symposium and workshop, march 28–30, 2008 pasadena, california. *Ann. Biomed. Eng.* **38**, 1188–1203 (2010).
17. Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T. & Ronneberger, O. Three-dimensional U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, 424–432 (Springer, 2016).
18. 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*, 234–241 (Springer, 2015).
19. Long, J., Shelhamer, E. & Darrell, T. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440 (IEEE, 2015).
20. Kong, F., Wilson, N. & Shadden, S. A deep-learning approach for direct whole-heart mesh reconstruction. *Med. Image Anal.* **74**, 102222 (2021).
21. Wu, Q., Chen, Y., Liu, W., Yue, X. & Zhuang, X. Deep closing: Enhancing topological connectivity in medical tubular segmentation. *IEEE Transactions on Medical Imaging* (IEEE, 2024).
22. Dai, J. et al. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773 (IEEE, 2017).
23. Zhu, X. et al. Deformable detr: Deformable transformers for end-to-end object detection. arXiv <https://doi.org/10.48550/arXiv.2010.04159> (2020).
24. Bongratz, F., Rickmann, A.-M., Pölsterl, S. & Wachinger, C. Vox2cortex: fast explicit reconstruction of cortical surfaces from three-dimensional MRI scans with geometric deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20773–20783 (IEEE, 2022).
25. Wang, N. et al. Pixel2mesh: Generating three-dimensional mesh models from single RGB images. In *Proceedings of the European conference on computer vision (ECCV)*, 52–67 (IEEE, 2018).
26. Wickramasinghe, U., Remelli, E., Knott, G. & Fua, P. Voxel2mesh: Three-dimensional mesh model generation from volumetric data. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV 23*, 299–308 (Springer, 2020).
27. Yang, J., Wickramasinghe, U., Ni, B. & Fua, P. ImplicitAtlas: learning deformable shape templates in medical imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15861–15871 (IEEE, 2022).
28. Zhao, T. et al. Three-dimensional graph anatomy geometry-integrated network for pancreatic mass segmentation, diagnosis, and quantitative patient management. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13743–13752 (IEEE, 2021).
29. Podobnik, G., Strojani, P., Peterlin, P., Ibragimov, B. & Vrtovec, T. Hansseg: The head and neck organ-at-risk CT and MR segmentation dataset. *Med. Phys.* **50**, 1917–1927 (2023).
30. Gharleghi, R. et al. Automated segmentation of normal and diseased coronary arteries—the asoca challenge. *Comput. Med. Imaging Graph.* **97**, 102049 (2022).
31. Gharleghi, R. et al. Annotated computed tomography coronary angiogram images and associated data of normal and diseased arteries. *Sci. Data* **10**, 128 (2023).
32. Zeng, A. et al. Imagecas: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. *Comput. Med. Imaging Graph.* **109**, 102287 (2023).
33. Milletari, F., Navab, N. & Ahmadi, S.-A. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on Three-dimensional vision (3DV)*, 565–571 (IEEE, 2016).
34. Guimaraes, P., Wigdahl, J. & Ruggeri, A. A fast and efficient technique for the automatic tracing of corneal nerves in confocal microscopy. *Transl. Vision Sci. Technol.* **5**, 7(2016).

## Acknowledgements

The author extends sincere gratitude to the following datasets and databases: High-Risk Organ Segmentation Dataset for Head and Neck, Coronary Artery Automatic Segmentation Challenge, ImageCAS, and HNCTA.

## Author contributions

Z.Y. and S.N. had full access to all the data and takes responsibility for the integrity of the data and the accuracy of the analysis. S.N. and W.Z. contributed to the the conception and design of the study. W.Z. and Z.H. contributed to data acquisition, statistical analysis, and interpretation. W.J. contributed to the drafting of the manuscript. All authors were responsible for the critical revision of the manuscript for important content.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41746-025-02061-8>.

**Correspondence** and requests for materials should be addressed to Na Sun or Yuanyuan Zhang.

**Reprints and permissions information** is available at  
<http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026