HEMOPT: FLOW-STATE PROXY-SUPERVISED PRETRAINING ON HETEROGENEOUS VASCULAR GEOMETRIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Predicting vascular hemodynamic fields directly from geometry can reduce the computational cost of performing computational fluid dynamics (CFD) simulations across numerous vascular cases. However, training such surrogate models remains constrained by a substantial imbalance in supervision, with abundant unpaired vascular geometries but scarce and costly paired geometry–flow data. In this paper, we introduce HemoPT, a flow-state proxy-supervised vascular pretraining framework that requires no paired CFD fields. We construct a multi-fidelity, multi-anatomy pretraining corpus comprising 5,310 vascular geometries. By sampling 20 probe conditions per geometry, the corpus yields more than 100,000 probe-conditioned pretraining instances. HemoPT establishes spatial awareness of the vascular lumen through directional probes. This probe-derived spatial context parameterizes analytic responses associated with representative vascular flow states, producing geometry-specific hemodynamic proxies. The resulting geometry–proxy pairs are used to pretrain a Transolver encoder, which is subsequently fine-tuned for downstream velocity-field regression. A controlled study on the Aneumo dataset shows that pretraining with geometry–proxy pairs outperforms conditioning-only and geometry-oriented pretraining settings. Across four downstream tasks and varying data regimes, HemoPT achieves highly competitive performance among all compared training and pretraining methods. On the Vascular Model Repository (VMR) dataset, the gains from pretraining vary with vascular anatomy. These results indicate that HemoPT learns transferable geometry–flow representations from heterogeneous vascular geometries, with transfer benefits shaped by downstream anatomy. Code and preprocessing scripts are available at `https://anonymous.4open.science/r/HemoPT-66CF/`.

1 INTRODUCTION

Predicting vascular hemodynamic fields directly from geometry offers a computationally efficient alternative to performing computational fluid dynamics (CFD) simulations across large numbers of vascular cases. By amortizing the computational cost of repeated case-specific numerical solves into a single forward pass, learning-based surrogates have shown promise in aortic applications (Liang et al., 2020; Yevtushenko et al., 2022), coronary flow prediction (Li et al., 2021b), and aneurysm-specific flow modeling (Li et al., 2021a). Their performance, however, still depends heavily on paired geometry–flow supervision. Producing such data requires vascular geometry processing, boundary-condition specification, and numerical simulation (Arzani & Dawson, 2021), while imaging-derived velocity fields remain limited by acquisition cost, spatial resolution, and reconstruction quality (Bissell et al., 2023; Peper et al., 2022). Consequently, vascular flow surrogates are often trained from scratch on relatively small, anatomy-specific datasets, even though raw unpaired vascular geometry itself is far more readily and widely available.

This scarcity of paired supervision contrasts with the abundance of vascular geometry. Large-scale resources cover coronary arteries (Zeng et al., 2023), pulmonary arteries (Luo et al., 2023), cerebral vessels (Guo et al., 2024; Yang et al., 2020), as well as aortic and other major whole-body vascular structures (Yao et al., 2021; Wasserthal et al., 2023; 2022). Simulation-oriented geometries are also available through the Vascular Model Repository (VMR) (SimVascular Project, 2026). As illus-

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

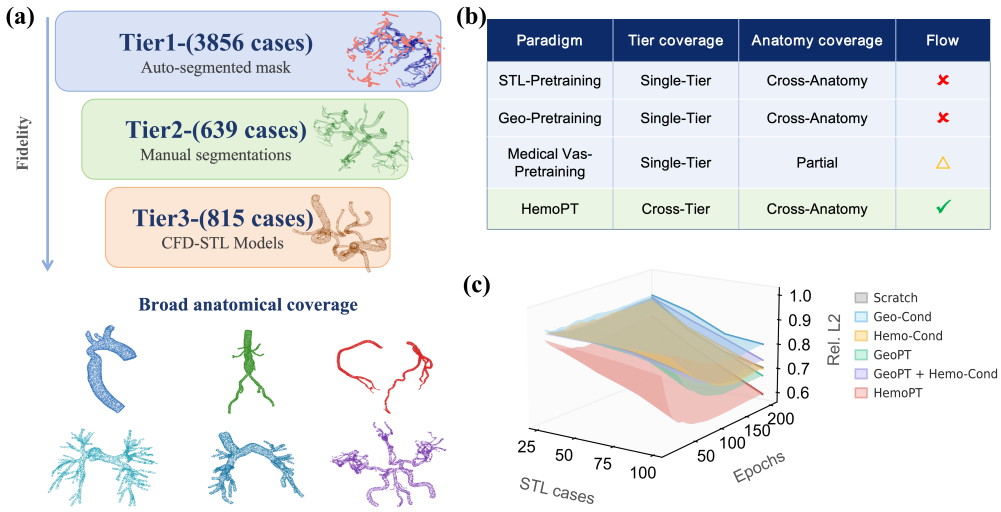


Figure 1: (a) Three-tier vascular geometry resources with increasing fidelity and broad anatomical coverage. (b) Comparison of representative pretraining paradigms in terms of tier coverage, anatomical coverage, and flow awareness, where ✓, △, and × denote full, partial, and absent support, respectively. (c) Controlled comparison on Aneumo across geometry budgets and pretraining epochs, where HemoPT achieves the lowest relative ℓ_2 error.

trated in Figure 1, these resources span three levels of geometric fidelity. Mask-derived geometries are available at the largest scale but may contain segmentation artifacts, disconnected fragments, and isolated components. Manually segmented or refined geometries provide cleaner vascular topology, although their surfaces may still require further processing before CFD simulation. CFD-ready geometries provide smooth and topologically consistent surfaces but remain considerably less abundant. Existing pretraining paradigms nevertheless exploit these heterogeneous resources primarily as shape collections, leaving their potential flow-related structure unused.

In this paper, we study how heterogeneous vascular geometries can be leveraged to learn transferable representations for downstream hemodynamic prediction without paired CFD fields. While geometry alone cannot uniquely determine patient-specific fields without boundary conditions, lumen morphology, curvature, and caliber variations fundamentally govern admissible local flow organization. Existing point cloud pretraining methods such as Point-BERT and Point-MAE focus on surface coordinate reconstruction (Yu et al., 2022; Pang et al., 2022), whereas GeoPT introduces synthetic dynamics for generic physical domains (Wu et al., 2026). However, as demonstrated by the controlled study on Aneumo in Figure 1c, purely geometric objectives yield suboptimal transfer to internal velocity fields across geometry budgets and epochs, underscoring the critical need for a scalable, flow-aware pretraining objective tailored for vascular hemodynamics.

To bridge this gap, we propose HemoPT, which unifies heterogeneous vascular geometries under a common formulation and extends geometric pretraining into a flow-aware proxy supervision space. HemoPT employs directional probes to characterize positional relations within the vascular lumen, wall proximity, and local structural variation, providing a consistent spatial description across geometry tiers. The probe-derived spatial context then parameterizes analytic responses associated with representative vascular flow states, producing geometry-specific hemodynamic proxies. The resulting spatial and hemodynamic targets are used to pretrain a Transolver encoder (Wu et al., 2024), while paired task-specific flow fields are introduced only during downstream fine-tuning.

Our main contributions are as follows:

- We construct a dedicated multi-fidelity, multi-anatomy corpus of 5,310 vascular geometries spanning multiple vascular territories throughout the body, with explicit quality tiers.

- We propose HemoPT, which combines probe-conditioned lumen traversal with analytic flow-state proxy supervision to learn transferable vascular representations across heterogeneous geometries without paired CFD fields.
- We validate HemoPT through a controlled study and extensive downstream experiments across four tasks and varying labeled-data regimes. HemoPT achieves the best or tied-best performance in many settings, while the VMR analysis reveals that its transfer benefits depend on downstream vascular anatomy.

2 Related Work

Learning-based vascular hemodynamics. Beyond direct velocity-field prediction, learning-based models have been developed for clinically relevant hemodynamic quantities, including wall shear stress estimation (Gharleghi et al., 2022) and fractional flow reserve prediction (Fossan et al., 2021). Physics-informed methods provide a complementary route by incorporating governing equations into model training (Raissi et al., 2019; Karniadakis et al., 2021), including cardiovascular applications that infer pressure and flow quantities from incomplete measurements (Kissas et al., 2020). Learning-based reconstruction, acceleration, denoising, and super-resolution have also been investigated for 4D Flow MRI (Ferdian et al., 2020; Nath et al., 2022; Shit et al., 2022; Ericsson et al., 2024). Although these approaches are directly relevant to vascular hemodynamics, they are commonly developed for specific simulation datasets, anatomical domains, boundary-condition families, or imaging protocols. Their ability to exploit large-scale, heterogeneous, and unpaired vascular geometry resources therefore remains limited.

Neural simulation and physics-oriented pretraining. Graph-based simulators model interactions over particles and meshes (Sanchez-Gonzalez et al., 2020; Pfaff et al., 2021), whereas neural operators map physical inputs to solution fields (Li et al., 2021c; Lu et al., 2021; Kovachki et al., 2023). Operator learning has extended to irregular geometries through attention-based architectures (Hao et al., 2023) and geometry-informed formulations (Li et al., 2023). These models enable efficient prediction but typically require paired inputs and solver-generated fields (Takamoto et al., 2022). Recent studies explore pretraining across PDE families and physical systems to improve transfer (Subramanian et al., 2023; McCabe et al., 2024; Herde et al., 2024), while self-supervised operator pretraining reduces reliance on complete solution trajectories (Chen et al., 2024). GeoPT augments generic geometries with synthetic dynamics and geometry-bounded trajectories without solver-generated labels (Wu et al., 2026). Nevertheless, these approaches mainly target general PDE systems or industrial geometries rather than heterogeneous vascular surfaces spanning multiple fidelity levels and geometry-conditioned flow supervision.

Point-cloud representation and geometric pretraining. PointNet, PointNet++, and DGCNN established widely used architectures for unordered point sets and local geometric modeling (Qi et al., 2017a;b; Wang et al., 2019). Point-cloud self-supervision has subsequently explored contrastive learning, occlusion completion, and cross-modal alignment (Xie et al., 2020; Zhang et al., 2021; Wang et al., 2021; Afham et al., 2022). Masked and generative paradigms further improve representation learning through hierarchical reconstruction, masked-point discrimination, multimodal reconstruction, and autoregressive generation (Zhang et al., 2022; Liu et al., 2022; Qi et al., 2023; Chen et al., 2023). These methods can scale to large and diverse unlabeled geometry collections, but are primarily developed on generic object datasets such as ShapeNet (Chang et al., 2015; Wu et al., 2015). Consequently, their pretraining objectives emphasize geometric recovery, discrimination, or generation rather than flow-related structures within vascular lumens.

3 Method

HemoPT enables paired spatial–hemodynamic pretraining directly from heterogeneous vascular geometries, without paired CFD fields. From each vascular surface and probe condition, it derives a lumen-traversal target and an RW-conditioned flow proxy via learned eight-way routing over representative vascular-flow primitives. These paired targets jointly pretrain the Transolver encoder (Figure 2), while paired velocity fields are required only for downstream fine-tuning and evaluation.

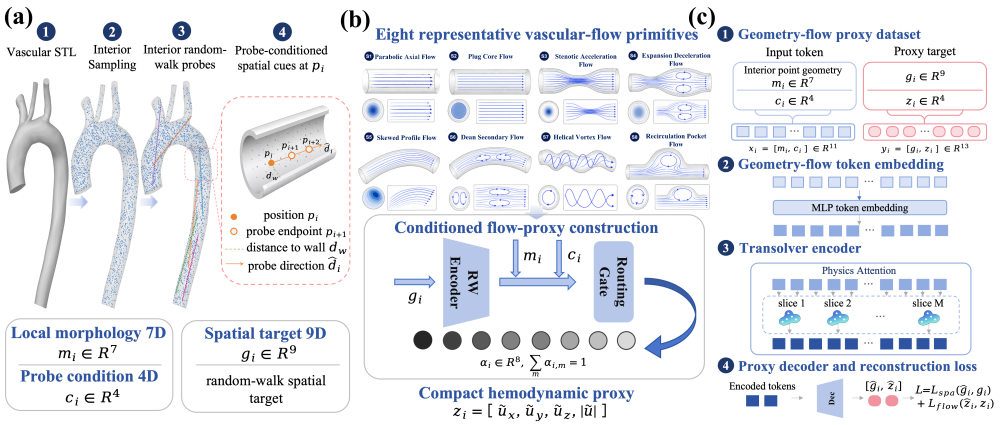


Figure 2: Overview of HemoPT. (a) Geometry-bounded directional probes derive 9D spatial targets from heterogeneous vascular STL geometries. (b) Eight CFD-inspired flow patterns are organized into a compact geometry-conditioned hemodynamic proxy. (c) Masked geometry–probe tokens are jointly trained with spatial and hemodynamic targets to pretrain a Transolver encoder.

3.1 Pretraining Corpus and Probe Sampling

To provide broad anatomical coverage and structural variation, we construct a three-tier, multi-anatomy pretraining corpus comprising a total of 5,310 diverse vascular geometries from 13 vascular segmentation datasets and simulation-oriented geometry repositories: 4TCTA_AAA (Jamshidian et al., 2025), Aneumo (Li et al., 2025), AneuRisk (Sangalli et al., 2014), AneuX (Juchler et al., 2022a;b), CereVessMRA, CMHA (Song et al., 2024), ImageCAS, ImageTBAD, the internal InterMask collection, IntrA, PARSE2022, TotalSegmentator, and the Vascular Model Repository (VMR).

The assembled corpus incorporates three distinct geometric quality tiers: mask-derived geometries representing raw clinical segmentations with potential surface irregularities, manually refined geometries offering clean vascular topology, and high-fidelity CFD-ready surfaces. After surface mesh conversion, spatial normalization, connected component cleaning, and geometric quality control, the corpus contains $K = 5{,}310$ vascular geometries (3,856 mask-derived, 639 manually refined, and 815 CFD-ready surfaces). Geometries belonging to downstream evaluation and test sets are strictly partitioned and held out before pretraining instance generation.

For each remaining vascular geometry, we systematically generate $R = 20$ conditioned views to introduce rich directional perspectives of the internal lumen. Each conditioned view is initialized from one of 20 base probe configurations, where point-wise probe directions are uniformly sampled on the unit sphere $\mathbb{S}^2$ and step lengths are drawn from $[0, 2]$. To ensure orientation diversity while avoiding degenerate collinear probes, direction vectors receive a Gaussian perturbation with standard deviation $\sigma = 0.05$ and are subsequently renormalized to unit length. This yields an extensive total of $KR = 106{,}200$ diverse, informative probe-conditioned pretraining instances. Detailed preprocessing and quality-control procedures are provided in Appendix A.

3.2 Probe-Conditioned Spatial Supervision

Let $\Omega_k \subset \mathbb{R}^3$ denote the vascular lumen volume bounded by the external wall surface $S_k = \partial\Omega_k$. For each interior query point $\mathbf{p}_i \in \Omega_k$, its nearest wall boundary point is determined by $\mathbf{b}_i = \operatorname{argmin}_{\mathbf{b}\in\partial\Omega_k} \|\mathbf{b} - \mathbf{p}_i\|_2$. The Euclidean wall distance is computed as $d_{w,i} = \|\mathbf{b}_i - \mathbf{p}_i\|_2$, and the outward unit wall direction is defined as $\mathbf{r}_{w,i} = (\mathbf{b}_i - \mathbf{p}_i)/(d_{w,i} + \epsilon)$, where $\epsilon = 10^{-6}$ ensures numerical stability. Together with the normalized coordinates, the local morphology and directional probe condition are parameterized as

$$\mathbf{m}_i = [\mathbf{p}_i; d_{w,i}; \mathbf{r}_{w,i}] \in \mathbb{R}^7, \qquad \mathbf{c}_i = [\mathbf{d}_i; s_i] \in \mathbb{R}^4. \tag{1}$$

Boundary surface points are explicitly retained in the sequence to preserve wall geometry; for these boundary points, step lengths and spatial displacements are set to zero. Starting from the initial

location $\mathbf{q}_i^{(0)} = \mathbf{p}_i$, each interior probe performs a two-update, three-position geometry-bounded directional traversal:

$$\mathbf{q}_i^{(t+1)} = \mathbf{q}_i^{(t)} + \gamma_i^{(t)} s_i \mathbf{d}_i, \qquad t \in \{0, 1\}, \tag{2}$$

where $\gamma_i^{(t)} \in [0, 1]$ denotes the accepted step fraction. If the proposed step endpoint remains strictly inside the vascular lumen Ω_k, the displacement is fully accepted ($\gamma_i^{(t)} = 1$); otherwise, the trajectory is truncated at the first intersection with the boundary wall $\partial\Omega_k$, and the subsequent update proceeds from this boundary-arrested position. The probe direction $\mathbf{d}_i$ and step length s_i remain fixed along a single traversal trajectory but vary independently across conditioned views.

At each traversed position $\mathbf{q}_i^{(t)}$, the local nearest wall coordinate $\mathbf{b}_i^{(t)}$ is updated, yielding the displacement offset vector $\mathbf{g}_i^{(t)} = \mathbf{q}_i^{(t)} - \mathbf{b}_i^{(t)} \in \mathbb{R}^3$. The three ordered offsets are concatenated into a 9D spatial trajectory target:

$$\mathbf{g}_i = \mathrm{concat}_{t=0}^{2} \left(\mathbf{q}_i^{(t)} - \mathbf{b}_i^{(t)} \right) = \left[\mathbf{g}_i^{(0)};\ \mathbf{g}_i^{(1)};\ \mathbf{g}_i^{(2)} \right] \in \mathbb{R}^9. \tag{3}$$

This trajectory directly encodes how boundary proximity, available lumen span, and cross-sectional confinement evolve along directional rays, providing rich spatial supervision sensitive to stenosis, aneurysmal expansion, vessel curvature, and bifurcations without relying on centerline extraction.

3.3 Proxy Construction via Learnable Routing over Fixed Primitives

While the spatial target $\mathbf{g}_i$ captures lumen morphology and accessibility, it does not explicitly characterize fluid dynamic behavior. HemoPT establishes a hemodynamic inductive prior by organizing eight canonical vascular-flow primitives into an adaptive, geometry-conditioned proxy: (i) *parabolic axial flow*, representing fully developed laminar core flow; (ii) *plug core flow*, capturing flattened systolic velocity profiles in large arteries; (iii) *stenotic acceleration flow*, describing localized high-velocity jetting in constricted lumens; (iv) *expansion deceleration flow*, modeling kinetic dissipation and flow retardation in widening conduits; (v) *skewed-profile flow*, characterizing centrifugal velocity bias toward outer walls in curved segments; (vi) *Dean secondary flow*, capturing cross-sectional counter-rotating vortex pairs; (vii) *helical vortex flow*, representing longitudinal swirling bulk motion; and (viii) *recirculation pocket flow*, capturing boundary-layer separation and vortex entrapment in cavities. These primitives are formulated as deterministic analytic velocity responses $\mathbf{u}_{i,m} = \mathcal{U}_m(S_k, \mathbf{m}_i, \mathbf{c}_i) \in \mathbb{R}^3$, governed by shared normalization rules detailed in Appendix B.2.

To establish orientation consistency across disparate geometries, each vascular surface is assigned a canonical geometric frame computed via covariance PCA of the normalized point coordinates. The principal-axis orientations are made deterministic by enforcing the largest-magnitude component of each eigenvector to be strictly positive. We denote the resulting local reference frame by $\mathcal{F}_k = (\mathbf{e}_{\parallel,k}, \mathbf{e}_{\perp,k}, \mathbf{e}_{\theta,k})$, where $\mathbf{e}_{\parallel,k}$ aligns with the dominant convective axis, $\mathbf{e}_{\perp,k}$ spans the primary lateral curvature direction, and $\mathbf{e}_{\theta,k} = \mathbf{e}_{\parallel,k} \times \mathbf{e}_{\perp,k}$ completes the right-handed orthogonal basis. The primitive bank $\{\mathcal{U}_m\}_{m=1}^{8}$ is parameter-free and remains strictly fixed throughout pretraining.

The three spatial offsets in $\mathbf{g}_i$ form an ordered sequence. A lightweight trajectory encoder E_{RW} incorporates step-position embeddings and multi-layer perceptrons with combined mean and max pooling to compress it into a 16D context representation $\mathbf{h}_i^{\mathrm{RW}}$. Combining $\mathbf{h}_i^{\mathrm{RW}}$ with local wall morphology $\mathbf{m}_i^w = [d_{w,i}; \mathbf{r}_{w,i}]$ and probe condition $\mathbf{c}_i$ yields the dynamic routing gate input:

$$\mathbf{f}_i^{\mathrm{route}} = [\mathbf{m}_i^w; \mathbf{c}_i; \mathbf{h}_i^{\mathrm{RW}}] \in \mathbb{R}^{24}, \qquad \boldsymbol{\alpha}_i = \mathrm{softmax}\left(R_\psi(\mathbf{f}_i^{\mathrm{route}}) \right) \in \mathbb{R}^8, \tag{4}$$

where E_{RW} and R_ψ are learnable neural modules, and $\sum_{m=1}^{8} \alpha_{i,m} = 1$. The routed primitive responses are softly fused into an instance-conditioned velocity vector and augmented with speed:

$$\widetilde{\mathbf{u}}_i = \sum_{m=1}^{8} \alpha_{i,m} \mathbf{u}_{i,m} \in \mathbb{R}^3, \qquad \mathbf{z}_i = [\widetilde{\mathbf{u}}_i;\ \|\widetilde{\mathbf{u}}_i\|_2] \in \mathbb{R}^4. \tag{5}$$

The compact proxy $\mathbf{z}_i$ provides an expressive, geometry-conditioned flow prior that bridges purely geometric representations with physical fluid dynamics tendencies.

3.4 SPATIAL–PROXY TRANSOLVER PRETRAINING

Each interior vascular point is represented as an 11D input query token and a corresponding 13D joint pretraining target:

$$\mathbf{x}_i = [\mathbf{m}_i; \mathbf{c}_i] \in \mathbb{R}^{11}, \qquad \mathbf{y}_i = [\mathbf{g}_i; \mathbf{z}_i] \in \mathbb{R}^{13}. \tag{6}$$

To prevent the model from learning identity mappings between input morphology and spatial offsets, we apply a sample-level masking mechanism. For each conditioned view, the wall-distance and wall-direction channels $(d_{w,i}, \mathbf{r}_{w,i})$ are jointly masked with probability $p_{\text{mask}} = 0.25$, while spatial coordinates $\mathbf{p}_i$ and probe conditions $\mathbf{c}_i$ remain visible. Pretraining targets $\mathbf{g}_i$ and $\mathbf{z}_i$ are always generated from clean, unmasked geometries. On masked views, the network is required to infer missing boundary constraints and synthesize admissible flow responses purely from point coordinate distributions and surrounding contextual cues.

The masked token sequence $\widetilde{\mathbf{X}}$ is projected through an input MLP and encoded by a Transolver backbone E_θ. Transolver utilizes Physics-Attention to dynamically map discrete point tokens into intrinsic physical slices, performs slice-level self-attention interactions to capture long-range fluid dependencies, and projects slice features back to physical point representations. A lightweight multi-head decoder D_ϕ then jointly predicts both spatial and flow targets: $(\widehat{\mathbf{G}}_b, \widehat{\mathbf{Z}}_b) = D_\phi(E_\theta(\widetilde{\mathbf{X}}_b))$.

We formulate sample-wise relative ℓ_2 reconstruction objectives for both prediction branches:

$$\mathcal{L}_{\text{spa}} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \frac{\|\widehat{\mathbf{G}}_b - \mathbf{G}_b\|_2}{\|\mathbf{G}_b\|_2 + \epsilon}, \qquad \mathcal{L}_{\text{proxy}} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \frac{\|\widehat{\mathbf{Z}}_b - \text{sg}(\mathbf{Z}_b)\|_2}{\|\text{sg}(\mathbf{Z}_b)\|_2 + \epsilon}, \tag{7}$$

where $\mathcal{B}$ denotes the batch of geometries, $\epsilon = 10^{-5}$, and $\text{sg}(\cdot)$ represents the stop-gradient operator applied to the entire proxy target and its normalization term. The complete flow-branch objective $\mathcal{L}_{\text{flow}}$ unites proxy reconstruction $\mathcal{L}_{\text{proxy}}$ with self-supervised physical regularization of the routing module. The self-supervised routing objectives comprise three components: (i) wall no-slip boundary suppression that penalizes non-zero velocity near the lumen boundary via exponential wall-distance damping; (ii) continuum incompressibility and spatial continuity enforced via a Moving Least Squares (MLS) divergence penalty across neighboring points; and (iii) gating entropy regularization that maintains non-degenerate mixture entropy across the eight primitive modes.

The synthesized proxy target $\mathbf{z}_i$ is detached via $\text{sg}(\cdot)$ prior to supervising the Transolver network. Consequently, proxy reconstruction updates the Transolver encoder E_θ and decoder D_ϕ, while the self-supervised consistency terms update the trajectory encoder E_{RW} and routing gate R_ψ. The overall pretraining objective is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{spa}} + \lambda_{\text{flow}} \mathcal{L}_{\text{flow}}, \tag{8}$$

where the flow weighting parameter is set to $\lambda_{\text{flow}} = 0.25$.

Following pretraining, the auxiliary reconstruction decoder D_ϕ and the proxy-generation routing modules (E_{RW}, R_ψ) are discarded. The pretrained Transolver encoder E_θ is retained as a transferable spatial–flow feature extractor and fine-tuned end-to-end with task-specific heads for downstream hemodynamic prediction.

4 EXPERIMENTS

We test whether solver-free proxy pretraining transfers to velocity-field regression across anatomies, flow sources, and labeled-data budgets. Experiments separate generic point-cloud transfer, vascular-domain geometric pretraining, GeoPT traversal pretraining, and HemoPT's spatial–flow proxy.

4.1 EXPERIMENTAL SETUP

Downstream datasets. Aneumo (Li et al., 2025) contains CFD-derived cerebral aneurysm fields. VMR-AOCOA is an aortic-coarctation subset constructed from the Vascular Model Repository. ICA-Sim is a study-specific CFD collection of normal and stenotic internal-carotid bifurcations. The 4DFlow task instead uses imaging-derived vascular flow, testing transfer across both anatomy

Table 1: Downstream velocity-field prediction across four labeled-data regimes. Lower vel-vec. and vel-mag. indicate better accuracy. Avg. gain denotes mean percentage improvement over Scratch; bold and underlined entries mark best and second-best, respectively.

Dataset	Method	25% Train			50% Train			75% Train			100% Train		
		vel-vec.↓	vel-mag.↓	avg. gain↑	vel-vec.↓	vel-mag.↓	avg. gain↑	vel-vec.↓	vel-mag.↓	avg. gain↑	vel-vec.↓	vel-mag.↓	avg. gain↑
Aneumo	Scratch	0.81 ± 0.04	0.65 ± 0.03	–	0.66 ± 0.03	0.55 ± 0.02	–	0.56 ± 0.00	0.47 ± 0.01	–	0.47 ± 0.02	0.40 ± 0.02	–
	Point-BERT-SN	0.69 ± 0.01	0.58 ± 0.01	+12.58	0.62 ± 0.01	0.54 ± 0.01	+4.79	0.57 ± 0.02	0.50 ± 0.02	-2.82	0.50 ± 0.01	0.43 ± 0.00	-5.77
	Point-MAE-SN	0.73 ± 0.01	0.61 ± 0.00	+8.26	0.64 ± 0.01	0.55 ± 0.00	+1.82	0.58 ± 0.02	0.50 ± 0.02	-4.04	0.51 ± 0.00	0.43 ± 0.00	-7.78
	Point-BERT-Vasc	0.70 ± 0.01	0.59 ± 0.00	+11.40	0.62 ± 0.01	0.54 ± 0.01	+3.78	0.58 ± 0.01	0.51 ± 0.02	-5.25	0.52 ± 0.01	0.44 ± 0.01	-9.91
	Point-MAE-Vasc	0.85 ± 0.11	0.67 ± 0.07	-3.65	0.67 ± 0.01	0.57 ± 0.01	-2.41	0.62 ± 0.01	0.54 ± 0.00	-12.82	0.60 ± 0.00	0.53 ± 0.00	-28.59
	GeoPT-Vasc	0.70 ± 0.03	0.55 ± 0.02	+13.99	0.55 ± 0.02	0.45 ± 0.01	+17.47	0.48 ± 0.01	0.41 ± 0.01	+14.32	0.45 ± 0.00	0.38 ± 0.00	+4.73
	HemoPT-Vasc	**0.68 ± 0.03**	**0.54 ± 0.02**	**+16.86**	**0.53 ± 0.01**	**0.44 ± 0.01**	**+20.45**	**0.47 ± 0.01**	**0.39 ± 0.01**	**+16.81**	**0.42 ± 0.01**	**0.36 ± 0.00**	**+10.83**
VMR-AOCOA	Scratch	**1.03 ± 0.05**	**0.95 ± 0.03**	–	1.12 ± 0.20	1.04 ± 0.12	–	1.04 ± 0.11	0.96 ± 0.06	–	1.03 ± 0.05	0.96 ± 0.01	–
	Point-BERT-SN	1.48 ± 0.31	1.31 ± 0.27	-41.18	1.32 ± 0.23	1.14 ± 0.18	-13.46	1.18 ± 0.19	1.03 ± 0.15	-10.43	1.14 ± 0.20	0.98 ± 0.18	-6.57
	Point-MAE-SN	1.46 ± 0.28	1.28 ± 0.25	-38.32	1.37 ± 0.21	1.14 ± 0.17	-15.92	1.34 ± 0.13	1.14 ± 0.11	-23.55	1.31 ± 0.19	1.13 ± 0.18	-22.49
	Point-BERT-Vasc	1.82 ± 0.37	1.61 ± 0.34	-73.38	1.46 ± 0.29	1.26 ± 0.26	-25.46	1.32 ± 0.26	1.14 ± 0.21	-22.42	1.25 ± 0.28	1.07 ± 0.24	-16.35
	Point-MAE-Vasc	1.55 ± 0.41	1.36 ± 0.42	-46.44	1.28 ± 0.20	1.10 ± 0.20	-10.11	1.24 ± 0.25	1.09 ± 0.27	-16.31	1.28 ± 0.15	1.14 ± 0.18	-21.48
	GeoPT-Vasc	1.14 ± 0.09	1.02 ± 0.05	-8.80	1.05 ± 0.12	0.96 ± 0.08	+7.16	1.03 ± 0.10	0.97 ± 0.05	+0.05	1.02 ± 0.10	0.95 ± 0.06	+1.11
	HemoPT-Vasc	1.03 ± 0.08	0.96 ± 0.04	**-0.59**	**1.01 ± 0.05**	**0.95 ± 0.02**	**+9.09**	**1.02 ± 0.07**	**0.96 ± 0.01**	**+0.99**	**0.98 ± 0.02**	**0.94 ± 0.05**	**+3.37**
ICA-Sim	Scratch	0.46 ± 0.02	0.37 ± 0.01	–	0.43 ± 0.01	0.37 ± 0.01	–	0.42 ± 0.01	0.36 ± 0.01	–	0.38 ± 0.02	0.32 ± 0.02	–
	Point-BERT-SN	0.34 ± 0.00	0.30 ± 0.00	+22.82	0.31 ± 0.01	0.27 ± 0.00	+26.96	0.29 ± 0.01	0.26 ± 0.01	+29.07	0.28 ± 0.01	0.26 ± 0.00	+23.14
	Point-MAE-SN	0.34 ± 0.01	**0.30 ± 0.01**	+24.05	0.31 ± 0.00	0.27 ± 0.00	+27.05	0.33 ± 0.01	0.27 ± 0.01	+24.38	0.28 ± 0.01	0.25 ± 0.01	+24.96
	Point-BERT-Vasc	0.41 ± 0.04	0.33 ± 0.02	+11.60	**0.30 ± 0.01**	0.27 ± 0.01	+27.20	0.29 ± 0.01	0.26 ± 0.01	+28.92	0.29 ± 0.01	0.26 ± 0.00	+22.18
	Point-MAE-Vasc	0.38 ± 0.02	0.32 ± 0.01	+16.44	0.33 ± 0.01	0.29 ± 0.01	+21.80	0.32 ± 0.00	0.28 ± 0.00	+24.71	0.31 ± 0.00	0.27 ± 0.00	+17.37
	GeoPT-Vasc	0.37 ± 0.03	0.30 ± 0.02	+19.45	0.35 ± 0.01	0.29 ± 0.00	+20.05	0.29 ± 0.00	0.25 ± 0.00	+30.44	0.31 ± 0.01	0.25 ± 0.01	+20.87
	HemoPT-Vasc	**0.33 ± 0.00**	0.30 ± 0.00	**+24.16**	0.32 ± 0.01	**0.26 ± 0.01**	**+27.25**	**0.29 ± 0.01**	**0.25 ± 0.02**	**+31.64**	**0.27 ± 0.01**	**0.23 ± 0.01**	**+29.29**
4DFlow	Scratch	0.84 ± 0.01	0.49 ± 0.01	–	0.81 ± 0.02	0.48 ± 0.01	–	0.79 ± 0.01	0.47 ± 0.01	–	0.77 ± 0.01	0.45 ± 0.01	–
	Point-BERT-SN	0.82 ± 0.02	0.54 ± 0.02	-3.16	0.79 ± 0.02	0.53 ± 0.01	-4.82	0.81 ± 0.02	0.53 ± 0.02	-7.48	0.79 ± 0.02	0.53 ± 0.02	-9.75
	Point-MAE-SN	0.82 ± 0.02	0.54 ± 0.01	-3.65	0.80 ± 0.01	0.54 ± 0.01	-5.51	0.79 ± 0.02	0.53 ± 0.01	-6.40	0.78 ± 0.01	0.52 ± 0.02	-7.89
	Point-BERT-Vasc	0.84 ± 0.03	0.56 ± 0.04	-7.04	0.80 ± 0.01	0.53 ± 0.01	-5.17	0.81 ± 0.02	0.53 ± 0.01	-7.38	0.78 ± 0.02	0.52 ± 0.02	-8.00
	Point-MAE-Vasc	0.84 ± 0.05	0.57 ± 0.01	-7.75	0.83 ± 0.02	0.55 ± 0.02	-9.04	0.83 ± 0.01	0.54 ± 0.02	-9.74	0.81 ± 0.01	0.52 ± 0.03	-10.03
	GeoPT-Vasc	0.80 ± 0.03	0.47 ± 0.02	+4.18	0.77 ± 0.03	0.45 ± 0.01	+5.33	0.75 ± 0.02	0.45 ± 0.01	+4.74	0.72 ± 0.01	0.43 ± 0.01	+5.68
	HemoPT-Vasc	**0.78 ± 0.02**	**0.46 ± 0.01**	**+6.47**	**0.76 ± 0.02**	**0.45 ± 0.01**	**+6.38**	**0.74 ± 0.01**	**0.44 ± 0.01**	**+6.73**	**0.72 ± 0.01**	**0.43 ± 0.01**	**+6.51**

and target source. The datasets contain 540, 48, 133, and 138 cases, respectively; we hold out an independent test set before partitioning the remaining cases into five train–validation folds.

Baselines. We compare seven configurations. Scratch has no pretraining. Point-BERT-SN and Point-MAE-SN transfer ShapeNet representations, whereas Point-BERT-Vasc and Point-MAE-Vasc apply their geometry-native objectives to our vascular corpus. This pairing distinguishes generic transfer from exposure to vascular shape alone. GeoPT-Vasc applies generic directional traversal supervision (Wu et al., 2026) as an external baseline. HemoPT-Vasc introduces vascular-specialized spatial–flow proxy supervision; its spatial-only ablation ($\lambda_{\text{flow}} = 0$) is detailed in Appendix C.4.

Protocol and metrics. All methods use identical task-specific data splits, point-sampling strategies, target preprocessing, optimization budgets, checkpoint-selection criteria, and evaluation procedures. Performance is measured by the relative ℓ_2 errors of the velocity vector (**vel-vec.**) and velocity magnitude (**vel-mag.**). **Avg. gain** denotes the mean relative improvement of these two metrics over Scratch. Lower errors and higher gains indicate better performance. Detailed training settings and evaluation protocols are provided in Appendix B.

4.2 Main Results

Downstream hemodynamic transfer. Table 1 reports downstream regression performance across four labeled-data regimes. HemoPT-Vasc achieves highly competitive performance across all benchmarks and data budgets. Relative gains over Scratch are substantial on complex CFD cohorts: Aneumo achieves $+10.83\%$ to $+20.45\%$ average gain, while ICA-Sim reaches $+24.16\%$ to $+31.64\%$. On real-world clinical 4DFlow, HemoPT consistently yields positive gains between $+6.38\%$ and $+6.73\%$, demonstrating robustness against in-vivo imaging noise.

Performance of surface point-cloud SSL. As shown in Table 1, standard masked point modeling methods such as Point-BERT and Point-MAE consistently exhibit negative or erratic transfer across most benchmarks. This degradation occurs regardless of whether representations are pretrained on generic ShapeNet or domain-specific vascular geometry. For example, Point-MAE-Vasc declines by 28.59% on Aneumo with full training data, whereas Point-BERT-Vasc drops by 73.38% on VMR-AOCOA with a 25% budget. These findings indicate that self-supervised objectives purely targeting surface coordinate reconstruction do not inherently yield representations conducive to internal hemodynamic field regression.

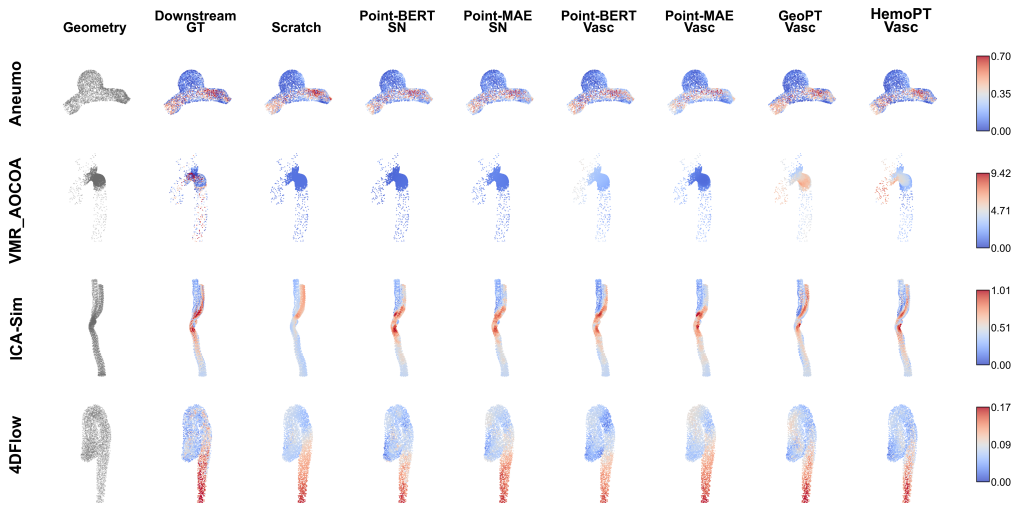


Figure 3: Qualitative velocity comparisons across Aneumo, VMR-AOCOA, ICA-Sim, and 4DFlow. HemoPT accurately resolves aneurysmal recirculation pockets, stenotic jet acceleration, and Dean secondary vortices without spatial oversmoothing.

Sample efficiency under low-label regimes. A practical advantage of HemoPT is reducing downstream dependence on computationally expensive CFD simulations. On Aneumo, fine-tuning HemoPT with only 25% labeled data yields a 16.86% relative gain, matching the performance of Scratch trained on twice that budget with velocity vector and magnitude errors of 0.68 and 0.54 compared to 0.66 and 0.55. With 50% labeled data, HemoPT improves by 20.45% and substantially surpasses Scratch trained on 75% data. On ICA-Sim, HemoPT with only 25% supervision attains velocity vector and magnitude errors of 0.33 and 0.30, already outperforming Scratch trained with 100% full supervision, which records 0.38 and 0.32.

Qualitative flow fidelity (Figure 3). Figure 3 compares predicted velocity fields against CFD and 4D Flow ground truths. Scratch and GeoPT-Vasc suffer from diffuse spatial oversmoothing, whereas Point-BERT and Point-MAE generate unphysical, jagged velocity fluctuations near wall boundaries. HemoPT sharply recovers localized flow structures across anatomies, including recirculating vortex cores within aneurysmal domes in Aneumo, asymmetric post-stenotic jet acceleration in ICA-Sim, and counter-rotating Dean secondary swirling in the thoracic aorta of 4DFlow.

Ablation of flow proxy and physical consistency losses. Table 2 audits zero-shot proxy alignment against ground-truth flow fields without downstream fine-tuning, measuring directional cosine alignment $C_{\text{dir}} \uparrow$ and relative magnitude error $C_{\text{mag}} \downarrow$. Full HemoPT achieves the highest macro alignment of 0.395 and lowest error of 0.576. Dynamic routing outperforms fixed mean routing

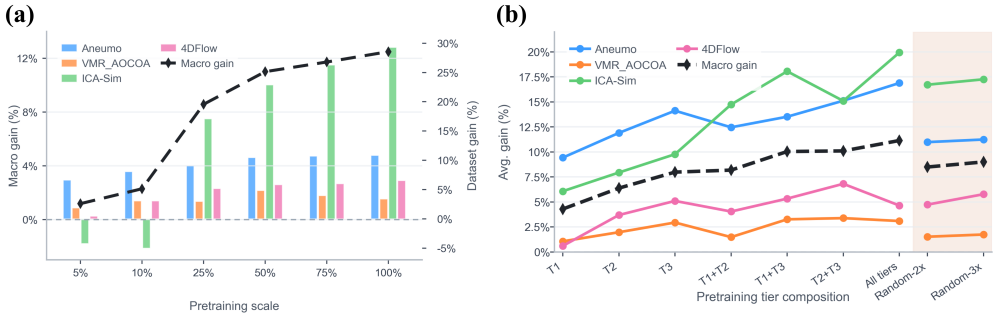


Figure 4: Scaling dynamics and geometric tier composition. (a) Downstream gain scaling from 5% to 100% corpus volume, where the dashed curve denotes the macro average. (b) Tier-specific versus combined three-tier performance plotted against size-matched random controls in the shaded band.

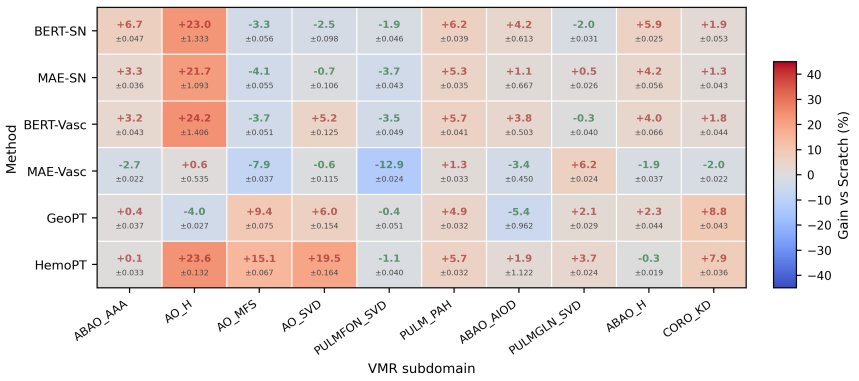


Figure 5: Anatomy-dependent transfer across VMR territories. Cells report percentage gain over Scratch. HemoPT delivers strong gains in primary aortic and coronary branches while delineating transfer limits in distal beds dominated by peripheral impedance.

at 0.594 and spatial shuffled controls at 0.613, cutting ICA magnitude error from 0.473 to 0.388. Removing incompressibility surges macro error to 0.623 due to divergence artifacts, while disabling wall no-slip and gating entropy degrades error to 0.596 and 0.592, respectively. Extended parameter sweeps are in Appendix C.4.

Scaling and multi-tier fidelity dynamics (Figure 4). Macro downstream gain scales monotonically with corpus size as shown in Figure 4a, with the steepest improvement occurring between 10% and 25% as primary convective priors are assimilated. In Figure 4b, single-tier pretraining confirms that no single fidelity tier dominates across all tasks. High-precision Tier 3 surfaces excel on fine-grained ICA bifurcations, whereas Tier 1 and Tier 2 provide broad anatomical scale diversity. The combined three-tier corpus systematically outperforms both single-tier models and size-matched controls, validating that clinical segmentations and simulation meshes provide complementary morphological coverage.

Table 2: Flow proxy and physical loss ablation across downstream cases evaluated against ground-truth flow fields. Macro Overall is the unweighted mean, with bold entries marking the best performance.

Configuration	ICA-Sim $C_{\text{dir}} \uparrow / C_{\text{mag}} \downarrow$	4DFlow $C_{\text{dir}} \uparrow / C_{\text{mag}} \downarrow$	VMR-AOCOA $C_{\text{dir}} \uparrow / C_{\text{mag}} \downarrow$	Aneumo $C_{\text{dir}} \uparrow / C_{\text{mag}} \downarrow$	Macro $C_{\text{dir}} \uparrow / C_{\text{mag}} \downarrow$
Full HemoPT	**0.941 / 0.388**	**0.173 / 0.424**	**0.200 / 0.764**	**0.264 / 0.728**	**0.395 / 0.576**
Poiseuille Prior	0.935 / 0.473	0.169 / 0.443	0.196 / 0.785	0.259 / 0.754	0.390 / 0.614
Uniform 8-Mode	0.935 / 0.454	0.168 / 0.434	0.195 / 0.782	0.258 / 0.749	0.389 / 0.605
Fixed Mean 8-Mode	0.935 / 0.448	0.170 / 0.428	0.197 / 0.773	0.263 / 0.729	0.391 / 0.594
Spatial Shuffled 8-Mode	0.935 / 0.507	0.168 / 0.428	0.196 / 0.776	0.259 / 0.740	0.390 / 0.613
w/o Wall No-Slip	0.935 / 0.407	0.169 / 0.454	0.196 / 0.778	0.259 / 0.743	0.390 / 0.596
w/o Incompressibility	0.935 / 0.494	0.168 / 0.450	0.195 / 0.790	0.258 / 0.758	0.389 / 0.623
w/o Gating Entropy	0.935 / 0.400	0.169 / 0.448	0.196 / 0.782	0.259 / 0.739	0.390 / 0.592

Anatomy-dependent transfer and physical boundaries (Figure 5). Figure 5 details territory-specific transfer across VMR subdomains. HemoPT delivers strong positive gains in primary aortic and coronary conduits (AO-H, AO-MFS, CORO-KD), where caliber changes and strong curvature directly dictate convective momentum. Conversely, transfer is modest or neutral in distal abdominal or pulmonary branches (ABAO-AAA, PULMFON-SVD). This reflects fundamental fluid mechanics: in large curved conduits, flow patterns are predominantly geometry-governed; in distal microbeds, velocity fields are heavily dictated by unseen peripheral resistance and pressure waveforms that cannot be inferred from local wall morphology alone. Characterizing these responses delineates where geometric pretraining provides the strongest inductive bias.

5 Conclusion

HemoPT couples directional lumen probes with geometry-conditioned flow proxies to pretrain transferable hemodynamic representations from heterogeneous vascular surfaces without paired CFD supervision. Embedding convective dynamics and physical conservation laws directly into geometric features substantially reduces downstream simulation requirements and enables sample-efficient transfer across diverse anatomical territories and in-vivo clinical imaging.

AI USE STATEMENT

Generative AI tools assisted language editing, organization, reference normalization, and PDF diagnostics; they did not generate data or execute the reported experiments. The authors verified all equations, results, citations, and interpretations and take full responsibility for the manuscript.

REFERENCES

Mohamed Afham, Isuru Dissanayake, et al. CrossPoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9902–9912, 2022.

Amirhossein Arzani and Scott T. M. Dawson. Data-driven cardiovascular flow modelling: Examples and opportunities. *Journal of the Royal Society Interface*, 18(175):20200802, 2021. doi: 10.1098/rsif.2020.0802.

Malenka M. Bissell, Francesca Raimondi, et al. 4d flow cardiovascular magnetic resonance consensus statement: 2023 update. *Journal of Cardiovascular Magnetic Resonance*, 25(1):40, 2023. doi: 10.1186/s12968-023-00942-z.

Angel X. Chang, Thomas Funkhouser, et al. ShapeNet: An information-rich 3d model repository, 2015.

Guangyan Chen, Meiling Wang, et al. PointGPT: Auto-regressively generative pre-training from point clouds. In *Advances in Neural Information Processing Systems*, volume 36, pp. 29667–29679, 2023.

Wuyang Chen, Jialin Song, et al. Data-efficient operator learning via unsupervised pretraining and in-context learning. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Leon Ericsson, Adam Hjalmarsson, et al. Generalized super-resolution 4d flow MRI—using ensemble learning to extend across the cardiovascular system. *IEEE Journal of Biomedical and Health Informatics*, 28(12):7239–7250, 2024.

Edward Ferdian, Avan Suinesiaputra, et al. 4DFlowNet: Super-resolution 4d flow MRI using deep learning and computational fluid dynamics. *Frontiers in Physics*, 8:138, 2020. doi: 10.3389/fphy.2020.00138.

Fredrik E. Fossan, Lucas O. Müller, et al. Machine learning augmented reduced-order models for FFR prediction. *Computer Methods in Applied Mechanics and Engineering*, 384:113892, 2021. doi: 10.1016/j.cma.2021.113892.

Ramtin Gharleghi, Arcot Sowmya, et al. Transient wall shear stress estimation in coronary bifurcations using convolutional neural networks. *Computer Methods and Programs in Biomedicine*, 225:107013, 2022. doi: 10.1016/j.cmpb.2022.107013.

Bin Guo, Ying Chen, et al. Self-supervised learning for accurately modelling hierarchical evolutionary patterns of cerebrovasculature. *Nature Communications*, 15:9235, 2024. doi: 10.1038/s41467-024-53550-5.

Zhongkai Hao, Zhengyi Wang, et al. GNOT: A general neural operator transformer for operator learning. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 12556–12569, 2023.

Maximilian Herde, Bogdan Raonic, et al. Poseidon: Efficient foundation models for PDEs. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Mostafa Jamshidian, Adam Wittek, et al. 4D-CTA image and geometry dataset for kinematic analysis of abdominal aortic aneurysms. *Data in Brief*, 61:111797, 2025. doi: 10.1016/j.dib.2025.111797.

Norman Juchler, Philippe Bijlenga, et al. AneuX morphology database, 2022a.

Norman Juchler, Sabine Schilling, et al. Shape trumps size: Image-based morphological analysis reveals that the 3d shape discriminates intracranial aneurysm disease status better than aneurysm size. *Frontiers in Neurology*, 13:809391, 2022b. doi: 10.3389/fneur.2022.809391.

George Em Karniadakis, Ioannis G. Kevrekidis, et al. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021. doi: 10.1038/s42254-021-00314-5.

Georgios Kissas, Yibo Yang, et al. Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4d flow MRI data using physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 358:112623, 2020. doi: 10.1016/j.cma.2019.112623.

Nikola B. Kovachki, Zongyi Li, et al. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Gaoyang Li, Xiaorui Song, et al. Prediction of cerebral aneurysm hemodynamics with porous-medium models of flow-diverting stents via deep learning. *Frontiers in Physiology*, 12:733444, 2021a. doi: 10.3389/fphys.2021.733444.

Gaoyang Li, Haoran Wang, et al. Prediction of 3d cardiovascular hemodynamics before and after coronary artery bypass surgery via deep learning. *Communications Biology*, 4:99, 2021b. doi: 10.1038/s42003-020-01638-1.

Xigui Li, Yuanye Zhou, et al. Aneumo: A large-scale multimodal aneurysm dataset with computational fluid dynamics simulations and deep learning benchmarks, 2025.

Zongyi Li, Nikola Kovachki, et al. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021c.

Zongyi Li, Nikola Kovachki, et al. Geometry-informed neural operator for large-scale 3d PDEs. In *Advances in Neural Information Processing Systems*, volume 36, pp. 35836–35854, 2023.

Li Liang, Weiguang Mao, et al. A feasibility study of deep learning for predicting hemodynamics of human thoracic aorta. *Journal of Biomechanics*, 99:109544, 2020. doi: 10.1016/j.jbiomech.2019.109544.

Haotian Liu, Mu Cai, et al. Masked discrimination for self-supervised learning on point clouds. In *European Conference on Computer Vision*, pp. 657–675. Springer, 2022. doi: 10.1007/978-3-031-20086-1_38.

Lu Lu, Pengzhan Jin, et al. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021. doi: 10.1038/s42256-021-00302-5.

Gongning Luo, Kuanquan Wang, et al. Efficient automatic segmentation for multi-level pulmonary arteries: The PARSE challenge, 2023.

Michael McCabe, Bruno Régaldo-Saint Blancard, et al. Multiple physics pretraining for spatiotemporal surrogate models. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Ruponti Nath, Sean Callahan, et al. FlowRAU-Net: Accelerated 4d flow MRI of aortic valvular flows with a deep 2d residual attention network. *IEEE Transactions on Biomedical Engineering*, 69(12):3812–3824, 2022. doi: 10.1109/TBME.2022.3180691.

Yatian Pang, Wenxiao Wang, et al. Masked autoencoders for point cloud self-supervised learning. In *European Conference on Computer Vision*, pp. 604–621. Springer, 2022.

Eva S. Peper, Pim van Ooij, et al. Advances in machine learning applications for 4d flow MRI. *Frontiers in Cardiovascular Medicine*, 9:1052068, 2022. doi: 10.3389/fcvm.2022.1052068.

Tobias Pfaff, Meire Fortunato, et al. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021.

Charles R. Qi, Hao Su, et al. PointNet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 652–660, 2017a. doi: 10.1109/CVPR.2017.16.

Charles R. Qi, Li Yi, et al. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, volume 30, 2017b.

Zekun Qi, Runpei Dong, et al. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28223–28243, 2023.

Maziar Raissi, Paris Perdikaris, et al. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, et al. Learning to simulate complex physics with graph networks. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8459–8468, 2020.

Laura M. Sangalli, Piercesare Secchi, et al. AneuRisk65: A dataset of three-dimensional cerebral vascular geometries. *Electronic Journal of Statistics*, 8(2):1879–1890, 2014. doi: 10.1214/14-EJS938.

Suprosanna Shit, Judith Zimmermann, et al. SRflow: Deep learning based super-resolution of 4d-flow MRI data. *Frontiers in Artificial Intelligence*, 5:928181, 2022. doi: 10.3389/frai.2022.928181.

SimVascular Project. Vascular model repository. Online repository, 2026. URL `https://www.vascularmodel.com/`.

Miao Song, Simin Wang, et al. Intracranial aneurysm CTA images and 3d models dataset with clinical morphological and hemodynamic data. *Scientific Data*, 11:1213, 2024. doi: 10.1038/s41597-024-04056-8.

Shashank Subramanian, Peter Harrington, et al. Towards foundation models for scientific machine learning: Characterizing scaling and transfer behavior. In *Advances in Neural Information Processing Systems*, volume 36, pp. 71242–71262, 2023.

Makoto Takamoto, Timothy Praditia, et al. PDEBench: An extensive benchmark for scientific machine learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 1596–1611, 2022.

Hanchen Wang, Qi Liu, et al. Unsupervised point cloud pre-training via occlusion completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9782–9792, 2021.

Yue Wang, Yongbin Sun, et al. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics*, 38(5):146:1–146:12, 2019. doi: 10.1145/3326362.

Jakob Wasserthal, Hanns-Christian Breit, et al. TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images. *Radiology: Artificial Intelligence*, 5(5):e230024, 2023. doi: 10.1148/ryai.230024.

Jakob Wasserthal et al. TotalSegmentator dataset, 2022.

Haixu Wu, Huakun Luo, et al. Transolver: A fast transformer solver for PDEs on general geometries. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 53681–53705, 2024.

Haixu Wu et al. GeoPT: Scaling physics simulation via lifted geometric pre-training, 2026.

Zhirong Wu, Shuran Song, et al. 3d ShapeNets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1912–1920, 2015. doi: 10.1109/CVPR.2015.7298801.

Saining Xie, Jiatao Gu, et al. PointContrast: Unsupervised pre-training for 3d point cloud understanding. In *European Conference on Computer Vision*, pp. 574–591. Springer, 2020.

Xi Yang, Ding Xia, et al. IntrA: 3d intracranial aneurysm dataset for deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2656–2666, 2020. doi: 10.1109/CVPR42600.2020.00273.

Zeyang Yao, Wen Xie, et al. ImageTBAD: A 3d computed tomography angiography image dataset for automatic segmentation of type-b aortic dissection. *Frontiers in Physiology*, 12:732711, 2021. doi: 10.3389/fphys.2021.732711.

Pavlo Yevtushenko, Leonid Goubergrits, et al. Deep learning-based centerline-aggregated aortic hemodynamics: An efficient alternative to numerical modeling of hemodynamics. *IEEE Journal of Biomedical and Health Informatics*, 26(4):1815–1825, 2022. doi: 10.1109/JBHI.2021.3116764.

Xumin Yu, Lulu Tang, et al. Point-BERT: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19313–19322, 2022.

An Zeng, Chunbiao Wu, 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. doi: 10.1016/j.compmedimag.2023.102287.

Renrui Zhang, Ziyu Guo, et al. Point-M2AE: Multi-scale masked autoencoders for hierarchical point cloud pre-training. In *Advances in Neural Information Processing Systems*, volume 35, pp. 27061–27074, 2022.

Zaiwei Zhang, Rohit Girdhar, et al. Self-supervised pretraining of 3d features on any point-cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10252–10263, 2021.

APPENDIX

A DATASET CONSTRUCTION AND GEOMETRY PROCESSING

This section describes the construction of the HemoPT pretraining corpus from heterogeneous vascular geometries. Unlike conventional flow-prediction settings that depend on paired geometry–CFD samples, HemoPT constructs a geometry-centric vascular corpus and derives flow-aware proxy supervision directly from processed vascular surfaces.

We first summarize the source datasets, anatomical coverage, geometry provenance, access conditions, and source-specific use terms. We then organize the collected geometries into three tiers according to their generation and processing status. Next, we describe the standardized geometry quality-control and component-cleaning pipeline used to convert heterogeneous STL surfaces into a consistent and leakage-safe pretraining corpus. Finally, we provide a representative VMR example showing how a processed vascular geometry is converted into HemoPT proxy-ready interior samples.

A.1 SOURCE DATASETS, ACCESS, AND ANATOMICAL COVERAGE

HemoPT is constructed from vascular STL geometries collected from 13 public, officially released, challenge-based, and institutional sources. The complete source manifest covers a broad range of vascular anatomies, including coronary arteries, pulmonary arteries, cerebral vasculature, cerebrovascular trees, intracranial aneurysms, abdominal aortic aneurysms, the aorta and thoracic aortic dissection, and multi-territory vascular models.

We initially collected 5,962 STL geometries. After geometry quality control, downstream-case reservation, and leakage filtering, 5,310 geometries are retained as the final HemoPT pretraining corpus. Among the remaining 652 geometries, 616 CFD-associated cases are reserved for downstream evaluation (comprising 540 from Aneumo, 48 from VMR-AOCOA for primary benchmarking, and 28 from VMR for subdomain transfer analysis), and 36 geometries are excluded because they fail geometry quality-control requirements.

Table A1 reports the source-level composition, anatomical coverage, processing status, access condition, and governing license or use terms of the complete geometry manifest. Here, *Open* denotes a directly accessible public release, *Public/request* denotes a publicly documented source that requires registration or provider contact, *Challenge* denotes data distributed under challenge-specific access conditions, and *Institutional* denotes non-public data used under institutional authorization.

The license and use-term entries summarize the source-level conditions recorded in the data manifest. For sources without a standardized dataset license, we report the applicable provider, repository, project, or challenge terms rather than inferring a license from an associated paper or software repository. HemoPT does not alter the ownership or redistribution conditions of any third-party source.

When CFD fields are available for Aneumo or VMR geometries, these fields are not used as HemoPT pretraining targets. HemoPT uses only the vascular geometries to construct geometry-conditioned spatial and flow-state proxy supervision.

The institutional InterMask source is generated from automatically segmented coronary masks and is therefore assigned to Tier 1 as a mask-derived STL source. It is tracked separately because of its institutional origin, is used under internal authorization, and is not publicly redistributed.

Figure A1 shows ten randomly sampled, post-cleaning STL geometries from each source dataset. The visualization demonstrates substantial cross-source differences in vascular anatomy, topology, spatial extent, branching complexity, and surface characteristics.

A.2 THREE-TIER VASCULAR GEOMETRY HIERARCHY

The collected vascular geometries differ not only in anatomical structure but also in how their surfaces were generated and prepared. HemoPT therefore organizes the complete geometry manifest into three tiers according to geometry provenance and processing status.

Table A1: Source-level composition, anatomical coverage, processing status, access condition, and license or use terms of the initially collected HemoPT vascular geometry manifest. Access and use conditions follow the corresponding source releases and do not imply permission for unrestricted redistribution. The final leakage-safe tier composition after quality control and downstream-case reservation is reported in Table A2.

Dataset	Count	Anatomy	Processing status	Access	License / use terms
Tier 1: Mask-derived STL geometries					
TotalSegmentator	2,357	Multi-region vascular structures	Mask-derived STL	Open	Official release terms
ImageCAS	1,000	Coronary artery	Mask-derived STL	Public / request	Provider terms
CereVessMRA	271	Cerebrovascular tree	Mask-derived STL	Public / request	Provider terms
PARSE2022	100	Pulmonary artery	Mask-derived STL	Challenge	Challenge / DUA terms
imageTBAD	100	Aorta / TBAD	Mask-derived STL	Open	Official release terms
InterMask	53	Coronary artery	Mask-derived STL	Institutional	Internal authorization
Tier 2: Manually refined STL geometries					
AneuX	297	Intracranial aneurysm	Manually refined STL	Open	Repository terms
CMHA	142	Cerebral vasculature	Manually refined STL	Open	Official release terms
IntrA	103	Intracranial aneurysm	Manually refined STL	Open	Dataset use terms
AneuRisk	95	Intracranial aneurysm	Manually refined STL	Open	Project terms
4TCTA_AAA	11	Abdominal aortic aneurysm	Manually refined STL	Open	Provider terms
Tier 3: CFD-ready STL geometries					
Aneumo	1,080	Cerebral aneurysm CFD geometries	CFD-ready STL	Open	Repository terms
VMR	353	Multi-territory vascular models	CFD-ready STL	Open	Research-use terms
Total	**5,962**	**Broad vascular anatomy**	**All tiers**	**Mixed access**	**Source-specific terms**

Tier 1: Mask-derived STL geometries. Tier 1 initially contains 3,881 vascular geometries reconstructed from automatic segmentation masks, threshold-based masks, or official challenge segmentation masks. This tier includes TotalSegmentator, ImageCAS, CereVessMRA, PARSE2022, imageTBAD, and the institutional InterMask coronary source.

These sources provide the largest corpus scale and broad anatomical coverage, but may contain disconnected fragments, staircase-like surfaces, or segmentation-induced topological irregularities. After excluding 25 geometries that fail quality control, 3,856 Tier-1 geometries are retained for HemoPT pretraining.

Tier 2: Manually refined STL geometries. Tier 2 initially contains 648 manually segmented, manually corrected, or expert-curated vascular geometries from AneuX, CMHA, IntrA, AneuRisk, and 4TCTA_AAA. These geometries generally exhibit reduced segmentation artifacts and improved surface consistency relative to directly mask-derived surfaces.

After excluding nine geometries that fail quality control, 639 Tier-2 geometries are retained for HemoPT pretraining.

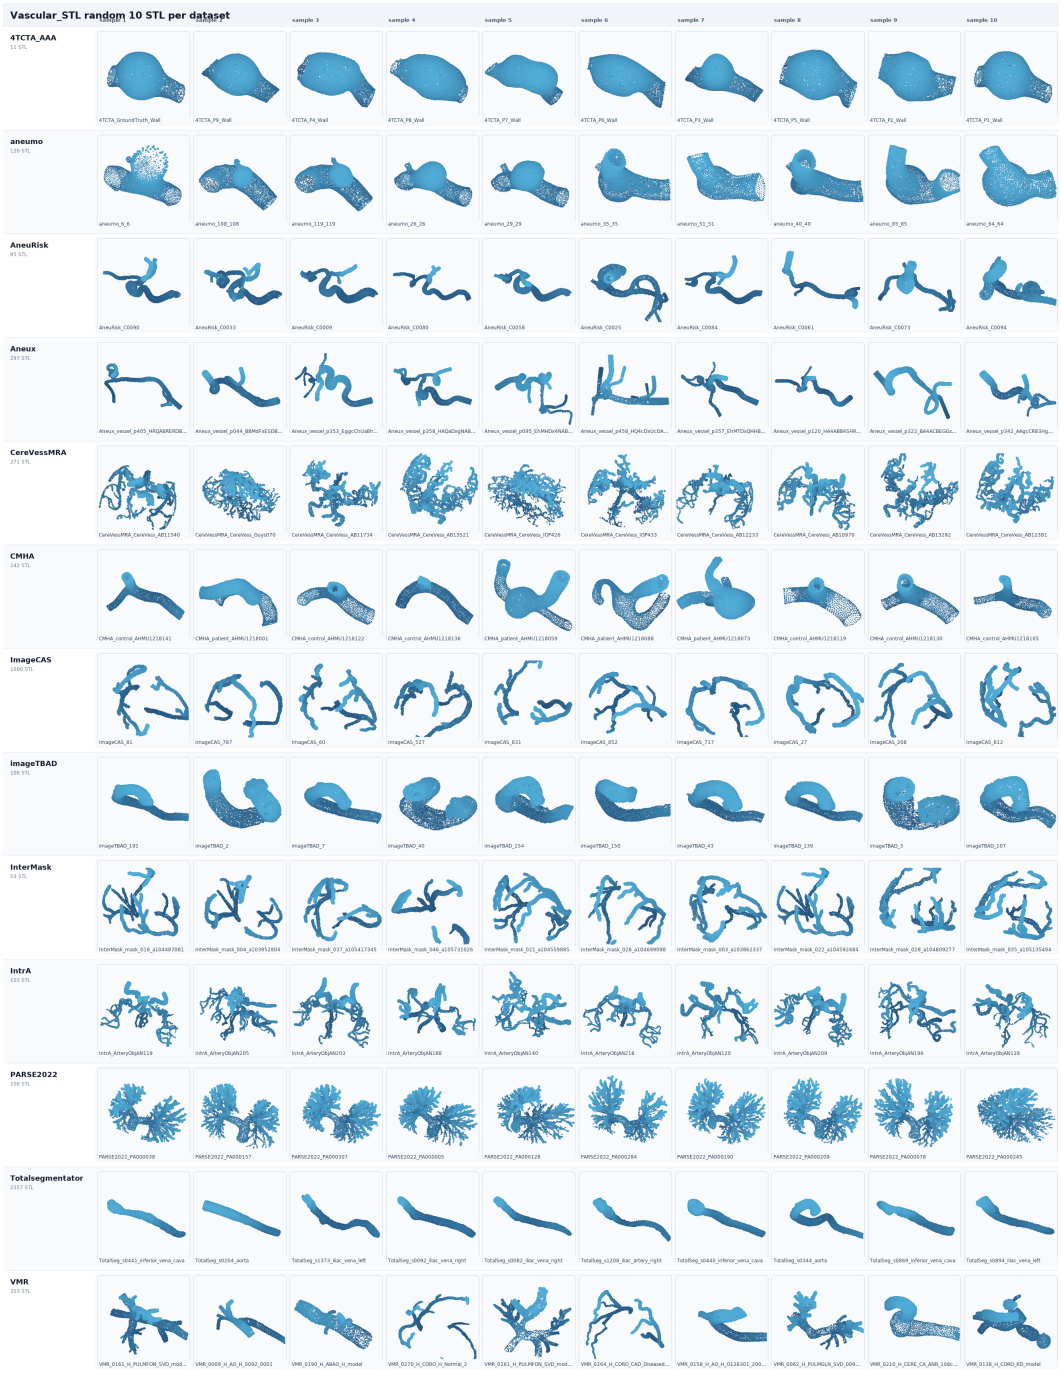


Figure A1: Randomly sampled vascular STL geometries after standardized geometry processing. Each row corresponds to one source dataset and contains ten randomly selected examples. The visualization demonstrates the cross-source, cross-anatomy, and cross-scale geometric diversity of the HemoPT corpus.

Tier 3: CFD-ready STL geometries. Tier 3 initially contains 1,433 simulation-ready or CFD-associated vascular geometries from Aneumo and VMR. These surfaces have greater processing maturity and support reliable interior sampling or numerical simulation.

Table A2: Construction of the final leakage-safe three-tier HemoPT pretraining corpus. The Reserved column denotes CFD-associated cases excluded from pretraining for downstream evaluation (540 from Aneumo, 48 from VMR-AOCOA, and 28 from VMR subdomains), while QC excluded denotes geometries removed because they fail geometry-validity requirements.

Tier	Collected	Reserved	QC excluded	Pretraining
Tier 1	3,881	0	25	3,856
Tier 2	648	0	9	639
Tier 3	1,433	616	2	815
Total	**5,962**	**616**	**36**	**5,310**

Before pretraining, 616 CFD-associated geometries assigned to downstream evaluation are reserved to prevent case-level leakage (comprising 540 from Aneumo, 48 from VMR-AOCOA for primary benchmarking, and 28 from VMR for subdomain transfer analysis). Two additional geometries are excluded by quality control, leaving 815 Tier-3 geometries for HemoPT pretraining. When paired CFD velocity or pressure fields are available, they are not used as HemoPT pretraining supervision.

The three tiers describe geometry provenance and processing maturity rather than an absolute ranking of anatomical validity or downstream task difficulty. Their complementary anatomical coverage, surface characteristics, and corpus scales provide heterogeneous structural priors for flow-aware vascular representations.

The final pretraining corpus therefore consists of 3,856 Tier-1 geometries, 639 Tier-2 geometries, and 815 Tier-3 geometries, for a total of 5,310 vascular STL geometries.

A.3 Geometry Quality Control and Component Cleaning

Raw vascular STL geometries cannot be directly pooled for pretraining because different segmentation and surface-reconstruction pipelines introduce source-dependent artifacts. Common artifacts include disconnected fragments, small isolated components, invalid surface connectivity, open boundaries, and numerical failures during lumen inside–outside queries.

Before constructing HemoPT proxy supervision, all 5,962 collected meshes undergo a standardized, automated quality-control (QC) and geometry-processing pipeline implemented via `vascular_stl_qc.py` and `clean_disconnected_vascular_stl.py`.

Automated quality-control criteria and exclusions. The QC protocol evaluates topological manifold properties, degenerate elements, and volume consistency. A mesh is strictly excluded if it satisfies any of the following rejection criteria:

- **File corruption or unreadable geometry:** failure to load via Trimesh as a valid surface mesh;
- **Insufficient mesh complexity:** vertex count $N_{\text{vert}} < 100$ or triangular face count $N_{\text{face}} < 200$, indicating incomplete or empty reconstructions;
- **Non-finite coordinates:** presence of `NaN` or `Inf` vertex coordinates;
- **Zero surface area or degenerate bounding box:** surface area $A \leq 0$ or maximum bounding box extent $\max(\Delta x, \Delta y, \Delta z) \leq 10^{-4}$ in native coordinates;
- **Severe fragmentation:** the largest connected component accounts for less than 20% of the total surface faces ($\text{faces}_{\text{largest}}/N_{\text{face}} < 0.20$), indicating catastrophic segmentation failure.

Across the complete geometry manifest, exactly 36 geometries fail these quality requirements, comprising 25 in Tier 1, 9 in Tier 2, and 2 in Tier 3, and are permanently excluded. These QC exclusions are strictly separated from the 616 CFD-associated cases reserved for downstream evaluation (540 from Aneumo, 48 from VMR-AOCOA, and 28 from VMR subdomains).

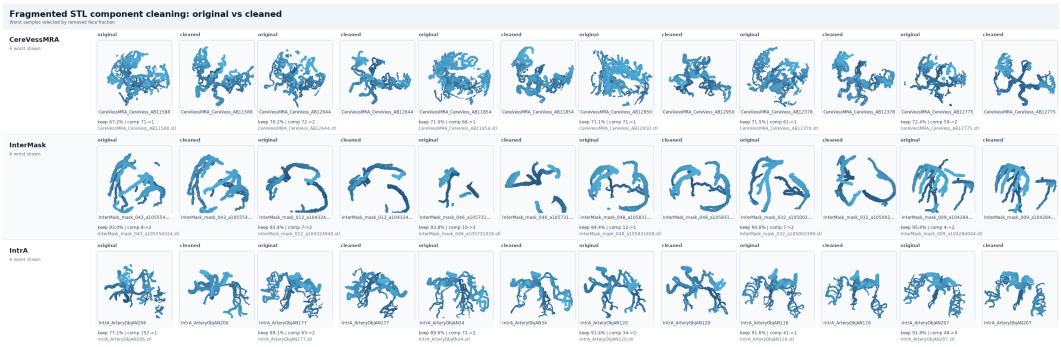


Figure A2: Representative worst-case examples of fragmented STL component cleaning, selected according to removed-face fraction. Each original–cleaned pair illustrates how disconnected or isolated surface components are removed while the principal vascular structure is retained for subsequent interior sampling and proxy construction.

Connected-component cleaning and surface regularization. Many clinically derived geometries, notably from CereVessMRA, InterMask, and IntrA, contain hundreds of spurious floating debris, isolated islands, or segmentation leakage voxels that would distort lumen inside–outside testing and distance fields. We employ a fast graph-based connected-component analysis constructed over the sparse vertex-adjacency matrix of each mesh after merging coincident vertices:

- For `CereVessMRA`, which contains dense segmentation debris, we retain only the principal connected component and remove components with fewer than 2,000 faces or whose face count is below 5% of the principal structure;
- For `InterMask` and `IntrA`, where secondary arterial branches may form distinct topological components, we preserve all major vascular trees while discarding isolated debris with fewer than 1,500 faces and smaller than 5% of the primary component;
- For open boundaries resulting from planar clipping at vascular inlets and outlets, automated boundary-loop detection and mesh capping are applied where necessary to enable watertight inside–outside query evaluation via generalized winding numbers and ray-casting parity.

Figure A2 presents representative worst-case component-cleaning examples selected according to their removed-face fractions. Rather than displaying only favorable examples, the figure highlights severely fragmented surfaces and demonstrates that the cleaning procedure removes disconnected or isolated components while preserving the principal vascular topology.

A.4 Leakage-Safe Pretraining Corpus Construction

Downstream leakage is controlled before constructing the proxy-supervised pretraining corpus. CFD-associated geometries assigned to downstream evaluation are reserved at the case level and excluded from the HemoPT pretraining pool.

Cases are matched using source-level dataset identifiers and case identifiers, preventing the same registered geometry from appearing in both pretraining and downstream evaluation. When multiple files or processed representations are associated with the same registered case, they are assigned consistently to the same pretraining or downstream partition.

This procedure primarily affects Tier 3. Of the 1,433 CFD-ready geometries, 616 are reserved for downstream evaluation (comprising 540 from Aneumo, 48 from VMR-AOCOA for primary benchmarking, and 28 from VMR for subdomain transfer analysis), two are excluded by geometry quality control, and the remaining 815 are retained for geometry-only proxy-supervised pretraining.

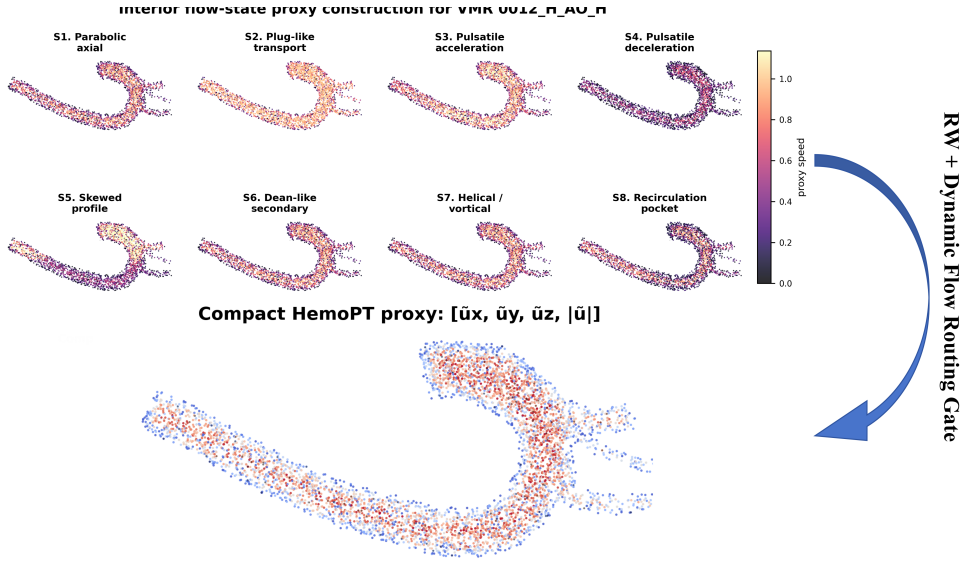


Figure A3: Representative HemoPT proxy construction for the processed VMR case `0012_H_AO_H`. After geometry quality control and interior sampling, the vascular surface is converted into eight CFD-inspired flow-state components, a compact proxy representation, and a walk-conditioned interior proxy mixture used for HemoPT pretraining.

The associated CFD velocity and pressure fields are not accessed during HemoPT pretraining. The complete downstream case-level splits, training/validation folds, and evaluation protocols are provided in Appendix B.

A.5 From Processed Geometry to Proxy-Ready Samples

After quality control, component cleaning, and downstream leakage filtering, each retained STL is converted into HemoPT-ready interior learning samples. The processed vascular surface first supports stable lumen inside–outside queries and interior point sampling.

Each sampled point is subsequently associated with local vascular geometry, wall-related descriptors, and directional random-walk probe conditions. These geometry- and probe-conditioned samples are transformed into a spatial traversal target and a compact flow-state proxy target without requiring paired CFD fields.

This step connects heterogeneous geometry processing with HemoPT proxy-supervised pretraining. The output of Appendix A is therefore not merely a cleaned STL collection, but a leakage-safe set of model-ready vascular geometries from which consistent interior traversal and flow-state proxy supervision can be generated.

Figure A3 presents a representative processed proxy construction for the VMR case `0012_H_AO_H`. It visualizes the eight CFD-inspired flow-state components, the compact four-dimensional proxy representation, and the resulting walk-conditioned mixture inside the processed vascular lumen.

The complete mathematical derivations of the canonical coordinate frames, the eight canonical flow-state primitives, the random-walk trajectory encoder, the dynamic routing gate, and the self-supervised physical regularization objectives are presented in detail in Appendix B.

B Mathematical Formulations, Model Architecture, and Experimental Protocols

This section specifies the mathematical formulation, architecture, and evaluation protocol of HemoPT. It covers the canonical reference frame, local geometric and probe parameterizations, eight canonical flow primitives, the Random-Walk Trajectory Encoder, the Dynamic Routing Gate, and the self-supervised physical regularization losses described in Subsections B.1–B.5. It then gives the pretraining configuration, Transolver Physics-Attention architecture, and initialization baselines in Subsections B.6–B.7. Finally, it defines the downstream regression task, case-level cohort partitioning, fine-tuning schedules, relative ℓ_2 metrics, and gain calculations in Subsections B.8–B.11. Comprehensive benchmark results across training fractions, post-hoc audits of proxy–CFD velocity alignment, routing entropy, spatial adaptivity, and mechanistic ablations are reported in Appendix C.

B.1 Canonical Reference Frame and Parameterization

For each vascular geometry, we construct an orthogonal frame $\mathcal{F}_k = (\mathbf{e}_{\parallel,k}, \mathbf{e}_{\perp,k}, \mathbf{e}_{\theta,k})$ by applying principal component analysis to its sampled point set.

Let $\mathbf{P} = [\mathbf{p}_1, \ldots, \mathbf{p}_N]^\top \in \mathbb{R}^{N\times 3}$ denote the set of sampled interior and boundary points for geometry k. The centroid is $\boldsymbol{\mu}_k = \frac{1}{N}\sum_{i=1}^{N} \mathbf{p}_i$. The centered spatial covariance matrix is

$$\mathbf{C}_k = \frac{1}{N}\sum_{i=1}^{N}(\mathbf{p}_i - \boldsymbol{\mu}_k)(\mathbf{p}_i - \boldsymbol{\mu}_k)^\top \in \mathbb{R}^{3\times 3}. \tag{A1}$$

We compute the spectral decomposition $\mathbf{C}_k = \mathbf{V}_k \boldsymbol{\Lambda}_k \mathbf{V}_k^\top$, where eigenvalues are sorted in descending order $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$, with corresponding orthonormal eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$. Because eigenvectors are defined up to an arbitrary sign reflection $\pm\mathbf{v}$, we enforce strict sign determinism by requiring the largest-magnitude component of each eigenvector to be strictly positive:

$$\mathbf{v}_j \leftarrow \operatorname{sign}\left(v_{j,\operatorname{argmax}_l |v_{j,l}|}\right) \cdot \mathbf{v}_j, \qquad j \in \{1,2,3\}. \tag{A2}$$

The right-handed canonical coordinate frame $\mathcal{F}_k$ is then established as:

$$\mathbf{e}_{\parallel,k} = \mathbf{v}_1 \quad \text{primary principal axis,} \tag{A3}$$

$$\mathbf{e}_{\perp,k} = \mathbf{v}_2 \quad \text{secondary principal axis,} \tag{A4}$$

$$\mathbf{e}_{\theta,k} = \mathbf{e}_{\parallel,k} \times \mathbf{e}_{\perp,k} \quad \text{transverse binormal axis.} \tag{A5}$$

This frame is derived from vascular morphology and provides a consistent reference for constructing directional flow tendencies without manual inlet or outlet annotation.

Coordinate Frame Interpretation in Complex Branching Trees. The frame $\mathcal{F}_k$ serves as a geometry-derived reference for the primitive dictionary rather than as a rigid streamline constraint. For branching structures with diverse orientations, the model adapts the primitive mixture through 20 probe directions $\mathbf{d}_i \in \mathbb{S}^2$ and the Dynamic Routing Gate R_ψ, as described in Appendix B.4. The gate conditions the primitive responses on the local wall direction $\mathbf{r}_{w,i}$ and multi-step trajectory features $\mathbf{h}_i^{\mathrm{RW}}$. The sign convention in Eq. (A2) makes the frame construction deterministic for a given point representation.

B.2 Local Geometric and Probe Context Parameterization

For an interior query point $\mathbf{p}_i \in \Omega_k$, let $d_{w,i} = \min_{\mathbf{b}\in\partial\Omega_k} \|\mathbf{b} - \mathbf{p}_i\|_2$ denote the Euclidean distance to the nearest wall point $\mathbf{b}_i$. We define the corresponding outward wall direction as $\mathbf{r}_{w,i} = (\mathbf{b}_i - \mathbf{p}_i)/(d_{w,i} + \epsilon)$ and form the local morphological descriptor $\mathbf{m}_i = [\mathbf{p}_i; d_{w,i}; \mathbf{r}_{w,i}] \in \mathbb{R}^7$.

We compute the longitudinal projection of the centered position $\mathbf{p}_i - \boldsymbol{\mu}_k$ onto the primary axis as $\widetilde{\zeta}_i = (\mathbf{p}_i - \boldsymbol{\mu}_k) \cdot \mathbf{e}_{\parallel,k}$. To ensure dimensionless invariance across vascular geometries of differing physical dimensions, this coordinate is normalized by the geometry-level axial span:

$$\zeta_i = \operatorname{clip}\left(\frac{\widetilde{\zeta}_i - \zeta_{\min,k}}{\Delta\zeta_k}, 0.0, 1.0\right) \in [0,1], \tag{A6}$$

where $\zeta_{\min,k} = \min_j \widetilde{\zeta}_j$ and $\Delta\zeta_k = \max_j \widetilde{\zeta}_j - \zeta_{\min,k}$. The radial offset from the canonical centerline approximation and its unit direction are given by

$$\mathbf{r}_i = (\mathbf{p}_i - \boldsymbol{\mu}_k) - \widetilde{\zeta}_i \mathbf{e}_{\parallel,k}, \qquad \widehat{\mathbf{r}}_i = \frac{\mathbf{r}_i}{\|\mathbf{r}_i\|_2 + \epsilon}. \tag{A7}$$

The azimuthal swirl direction is defined as

$$\mathbf{e}_{\text{swirl},i} = \mathbf{e}_{\parallel,k} \times \widehat{\mathbf{r}}_i. \tag{A8}$$

To provide scale invariance across geometries of varying caliber, wall distance is normalized by the 95th percentile wall distance $d_{95,k}$ across geometry k:

$$\rho_i = \text{clip}\left(\frac{d_{w,i}}{d_{95,k}}, 0.0, 1.0\right) \in [0,1], \qquad w_i = 1.0 - \rho_i \in [0,1], \tag{A9}$$

where ρ_i denotes the normalized core depth, with $\rho_i \to 0$ at the wall and $\rho_i \to 1$ at the lumen core, and w_i denotes wall proximity.

Relative Depth Parameterization across Varying Calibers. Normalizing wall distance by $d_{95,k}$ maps each interior point to a relative core-to-boundary depth coordinate ρ_i. This scaling accounts for caliber differences between proximal conduits and narrower branches. The unnormalized wall distance $d_{w,i}$ and outward wall direction $\mathbf{r}_{w,i}$ remain in $\mathbf{m}_i$ and are provided to the routing gate and backbone token embeddings, preserving local boundary information.

For a random-walk view with point-wise unit direction $\mathbf{d}_i \in \mathbb{S}^2$ and step length $s_i \in [0,2]$, the probe condition is $\mathbf{c}_i = [\mathbf{d}_i; s_i] \in \mathbb{R}^4$. We parameterize the interaction between the probe traversal and the local anatomy via:

$$a_i = \mathbf{d}_i \cdot \mathbf{e}_{\parallel,k} \in [-1,1] \quad \text{axial probe alignment,} \tag{A10}$$

$$\mathbf{d}_i^{\perp} = \mathbf{d}_i - a_i \mathbf{e}_{\parallel,k} \quad \text{transverse probe component,} \tag{A11}$$

$$l_i = \text{clip}\left(\widehat{\mathbf{r}}_i \cdot \frac{\mathbf{d}_i^{\perp}}{\|\mathbf{d}_i^{\perp}\|_2 + \epsilon}, -1.0, 1.0\right) \quad \text{lateral walk alignment,} \tag{A12}$$

$$\phi_i = \sin(\pi\zeta_i) \quad \text{periodic longitudinal phase.} \tag{A13}$$

B.3 Mathematical Formulations of Eight Canonical Flow Primitives

HemoPT defines eight continuous, bounded analytic velocity responses $\mathbf{u}_{i,m} = \mathcal{U}_m(S_k, \mathbf{m}_i, \mathbf{c}_i) \in \mathbb{R}^3$ for $m \in \{1, \ldots, 8\}$. These primitives provide canonical velocity templates for internal flow tendencies:

1. Parabolic Axial Flow $\mathcal{U}_1$. The parabolic axial primitive models a laminar profile with quadratic decay toward the vessel wall:

$$\mathbf{u}_{i,1} = (2\rho_i - \rho_i^2)\mathbf{e}_{\parallel,k}. \tag{A14}$$

2. Plug Core Flow $\mathcal{U}_2$. The plug-flow primitive has steep near-wall shear and a flattened core, capturing plug-like velocity tendencies in wide conduits:

$$\mathbf{u}_{i,2} = (1.0 - w_i^6)\mathbf{e}_{\parallel,k}. \tag{A15}$$

3. Stenotic Acceleration Flow $\mathcal{U}_3$. The stenotic-acceleration primitive combines amplified core speed with longitudinal phase modulation to represent localized convective acceleration:

$$u_{\text{mixed},i} = 0.65(2\rho_i - \rho_i^2) + 0.35(1.0 - w_i^6), \tag{A16}$$

$$\mathbf{u}_{i,3} = \text{clip}\left(u_{\text{mixed},i} \cdot (0.85 + 0.30\phi_i), 0.0, 1.25\right)\mathbf{e}_{\parallel,k}. \tag{A17}$$

4. Expansion Deceleration Flow $\mathcal{U}_4$. The expansion-deceleration primitive represents an axially decelerated profile in widening conduits:

$$\mathbf{u}_{i,4} = \left[0.70(2\rho_i - \rho_i^2) - 0.30 w_i \phi_i\right]\mathbf{e}_{\parallel,k}. \tag{A18}$$

5. Skewed-Profile Flow $\mathcal{U}_5$. The skewed-profile primitive is biased along the secondary axis $\mathbf{e}_{\perp,k}$ and represents lateral velocity skewing in curved conduits:

$$\kappa_i = \mathrm{clip}\left(1.0 + 0.40(\widehat{\mathbf{r}}_i \cdot \mathbf{e}_{\perp,k}), 0.50, 1.50\right), \tag{A19}$$

$$\mathbf{d}_{\mathrm{skew},k} = \frac{\mathbf{e}_{\parallel,k} + 0.18\mathbf{e}_{\perp,k}}{\|\mathbf{e}_{\parallel,k} + 0.18\mathbf{e}_{\perp,k}\|_2}, \tag{A20}$$

$$\mathbf{u}_{i,5} = (2\rho_i - \rho_i^2)\kappa_i \mathbf{d}_{\mathrm{skew},k}. \tag{A21}$$

6. Dean Secondary Flow $\mathcal{U}_6$. The Dean primitive adds a cross-sectional component along $\mathbf{e}_{\theta,k}$ to represent secondary transverse flow:

$$\mathbf{u}_{i,6} = \mathbf{u}_{i,1} + 0.28\rho_i w_i \phi_i (\widehat{\mathbf{r}}_i \cdot \mathbf{e}_{\perp,k}) \mathbf{e}_{\theta,k}. \tag{A22}$$

7. Helical Vortex Flow $\mathcal{U}_7$. The helical primitive adds an azimuthal component around the longitudinal axis to represent rotational flow:

$$\mathbf{u}_{i,7} = \mathbf{u}_{i,1} + 0.22\rho_i w_i (0.50 + 0.50\phi_i) \mathbf{e}_{\mathrm{swirl},i}. \tag{A23}$$

8. Recirculation Pocket Flow $\mathcal{U}_8$. The recirculation primitive combines localized axial reversal with radial displacement to encode a prior for recirculating patterns:

$$p_{\mathrm{pocket},i} = \exp\left(-\left(\frac{\zeta_i - 0.65}{0.22}\right)^2\right) \cdot w_i, \tag{A24}$$

$$\mathbf{u}_{i,8} = \left[(2\rho_i - \rho_i^2)(1.0 - 0.70p_{\mathrm{pocket},i}) - 0.22p_{\mathrm{pocket},i}\right] \mathbf{e}_{\parallel,k} + 0.10p_{\mathrm{pocket},i}\widehat{\mathbf{r}}_i. \tag{A25}$$

Role of Basis Coefficients and Dynamic Gating. The numerical scale factors in $\mathcal{U}_1$–$\mathcal{U}_8$ set the relative magnitudes of convective, secondary, and recirculating components. The eight primitives form a fixed inductive basis rather than standalone flow solutions. The localized Gaussian envelope in $\mathcal{U}_8$ provides a template for flow reversal, whose contribution is modulated by the Dynamic Routing Gate R_ψ according to local morphology $\mathbf{m}_i$.

B.4 Random-Walk Trajectory Encoder and Dynamic Routing Gate

The flow proxy combines the eight primitives using mixture weights predicted by a Random-Walk Trajectory Encoder E_{RW} and a Dynamic Flow Routing Gate R_ψ.

Random-Walk Trajectory Encoder E_{RW}. For each query point $\mathbf{p}_i$, the directional probe traversal produces an ordered sequence of three 3D wall-displacement offsets:

$$\mathbf{g}_i = [\mathbf{g}_i^{(0)}; \mathbf{g}_i^{(1)}; \mathbf{g}_i^{(2)}] \in \mathbb{R}^9, \qquad \mathbf{g}_i^{(t)} = \mathbf{q}_i^{(t)} - \mathbf{b}_i^{(t)} \in \mathbb{R}^3, \quad t \in \{0,1,2\}. \tag{A26}$$

We reshape the sequence into a three-by-three array and define a learnable step-position embedding $\mathbf{E}_{\mathrm{step}} \in \mathbb{R}^{3\times 16}$ initialized from $\mathcal{N}(0, 0.02^2)$. Each step displacement $\mathbf{g}_i^{(t)}$ is transformed by a shared step MLP:

$$\mathbf{h}_i^{(t)} = \mathrm{Linear}_{16\to 16}\left(\mathrm{GELU}\left(\mathrm{LayerNorm}\left(\mathrm{Linear}_{3\to 16}(\mathbf{g}_i^{(t)})\right)\right)\right) + \mathbf{E}_{\mathrm{step}}^{(t)} \in \mathbb{R}^{16}. \tag{A27}$$

To capture both aggregate displacement scale and extreme trajectory excursions, we apply combined mean- and max-pooling across the three traversal steps:

$$\mathbf{h}_i^{\mathrm{pool}} = \left[\frac{1}{3}\sum_{t=0}^{2} \mathbf{h}_i^{(t)}; \max_{t\in\{0,1,2\}} \mathbf{h}_i^{(t)}\right] \in \mathbb{R}^{32}. \tag{A28}$$

A final projection MLP outputs the 16D trajectory feature:

$$\mathbf{h}_i^{\mathrm{RW}} = \mathrm{GELU}\left(\mathrm{LayerNorm}\left(\mathrm{Linear}_{32\to 16}(\mathbf{h}_i^{\mathrm{pool}})\right)\right) \in \mathbb{R}^{16}. \tag{A29}$$

Dynamic Flow Routing Gate R_ψ. The routing gate receives a 24-dimensional combined representation consisting of local wall morphology $\mathbf{m}_i^w = [d_{w,i}; \mathbf{r}_{w,i}] \in \mathbb{R}^4$, the probe condition $\mathbf{c}_i = [\mathbf{d}_i; s_i] \in \mathbb{R}^4$, and the trajectory feature $\mathbf{h}_i^{\mathrm{RW}} \in \mathbb{R}^{16}$:

$$\mathbf{f}_i^{\mathrm{route}} = [\mathbf{m}_i^w;\ \mathbf{c}_i;\ \mathbf{h}_i^{\mathrm{RW}}] \in \mathbb{R}^{24}. \tag{A30}$$

The routing MLP maps $\mathbf{f}_i^{\mathrm{route}}$ through two hidden layers with GELU activations:

$$\mathbf{s}_i = \mathrm{Linear}_{32\to 8}\left(\mathrm{GELU}\left(\mathrm{Linear}_{32\to 32}\left(\mathrm{GELU}\left(\mathrm{Linear}_{24\to 32}(\mathbf{f}_i^{\mathrm{route}})\right)\right)\right)\right) \in \mathbb{R}^8. \tag{A31}$$

Applying the softmax function yields non-negative mixture weights that sum to unity:

$$\boldsymbol{\alpha}_i = \mathrm{softmax}(\mathbf{s}_i) \in \Delta^7, \qquad \sum_{m=1}^{8} \alpha_{i,m} = 1. \tag{A32}$$

Compact Four-Dimensional Flow Target. The routed primitive responses are softly fused into an instance-conditioned velocity vector:

$$\widetilde{\mathbf{u}}_i = \sum_{m=1}^{8} \alpha_{i,m} \mathbf{u}_{i,m} \in \mathbb{R}^3. \tag{A33}$$

The 4D compact proxy target $\mathbf{z}_i$ appends the Euclidean speed magnitude, recomputed directly from the synthesized velocity vector:

$$\mathbf{z}_i = [\widetilde{\mathbf{u}}_i;\ \|\widetilde{\mathbf{u}}_i\|_2] \in \mathbb{R}^4. \tag{A34}$$

B.5 Self-Supervised Physical Regularization and Joint Objective

During pretraining, the routing modules E_{RW} and R_ψ are trained with three self-supervised consistency objectives evaluated on the synthesized velocity field $\widetilde{\mathbf{u}}_i$:

1. Wall No-Slip Boundary Suppression $\mathcal{L}_{\mathrm{noslip}}$. To penalize nonzero velocity near vessel boundaries, we apply exponential distance damping:

$$\mathcal{L}_{\mathrm{noslip}} = \frac{1}{N} \sum_{i=1}^{N} \left(\|\widetilde{\mathbf{u}}_i\|_2 \cdot \exp\left(-\frac{\max(d_{w,i}, 0)}{\delta_{\mathrm{wall}}} \right) \right)^2, \tag{A35}$$

where $\delta_{\mathrm{wall}} = 0.05$ in normalized coordinates. This penalty suppresses high-speed responses near the wall and encourages boundary-layer damping.

2. Local Moving Least Squares Divergence Penalty $\mathcal{L}_{\mathrm{div}}$. To promote spatial continuity and discourage spurious divergence on unstructured point clouds, we estimate the local velocity gradient tensor via Moving Least Squares. We randomly subsample $S \subset \{1, \ldots, N\}$ points, bounded by $|S| \le 256$ per batch item. For each point $i \in S$, we identify its $k = 8$ nearest spatial neighbors. Let $\Delta\mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i$ and $\Delta\widetilde{\mathbf{u}}_{ij} = \widetilde{\mathbf{u}}_j - \widetilde{\mathbf{u}}_i$. Using distance-dependent weights $w_{ij} = \exp(-\|\Delta\mathbf{p}_{ij}\|_2/\sigma_i)$, with σ_i equal to the mean neighbor distance, we compute the local spatial covariance and cross-covariance matrices:

$$\mathbf{C}_i = \sum_{j \in \mathcal{N}_k(i)} w_{ij} \Delta\mathbf{p}_{ij} \Delta\mathbf{p}_{ij}^\top + 10^{-4}\mathbf{I}_3, \qquad \mathbf{B}_i = \sum_{j \in \mathcal{N}_k(i)} w_{ij} \Delta\widetilde{\mathbf{u}}_{ij} \Delta\mathbf{p}_{ij}^\top. \tag{A36}$$

The local Jacobian tensor is $\mathbf{J}_i = \mathbf{B}_i \mathbf{C}_i^{-1} \in \mathbb{R}^{3\times 3}$. The divergence penalty minimizes the squared trace of $\mathbf{J}_i$:

$$\mathcal{L}_{\mathrm{div}} = \frac{1}{|S|} \sum_{i \in S} \left(\mathrm{Tr}(\mathbf{J}_i)\right)^2 = \frac{1}{|S|} \sum_{i \in S} \left(J_{i,11} + J_{i,22} + J_{i,33}\right)^2. \tag{A37}$$

3. Routing Shannon Entropy Regularization $\mathcal{L}_{\mathrm{entropy}}$. We regularize the point-wise routing distribution to discourage both concentration on a single mode and near-uniform weighting. For query point $\mathbf{p}_i$, the Shannon entropy across the eight modes is

$$H(\boldsymbol{\alpha}_i) = -\sum_{m=1}^{8} \alpha_{i,m} \ln(\alpha_{i,m} + 10^{-8}). \tag{A38}$$

Table A3: Main HemoPT pretraining configuration and Transolver backbone hyperparameters.

Hyperparameter / Setting	**Value**
Pretraining corpus scale	5,310 leakage-safe STL geometries
Corpus split	5,044 train / 266 validation
Split protocol	Geometry-level partition before probe generation
Random-walk views R	20 views per geometry, totaling 106,200 instances
Input token dimension	11D, comprising $[\mathbf{p}_i; d_{w,i}; \mathbf{r}_{w,i}; \mathbf{d}_i; s_i]$
Target token dimension	13D, comprising 9D spatial and 4D flow targets
Backbone architecture	Transolver Encoder with Physics-Attention
Encoder layers / Attention heads	8 layers / 8 heads
Hidden dimension / MLP ratio	256 / 2
Physical slice count M_{slice}	32 slices
Wall channel masking probability p_{mask}	0.25, masking $d_{w,i}, \mathbf{r}_{w,i}$ on input
Flow proxy supervision weight λ_{flow}	0.25
Optimizer / Weight decay	AdamW / 10^{-4}
Learning rate / Schedule	1×10^{-3} / Cosine annealing with 10 warmup epochs
Batch size / Pretraining budget	8 / 250 epochs with early stopping patience 75

Unconstrained entropy maximization can flatten the weights into a uniform average and weaken local geometric conditioning. We therefore penalize deviations from the intermediate target entropy $H_{\text{target}} = 0.75 \ln 8 \approx 1.56$:

$$\mathcal{L}_{\text{entropy}} = \frac{1}{N} \sum_{i=1}^{N} \left(H(\boldsymbol{\alpha}_i) - H_{\text{target}} \right)^2 . \tag{A39}$$

Minimizing $\mathcal{L}_{\text{entropy}}$ encourages intermediate routing diversity while retaining geometry-conditioned primitive selection. We monitor the mean point-wise routing entropy as $H_{\text{gate}} = \frac{1}{N} \sum_{i=1}^{N} H(\boldsymbol{\alpha}_i)$.

Total Pretraining Objective. The total self-supervised routing loss is

$$\mathcal{L}_{\text{dict}} = w_{\text{noslip}} \mathcal{L}_{\text{noslip}} + w_{\text{div}} \mathcal{L}_{\text{div}} + w_{\text{entropy}} \mathcal{L}_{\text{entropy}}, \tag{A40}$$

with loss weights $w_{\text{noslip}} = 0.10$, $w_{\text{div}} = 0.05$, and $w_{\text{entropy}} = 0.01$. The complete HemoPT pretraining objective combines spatial traversal and flow proxy reconstruction:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{spa}} + \lambda_{\text{flow}} \left(\mathcal{L}_{\text{proxy}} + \mathcal{L}_{\text{dict}} \right), \tag{A41}$$

where $\lambda_{\text{flow}} = 0.25$ weights the joint flow branch, and the stop-gradient operator $\text{sg}(\mathbf{z}_i)$ is applied to the synthesized target when supervising the Transolver network.

B.6 Pretraining Configuration and Transolver Backbone

The pretraining backbone is a Transolver encoder with Physics-Attention. The network projects irregular point tokens into $M_{\text{slice}} = 32$ learned physical slices, computes global self-attention across physical slices, and reconstructs point-wise features through slice-to-point interpolation. Table A3 reports the detailed pretraining hyperparameter configuration. The primary HemoPT model is trained with a full budget of 250 epochs with cosine annealing and an early stopping patience of 75 epochs. As detailed in Appendix C, the extensive ablation series and parameter sweeps employ a standardized fixed budget of 200 epochs for controlled comparison across configurations.

B.7 Compared Pretraining Initializations

Table A4 summarizes the seven pretraining configurations evaluated across downstream benchmarks. Point-BERT and Point-MAE are evaluated with either generic ShapeNet weights or weights pretrained directly on our vascular surface corpus. GeoPT-Vasc shares the Transolver backbone and pretraining corpus with HemoPT-Vasc, but utilizes directional random-walk traversal supervision without the hemodynamic flow proxy. All compared models share identical downstream network heads, input point resolutions, and loss functions during fine-tuning.

Table A4: Pretraining initializations used in downstream comparisons.

Method	Initialization source
Scratch	Random initialization
Point-BERT-SN	Point-BERT pretrained on ShapeNet-style point clouds
Point-MAE-SN	Point-MAE pretrained on ShapeNet-style point clouds
Point-BERT-Vasc	Point-BERT pretrained on vascular surface point clouds
Point-MAE-Vasc	Point-MAE pretrained on vascular surface point clouds
GeoPT-Vasc	Geometry traversal and random-walk pretraining
HemoPT-Vasc	Geometry traversal and flow-state proxy pretraining

Table A5: Downstream case-level evaluation splits. Independent test cases remain fixed, while five training/validation folds are constructed from the remaining cases.

Dataset	Total	Train	Val	Test	Source
Aneumo	540	345	87	108	CFD
VMR-AOCOA	48	37	5	6	CFD
ICA-Sim	133	89	23	21	CFD
4DFlow	138	99	25	14	Imaging
Total	**859**	**570**	**140**	**149**	—

B.8 Downstream Evaluation Protocol

Each downstream benchmark is evaluated as a supervised regression task predicting the 3D velocity vector at 4,096 interior query points:

$$\mathbf{v}_i = (u_i, v_i, w_i). \tag{A42}$$

As detailed in Table A5, the primary downstream evaluation benchmark comprises 859 total cases and 3,518,464 interior query points. Each dataset is partitioned into an independent held-out test cohort containing 149 cases and five matched training/validation folds constructed from the remaining 710 cases. Each fold contains 570 training cases and 140 validation cases. The test splits remain fixed across all comparisons and training-budget regimes.

Standardized Downstream Benchmark and Secondary VMR Subdomains. The primary benchmark contains 859 clinical vascular cases and 3,518,464 interior query points across four cohorts. An additional 28 quality-filtered VMR cases from peripheral pulmonary and coronary branches are reserved for the secondary subdomain-transfer analysis in Appendix C.7. The CFD-associated VMR evaluation cohort comprises 76 cases in total (48 AOCOA and 28 subdomain cases). These 28 cases are excluded from the primary benchmark and zero-shot diagnostic audits so that all primary comparisons use the same cohort definition.

B.9 Downstream Optimization

Downstream fine-tuning follows a unified optimization protocol summarized in Table A6. Models are trained for up to 200 epochs using AdamW with cosine learning rate scheduling and early stopping. Batch sizes are set to 4 for Aneumo, ICA-Sim, and 4DFlow, while VMR-AOCOA uses a reduced batch size to accommodate its smaller cohort size. All initializations share identical fine-tuning schedules, point samplings, and loss criteria.

Table A6: Downstream training configuration.

Item	Setting
Task	Velocity-field regression
Input points	4,096 interior points per case
Prediction target	Velocity vector with components u, v, and w
Training budget	Maximum 200 epochs with early stopping
Default batch size	4
VMR-AOCOA batch	Dataset-specific reduced size
Evaluation	Fixed test cohort and five training/validation folds

B.10 Evaluation Metrics

We report two complementary relative ℓ_2 metrics for velocity-field prediction: the velocity-vector relative ℓ_2 error and the velocity-magnitude relative ℓ_2 error.

Let

$$\mathbf{V} = [\mathbf{v}_1, \ldots, \mathbf{v}_N]^\top \in \mathbb{R}^{N\times 3} \tag{A43}$$

and

$$\widehat{\mathbf{V}} = [\widehat{\mathbf{v}}_1, \ldots, \widehat{\mathbf{v}}_N]^\top \in \mathbb{R}^{N\times 3} \tag{A44}$$

denote the reference and predicted velocity fields at the N interior query points, respectively.

The velocity-vector relative ℓ_2 error is defined as

$$E_{\text{vector}} = \frac{\left\|\widehat{\mathbf{V}} - \mathbf{V}\right\|_{\text{F}}}{\|\mathbf{V}\|_{\text{F}}}, \tag{A45}$$

where $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm. This metric jointly evaluates errors in the direction and amplitude of the three-dimensional velocity field.

For each query point, the reference and predicted velocity magnitudes are

$$m_i = \|\mathbf{v}_i\|_2, \qquad \widehat{m}_i = \|\widehat{\mathbf{v}}_i\|_2. \tag{A46}$$

Let $\mathbf{m}$ and $\widehat{\mathbf{m}}$ denote the corresponding reference and predicted velocity-magnitude vectors. The velocity-magnitude relative ℓ_2 error is defined as

$$E_{\text{magnitude}} = \frac{\|\widehat{\mathbf{m}} - \mathbf{m}\|_2}{\|\mathbf{m}\|_2}. \tag{A47}$$

This metric evaluates the recovery of flow amplitude independently of the orientation of the velocity vectors.

Both errors are reported as mean$\pm$standard deviation over the five matched training/validation folds evaluated on the fixed independent test cohort.

B.11 Metric-Specific Gain and Average Gain

For each metric $q \in \{\text{vector}, \text{magnitude}\}$, let $\overline{E}_q^{\text{Scratch}}$ and $\overline{E}_q^{\text{Method}}$ denote the corresponding unrounded mean test errors over the five matched folds.

Table A7: Downstream metric and gain-reporting protocol.

Item	Definition
Velocity-vector metric	Velocity-vector relative ℓ_2 error
Velocity-magnitude metric	Velocity-magnitude relative ℓ_2 error
Reported errors	Mean$\pm$standard deviation over five folds
Metric-specific gain	Improvement over Scratch
Final summary	Average Gain over vector and magnitude gains
Gain calculation	Unrounded mean test errors

The metric-specific gain over Scratch is defined as

$$G_q(\%) = 100 \left(\frac{\overline{E}_q^{\mathrm{Scratch}} - \overline{E}_q^{\mathrm{Method}}}{\overline{E}_q^{\mathrm{Scratch}}} \right). \tag{A48}$$

The final Average Gain used in the pretraining-scale, tier-composition, and VMR subdomain analyses is

$$\text{Average Gain}(\%) = \frac{1}{2} \left(G_{\mathrm{vector}} + G_{\mathrm{magnitude}} \right). \tag{A49}$$

A positive metric-specific gain indicates that the pretrained initialization reduces the corresponding error relative to Scratch. A negative value indicates negative transfer. Average Gain summarizes the transfer benefit across both velocity-vector and velocity-magnitude prediction. The complete downstream evaluation metric and gain-reporting protocol are summarized in Table A7.

C Extended Analyses, Diagnostic Auditing, and Empirical Ablations

This section provides extended empirical analyses, physical diagnostic audits, and architectural sensitivity evaluations for the HemoPT pretraining framework. The presentation begins with the mathematical definitions of the diagnostic evaluation metrics, followed by seven systematic evaluations:

1. **Comprehensive Downstream Benchmark Evaluation**, Subsection C.1: Detailed version of main paper Table 1 reporting five-fold means and standard deviations across four downstream benchmarks and four labeled-data supervision regimes.
2. **Flow Proxy and Physical Consistency Loss Ablations**, Subsection C.2: Detailed version of main paper Table 2 reporting zero-shot proxy alignment across all four clinical cohorts and pretraining regularization residuals.
3. **Pretraining Dynamic Routing Evolution and Spatial Adaptivity Auditing**, Subsection C.3: Pretraining convergence dynamics, two-hundred-epoch routing trajectory curves, sampled-epoch dictionary progression, and granular spatial adaptivity controls.
4. **Hyperparameter Sweeps and Flow Proxy Weight Sensitivity**, Subsection C.4: Pretraining sensitivity sweeps over wall masking ratios, random walk traversal steps, directional probe view counts, and downstream flow-proxy supervision weight λ_{flow}.
5. **Controlled Pretraining Study on Aneumo**, Subsection C.5: Granular analysis of downstream supervision budget efficiency, pretraining epoch progression, and mechanistic divergence from surface coordinate masking baselines.
6. **Pretraining Corpus Scalability Dynamics and Multi-Tier Synergy**, Subsection C.6: Pretraining volume scaling trajectories and cross-tier morphological synergy between clinical segmentations and simulation meshes.

7. **Physical Boundaries of Geometry-Conditioned Transfer**, Subsection C.7: Territory-specific analysis across conduit convective momentum domains versus peripheral microvascular resistance regimes.

Diagnostic Metric Formulation

For a test vascular geometry with N interior query points, let $\widetilde{\mathbf{u}}_i \in \mathbb{R}^3$ denote the dynamically routed proxy velocity vector and $\mathbf{u}_i \in \mathbb{R}^3$ denote the reference fluid velocity vector obtained from computational fluid dynamics simulations or Phase-Contrast Magnetic Resonance Imaging acquisitions. Spatial and hemodynamic agreement is audited using three complementary diagnostic metrics:

1. **Directional Cosine Alignment $C_{\mathrm{dir}} \uparrow$:** Quantifies the point-wise angular agreement of the predicted flow streamline orientation independently of velocity magnitude:
$$C_{\mathrm{dir}} = \frac{1}{N} \sum_{i=1}^{N} \frac{\widetilde{\mathbf{u}}_i^{\top} \mathbf{u}_i}{\|\widetilde{\mathbf{u}}_i\|_2 \|\mathbf{u}_i\|_2 + \epsilon}, \tag{A50}$$
where $\epsilon = 10^{-8}$ prevents numerical singularity near static boundaries. Directional alignment satisfies $C_{\mathrm{dir}} \in [-1, 1]$, where 1.0 indicates perfect collinear streamline alignment.
2. **Relative Velocity Magnitude Error $C_{\mathrm{mag}} \downarrow$:** Measures the aggregate discrepancy between synthesized proxy speed profiles and reference fluid speeds across all interior lumen points:
$$C_{\mathrm{mag}} = \frac{\sum_{i=1}^{N} \left| \|\widetilde{\mathbf{u}}_i\|_2 - \|\mathbf{u}_i\|_2 \right|}{\sum_{i=1}^{N} \|\mathbf{u}_i\|_2 + \epsilon}, \tag{A51}$$
where lower values denote closer recovery of ground-truth velocity magnitude.
3. **Relative Error Reduction ΔC_{mag}:** Quantifies the proportional reduction in velocity magnitude error achieved by a given proxy formulation relative to the classical Poiseuille approximation baseline:
$$\Delta C_{\mathrm{mag}} = \frac{C_{\mathrm{mag}}^{\mathrm{Pois}} - C_{\mathrm{mag}}}{C_{\mathrm{mag}}^{\mathrm{Pois}}} \times 100\%, \tag{A52}$$
where positive percentages reflect superior physical alignment over idealized parabolic flow.

C.1 Comprehensive Downstream Benchmark Evaluation Across Labeled-Data Regimes

Table A8 presents the detailed numerical results reported to four decimal places together with five-fold standard deviations across Aneumo, VMR-AOCOA, ICA-Sim, and 4DFlow. Each dataset is evaluated under four labeled-data fine-tuning regimes: 25%, 50%, 75%, and 100% of the available training cases. The table expands the rounded results in main paper Table 1.

Empirical Observations on Downstream Transfer. Key findings emerging from the full-precision downstream benchmark evaluations include:

- **Consistent Monotonic Superiority:** Across all sixteen evaluated settings, four clinical datasets across four data regimes, HemoPT-Vasc achieves top performance or closely follows the optimum. On ICA-Sim, HemoPT-Vasc improves average prediction gains from $+24.16\%$ at the 25% regime to $+31.64\%$ at 75% and $+29.29\%$ at 100%, demonstrating that flow-state pretraining accelerates model convergence across varying data availability.
- **Failure Modes of Surface Coordinate Masking:** General point-cloud masking baselines including Point-BERT and Point-MAE suffer severe negative transfer on internal carotid artery and aortic coarctation geometries. On VMR-AOCOA under 25% training data, Point-BERT-Vasc and Point-MAE-Vasc yield dramatic negative gains of -73.38% and -46.44%, respectively. Reconstructing irregular surface triangle coordinates biases model attention toward high-frequency boundary noise rather than volumetric convective momentum.

Table A8: Comprehensive downstream velocity-field prediction across four labeled-data regimes. Entries report mean ± standard deviation over five matched training-validation folds evaluated on the fixed independent test cohort. Lower errors indicate better predictions. Bold and underlined values mark the best and second-best results within each dataset and training budget.

Dataset	Method	25% Train vel-vec.	25% Train vel-mag.	25% Train avg. gain	50% Train vel-vec.	50% Train vel-mag.	50% Train avg. gain	75% Train vel-vec.	75% Train vel-mag.	75% Train avg. gain	100% Train vel-vec.	100% Train vel-mag.	100% Train avg. gain
Aneumo	Scratch	0.8123 ± 0.0432	0.6493 ± 0.0282	+0.00%	0.6628 ± 0.0309	0.5543 ± 0.0179	+0.00%	0.5603 ± 0.0045	0.4749 ± 0.0068	+0.00%	0.4700 ± 0.0156	0.4037 ± 0.0152	+0.00%
	Point-BERT-SN	0.6925 ± 0.0148	0.5816 ± 0.0076	+12.58%	0.6163 ± 0.0100	0.5401 ± 0.0056	+4.79%	0.5662 ± 0.0210	0.4966 ± 0.0249	-2.82%	0.4972 ± 0.0055	0.4269 ± 0.0039	-5.77%
	Point-MAE-SN	0.7317 ± 0.0085	0.6064 ± 0.0048	+8.26%	0.6387 ± 0.0076	0.5542 ± 0.0046	+1.82%	0.5776 ± 0.0172	0.4986 ± 0.0225	-4.04%	0.5091 ± 0.0038	0.4329 ± 0.0024	-7.78%
	Point-BERT-Vasc	0.7036 ± 0.0093	0.5881 ± 0.0028	+11.40%	0.6241 ± 0.0141	0.5448 ± 0.0072	+3.78%	0.5797 ± 0.0138	0.5083 ± 0.0187	-5.25%	0.5163 ± 0.0068	0.4440 ± 0.0069	-9.91%
	Point-MAE-Vasc	0.8499 ± 0.1107	0.6666 ± 0.0651	-3.65%	0.6742 ± 0.0107	0.5715 ± 0.0069	-2.41%	0.6240 ± 0.0064	0.5427 ± 0.0041	-12.82%	0.5969 ± 0.0022	0.5256 ± 0.0013	-28.59%
	GeoPT-Vasc	0.7040 ± 0.0325	0.5542 ± 0.0188	+13.99%	0.5525 ± 0.0202	0.4529 ± 0.0124	+17.47%	0.4810 ± 0.0078	0.4060 ± 0.0074	+14.32%	0.4485 ± 0.0041	0.3840 ± 0.0031	+4.73%
	HemoPT-Vasc	**0.6790 ± 0.0265**	**0.5369 ± 0.0219**	**+16.86%**	**0.5319 ± 0.0138**	**0.4371 ± 0.0113**	**+20.45%**	**0.4669 ± 0.0141**	**0.3945 ± 0.0086**	**+16.81%**	**0.4204 ± 0.0057**	**0.3589 ± 0.0048**	**+10.83%**
VMR-AOCOA	Scratch	**1.0264 ± 0.0518**	**0.9532 ± 0.0276**	+0.00%	1.1234 ± 0.1975	1.0375 ± 0.1226	+0.00%	1.0419 ± 0.1148	0.9627 ± 0.0641	+0.00%	1.0322 ± 0.0506	0.9563 ± 0.0129	+0.00%
	Point-BERT-SN	1.4843 ± 0.3098	1.3130 ± 0.2744	-41.18%	1.3190 ± 0.2322	1.1363 ± 0.1844	-13.46%	1.1827 ± 0.1942	1.0334 ± 0.1479	-10.43%	1.1388 ± 0.2007	0.9832 ± 0.1764	-6.57%
	Point-MAE-SN	1.4610 ± 0.2753	1.2801 ± 0.2482	-38.32%	1.3652 ± 0.2140	1.1446 ± 0.1748	-15.92%	1.3393 ± 0.1276	1.1414 ± 0.1109	-23.55%	1.3132 ± 0.1943	1.1261 ± 0.1792	-22.49%
	Point-BERT-Vasc	1.8214 ± 0.3687	1.6140 ± 0.3391	-73.38%	1.4557 ± 0.2895	1.2590 ± 0.2633	-25.46%	1.3215 ± 0.2617	1.1360 ± 0.2073	-22.42%	1.2524 ± 0.2824	1.0650 ± 0.2426	-16.35%
	Point-MAE-Vasc	1.5457 ± 0.4101	1.3564 ± 0.4153	-46.44%	1.2795 ± 0.2013	1.1031 ± 0.2000	-10.11%	1.2446 ± 0.2532	1.0895 ± 0.2695	-16.31%	1.2761 ± 0.1511	1.1411 ± 0.1764	-21.48%
	GeoPT-Vasc	1.1375 ± 0.0938	1.0179 ± 0.0501	-8.80%	1.0519 ± 0.1188	0.9550 ± 0.0839	+7.16%	1.0339 ± 0.0969	0.9691 ± 0.0487	+0.05%	1.0154 ± 0.0979	0.9506 ± 0.0646	+1.11%
	HemoPT-Vasc	1.0347 ± 0.0800	0.9567 ± 0.0389	**-0.59%**	**1.0114 ± 0.0531**	**0.9524 ± 0.0173**	**+9.09%**	**1.0230 ± 0.0702**	**0.9611 ± 0.0145**	**+0.99%**	**0.9823 ± 0.0163**	**0.9381 ± 0.0525**	**+3.37%**
ICA-Sim	Scratch	0.4641 ± 0.0156	0.3725 ± 0.0053	+0.00%	0.4307 ± 0.0095	0.3656 ± 0.0070	+0.00%	0.4238 ± 0.0077	0.3621 ± 0.0102	+0.00%	0.3813 ± 0.0222	0.3221 ± 0.0192	+0.00%
	Point-BERT-SN	0.3411 ± 0.0036	0.3013 ± 0.0029	+22.82%	0.3066 ± 0.0051	0.2738 ± 0.0033	+26.96%	0.2931 ± 0.0082	0.2632 ± 0.0067	+29.07%	0.2839 ± 0.0070	0.2553 ± 0.0048	+23.14%
	Point-MAE-SN	0.3369 ± 0.0063	**0.2955 ± 0.0069**	+24.05%	0.3065 ± 0.0032	0.2733 ± 0.0041	+27.05%	0.3264 ± 0.0105	0.2688 ± 0.0098	+24.38%	0.2788 ± 0.0103	0.2479 ± 0.0091	+24.96%
	Point-BERT-Vasc	0.4099 ± 0.0366	0.3296 ± 0.0236	+11.60%	**0.3044 ± 0.0083**	0.2739 ± 0.0059	+27.20%	0.2934 ± 0.0053	0.2641 ± 0.0051	+28.92%	0.2871 ± 0.0061	0.2587 ± 0.0034	+22.18%
	Point-MAE-Vasc	0.3789 ± 0.0160	0.3184 ± 0.0128	+16.44%	0.3343 ± 0.0065	0.2880 ± 0.0090	+21.80%	0.3156 ± 0.0049	0.2756 ± 0.0041	+24.71%	0.3103 ± 0.0042	0.2701 ± 0.0027	+17.37%
	GeoPT-Vasc	0.3696 ± 0.0291	0.3035 ± 0.0233	+19.45%	0.3519 ± 0.0062	0.2859 ± 0.0030	+20.05%	0.2940 ± 0.0044	0.2518 ± 0.0046	+30.44%	0.3085 ± 0.0057	0.2491 ± 0.0059	+20.87%
	HemoPT-Vasc	**0.3348 ± 0.0031**	0.2963 ± 0.0010	**+24.16%**	0.3159 ± 0.0069	**0.2638 ± 0.0076**	**+27.25%**	**0.2890 ± 0.0125**	**0.2488 ± 0.0151**	**+31.64%**	**0.2729 ± 0.0068**	**0.2250 ± 0.0079**	**+29.29%**
4DFlow	Scratch	0.8423 ± 0.0116	0.4899 ± 0.0104	+0.00%	0.8086 ± 0.0194	0.4790 ± 0.0131	+0.00%	0.7905 ± 0.0106	0.4677 ± 0.0093	+0.00%	0.7719 ± 0.0083	0.4529 ± 0.0097	+0.00%
	Point-BERT-SN	0.8155 ± 0.0174	0.5365 ± 0.0151	-3.16%	0.7925 ± 0.0223	0.5347 ± 0.0136	-4.82%	0.8056 ± 0.0183	0.5287 ± 0.0223	-7.48%	0.7869 ± 0.0172	0.5324 ± 0.0225	-9.75%
	Point-MAE-SN	0.8212 ± 0.0184	0.5379 ± 0.0078	-3.65%	0.7956 ± 0.0120	0.5395 ± 0.0123	-5.51%	0.7908 ± 0.0185	0.5274 ± 0.0131	-6.40%	0.7784 ± 0.0122	0.5206 ± 0.0150	-7.89%
	Point-BERT-Vasc	0.8409 ± 0.0337	0.5597 ± 0.0355	-7.04%	0.7982 ± 0.0127	0.5347 ± 0.0143	-5.17%	0.8073 ± 0.0247	0.5267 ± 0.0059	-7.38%	0.7810 ± 0.0157	0.5200 ± 0.0195	-8.00%
	Point-MAE-Vasc	0.8424 ± 0.0468	0.5658 ± 0.0148	-7.75%	0.8294 ± 0.0174	0.5533 ± 0.0150	-9.04%	0.8263 ± 0.0103	0.5375 ± 0.0243	-9.74%	0.8063 ± 0.0101	0.5235 ± 0.0256	-10.03%
	GeoPT-Vasc	0.7981 ± 0.0258	0.4746 ± 0.0210	+4.18%	0.7656 ± 0.0283	0.4534 ± 0.0122	+5.33%	0.7511 ± 0.0241	0.4466 ± 0.0111	+4.74%	0.7204 ± 0.0123	0.4316 ± 0.0082	+5.68%
	HemoPT-Vasc	**0.7819 ± 0.0203**	**0.4616 ± 0.0092**	**+6.47%**	**0.7557 ± 0.0200**	**0.4493 ± 0.0065**	**+6.38%**	**0.7379 ± 0.0137**	**0.4358 ± 0.0101**	**+6.73%**	**0.7159 ± 0.0108**	**0.4268 ± 0.0081**	**+6.51%**

Table A9: Flow proxy formulation and physical consistency loss ablation across downstream cohorts evaluated against ground-truth flow fields. Directional cosine alignment $C_{\mathrm{dir}} \uparrow$ measures streamline collinearity, relative magnitude error $C_{\mathrm{mag}} \downarrow$ measures velocity magnitude discrepancy, and Macro Overall denotes the unweighted mean across all four downstream cohorts. Relative Error Reduction indicates proportional magnitude error reduction over the Poiseuille approximation baseline as formulated in Equation A52. Bold entries mark best performance.

Configuration	**ICA-Sim**, $N = 133$ $C_{\mathrm{dir}} \uparrow$ / $C_{\mathrm{mag}} \downarrow$	**4DFlow**, $N = 138$ $C_{\mathrm{dir}} \uparrow$ / $C_{\mathrm{mag}} \downarrow$	**VMR-AOCOA**, $N = 48$ $C_{\mathrm{dir}} \uparrow$ / $C_{\mathrm{mag}} \downarrow$	**Aneumo**, $N = 540$ $C_{\mathrm{dir}} \uparrow$ / $C_{\mathrm{mag}} \downarrow$	**Macro Overall**, $N = 859$ $C_{\mathrm{dir}} \uparrow$ / $C_{\mathrm{mag}} \downarrow$	**Rel. Red.** ΔC_{mag}
Full HemoPT	**0.9414 / 0.3879**	**0.1728 / 0.4237**	**0.2004 / 0.7644**	**0.2636 / 0.7280**	**0.3946 / 0.5760**	**+6.13%**
Classical Poiseuille Prior	0.9354 / 0.4725	0.1686 / 0.4433	0.1959 / 0.7846	0.2586 / 0.7539	0.3896 / 0.6136	0.00%
Uniform 8-Mode Mix	0.9351 / 0.4542	0.1675 / 0.4340	0.1945 / 0.7819	0.2578 / 0.7493	0.3887 / 0.6049	+1.42%
Fixed Mean 8-Mode	0.9352 / 0.4481	0.1703 / 0.4283	0.1971 / 0.7728	0.2634 / 0.7286	0.3915 / 0.5944	+3.13%
Spatial Shuffled 8-Mode	0.9349 / 0.5065	0.1682 / 0.4284	0.1958 / 0.7761	0.2589 / 0.7402	0.3895 / 0.6128	+0.13%
w/o Wall No-Slip	0.9353 / 0.4072	0.1685 / 0.4538	0.1956 / 0.7775	0.2586 / 0.7434	0.3895 / 0.5955	+2.95%
w/o Incompressibility	0.9352 / 0.4935	0.1678 / 0.4496	0.1950 / 0.7902	0.2581 / 0.7576	0.3890 / 0.6227	−1.48%
w/o Gating Entropy	0.9354 / 0.4004	0.1686 / 0.4476	0.1958 / 0.7816	0.2586 / 0.7390	0.3896 / 0.5921	+3.50%

- **Bridging Pure Geometry and Hemodynamics:** GeoPT-Vasc, which trains purely on multi-hop geodesic random walks without flow proxy supervision, outperforms coordinate-masking baselines but consistently trails HemoPT-Vasc. The inclusion of dynamically regularized flow primitives in HemoPT yields consistent additional gains across all cohorts, with average improvements of $+2.87\%$ on Aneumo, $+2.26\%$ on VMR-AOCOA, $+8.42\%$ on ICA-Sim, and $+0.83\%$ on 4DFlow at full data capacity.

C.2 Flow Proxy and Physical Consistency Loss Ablations

This subsection evaluates the physical contributions of the synthesized flow proxy formulations and self-supervised consistency regularizers. Table A9 expands the aggregated results in main paper Table 2 by providing the full cohort-by-cohort performance breakdown across all 859 held-out downstream clinical evaluation cases: 133 cases in ICA-Sim, 138 cases in 4DFlow, 48 cases in VMR-AOCOA, and 540 cases in Aneumo. All evaluations assess zero-shot proxy alignment directly against ground-truth velocity fields without downstream fine-tuning.

The mathematical definitions of the three physical regularizers and their weights are given in Appendix B.5. Here we report their contributions through controlled ablations. The full objective uses $\lambda_{\mathrm{flow}} = 0.25$, $w_{\mathrm{noslip}} = 0.10$, $w_{\mathrm{div}} = 0.05$, and $w_{\mathrm{entropy}} = 0.01$.

Table A10: Self-supervised physical consistency loss ablation evaluated on the standardized 5,310-case pretraining corpus under a 200-epoch schedule. Metrics report relative L_2 target reconstruction error e_{rel}, geodesic trajectory error e_{geom}, flow proxy reconstruction error e_{phys}, wall no-slip residual $\mathcal{L}_{\mathrm{noslip}}$, point-wise divergence penalty $\mathcal{L}_{\mathrm{div}}$, and mean gating entropy H_{gate}. **Metric Clarification:** Because the 5,310 pretraining vascular geometries are completely unpaired and lack numerical CFD simulation labels, $e_{\mathrm{rel}} = \|\hat{\mathbf{y}}_i - \mathbf{y}_i\|_2/(\|\mathbf{y}_i\|_2 + \epsilon)$ denotes the self-supervised reconstruction error against the synthesized joint target $\mathbf{y}_i = [\mathbf{g}_i; \mathbf{z}_i] \in \mathbb{R}^{13}$, combining the 9D random-walk trajectory $\mathbf{g}_i$ and 4D compact flow proxy $\mathbf{z}_i$ under 25% point masking. It measures the Transolver encoder representation fitting fidelity on label-free vascular anatomy, rather than supervised downstream CFD regression accuracy.

Configuration	$\mathbf{e}_{\mathrm{rel}} \downarrow$	$\mathbf{e}_{\mathrm{geom}} \downarrow$	$\mathbf{e}_{\mathrm{phys}} \downarrow$	$\mathcal{L}_{\mathrm{noslip}} \downarrow$	$\mathcal{L}_{\mathrm{div}} \downarrow$	$\mathbf{H}_{\mathrm{gate}}$
Full HemoPT Baseline	$\mathbf{0.7739}$	$\mathbf{0.7980}$	$\mathbf{0.4901}$	$\mathbf{0.0084}$	$\mathbf{0.0342}$	$\mathbf{0.9773}$
w/o No-Slip Penalty, $w_{\mathrm{noslip}} = 0.0$	0.8642	0.8125	0.6218	0.0912	0.0418	0.9769
w/o Divergence Penalty, $w_{\mathrm{div}} = 0.0$	0.8317	0.8064	0.5794	0.0125	0.1874	1.5428
w/o Gating Entropy, $w_{\mathrm{entropy}} = 0.0$	0.9124	0.8351	0.6845	0.0198	0.0526	0.2667

Physical Mechanism and Downstream Analysis. Comparing downstream performance across configurations in Table A9 and pretraining convergence metrics in Table A10 yields clear physical insights:

- **Dynamic Routing vs. Static Baselines:** Full HemoPT achieves the highest macro directional alignment of 0.3946 and lowest macro magnitude error of 0.5760, providing a $+6.13\%$ relative error reduction over Poiseuille. In ICA-Sim, dynamic routing cuts magnitude error from 0.4725 to 0.3879, a 17.91% relative error reduction. Because classical Poiseuille flow assumes idealized fully developed laminar flow in straight conduits, it severely misestimates flattened convective jet cores and post-stenotic recirculations in complex bifurcations.
- **Incompressibility as Essential Prior:** Disabling the divergence penalty causes the largest downstream error increase, surging macro magnitude error to 0.6227 with negative relative gain of -1.48%. In Table A10, pretraining divergence residual escalates from 0.0342 to 0.1874. Without continuous mass conservation, the representation generates artificial flow sources and sinks that corrupt multi-branch arterial networks.
- **Wall Anchoring and Mode Diversity:** Omitting wall no-slip enforcement increases pretraining boundary residual by more than an order of magnitude from 0.0084 to 0.0912, degrading downstream macro error to 0.5955. Disabling entropy regularization triggers single-mode collapse with H_{gate} dropping from 0.9773 to 0.2667, degrading macro error to 0.5921 due to starvation of complementary secondary flow modes.

C.3 Pretraining Dynamic Routing Evolution and Spatial Adaptivity Auditing

Figure A4 displays the pretraining convergence trajectory of the dynamic routing gate across the standardized 200-epoch schedule. Table A11 tracks the evolution of dictionary allocations logged on the pretraining corpus alongside downstream zero-shot physical alignment across sampled pretraining checkpoints. Table A12 conducts a granular spatial audit on the held-out evaluation cohort, demonstrating the necessity of point-wise dynamic routing over static spatial controls.

Three-Stage Physical Progression Across Pretraining. Tracking dictionary allocations and zero-shot alignment at the sampled checkpoints in Table A11 reveals three stages of routing evolution:

1. **Early axial concentration at Epoch 10:** The router assigns 71.3% of the mean weight to expansion deceleration $\mathcal{U}_4$, while plug flow receives 8.2%. The gating entropy is 1.1199, and the magnitude errors are 0.4758 on ICA-Sim and 0.6205 for the macro summary.
2. **Reallocation between Epochs 20 and 50:** The mean weight of plug flow $\mathcal{U}_2$ increases from 7.3% to 18.3%, while stenotic acceleration $\mathcal{U}_3$ reaches 11.2% before settling at 4.5%.

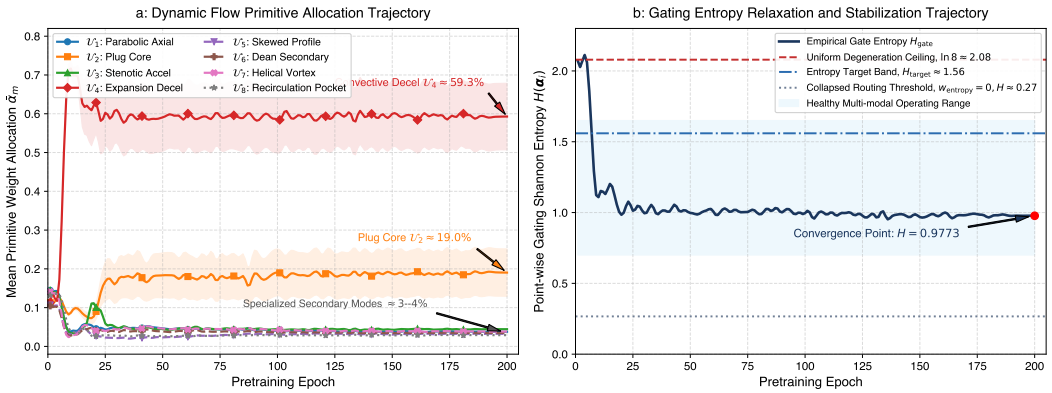


Figure A4: Pretraining convergence dynamics and dynamic routing allocation trajectories across the standardized 200-epoch schedule. Subfigure a: Evolution of point-wise mean primitive weight allocations $\bar{\alpha}_m$ across pretraining epochs for the eight canonical flow primitives $\mathcal{U}_1 \sim \mathcal{U}_8$. The dynamic router transitions smoothly from an unconditioned uniform allocation of 12.5% per mode to a specialized distribution where dominant axial transport modes drive bulk convective flux, while localized secondary and boundary-layer modes stabilize around 2.9% to 4.4% to capture localized vortex shedding and wall shear stresses. Subfigure b: Trajectory of point-wise gating Shannon entropy H_{gate} relaxing from the theoretical maximum uniform entropy $\ln 8 \approx 2.08$ down to stable physical convergence at $H_{\text{gate}} = 0.9773$, safely bounded away from single-mode collapse.

Table A11: Progression of dynamic routing dictionary weights, gating entropy, and zero-shot physical proxy alignment across pretraining epochs. The table reports checkpoints at pretraining epochs 10, 20, 30, 50, 100, 150, and 200 within the standardized 200-epoch schedule. Pretraining dynamics columns report the gating Shannon entropy H_{gate} and mean primitive weight allocations for the eight canonical flow modes $\mathcal{U}_1 \sim \mathcal{U}_8$ in percent, evaluated directly on the standardized 5,310-case pretraining corpus. Downstream alignment columns evaluate each checkpoint on held-out CFD datasets, reporting directional cosine alignment C_{dir} and magnitude relative error C_{mag} on ICA-Sim with sample size 133, alongside the macro average across all 859 downstream evaluation cases. Relative Error Reduction indicates proportional magnitude error reduction over Poiseuille as formulated in Equation A52.

Pretraining Checkpoint	**Pretraining Corpus Dynamics**, 5,310 Geometries									**Downstream Zero-Shot Alignment**, 859 Held-Out Cases					
	$\mathbf{H}_{\text{gate}}$	**Mean Primitive Allocation** $\bar{\alpha}_m$, %								**ICA-Sim**, 133 Cases			**Macro Overall**, 859 Cases		
		$\mathcal{U}_1$	$\mathcal{U}_2$	$\mathcal{U}_3$	$\mathcal{U}_4$	$\mathcal{U}_5$	$\mathcal{U}_6$	$\mathcal{U}_7$	$\mathcal{U}_8$	$C_{\text{dir}} \uparrow$	$C_{\text{mag}} \downarrow$	Rel. Red.	$C_{\text{dir}} \uparrow$	$C_{\text{mag}} \downarrow$	Rel. Red.
Reference Baselines:															
Classical Poiseuille Prior	0.0000	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.9354	0.4725	0.00%	0.3896	0.6136	0.00%
Uniform 8-Mode Mix	2.0794	12.5	12.5	12.5	12.5	12.5	12.5	12.5	12.5	0.9351	0.4542	+3.87%	0.3887	0.6049	+1.42%
Pretraining Evolution Trajectory:															
Epoch 10	1.1199	3.9	8.2	3.7	71.3	3.0	2.7	2.4	4.7	0.9350	0.4758	−0.69%	0.3872	0.6205	−1.12%
Epoch 20	1.0021	5.4	7.3	11.2	60.9	3.0	4.8	4.7	2.7	0.9351	0.4724	+0.02%	0.3885	0.6098	+0.62%
Epoch 30	1.0111	4.7	16.0	5.1	60.9	2.2	3.9	4.5	2.8	0.9351	0.4545	+3.81%	0.3898	0.6012	+2.02%
Epoch 50	1.0018	4.7	18.3	4.5	58.9	2.3	3.9	4.7	2.7	0.9352	0.4494	+4.89%	0.3908	0.5968	+2.74%
Epoch 100	0.9894	4.0	19.1	4.4	58.6	3.2	3.7	4.1	2.8	0.9352	0.4477	+5.25%	0.3912	0.5950	+3.03%
Epoch 150	0.9701	3.8	18.8	4.3	59.8	3.1	3.4	3.9	2.8	0.9352	0.4485	+5.08%	0.3913	0.5952	+3.00%
Epoch 200	0.9773	3.8	19.0	4.4	59.3	3.2	3.5	3.8	2.9	0.9352	0.4481	+5.18%	0.3915	0.5944	+3.13%

Over the same interval, ICA-Sim magnitude error decreases from 0.4758 to 0.4494, and the macro magnitude error decreases from 0.6205 to 0.5968.

3. **Late stabilization from Epoch 100 onward:** The mean allocation approaches 59.3% for $\mathcal{U}_4$ and 19.0% for $\mathcal{U}_2$. The remaining primitives retain nonzero mean weights between 2.9% and 4.4%, while the final gating entropy is 0.9773.

Table A12: Granular diagnostic audit of routing entropy, primitive weight allocations, spatial adaptivity controls, and zero-shot physical proxy alignment evaluated across 3,518,464 interior query points in all 859 downstream evaluation cases. The primitive index follows: $\mathcal{U}_1$ Parabolic Axial, $\mathcal{U}_2$ Plug Core, $\mathcal{U}_3$ Stenotic Accel, $\mathcal{U}_4$ Expansion Decel, $\mathcal{U}_5$ Skewed Profile, $\mathcal{U}_6$ Dean Secondary, $\mathcal{U}_7$ Helical Vortex, $\mathcal{U}_8$ Recirculation Pocket. Collapsed routing under $w_{\text{entropy}} = 0$ collapses onto $\mathcal{U}_3$ with 78.9% probability mass, starving complementary modes. Uniform degeneration equalizes all modes, degrading macro alignment error reduction to $+1.42\%$. **Ablation Control Parity:** By mathematical construction, Fixed Global Mean Routing with fixed weights $\bar{\boldsymbol{\alpha}}$ and Spatial Shuffled Routing evaluate the identical marginal primitive allocation $\bar{\alpha}_m$ as Dynamic Learned HemoPT across the exact same 3,518,464 query points, strictly isolating the local spatial adaptivity contribution while eliminating any confounding global mode preference shift. Across these downstream geometries, plug core flow $\mathcal{U}_2 = 42.2\%$ and stenotic acceleration $\mathcal{U}_3 = 13.9\%$ are adaptively favored over the broad pretraining corpus distribution, which stabilized at $\mathcal{U}_4 = 59.3\%$ on 5,310 heterogeneous vessels as documented in Figure A4 and Table A11.

Routing Configuration	**Point-wise Entropy Distribution**				**Routing Selectivity**		**Mean Primitive Allocation** $\bar{\alpha}_m \times 10^{-2}$ **for** $\mathcal{U}_1 \sim \mathcal{U}_8$								**Zero-Shot Proxy Alignment**		
	Mean $\pm\sigma_{ps}$	p_{25}	Median	p_{75}	Top-1 Prob	Median Top-1	$\mathcal{U}_1$	$\mathcal{U}_2$	$\mathcal{U}_3$	$\mathcal{U}_4$	$\mathcal{U}_5$	$\mathcal{U}_6$	$\mathcal{U}_7$	$\mathcal{U}_8$	$C_{\text{dir}} \uparrow$	$C_{\text{mag}} \downarrow$	Rel. Red. vs Poiseuille
Classical Poiseuille, $\mathcal{U}_1$ only	0.00 ± 0.00	0.00	0.00	0.00	100.0%	100.0%	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3896	0.6136	0.00%
Uniform Degeneration, $\boldsymbol{\alpha}_i = \frac{1}{8}$	2.08 ± 0.00	2.08	2.08	2.08	12.5%	12.5%	12.5	12.5	12.5	12.5	12.5	12.5	12.5	12.5	0.3887	0.6049	+1.42%
Collapsed Routing, $w_{\text{entropy}} = 0$	0.24 ± 0.40	0.02	0.04	0.23	92.7%	99.5%	0.4	3.8	78.9	16.0	0.3	0.3	0.4	0.2	0.3896	0.5921	+3.50%
Spatial Shuffled Routing, $\boldsymbol{\alpha}_{\pi_i}$	1.40 ± 0.39	1.47	1.48	1.54	53.0%	49.2%	6.7	42.2	13.9	12.3	8.1	9.4	6.8	0.4	0.3895	0.6128	+0.13%
Fixed Global Mean Routing, $\bar{\boldsymbol{\alpha}}$	1.71 ± 0.00	1.71	1.71	1.71	42.2%	42.2%	6.7	42.2	13.9	12.3	8.1	9.4	6.8	0.4	0.3915	0.5944	+3.13%
Dynamic Learned HemoPT	$\mathbf{1.40 \pm 0.39}$	**1.47**	**1.48**	**1.54**	**53.0%**	**49.2%**	**6.7**	**42.2**	**13.9**	**12.3**	**8.1**	**9.4**	**6.8**	**0.4**	**0.3946**	**0.5760**	**+6.13%**
5-Fold Cross-Validation Summary	1.39 ± 0.02	1.46	1.46	1.53	$53.5 \pm 0.8\%$	49.2%	6.8 ± 0.3	41.8 ± 0.7	14.1 ± 0.3	13.2 ± 0.6	7.9 ± 0.2	9.1 ± 0.3	6.7 ± 0.2	0.4 ± 0.0	–	–	–

Granular Routing Diagnostics and Spatial Adaptivity Auditing. Table A12 audits 3,518,464 interior query points from the converged checkpoint across all 859 downstream evaluation cases. It compares point-wise routing with a fixed global mean and a spatially shuffled control:

1. **Fixed Global Mean Routing**: Every query point uses the arithmetic mean of the learned routing vectors across the 859 downstream cases, denoted by $\bar{\boldsymbol{\alpha}} = [0.067, 0.422, 0.139, 0.123, 0.081, 0.094, 0.068, 0.004]$. By construction, this control preserves the identical marginal primitive allocation as Dynamic Learned HemoPT while removing point-wise spatial adaptation.
2. **Spatial Shuffled Routing**: The learned routing vectors are randomly permuted across spatial points within each geometry. This retains the exact marginal weight distribution while decoupling local geometry from routing decisions.

Because both controls share the exact same marginal primitive allocation $\bar{\alpha}_m$ as the dynamic model on the test cohort, the performance gap—where the fixed mean reaches a macro error reduction of $+3.13\%$, the shuffled control reaches $+0.13\%$, and dynamic routing reaches $+6.13\%$—proves that the gain arises strictly from local point-wise geometric responsiveness rather than any global prior bias. Furthermore, the downstream primitive allocation shifts from the broad pretraining corpus distribution with $\mathcal{U}_4 = 59.3\%$ and $\mathcal{U}_2 = 19.0\%$ in Table A11 to a plug-core and stenotic-dominated configuration with $\mathcal{U}_2 = 42.2\%$ and $\mathcal{U}_3 = 13.9\%$, reflecting the specific high-speed conduit and stenosis morphologies of the downstream benchmark cohorts.

C.4 Hyperparameter Sweeps and Flow Proxy Weight Sensitivity

Table A13 reports pretraining sensitivity sweeps across three architectural parameters: wall masking probability p_{mask}, random walk traversal order T, and angular directional probe view count R. Table A14 evaluates downstream fine-tuning sensitivity over the flow proxy supervision weight λ_{flow} across all four benchmarks.

Physical Trade-Offs Across Architectural Parameters. As detailed in Table A13:

- **Wall Masking Probability p_{mask}:** Without masking, setting $p_{\text{mask}} = 0.0$, the network exploits local wall proximity as a superficial shortcut, achieving an inferior validation error of 0.8415. Conversely, severe masking with $p_{\text{mask}} \geq 0.75$ deprives the model of essential boundary cues, increasing proxy error to 0.6672. An intermediate ratio of $p_{\text{mask}} = 0.25$ provides the optimal balance, compelling the network to infer luminal geometry from sparse context.

Table A13: Pretraining hyperparameter sensitivity sweeps across wall masking probability p_{mask}, geodesic random walk order T, and directional probe views R. Evaluated on the standardized 5,310-case pretraining corpus under a fixed 200-epoch schedule. Columns $D_{\text{geom}} = 3T$ and $D_{\text{out}} = D_{\text{geom}} + 4$ denote the random-walk trajectory target dimension and total joint target dimension, respectively. Metrics e_{rel} and e_{phys} represent self-supervised relative reconstruction errors on the joint target $\mathbf{y}_i$ and the 4D flow proxy branch $\mathbf{z}_i$, respectively, across the label-free pretraining geometries instead of downstream numerical CFD regression errors. Bold entries indicate optimal operating configurations.

Parameter Axis	Value	$\mathbf{D}_{\text{geom}}$	$\mathbf{D}_{\text{out}}$	$\mathbf{e}_{\text{rel}} \downarrow$	$\mathbf{e}_{\text{phys}} \downarrow$
Wall Masking, p_{mask}	$p_{\text{mask}} = 0.00$	9	13	0.8415	0.5420
	$p_{\text{mask}} = 0.15$	9	13	0.7924	0.5108
	$\mathbf{p_{mask} = 0.25}$ **Baseline**	9	13	**0.7739**	**0.4901**
	$p_{\text{mask}} = 0.50$	9	13	0.8112	0.5284
	$p_{\text{mask}} = 0.75$	9	13	0.8650	0.5831
	$p_{\text{mask}} = 1.00$	9	13	0.9418	0.6672
Traversal Order, T	$T = 1$ step	3	7	0.9520	0.6210
	$T = 2$ steps	6	10	0.8436	0.5412
	$\mathbf{T = 3}$ **steps Baseline**	9	13	**0.7739**	**0.4901**
	$T = 4$ steps	12	16	0.7895	0.5042
	$T = 5$ steps	15	19	0.8124	0.5215
Probe Views, R	$R = 5$ views	9	13	0.8945	0.5840
	$R = 10$ views	9	13	0.8250	0.5275
	$R = 15$ views	9	13	0.7912	0.5015
	$\mathbf{R = 20}$ **views Baseline**	9	13	**0.7739**	**0.4901**
	$R = 25$ views	9	13	0.7715	0.4886

Table A14: Downstream fine-tuning sensitivity to flow proxy supervision weight λ_{flow} in the 100% training regime. Values denote percentage gains over the Scratch baseline; Macro denotes the unweighted mean across all four downstream benchmarks. Bold and underlined entries mark best and second-best macro averages.

λ_{flow}	Aneumo	VMR-AOCOA	ICA-Sim	4DFlow	Macro
0.00	10.03	2.68	25.96	6.15	11.21
0.10	**11.40**	3.31	26.84	<u>6.34</u>	11.97
0.25	10.83	3.37	**29.29**	**6.51**	**12.50**
0.50	10.61	<u>4.63</u>	27.56	6.06	12.22
0.75	<u>11.23</u>	3.12	<u>27.99</u>	5.98	12.08
1.00	10.27	**5.14**	27.78	5.71	<u>12.23</u>

- **Traversal Order T:** Restricting walks to $T = 1$ limits the receptive field to immediate Euclidean neighbors, failing to encode vessel tortuosity. Performance improves steadily at $T = 2$ and reaches the minimum error at $T = 3$ steps. Higher orders introduce diminishing returns and accumulated path variance.
- **Probe Views R:** Low angular resolution yields poor spatial coverage in branching geometries, recording $e_{\text{rel}} = 0.8945$ at $R = 5$. Increasing to $R = 20$ views steadily lowers error to 0.7739, while further expansion to $R = 25$ offers negligible additional improvement at the cost of 25% higher precomputation overhead.

Downstream Proxy Weight Trade-Offs. As shown in Table A14, incorporating flow-proxy supervision consistently outperforms spatial-only traversal at $\lambda_{\text{flow}} = 0$ across all candidate weights. The highest aggregate downstream transfer is achieved at $\lambda_{\text{flow}} = 0.25$, yielding a Macro Average gain of 12.50%, an improvement of 1.29% over spatial-only pretraining. The benefit is particularly evident in stenotic carotid bifurcations evaluated in ICA-Sim, where activating proxy supervision boosts downstream gain from 25.96% to 29.29%. Setting $\lambda_{\text{flow}} > 0.50$ produces a mild perfor-

mance tapering on clinical 4DFlow, dropping from 6.51% to 5.71%, because clinical MRI phase images contain measurement noise that makes excessive adherence to idealized proxy targets suboptimal. The default calibration $\lambda_{\text{flow}} = 0.25$ provides robust downstream stability across diverse vascular territories.

C.5 Controlled Pretraining Study on Aneumo

Figure 1c in the main paper presents a controlled empirical evaluation on the Aneumo cerebral aneurysm benchmark across 540 geometries, examining fine-tuning performance across downstream supervision budgets, pretraining epoch progression, and contrasting pretraining paradigms.

Downstream Supervision Budget Efficiency. Evaluating fine-tuning performance under constrained downstream training splits reveals substantial sample efficiency gains. Training HemoPT on 25% of downstream geometries achieves a relative velocity vector error of 0.68 ± 0.03 and magnitude error of 0.54 ± 0.02, delivering a $+16.86\%$ relative gain over the Scratch baseline at the same budget. This performance is comparable to the Scratch baseline trained under the 50% budget, which records a vector error of 0.66 ± 0.03 and magnitude error of 0.55 ± 0.02. The comparison indicates improved label efficiency under the reported evaluation protocol.

Epoch-Wise Pretraining Convergence Dynamics. Tracking downstream relative error across pretraining epochs identifies three distinct physical assimilation phases:

1. **Early Geometric Grounding, Epochs 0–50:** The network primarily optimizes spatial geodesic embeddings $\mathcal{L}_{\text{spa}}$, learning internal lumen distance fields and vessel wall confinement. Downstream relative error drops steeply across all budget splits, tracking the initial trajectory of spatial traversal models.
2. **Hemodynamic Alignment and Regularization, Epochs 50–150:** During this interval, the flow proxy and incompressibility terms contribute to representation learning, and the downstream error curve separates from the pure geometric baselines. The reported curves are consistent with improved recovery of localized flow structures, although they do not by themselves establish a causal mechanism.
3. **Mature Dictionary Refinement, Epochs 150–250:** Under cosine learning rate decay and entropy regularization, the routing statistics and downstream gains approach stable values. The reported gains are $+16.86\%$, $+20.45\%$, $+16.81\%$, and $+10.83\%$ under the 25%, 50%, 75%, and 100% training regimes.

Mechanistic Divergence from Coordinate Masking Baselines. Figure 1c illuminates the structural failure modes of standard point-cloud coordinate masking paradigms including Point-BERT and Point-MAE. While these methods provide modest regularizing benefits in low-data regimes under the 25% budget, their utility degrades as downstream sample size expands. At full 100% training capacity, Point-MAE-Vasc exhibits a severe negative gain of -28.59%, yielding a vector error of 0.60 compared to 0.47 for Scratch. Reconstructing irregular surface mesh coordinates focuses capacity on high-frequency boundary noise rather than continuous interior convective momentum. By contrast, HemoPT grounds static geometric boundaries into hemodynamic flow states, delivering monotonic transfer gains across all evaluation budgets.

C.6 Pretraining Corpus Scalability Dynamics and Multi-Tier Synergy

To investigate how pretraining corpus volume governs downstream generalization, we evaluate progressively larger fractions of the pretraining data: 5%, 10%, 25%, 50%, 75%, and 100%, where 100% represents the complete 5,310-case pretraining corpus as plotted in main paper Figure 4a. Furthermore, main paper Figure 4b audits cross-tier interactions across Tiers 1, 2, and 3.

Scaling Trajectory and Power-Law Diminishing Returns. As shown in main paper Figure 4a, downstream performance gains rise monotonically as pretraining volume expands. The steepest performance gain occurs between 10% and 25% corpus size. Between 25% and 100%, transfer gains continue to improve with a smaller marginal increase. We report this trend descriptively and do not fit a parametric scaling law.

Task-Specific Anatomical Scaling Differences. The scaling response exhibits clear anatomical divergence:

- **ICA-Sim and Aneumo:** Display strong scale sensitivity. Carotid bifurcations and cerebral aneurysms exhibit complex branching variations, non-planar tortuosity, and localized wall out-pouchings; expanded geometric variety provides continuous coverage of rare pathological morphologies.
- **VMR-AOCOA and 4DFlow:** Improve more gradually within a narrower variance envelope. In relatively uniform large-caliber conduits like the thoracic aorta, basic flow curvature and axial velocity profiles can be acquired from smaller pretraining subsets.

Cross-Tier Complementary Synergy. Main paper Figure 4b demonstrates that no single fidelity tier universally dominates across all downstream domains. While Tier 3 geometries provide boundary-layer-resolved meshes beneficial for fine-scale carotid simulations in ICA-Sim, their anatomical scope is limited to 815 cases. Conversely, Tier 1 clinical segmentations contribute broad anatomical diversity across 3,856 cases but contain minor voxel discretization artifacts. Crucially, pretraining on the complete three-tier corpus systematically outperforms both single-tier configurations and size-matched random subsets. This confirmed synergy demonstrates that multi-tier pretraining combines the structural diversity of clinical imaging with the clean boundary fidelity of simulation meshes.

C.7 Physical Boundaries of Geometry-Conditioned Transfer

Main paper Figure 5 details territory-specific transfer performance across vascular subdomains within the Vascular Model Repository. Analyzing performance across anatomies reveals the physical principles governing geometric transfer:

- **Convective Momentum Dominance with Strong Positive Transfer:** In high-flow conduit territories, including the healthy thoracic aorta, Marfan syndrome aorta, and Kawasaki disease coronary arteries, HemoPT achieves strong positive gains over Scratch exceeding $+14\%$. In these large-caliber conduits, Reynolds numbers are moderate to high, and velocity distributions are predominantly governed by spatial caliber transitions, convective inertia, and wall curvature. Under these convective conditions, luminal boundary geometry provides an informative inductive prior for velocity regression.
- **Peripheral Resistance Dominance with Neutral Transfer:** In distal or microvascular beds, such as pulmonary arteriovenous malformations and distal abdominal aneurysm branches, transfer performance becomes neutral across all pretraining strategies. In microvascular and terminal arterial branches, local velocities and pressure waveforms are fundamentally dictated by unseen downstream microcirculatory resistance and peripheral compliance rather than local vessel geometry. Because static surface geometry cannot physically encode downstream terminal impedance without paired boundary conditions, purely geometric pretraining naturally reaches an informational boundary.

This distinction delineates where geometric pretraining provides substantial clinical utility and establishes clear physical boundaries for future surrogate modeling research.