

From Centerlines to Hemodynamics: Anisotropic RBF Decoders for Coronary Arteries

Anonymous authors

Paper under double-blind review

Abstract

Accurate and rapid estimation of hemodynamic metrics, such as pressure and wall shear stress (WSS), is important for assessing the severity of Coronary Artery Disease (CAD). Existing approaches, including invasive Fractional Flow Reserve (FFR) measurements and computationally expensive Computational Fluid Dynamics (CFD) simulations, face challenges in invasiveness, cost, and speed. We present a framework for fast, non-invasive coronary hemodynamics prediction. The model encodes 1D vessel centerlines together with inlet flow rate using a transformer-based encoder, and predicts continuous wall-based fields via an anisotropic Radial Basis Function (RBF) decoder aligned with vessel morphology. To support training and evaluation, we introduce two datasets with paired steady-state OpenFOAM simulations: (i) a synthetic benchmark of 4,200 single-vessel geometries with controlled anatomical variations, and (ii) a multi-vessel dataset derived from ImageCAS including 4,800 cases spanning both right and left coronary arteries, generated by randomly introducing stenoses and varying physiologically plausible flow rates. Across both datasets, our method achieves lower pressure and WSS errors than strong neural-operator baselines (GNOT, Transolver, and ONO) at a fraction of the computational cost of CFD. On the multi-vessel dataset, using 1,024 anisotropic RBF centers our model reduces the mean relative ℓ_2 error by 52% compared to the best neural-operator baseline, while at 128 centers it requires $13.8\times$ fewer FLOPs than GNOT and still outperforms all baselines. The single-vessel dataset is publicly available.¹

1 Introduction

Cardiovascular disease remains the leading cause of mortality worldwide, with coronary artery disease (CAD) accounting for a large fraction of deaths Virani & et al. (2021). Clinical decision-making in CAD depends on determining whether a coronary stenosis significantly impairs blood flow. The gold standard for this assessment is fractional flow reserve (FFR), defined as the ratio of distal to proximal coronary pressure during pharmacologically induced hyperemia Tonino et al. (2009). Despite its diagnostic value, routine FFR measurement is limited by the need for catheterization, pressure-wire instrumentation, and pharmacological intervention Pijls et al. (1996).

Patient-specific computational fluid dynamics (CFD) offers a non-invasive alternative by solving the incompressible Navier-Stokes equations on reconstructed vascular domains, accurately recovering pressure and wall shear stress (WSS) fields and enabling FFR estimation from imaging Min et al. (2015). However, generating high-quality meshes and solving three-dimensional flow fields typically requires hours of computation per case Sankaran et al. (2012), limiting applicability in time-sensitive clinical workflows and large-scale population studies.

Machine learning surrogates have been explored to reduce this cost. Physics-informed neural networks (PINNs) Raissi et al. (2019); Bafghi & Raissi (2023) embed governing equations into training but require retraining per geometry and struggle with multi-scale optimization difficulties Wang et al. (2020; 2023). Neural operators learn mappings between function spaces Kovachki et al. (2023) and avoid this limitation:

¹Dataset URL withheld for anonymous review.

DeepONet Lu et al. (2019) uses branch-trunk decomposition; FNO Li et al. (2020), Geo-FNO Li et al. (2022), and GINO Li et al. (2023) extend spectral learning to irregular geometries; and transformer-based operators including GNOT Hao et al. (2023), Transolver Wu et al. (2024), and ONO Xiao et al. (2023) further improve expressiveness through attention mechanisms. Geometric deep learning methods have also been applied directly to coronary artery surface meshes Nannini et al. (2025); Suk et al. (2024b;a); Kuang et al. (2024); Rygiel et al. (2023).

Despite this progress, a structural gap remains. Coronary arteries are tubular domains naturally parameterized by a one-dimensional centerline with an associated radius field, while clinically relevant outputs are defined on the vessel wall. Existing neural operator architectures operate on volumetric or surface discretizations and do not exploit this intrinsic low-dimensional structure, incurring unnecessary computational cost and failing to provide a unified framework for geometry-aware encoding with continuous, mesh-independent field reconstruction.

We address this gap by learning the nonlinear solution operator

$$\mathcal{G} : (\Gamma, r(s), q_{\text{in}}) \mapsto (p(\mathbf{x}), \tau_w(\mathbf{x})),$$

where Γ is the vessel centerline, $r(s)$ is the radius field, q_{in} is the inlet flow rate, and (p, τ_w) are spatially continuous pressure and WSS fields at arbitrary wall locations \mathbf{x} . Our Transformer-Anisotropic RBF Network encodes the 1D centerline via Fourier positional embeddings and a transformer encoder, conditions on inlet flow rate through Feature-wise Linear Modulation (FiLM), and decodes continuous fields as a weighted superposition of anisotropic Gaussian RBF kernels centered along the vessel. Each kernel carries a learned full-precision matrix, allowing orientation and scale to adapt to local vessel geometry. Evaluation at arbitrary query points is achieved by kernel aggregation at near-constant cost regardless of surface resolution.

Contributions

- **Geometry-aware operator formulation.** A reduced-order representation that exploits the intrinsic 1D centerline structure of tubular geometries for efficient hemodynamic prediction.
- **Mesh-free continuous field reconstruction.** An anisotropic RBF decoder enabling continuous pressure and WSS evaluation independent of surface discretization, with inference cost nearly invariant to query-point count.
- **Large-scale synthetic benchmark.** Two paired CFD datasets (4,200 single-vessel and 4,800 multi-vessel geometries) supporting reproducible evaluation; the single-vessel set is publicly released.
- **Improved accuracy–efficiency trade-off.** Lower ℓ_2 errors than GNOT, Transolver, and ONO on both datasets, with up to $13.8\times$ fewer FLOPs at matched RBF count.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the methodology. Section 4 presents the datasets. Section 5 reports experimental results. Section 6 discusses limitations and future directions.

2 Related Work

Computational hemodynamics and reduced-order modeling. Patient-specific CFD remains the reference standard for coronary hemodynamics, solving the incompressible Navier-Stokes equations on anatomically reconstructed domains to recover pressure, WSS, and FFR Min et al. (2015); Sankaran et al. (2012). Its computational expense, however, limits use in time-sensitive and large-scale settings.

Classical reduced-order modeling (ROM) addresses this through projection-based methods such as proper orthogonal decomposition (POD) and reduced basis approaches, which construct low-dimensional subspaces from high-fidelity simulations. While these provide substantial savings with physical interpretability, they rely on linear approximations and often require intrusive access to the governing equations. One-dimensional and lumped-parameter models offer further efficiency along vessel centerlines but cannot capture three-dimensional

effects such as secondary flows and localized WSS distributions. Hybrid data-driven ROM approaches have been proposed to reconstruct pressure and WSS from limited simulation data Morgan et al. (2023), and graph neural networks have been applied to learn reduced-order cardiovascular flow models Pegolotti et al. (2023), but generalization across complex geometries remains limited.

Machine learning for hemodynamic prediction. PINNs Raissi et al. (2019); Bafghi & Raissi (2023); Kissas et al. (2020) embed governing equations as training regularizers but solve individual PDE instances rather than amortized operators, and optimization for multi-scale flows is sensitive to loss balancing Wang et al. (2020). Geometric deep learning methods applied to coronary artery surface meshes Nannini et al. (2025); Suk et al. (2024b;a) and physics-informed self-supervised approaches for hemodynamic digital twins Kuang et al. (2024) have shown promise, but their cost scales with the number of discretization points. CenterlinePointNet++ Rygiel et al. (2023) also leverages centerline and surface point cloud representations for coronary pressure drop and vFFR estimation, but predicts only a scalar per vessel rather than spatially continuous wall fields. In a concurrent line of work, physics-informed approaches have also been applied to estimate coronary flow reserve directly from angiography Thakur et al. (2026).

Neural operators. Neural operators learn resolution-independent mappings between function spaces Kovachki et al. (2023), avoiding the per-instance retraining of PINNs. DeepONet Lu et al. (2019) uses branch-trunk decomposition; FNO Li et al. (2020), Geo-FNO Li et al. (2022), and GINO Li et al. (2023) extend spectral convolution to irregular geometries via learned coordinate transforms or graph representations; and transformer-based operators GNOT Hao et al. (2023), Transolver Wu et al. (2024) (recently extended to million-scale meshes as Transolver++ Luo et al. (2025)), and ONO Xiao et al. (2023) add long-range expressiveness through attention; and Universal Physics Transformers Alkin et al. (2024) provide a general encoder-approximator-decoder framework for scaling neural operators across discretization types. Despite these advances, all these families operate on volumetric or surface discretizations and do not exploit the intrinsic one-dimensional structure of tubular vascular geometries.

Mesh-free and kernel-based representations. RBF approximations have a long history in computational mechanics for mesh-free interpolation, constructing solutions as weighted superpositions of kernel functions centered at selected locations Fasshauer (2007). Anisotropic extensions using ellipsoidal support regions have been explored for partition-of-unity interpolation Cavoretto et al. (2018), and recent work has revisited learned anisotropic RBF kernels in regression settings Gerber & Lloyd (2026). Classical RBF methods, however, use fixed kernels and do not adapt to parameter-dependent solution behavior in an operator-learning context.

Gap. No existing approach combines geometry-aware encoding of the 1D centerline structure with continuous, mesh-independent field reconstruction in a unified operator framework. This is the gap our method addresses.

3 Methodology

Our proposed Transformer-Anisotropic RBF Network predicts continuous pressure and wall shear stress (WSS) fields on coronary artery walls from vessel geometry and inlet flow rate. The framework combines a transformer-based encoder for geometry-aware feature extraction with an anisotropic Radial Basis Function (RBF) decoder for continuous field reconstruction (Fig. 1).

3.1 Problem Formulation

Given a vessel centerline $\mathbf{C} = \{(\mathbf{x}_i, r_i)\}_{i=1}^M$ with coordinates $\mathbf{x}_i \in \mathbb{R}^3$, local radii $r_i \in \mathbb{R}^+$, and inlet flow rate $q \in \mathbb{R}$, we learn the operator

$$\mathcal{G} : (\mathbf{C}, q) \mapsto (P(\mathbf{x}), \text{WSS}(\mathbf{x})),$$

predicting pressure and WSS magnitude at arbitrary query points \mathbf{x} on the vessel wall. Figure 2 illustrates this mapping for single- and multi-vessel anatomies.

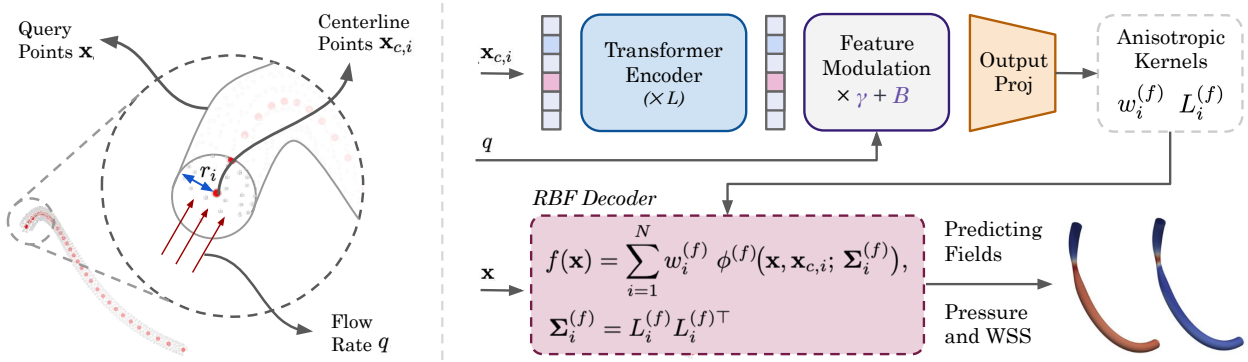


Figure 1: Architecture of the Transformer-Anisotropic RBF Network. The transformer encoder processes centerline geometry and flow rate, producing parameters for anisotropic RBF kernels aligned with vessel morphology to reconstruct continuous pressure and WSS fields.

3.2 Architecture Overview

The model has two main components: (1) a transformer encoder that processes centerline geometry and flow rate, and (2) an anisotropic RBF decoder that reconstructs continuous pressure and WSS fields on the wall surface.

Transformer Encoder with 4D Fourier. Each centerline point $\mathbf{c}_i = (x_i, y_i, z_i, r_i)$ is encoded with sinusoidal positional embeddings Tancik et al. (2020) applied independently to each of the four input coordinates at K exponentially spaced frequencies:

$$\mathbf{e}(\mathbf{c}) = \left[\sin(\omega_k c_j), \cos(\omega_k c_j) \right]_{\substack{j=1,\dots,4 \\ k=0,\dots,K-1}} \in \mathbb{R}^{8K}, \quad \omega_k = \pi \cdot 2^k. \quad (1)$$

The embedding is concatenated with the raw coordinates and linearly projected to the model dimension d :

$$\mathbf{h}_i^{(0)} = \mathbf{W}[\mathbf{e}(\mathbf{c}_i); \mathbf{c}_i] + \mathbf{b}, \quad \mathbf{W} \in \mathbb{R}^{d \times (8K+4)}. \quad (2)$$

Learned positional embeddings are added to the token sequence, which is then processed by a stack of transformer encoder layers Vaswani et al. (2017) to produce features $\{\mathbf{h}_i\}_{i=1}^M$.

Flow Rate Conditioning. After the transformer encoder, the inlet flow rate q modulates the encoded features via Feature-wise Linear Modulation (FiLM) Pérez et al. (2018). A two-layer MLP maps the scalar flow rate to scale and shift vectors:

$$\boldsymbol{\gamma}, \boldsymbol{\beta} = \text{MLP}(q), \quad \boldsymbol{\gamma}, \boldsymbol{\beta} \in \mathbb{R}^d, \quad (3)$$

$$\tilde{\mathbf{h}}_i = \boldsymbol{\gamma} \odot \mathbf{h}_i + \boldsymbol{\beta}, \quad (4)$$

where \odot denotes element-wise multiplication broadcast over all M centerline tokens.

Anisotropic RBF Decoder. A linear layer maps each modulated feature $\tilde{\mathbf{h}}_i$ to 14 parameters per centerline point: for each field $f \in \{\text{P}, \text{WSS}\}$, a scalar weight $w_i^{(f)}$ and six entries of a lower-triangular Cholesky factor $\mathbf{L}_i^{(f)} \in \mathbb{R}^{3 \times 3}$ (with positive diagonal enforced via squaring). The Cholesky factor defines a positive-definite precision matrix $\boldsymbol{\Sigma}_i^{(f)-1} = \mathbf{L}_i^{(f)} \mathbf{L}_i^{(f)\top}$, yielding an anisotropic Gaussian kernel centered at centerline coordinate $\mathbf{x}_{c,i}$:

$$\phi^{(f)}(\mathbf{x}, \mathbf{x}_{c,i}) = \exp\left(-(\mathbf{x} - \mathbf{x}_{c,i})^\top \boldsymbol{\Sigma}_i^{(f)-1} (\mathbf{x} - \mathbf{x}_{c,i})\right). \quad (5)$$

Given an arbitrary query point \mathbf{x} on the vessel wall, the predicted fields are obtained as weighted sums over all M kernels, one per centerline point:

$$\text{P}(\mathbf{x}) = \sum_{i=1}^M w_i^{(\text{P})} \phi^{(\text{P})}(\mathbf{x}, \mathbf{x}_{c,i}), \quad \text{WSS}(\mathbf{x}) = \sum_{i=1}^M w_i^{(\text{WSS})} \phi^{(\text{WSS})}(\mathbf{x}, \mathbf{x}_{c,i}). \quad (6)$$




Figure 2: Diverse coronary artery geometries from the single-vessel (top) and multi-vessel (bottom) datasets, colored by pressure. Multi-vessel cases include right coronary arteries (RCA, first three) and left coronary artery (LCA, last three) trees with varying branching complexity.

4 Dataset

We use two datasets in this work: a synthetic single-vessel dataset and a multi-vessel coronary artery dataset. Both datasets are generated with controlled anatomical variations and paired with steady-state flow simulations. Figure 2 shows representative geometries from both datasets.

4.1 Single-Vessel Dataset

To construct the single-vessel dataset, we generated synthetic coronary vessel geometries with controlled anatomical variations. The vessel length ranged from 40 mm to 70 mm, tapering ratios (defined as the ratio of outlet to inlet radius) varied between 0.6 and 0.8, and stenosis severity ranged from 30% to 70%. These parameters were used to generate vessel centerline representations consisting of coordinates and radius values, denoted as (x, y, z, r) .

Using these centerlines, we constructed 3D surface and volume meshes and performed steady-state flow simulations using OpenFOAM Jasak (2009). Physiologically relevant inlet pressures and flow rates were prescribed for each case. In total, 4,200 single-vessel geometries were generated along with their corresponding flow fields. The dataset was partitioned into 3,600 training, 400 validation, and 200 test cases, with fixed splits held constant across all experiments to ensure reproducibility.

4.2 Multi-Vessel Dataset

In addition to the single-vessel data, we constructed a multi-vessel coronary artery dataset based on anatomies from the ImageCAS dataset Zeng et al. (2023), which contains coronary artery segmentations from 1,000 patients. We extract both right coronary artery (RCA) and left coronary artery (LCA) centerlines from 141 patients, yielding distinct vessel trees that vary in branching topology and curvature. For each anatomy, we generate multiple simulation cases by (i) introducing synthetic stenoses at random locations with varying severities, and (ii) assigning different physiologically plausible inlet pressures and volumetric flow rates. This produces up to 10 variants per base geometry, each with different hemodynamic boundary conditions.

For each variant, 3D meshes were generated and steady-state flow simulations were performed using the same OpenFOAM pipeline as in the single-vessel dataset (simpleFoam solver, Newtonian fluid, no-slip walls; see Appendix A.1 for simulation details and Appendix A.2 for preprocessing). The final multi-vessel dataset consists of 4,000 training, 400 validation, and 400 test cases, with fixed splits across all experiments. Across all splits, the dataset contains 3,486 RCA and 1,314 LCA cases.

Table 1: Performance on synthetic single-vessel geometries (3,600 train / 400 val / 200 test). Relative ℓ_2 error (\downarrow) for pressure and WSS.

Method	Validation			Test		
	Pressure	WSS	Mean	Pressure	WSS	Mean
Low-Fidelity	0.639	0.488	0.563	0.635	0.488	0.562
GNOT Hao et al. (2023)	0.102	0.147	0.125	0.125	0.153	0.139
Transolver Wu et al. (2024)	0.228	0.334	0.281	0.245	0.336	0.291
ONO Xiao et al. (2023)	0.127	0.139	0.133	0.142	0.145	0.144
Ours (64 RBFs)	0.116	0.155	0.135	0.116	0.150	0.133
Ours (128 RBFs)	0.105	0.142	0.123	0.105	0.144	0.124
Ours (256 RBFs)	0.096	0.146	0.121	0.124	0.157	0.140
Ours (512 RBFs)	0.088	0.119	0.104	0.089	0.124	0.107

5 Experiments

We evaluate our method on both (i) synthetic single-vessel geometries and (ii) realistic multi-vessel coronary anatomies (RCA/LCA) to assess accuracy, robustness, and computational efficiency. Across all experiments, we compare against neural operator baselines, including GNOT Hao et al. (2023), Transolver Wu et al. (2024), and ONO Xiao et al. (2023). All models are trained with comparable parameter counts (~ 2.9 M) to predict both pressure and WSS fields (see Appendix A.4 for architecture details). All models were trained for 5,000 epochs on both single-vessel and multi-vessel data; we report results from the checkpoint with the lowest validation loss (see Appendix A.3). Evaluation metrics are defined in Appendix A.6.

5.1 Predictive Accuracy

Single-vessel results. Table 1 reports relative ℓ_2 errors on both validation (400 cases) and test (200 cases) splits for our model at 64, 128, 256, and 512 RBFs alongside the baselines. Our model obtains the lowest mean error on both splits at 512 RBFs, with particularly strong pressure predictions. Among the baselines, GNOT and ONO are competitive, while Transolver lags behind.

Multi-vessel results. We extend the evaluation to a multi-vessel coronary artery dataset derived from ImageCAS Zeng et al. (2023), containing both RCA and LCA anatomies. Because multi-vessel geometries vary widely in length and branching complexity, we study the effect of the number of RBF bases on this dataset. Table 2 reports results for our model at 128, 256, 512, and 1,024 RBFs alongside the baselines. All neural-operator baselines struggle on this more complex dataset (test mean $\ell_2 > 0.64$), while the 1D Poiseuille low-fidelity model fares even worse ($\ell_2 = 0.768$). Our model obtains substantially lower errors at every RBF count, with performance improving consistently as the number of bases increases. The best configuration (1,024 RBFs) reaches a test mean ℓ_2 of 0.309, roughly half the best baseline error. Figure 3 visualizes the pointwise absolute error on representative RCA and LCA test cases, confirming that our model produces substantially lower errors across both pressure and WSS fields. A qualitative comparison on single-vessel data is provided in Appendix A.9.

Effect of RBF size. Tables 1 and 2 show how the number of RBF bases affects accuracy. On single-vessel data, 512 RBFs achieves the best test performance (mean $\ell_2 = 0.107$), with a general trend of improvement from 64 to 512. The 256-RBF configuration is a slight exception (test mean 0.140 vs. 0.124 for 128 RBFs). On multi-vessel data, where vessel lengths and branching complexity vary more, accuracy improves consistently up to 1,024 bases (test mean $\ell_2 = 0.309$), suggesting that complex geometries benefit from higher spatial resolution in the RBF representation.

Figure 4 shows the per-case error distribution for all methods on the test set. On multi-vessel data, our model exhibits substantially lower median error and tighter interquartile range than all baselines for both

Table 2: Performance on multi-vessel RCA/LCA geometries (4,000 train / 400 val / 400 test). Relative ℓ_2 error (\downarrow) for pressure and WSS.

Method	Validation			Test		
	Pressure	WSS	Mean	Pressure	WSS	Mean
Low-Fidelity	0.852	0.659	0.756	0.852	0.683	0.768
GNOT Hao et al. (2023)	0.557	0.735	0.646	0.589	0.711	0.650
Transolver Wu et al. (2024)	0.553	0.730	0.642	0.584	0.708	0.646
ONO Xiao et al. (2023)	0.568	0.739	0.653	0.582	0.703	0.643
Ours (128 RBFs)	0.266	0.457	0.362	0.438	0.461	0.449
Ours (256 RBFs)	0.239	0.404	0.322	0.412	0.396	0.404
Ours (512 RBFs)	0.196	0.374	0.285	0.329	0.364	0.347
Ours (1024 RBFs)	0.202	0.368	0.285	0.266	0.352	0.309

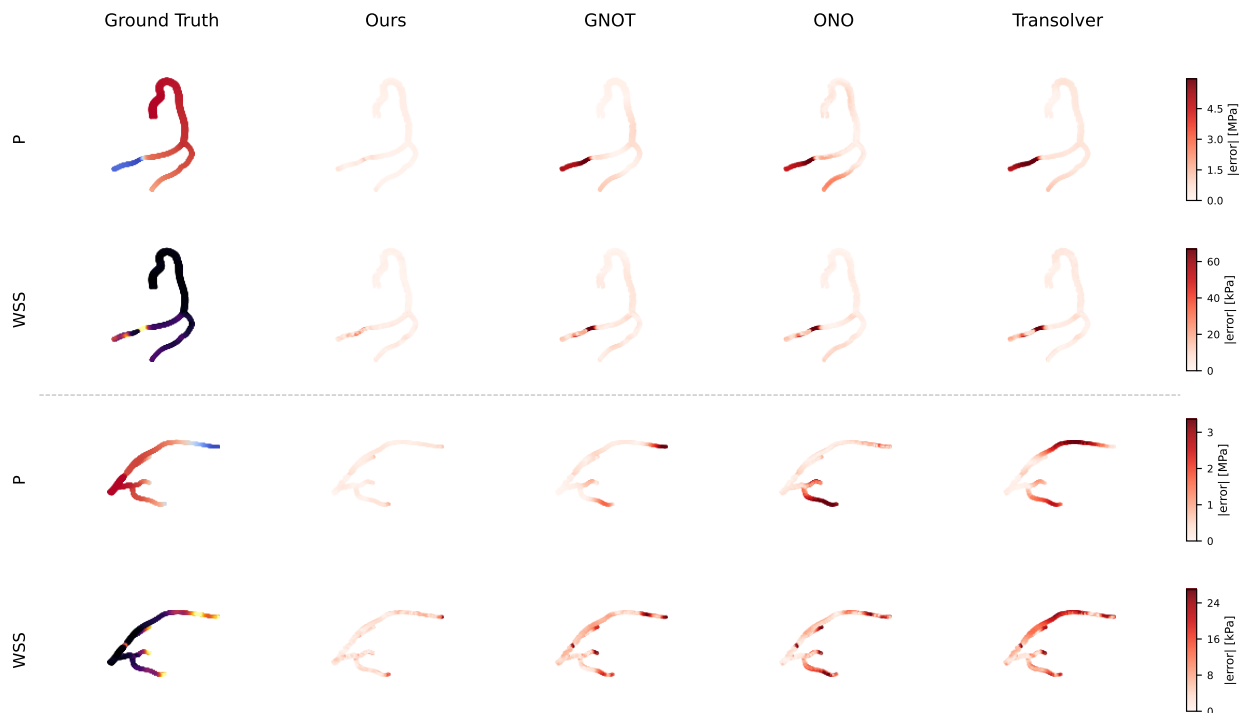


Figure 3: Pointwise absolute prediction error on two multi-vessel test geometries (top: RCA, bottom: LCA). The first column shows the ground-truth CFD solution; columns 2-5 display $|\hat{y} - y|$ for each model on a shared color scale. Our model (1024 RBFs) correctly captures the pressure drop and WSS elevation at stenosis locations, while the baseline neural operators fail to localize these regions, producing near-uniform error across the vessel wall.

pressure and WSS. An analysis of circumferential WSS reconstruction is provided in Appendix A.10, and training-set error comparisons are in Appendix A.11.

FFR evaluation. We compare predicted fractional flow reserve (FFR) against ground truth for all methods on the test set (Figure 5). Our model obtains the lowest FFR MAE on both datasets and the highest correlation with ground truth. The low-fidelity 1D Poiseuille baseline is essentially uncorrelated with the

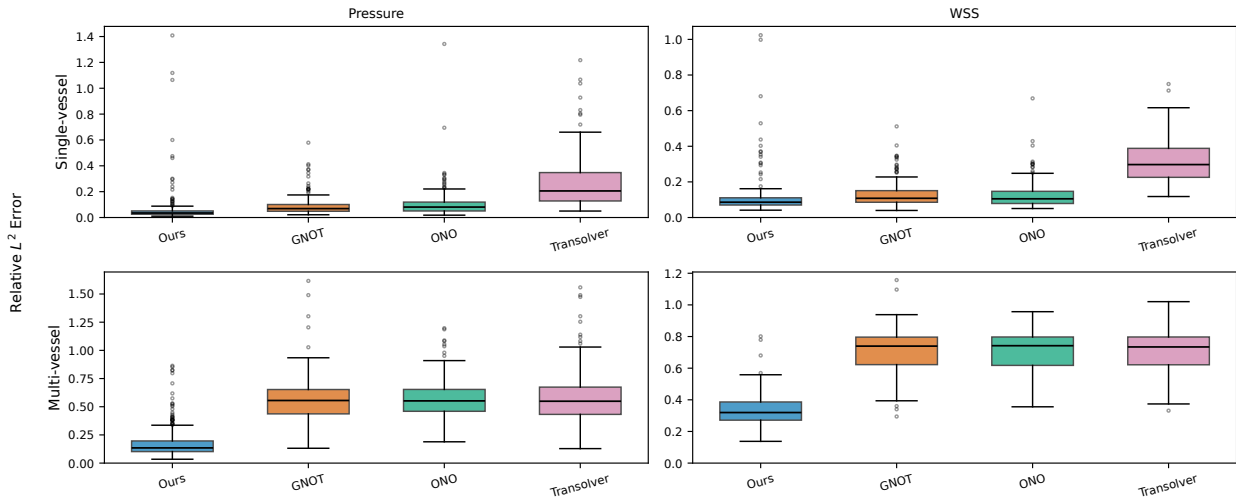


Figure 4: Per-case relative ℓ_2 error distributions on the test set. Each box shows the median, interquartile range, and outliers across all test cases. Our model achieves lower and more consistent errors than all baselines, particularly on multi-vessel data.

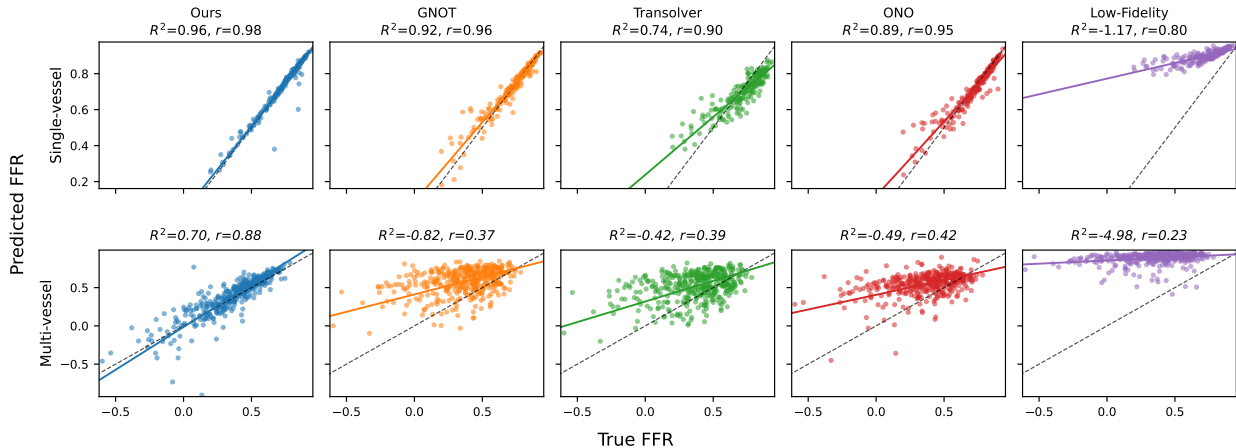


Figure 5: Predicted vs. true FFR on the test set. Top row: single-vessel; bottom row: multi-vessel. Columns show our model, three neural-operator baselines, and the 1D Poiseuille low-fidelity model. Our model obtains the highest R^2 and correlation on both datasets. The low-fidelity model systematically overestimates FFR (underestimates pressure drop), predicting most cases as non-significant, and is essentially uncorrelated with the ground truth on multi-vessel data.

ground truth on multi-vessel data. Detailed FFR metrics (MAE and classification accuracy) are reported in Appendix A.7.

5.2 Computational Efficiency

To analyze the relationship between computational cost and performance, we measure giga floating-point operations (GFLOPs) for a single forward pass under two controlled experiments: (i) fixing the number of RBFs and varying the number of query surface points, and (ii) fixing the query count and varying the number of RBFs. As shown in Figure 6, our method achieves better accuracy per FLOP than all baselines. At 128 RBF centers, our model requires only 0.53 GFLOPs, which is $13.8\times$ fewer than GNOT (7.31), $19.6\times$ fewer than ONO (10.37), and $8.7\times$ fewer than Transolver (4.59). At 64 centers, the cost drops further to just

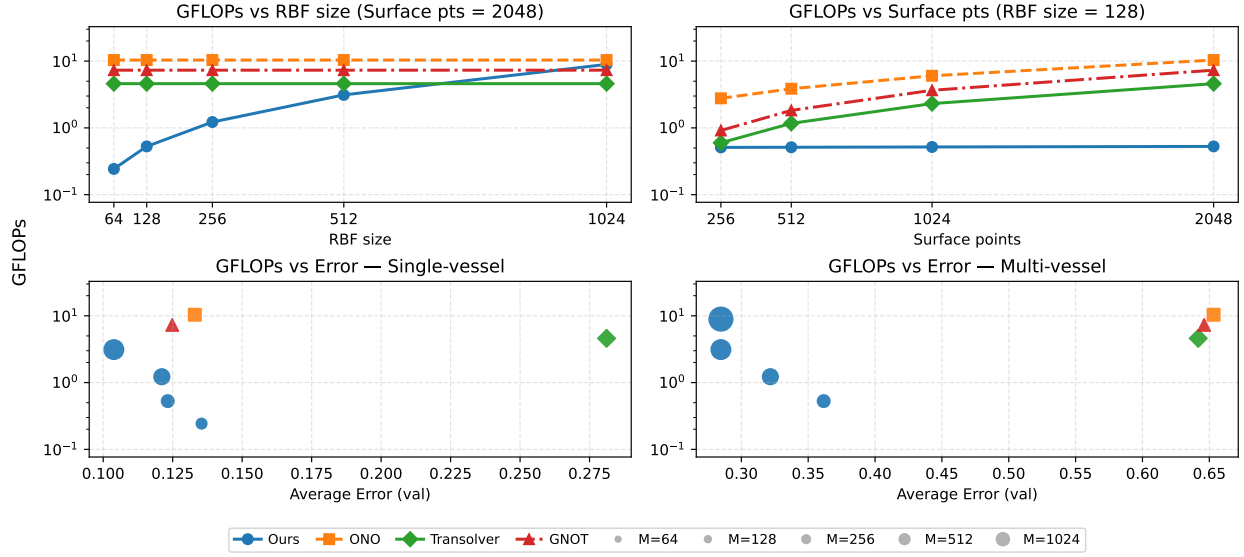


Figure 6: Computational cost vs. accuracy. (Left) GFLOPs increase with RBF count but remain comparable to baselines at 128 bases. (Middle) Our model’s cost is nearly invariant to the number of query points, unlike baselines that scale steeply. (Right) Our model achieves lower error at lower computational cost.

0.24 GFLOPs. At 512 centers the cost rises to 3.12 GFLOPs but remains $2.3\times$ lower than GNOT and $3.3\times$ lower than ONO. Crucially, the shallow RBF decoder makes inference cost nearly invariant to the number of query points (0.51-0.53 GFLOPs from 256 to 2,048 points), whereas baselines scale steeply (up to $8\times$ increase over the same range for GNOT and Transolver). Wall-clock timing benchmarks confirming these ratios are provided in Appendix A.8.

5.3 Ablation Study

We ablate four design choices on the single-vessel dataset using $M=128$ centerline points trained for 300 epochs. Starting from the default configuration (Table 3, row 1), each variant modifies exactly one component. Recall that the RBF decoder predicts, for each centerline point i and each field $f \in \{P, WSS\}$, a weight $w_i^{(f)}$ and a lower-triangular Cholesky factor $\mathbf{L}_i^{(f)}$ whose diagonal must be positive. The four axes we vary are as follows.

Conditioning mechanism. FiLM Pérez et al. (2018) applies a learned affine modulation $\tilde{\mathbf{h}}_i = \gamma(q) \odot \mathbf{h}_i + \beta(q)$ uniformly to all tokens after the encoder. The CLS alternative prepends a flow-rate token $\mathbf{t}_q = \phi_q(q)$ to the sequence, letting the transformer mix it via self-attention, and discards it before decoding.

Diagonal positivity. The diagonal entries of $\mathbf{L}_i^{(f)}$ are mapped to \mathbb{R}^+ via either $g(x) = x^2 + \epsilon$ (squared) or $g(x) = \log(1 + e^x) + \epsilon$ (softplus).

Kernel shape. The full anisotropic kernel uses the precision matrix $\Sigma_i^{-1} = \mathbf{L}_i \mathbf{L}_i^\top \in \mathbb{R}^{3 \times 3}$ (6 free parameters). The isotropic variant replaces it with a single scalar bandwidth σ_i :

$$\phi_{\text{iso}}(\mathbf{x}, \mathbf{x}_{c,i}) = \exp(-\sigma_i^2 \|\mathbf{x} - \mathbf{x}_{c,i}\|^2). \quad (7)$$

Kernel sharing. The default model learns independent kernel parameters for pressure and WSS ($2 \times 6=12$ covariance parameters per point). The shared variant uses a single set of 6 parameters for both fields.

The baseline configuration achieves the lowest mean error (0.162). Replacing anisotropic kernels with isotropic ones causes the largest WSS degradation ($0.210 \rightarrow 0.338$), consistent with the elongated geometry of coronary

Table 3: Design choice ablation on single-vessel data ($M=128$, 300 epochs). Each row modifies one component from the baseline. Metrics are relative ℓ_2 error at the best validation epoch.

Component	Variant	Pressure	WSS	Mean
<i>Baseline</i> (FiLM, squared, aniso, separate)		0.113	0.210	0.162
Conditioning	CLS token	0.112	0.224	0.168
Positivity	Softplus	0.137	0.338	0.237
Kernel shape	Isotropic	0.121	0.338	0.230
Kernel sharing	Shared (P & WSS)	0.122	0.253	0.187

arteries where directional bandwidth flexibility is beneficial. Softplus positivity enforcement produces a comparable WSS increase, suggesting that the sharper gradient of the squared mapping ($g'(x) = 2x$) is advantageous for learning the precision matrix entries. Sharing kernel parameters between pressure and WSS raises the mean error from 0.162 to 0.187, confirming that the two fields have distinct spatial correlation structures. FiLM and CLS conditioning yield similar pressure errors, but FiLM produces a lower WSS error (0.210 vs. 0.224), likely because the uniform affine modulation couples more directly with the subsequent linear decoder head.

We also test whether simply providing centerline information to an existing baseline is sufficient by augmenting GNOT with a centerline cross-attention branch. On single-vessel data, the added branch degrades performance (0.203 vs. 0.139), while on multi-vessel data it improves over the standard GNOT (0.596 vs. 0.650). Full results are in Appendix A.5.

6 Conclusion

We present a Transformer-Anisotropic RBF framework for predicting wall pressure and shear stress in coronary arteries from centerline geometry, hemodynamic descriptors, and inlet flow rate. The model provides high-fidelity, continuous predictions at low cost, obtaining lower ℓ_2 errors than neural operator baselines while requiring up to $13.8\times$ fewer FLOPs (at 128 RBF centers vs. GNOT, both evaluated with 2,048 query points). We also introduce a large-scale synthetic coronary hemodynamics dataset (4,200 single-vessel and 4,800 multi-vessel geometries with CFD-derived pressure and WSS), supporting robust training, reproducible benchmarking, and future work on multi-vessel anatomies.

Limitations and future work. All results reported in this paper are from single training runs; multi-seed evaluation is left for future work. The training data is entirely synthetic, generated from parametric vessel geometries with idealized boundary conditions (steady-state, Newtonian, rigid walls). While this controlled setup enables reproducible evaluation, it does not capture the full variability of patient-specific anatomies, pulsatile flow, or compliant vessel walls. In particular, coronary flow is inherently pulsatile and diastole-dominated; extending the framework to time-varying predictions (e.g., via time-conditioned FiLM or temporal kernel modulation) is an important direction. Bridging this domain gap is essential for clinical translation: future work should explore augmentation strategies such as physics-guided perturbations, anatomically informed generative models, and incorporation of patient-derived imaging data (e.g., from coronary CT angiography) to improve the diversity and realism of training distributions. In particular, learning-based approaches to synthetic data generation, such as using generative models to create realistic vessel geometries or to place and shape stenoses along the coronary tree, could greatly expand the training distribution beyond hand-crafted parametric rules. Validation against invasive FFR measurements and extension to patient-specific anatomies are important directions for clinical translation. Currently, our model does not provide uncertainty estimates. Incorporating both aleatoric and epistemic uncertainty through Bayesian approaches, ensembles, or probabilistic output layers would strengthen the reliability of predictions, particularly for cases outside the training distribution. Together, the framework and dataset provide a foundation for fast, anatomically aware hemodynamic surrogate modeling on synthetic coronary geometries, with potential for future extension to patient-specific clinical workflows.

Broader impact. This work is a research contribution to computational hemodynamics and is not intended for clinical use. All training and evaluation are performed on synthetic data; predictions have not been validated against clinical measurements. Any future clinical application would require rigorous regulatory approval and prospective validation.

References

- Benedikt Alkin, Andreas Fuerst, Simon Schmid, Lukas Gruber, Markus Holzleitner, and Johannes Brandstetter. Universal physics transformers: A framework for efficiently scaling neural operators. In *Neural Information Processing Systems*, 2024. URL <https://api.semanticscholar.org/CorpusID:267770281>.
- Reza Akbarian Bafghi and Maziar Raissi. Pinns-tf2: Fast and user-friendly physics-informed neural networks in tensorflow v2. *ArXiv*, abs/2311.03626, 2023. URL <https://api.semanticscholar.org/CorpusID:265043331>.
- Roberto Cavoretto, Alessandra De Rossi, and Emma Perracchione. An adaptive algorithm based on rbf-pu collocation for solving elliptic pdes. *ArXiv*, abs/1811.05193, 2018. URL <https://api.semanticscholar.org/CorpusID:53237218>.
- Gregory E. Fasshauer. Meshfree approximation methods with matlab. In *Interdisciplinary Mathematical Sciences*, 2007. URL <https://api.semanticscholar.org/CorpusID:27398882>.
- Luciano Gerber and Huw Lloyd. Revisiting chebyshev polynomial and anisotropic rbf models for tabular regression. *ArXiv*, abs/2602.22422, 2026. URL <https://api.semanticscholar.org/CorpusID:286082432>.
- Zhongkai Hao, Chengyang Ying, Zhengyi Wang, Hang Su, Yinpeng Dong, Songming Liu, Ze Cheng, Jun Zhu, and Jian Song. Gnot: A general neural operator transformer for operator learning. *ArXiv*, abs/2302.14376, 2023. URL <https://api.semanticscholar.org/CorpusID:257232579>.
- Hrvoje Jasak. Openfoam: Open source CFD in research and industry. *International Journal of Naval Architecture and Ocean Engineering*, 1(2):89–94, 2009. doi: 10.2478/IJNAOE-2013-0011. URL <https://www.sciencedirect.com/science/article/pii/S2092678216303879>. DOI ambiguity: many secondary sources cite 10.3744/JNAOE.2009.1.2.089; publisher-hosted page lists 10.2478/IJNAOE-2013-0011.
- Georgios Kissas, Yibo Yang, Enrui Hwuang, Weiwen R. Witschey, John A. Detre, and Paris Perdikaris. 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, Zong-Yi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *J. Mach. Learn. Res.*, 24:89:1–89:97, 2023. URL <https://api.semanticscholar.org/CorpusID:259149906>.
- Keying Kuang, Frances Dean, Jack B. Jedlicki, David Ouyang, Anthony Philippakis, David Sontag, and Ahmed Alaa. Med-real2sim: Non-invasive medical digital twins using physics-informed self-supervised learning. In *Neural Information Processing Systems*, 2024. URL <https://api.semanticscholar.org/CorpusID:268247645>.
- Zong-Yi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *ArXiv*, abs/2010.08895, 2020. URL <https://api.semanticscholar.org/CorpusID:224705257>.
- Zong-Yi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *J. Mach. Learn. Res.*, 24:388:1–388:26, 2022. URL <https://api.semanticscholar.org/CorpusID:250450893>.

- Zong-Yi Li, Nikola B. Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and Anima Anandkumar. Geometry-informed neural operator for large-scale 3d pdes. *ArXiv*, abs/2309.00583, 2023. URL <https://api.semanticscholar.org/CorpusID:261494027>.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218 – 229, 2019. URL <https://api.semanticscholar.org/CorpusID:233822586>.
- Huakun Luo, Haixu Wu, Hang Zhou, Lanxiang Xing, Yichen Di, Jianmin Wang, and Mingsheng Long. Transolver++: An accurate neural solver for pdes on million-scale geometries. *ArXiv*, abs/2502.02414, 2025. URL <https://api.semanticscholar.org/CorpusID:276106855>.
- James K. Min, Charles A. Taylor, Stephan Achenbach, Bon-Kwon Koo, Jonathon A Leipsic, Bjarne Linde Nørgaard, Nico H.J. Pijls, and Bernard de Bruyne. Noninvasive fractional flow reserve derived from coronary ct angiography: Clinical data and scientific principles. *JACC. Cardiovascular imaging*, 8 10: 1209–22, 2015. URL <https://api.semanticscholar.org/CorpusID:3037186>.
- Benjamin Morgan, Amal Roy Murali, George Preston, Yidnekachew Sima, Luis Alberto Marcelo Chamorro, Christos V. Bourantas, Ryo Torii, Anthony Mathur, Andreas Baumbach, Marc C. Jacob, Sergey Karabasov, and Rob Krams. A physics-based machine learning technique rapidly reconstructs the wall-shear stress and pressure fields in coronary arteries. *Frontiers in Cardiovascular Medicine*, 10, 2023. URL <https://api.semanticscholar.org/CorpusID:263279820>.
- Gabriele Nannini, Julian Suk, Patryk Rygiel, Simone Saitta, Luca Mariani, Riccardo Maragna, Andrea Baggiano, Gianluca Pontone, Jelmer M. Wolterink, and Alberto Redaelli. Learning hemodynamic scalar fields on coronary artery meshes: A benchmark of geometric deep learning models. *Computers in Biology and Medicine*, 195:110477, 2025. URL <https://api.semanticscholar.org/CorpusID:275570456>.
- Luca Pegolotti, Martin R. Pfaller, Natalia L. Rubio, Ke Ding, Rita Brugarolas Brufau, Eric F. Darve, and Alison L. Marsden. Learning reduced-order models for cardiovascular simulations with graph neural networks. *Computers in biology and medicine*, 168:107676, 2023. URL <https://api.semanticscholar.org/CorpusID:257496867>.
- Ethan Pérez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Nico H.J. Pijls, Bernard De Bruyne, Kathinka H. Peels, Pepijn H. van der Voort, Hans J.R.M. Bonnier, Jozef Bartunek, and Jacques J. Koolen. Measurement of fractional flow reserve to assess the functional severity of coronary-artery stenoses. *The New England journal of medicine*, 334 26:1703–8, 1996. URL <https://api.semanticscholar.org/CorpusID:34999507>.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 378:686–707, 2019. URL <https://api.semanticscholar.org/CorpusID:57379996>.
- Patryk Rygiel, Paweł Pluszka, Maciej Zięba, and Tomasz Konopczynski. Centerlinepointnet++: A new point cloud based architecture for coronary artery pressure drop and vffr estimation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. Springer, 2023. URL <https://api.semanticscholar.org/CorpusID:261049553>.
- Sethuraman Sankaran, Mahdi Esmaily Moghadam, Andrew M. Kahn, Elaine E. Tseng, Julius M. Guccione, and Alison L. Marsden. Patient-specific multiscale modeling of blood flow for coronary artery bypass graft surgery. *Annals of biomedical engineering*, 40:2228 – 2242, 2012. URL <https://api.semanticscholar.org/CorpusID:11594987>.

- Julian Suk, Baris Imre, and Jelmer M. Wolterink. Lab-gatr: geometric algebra transformers for large biomedical surface and volume meshes. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*. Springer, 2024a. URL <https://api.semanticscholar.org/CorpusID:268531880>.
- Julian Suk, Guido Nannini, Patryk Rygiel, Christoph Brune, Gianluca Pontone, Alberto Redaelli, and Jelmer M. Wolterink. Deep vectorised operators for pulsatile hemodynamics estimation in coronary arteries from a steady-state prior. *Computer Methods and Programs in Biomedicine*, 262:108958, 2024b. URL <https://api.semanticscholar.org/CorpusID:273375640>.
- Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Neural Information Processing Systems*, 2020. URL <https://api.semanticscholar.org/CorpusID:219791950>.
- Sukirt Thakur, Marcus Roper, Yang Zhou, Reza Akbarian Bafghi, Brahmajee. K. Nallamothu, C. Alberto Figueroa, Srinivas Paruchuri, Scott Burger, and Maziar Raissi. Punch: Physics-informed uncertainty-aware network for coronary hemodynamics. *ArXiv*, abs/2601.17192, 2026. URL <https://api.semanticscholar.org/CorpusID:285050399>.
- Pim A. L. Tonino, Bernard de Bruyne, Nico H.J. Pijls, Uwe Siebert, Fumiaki Ikeno, Marcel van ’t Veer, Volker Klauss, Ganesh Manoharan, Thomas Engstrøm, Keith G Oldroyd, Peter N Ver Lee, Philip A MacCarthy, and William F. Fearon. Fractional flow reserve versus angiography for guiding percutaneous coronary intervention. *The New England journal of medicine*, 360 3:213–24, 2009. URL <https://api.semanticscholar.org/CorpusID:18683765>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing Systems*, 2017. URL <https://api.semanticscholar.org/CorpusID:13756489>.
- Salim S. Virani and et al. Heart disease and stroke statistics—2021 update: A report from the american heart association. *Circulation*, 143(8):e254–e743, 2021. doi: 10.1161/CIR.0000000000000950.
- Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Karianne J. Bergen, Carla P. Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec, Tie-Yan Liu, Arjun K. Manrai, Debora S. Marks, Bharath Ramsundar, Le Song, Jimeng Sun, Jian Tang, Petar Velickovic, Max Welling, Linfeng Zhang, Connor W. Coley, Yoshua Bengio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*, 620:47–60, 2023. URL <https://api.semanticscholar.org/CorpusID:260384616>.
- Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why pinns fail to train: A neural tangent kernel perspective. *J. Comput. Phys.*, 449:110768, 2020. URL <https://api.semanticscholar.org/CorpusID:220845655>.
- Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries. *ArXiv*, abs/2402.02366, 2024. URL <https://api.semanticscholar.org/CorpusID:267411758>.
- Zipeng Xiao, Zhongkai Hao, Bokai Lin, Zhijie Deng, and Hang Su. Improved operator learning by orthogonal attention. *ArXiv*, abs/2310.12487, 2023. URL <https://api.semanticscholar.org/CorpusID:264306159>.
- An Zeng, Chunbiao Wu, Meiping Huang, Jian Zhuang, Shanshan Bi, Dan Pan, Najeeb Ullah, Kaleem Nawaz Khan, Tianchen Wang, Yiyu Shi, Xiawei Li, Guisen Lin, and Xiaowei Xu. 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. URL <https://api.semanticscholar.org/CorpusID:253264830>.

A Experimental Details

This appendix provides additional details supporting Section 5, covering CFD simulation setup, data preprocessing, training configuration, model architectures, a centerline input ablation for GNOT, evaluation metrics, FFR evaluation, wall-clock inference timing, and supplementary results.

A.1 CFD Simulation Setup

All flow simulations use OpenFOAM Jasak (2009) (version 11) with the `simpleFoam` steady-state incompressible solver. Volume meshes are generated with `blockMesh` followed by `snappyHexMesh` for body-fitted refinement. The fluid is modeled as Newtonian with kinematic viscosity $\nu = 3.77 \text{ mm}^2/\text{s}$. Inlet boundary conditions prescribe a fixed pressure, while the outlet uses a parabolic (Poiseuille-type) velocity profile. Walls are treated as no-slip boundaries. Each simulation runs for 4,000 iterations. Wall shear stress is computed via the `wallShearStress` function object. The output fields (kinematic pressure p/ρ and WSS magnitude) are sampled at the surface mesh vertices. Inlet pressures range from 10,000 to 16,000 mm^2/s^2 (kinematic) and maximum outlet velocities from 0.09 to 0.11 mm/s for both datasets.

A.2 Data Preprocessing

All inputs undergo a fixed preprocessing pipeline before being fed to the model. For each case, surface and centerline positions are centered by subtracting the mean surface position, so that all geometries share a common origin. Surface and centerline coordinates are min-max normalized to $[0, 1]$ using global minimum and maximum values (with a buffer of 1.0) computed across all data splits. Four scalar fields are standardized to zero mean and unit variance using statistics computed on the training split only: surface pressure, surface WSS magnitude, centerline radius, and centerline flowrate. Surface normals and centerline tangent vectors are not normalized. During training, 2,048 surface points are randomly sampled per case. For the RBF model, M centerline points are randomly sampled and sorted by index to preserve spatial ordering along the vessel.

A.3 Training Configuration

We optimize all models with AdamW (learning rate 10^{-3} , weight decay 5×10^{-4}) with batch size 128. We train all models for 5,000 epochs on both datasets. Single-vessel experiments were run on an NVIDIA Quadro RTX 8000 GPU; multi-vessel experiments were run on an NVIDIA A100 80GB PCIe GPU. On the multi-vessel dataset, training takes approximately 11 h for our RBF model (512 centers), 22 h (1,024 centers), 38 h for GNOT, 45 h for ONO, and 16 h for Transolver on the A100. The learning rate schedule uses a 1,000-epoch linear warmup from 0 to full LR, followed by cosine annealing to 0. The random seed is fixed to 1234.

For the single-vessel dataset we train on 3,600 cases, validate on 400, and test on 200; for the multi-vessel dataset we train on 4,000 cases, validate on 400, and test on 400. The training objective minimizes

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \|\hat{p}_i - p_i\|^2 + \frac{1}{N} \sum_{i=1}^N \|\widehat{\text{wss}}_i - \text{wss}_i\|^2, \quad (8)$$

where \hat{p}_i and $\widehat{\text{wss}}_i$ are predictions and p_i , wss_i are targets, all in z-score normalized space. We select the checkpoint with the lowest validation loss.

A.4 Model Configurations

All baseline models operate directly on surface points and receive only the scalar flow rate as a global condition; our RBF model instead processes centerline points through the transformer encoder and evaluates the learned kernels at the surface query locations. All baselines were reimplemented from the respective official codebases and adapted to our hemodynamic prediction task (surface-based regression of pressure and WSS with FiLM Pérez et al. (2018) flow-rate conditioning). Hyperparameters (depth, width, attention heads) were tuned so that all models have comparable parameter counts ($\sim 2.9\text{M}$).

Ours (RBF). The encoder receives M centerline points, each described by position and radius (x, y, z, r) ; the scalar flow rate is injected via FiLM after the encoder. The decoder evaluates the learned anisotropic RBF kernels at 2,048 surface query positions (x, y, z) . Architecture: transformer encoder with 5 layers, model dimension 256, feedforward dimension 512, 1 attention head, and dropout 0.1. Centerline tokens are built via Fourier embeddings (hidden dimension 48) projected to the model dimension with a linear layer; learned positional encodings are added over the token sequence. Total: 2.82-3.05M parameters depending on the number of RBF centers (128-1,024).

ONO Xiao et al. (2023). Receives 2,048 surface points (x, y, z) and their normals (n_x, n_y, n_z) as input; the scalar flow rate is injected via FiLM. 4 layers, model dimension 240, 1 attention head, dropout 0.1, with orthogonal projection dimension $\psi_{\text{dim}}=8$ and linear attention. Total: 2.91M parameters.

GNOT Hao et al. (2023). Receives 2,048 surface points (x, y, z) and their normals (n_x, n_y, n_z) as input; the scalar flow rate is broadcast to every point and embedded jointly. 4 layers, model dimension 126, 1 attention head, dropout 0.1, with 4 mixture-of-experts and heterogeneous cross-attention with normalized linear attention. Total: 2.97M parameters.

Transolver Wu et al. (2024). Receives 2,048 surface points (x, y, z) and their normals (n_x, n_y, n_z) as input; the scalar flow rate is injected via FiLM. 4 layers, model dimension 256, 1 attention head, dropout 0.1, with 32 physics-aware slices. Total: 2.93M parameters.

A.5 Centerline Input Ablation

To test whether simply providing centerline information to an existing baseline is sufficient, we ablate the input configuration of GNOT, which natively supports heterogeneous inputs on different grids via its cross-attention branches Hao et al. (2023). We evaluate three configurations: *normals-only* (surface positions, normals, and flow rate), *centerline-only* (surface positions, centerline positions and radii, and flow rate, no normals), and *both* (all inputs). ONO and Transolver operate exclusively on surface tokens through self-attention and have no mechanism to ingest a separate input sequence, so this ablation is limited to GNOT.

Table 4 reports the results. On single-vessel data, normals-only GNOT achieves the best test error (0.139), while adding the centerline branch degrades performance; the centerline-only variant reaches 0.173 and the combined variant 0.203. This suggests that, for simple tubular geometries, surface normals already encode sufficient geometric information, and that the additional centerline branch introduces optimization difficulty. On multi-vessel data, the trend reverses: the combined variant improves over normals-only GNOT (0.596 vs. 0.650), indicating that explicit centerline geometry becomes beneficial when vessel topology is complex.

Table 4: GNOT input ablation. Test relative ℓ_2 error for three input configurations.

Dataset	Input config	Params	Pressure	WSS	Mean
Single-vessel	Normals only	2.97M	0.125	0.153	0.139
	Centerline only	2.97M	0.120	0.227	0.173
	Both	3.15M	0.158	0.248	0.203
Multi-vessel	Normals only	2.97M	0.589	0.711	0.650
	Centerline only	2.97M	0.556	0.642	0.599
	Both	3.15M	0.540	0.651	0.596

A.6 Evaluation Metrics

Fractional Flow Reserve (FFR). FFR is a clinically important hemodynamic index used to assess the functional significance of coronary artery stenoses Pijls et al. (1996). Clinically, it is defined as the ratio of mean distal coronary pressure to mean aortic pressure during maximal hyperemia. In our steady-state CFD

simulations, we compute a pointwise pressure ratio at every surface point \mathbf{x} :

$$\text{FFR}(\mathbf{x}) = \frac{P(\mathbf{x})}{P_a}, \quad (9)$$

where $P(\mathbf{x}) = \tilde{p}(\mathbf{x})\rho/1000 + P_a$ converts the OpenFOAM kinematic pressure output \tilde{p} (in mm^2/s^2) to absolute pressure using blood density $\rho = 1.06 \text{ g/mL}$, and P_a is the prescribed inlet (aortic) pressure. We report the minimum FFR over the vessel surface, $\text{FFR}_{\min} = \min_{\mathbf{x}} \text{FFR}(\mathbf{x})$, as the per-case summary metric. An FFR value below 0.80 is typically considered indicative of flow-limiting stenosis and is used to guide revascularization decisions Tonino et al. (2009).

Low-fidelity baseline. We compute a 1D pressure profile along the vessel centerline using the Hagen-Poiseuille law. For each centerline segment with local radius $r(s)$, the axial pressure gradient is

$$\frac{dP}{ds} = \frac{8\mu Q}{\pi r(s)^4}, \quad (10)$$

where $\mu = 0.004 \text{ Pa}\cdot\text{s}$ is the dynamic viscosity and Q is the volumetric flow rate. The cumulative pressure drop is integrated along the centerline to obtain $P(s)$, from which FFR is computed as $(P_a - \Delta P(s))/P_a$. For multi-vessel cases, flow is split equally among downstream branches at each bifurcation.

A.7 FFR Evaluation

Tables 5 and 6 report FFR classification metrics on the test set for single-vessel and multi-vessel data, respectively.

Single-vessel (Table 5; threshold $\text{FFR} < 0.8$; 134 positive, 66 negative). Our model (512 RBFs) achieves the lowest MAE (0.013) with sensitivity 97.8% and specificity 90.9%. GNOT obtains comparable classification performance (sensitivity 98.5%, specificity 90.9%) but with double the MAE (0.026). The low-fidelity 1D Poiseuille baseline predicts nearly all cases as non-significant ($\text{FFR} \geq 0.8$), yielding only 1.5% sensitivity despite 100% specificity.

Multi-vessel (Table 6; threshold $\text{FFR} < 0.5$; 273 positive, 127 negative). All 400 multi-vessel test cases have ground-truth $\text{FFR} < 0.8$, reflecting the high prevalence of hemodynamically significant stenoses in the ImageCAS-derived dataset. Classification at the standard 0.8 threshold is therefore uninformative (all models achieve $\geq 97.8\%$ accuracy trivially). We instead evaluate at $\text{FFR} < 0.5$, which distinguishes severe from moderate stenoses and yields a more balanced class distribution. At this threshold, our model (1024 RBFs) achieves 86.0% accuracy with 82.4% sensitivity and 93.7% specificity, while all baselines perform near chance level (50 to 56% accuracy). The MAE of our model (0.084) is more than $2.5\times$ lower than the best baseline, confirming that our predictions are quantitatively closer to the ground truth across the full FFR range.

Table 5: FFR evaluation on single-vessel test data (threshold $\text{FFR} < 0.8$; 134 positive, 66 negative cases).

Method	MAE (\downarrow)	Sensitivity	Specificity	Accuracy	F1
Low-Fidelity	0.187	0.015	1.000	0.340	0.029
GNOT Hao et al. (2023)	0.026	0.985	0.909	0.960	0.971
Transolver Wu et al. (2024)	0.062	0.963	0.409	0.780	0.854
ONO Xiao et al. (2023)	0.031	0.993	0.788	0.925	0.947
Ours (512 RBFs)	0.013	0.978	0.909	0.955	0.967

Figure 7 shows a Bland-Altman analysis of FFR agreement for our best model on each dataset. On single-vessel data, the mean bias is near zero with narrow limits of agreement, confirming accurate FFR estimation. On multi-vessel data, the bias remains small but the limits of agreement are wider, reflecting the greater difficulty of the multi-vessel prediction task.

Table 6: FFR evaluation on multi-vessel test data (threshold FFR < 0.5; 273 positive, 127 negative cases). All test cases have FFR < 0.8; we use 0.5 to distinguish severe from moderate stenoses.

Method	MAE (\downarrow)	Sensitivity	Specificity	Accuracy	F1
Low-Fidelity	0.523	0.004	1.000	0.320	0.007
GNOT Hao et al. (2023)	0.249	0.330	0.866	0.500	0.474
Transolver Wu et al. (2024)	0.217	0.462	0.764	0.557	0.587
ONO Xiao et al. (2023)	0.218	0.359	0.866	0.520	0.505
Ours (1024 RBFs)	0.084	0.824	0.937	0.860	0.889




Figure 7: Bland-Altman plots of predicted vs. true FFR for our model on single-vessel (left, 512 RBFs) and multi-vessel (right, 1024 RBFs) test sets. Solid line: mean bias; dashed lines: ± 1.96 standard deviation limits of agreement.

A.8 Wall-Clock Inference Timing

Table 7 reports wall-clock inference latency on an NVIDIA A100 80GB GPU with 2,048 surface query points and batch size 32. Each model was run for 10 warm-up forward passes followed by 100 timed passes with explicit `cuda.synchronize()`; we report per-sample time. A batch size of 32 is used to ensure the GPU is compute-saturated, so that wall-clock differences reflect actual computational cost rather than kernel-launch overhead. At 128 RBFs, our model runs in 0.18 ms/sample, 14 \times faster than GNOT (2.54 ms), 12.8 \times faster than ONO (2.31 ms), and 4.5 \times faster than Transolver (0.81 ms), consistent with the theoretical FLOP ratios. Even at 1,024 RBFs (1.32 ms), our model remains faster than all baselines.

Table 7: Wall-clock inference time on an NVIDIA A100 GPU with 2,048 surface points and batch size 32.

Method	Params (M)	GFLOPs	Time (ms/sample)
GNOT Hao et al. (2023)	2.97	7.31	2.54
Transolver Wu et al. (2024)	2.93	4.59	0.81
ONO Xiao et al. (2023)	2.91	10.37	2.31
Ours (128 RBFs)	2.82	0.53	0.18
Ours (256 RBFs)	2.85	1.22	0.33
Ours (512 RBFs)	2.92	3.12	0.63
Ours (1024 RBFs)	3.05	8.93	1.32

A.9 Qualitative Single-Vessel Comparison

Figure 8 shows the pointwise absolute prediction error on two single-vessel test cases. On this simpler geometry, our model, GNOT, and ONO all achieve low errors, while Transolver lags behind on both examples.



Figure 8: Pointwise absolute prediction error on two single-vessel test geometries, following the same layout as Figure 3. Our model (128 RBFs) attains relative ℓ_2 errors of 3.51%/5.83% and 2.74%/5.18% (P/WSS). GNOT (3.15%/5.95%, 3.87%/5.83%) and ONO (5.29%/8.16%, 2.32%/5.11%) are competitive on single-vessel data, while Transolver shows higher error (11.54%/20.85%, 17.19%/21.65%).




Figure 9: Circumferential WSS profile at two axial stations of a single-vessel test case. Left: stenotic throat (minimum radius); right: healthy upstream section. Scatter points show individual surface samples; curves are smoothed fits. The anisotropic RBF decoder captures the circumferential WSS variation at both stations.

A.10 Circumferential WSS Reconstruction

A natural concern with centerline-anchored RBF kernels is whether they can capture circumferential variation in WSS, which is non-axisymmetric near stenoses. Figure 9 plots the predicted and ground-truth WSS as a function of circumferential angle θ at two axial stations of a single-vessel test case: (left) the stenotic throat, where WSS is elevated and varies around the circumference, and (right) a healthy upstream section. At both stations, the anisotropic RBF decoder closely tracks the ground-truth circumferential profile, confirming that the learned precision matrices adapt to capture non-axisymmetric WSS patterns despite the kernels being anchored on the 1D centerline.

A.11 Supplementary Results

Tables 8 and 9 report training-set errors at the final training epoch for single-vessel and multi-vessel datasets, respectively. Comparing with the test results in Table 1 and Table 2 reveals the generalization gap for each model. On single-vessel data, GNOT achieves the lowest training error (0.017) but degrades to 0.139 at test time ($\sim 8\times$ gap), indicating severe overfitting. ONO shows a similar pattern (0.033 \rightarrow 0.144, $\sim 4\times$ gap). Transolver does not converge well, with high training error (0.173) and correspondingly high test error (0.291). Our model (512 RBFs) reaches a training error of 0.030 with the smallest test-time degradation (0.107, $\sim 3.6\times$ gap), demonstrating stronger generalization.

On multi-vessel data, the baselines also converge to low training errors by the final epoch (GNOT 0.168, ONO 0.205, Transolver 0.168) but degrade catastrophically at test time (all > 0.64), with best-validation checkpoints selected at very early epochs out of 5,000. Our model scales consistently with the number of RBF bases, reaching a training error of 0.126 at 1,024 RBFs with a test error of 0.309 ($\sim 2.5\times$ gap), substantially outperforming all baselines.

Table 8: Single-vessel training set (3,600 cases). Relative ℓ_2 error at the final training epoch.

Method	Pressure	WSS	Mean
GNOT Hao et al. (2023)	0.011	0.023	0.017
Transolver Wu et al. (2024)	0.135	0.212	0.173
ONO Xiao et al. (2023)	0.022	0.043	0.033
Ours (64 RBFs)	0.037	0.096	0.067
Ours (128 RBFs)	0.027	0.067	0.047
Ours (256 RBFs)	0.023	0.059	0.041
Ours (512 RBFs)	0.019	0.041	0.030

Table 9: Multi-vessel training set (4,000 cases). Relative ℓ_2 error at the final training epoch.

Method	Pressure	WSS	Mean
GNOT Hao et al. (2023)	0.114	0.221	0.168
Transolver Wu et al. (2024)	0.116	0.219	0.168
ONO Xiao et al. (2023)	0.141	0.269	0.205
Ours (128 RBFs)	0.139	0.375	0.257
Ours (256 RBFs)	0.106	0.311	0.208
Ours (512 RBFs)	0.079	0.241	0.160
Ours (1024 RBFs)	0.062	0.191	0.126