



Learning hemodynamic scalar fields on coronary artery meshes: A benchmark of geometric deep learning models

Guido Nannini ^a^{*}, Julian Suk ^b, Patryk Rygiel ^b, Simone Saitta ^{c,d}, Luca Mariani ^a,
Riccardo Maranga ^e, Andrea Baggiano ^{e,f}, Gianluca Pontone ^{e,g}, Jelmer M. Wolterink ^b,
Alberto Redaelli ^a

^a Department of Electronics Information and Bioengineering, Politecnico di Milano, Milan, Italy

^b Department of Applied Mathematics and Technical Medical Center, University of Twente, Enschede, The Netherlands

^c Department of Biomedical Engineering and Physics, Amsterdam UMC, Amsterdam, The Netherlands

^d Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands

^e Department of Perioperative Cardiology and Cardiovascular Imaging, Centro Cardiologico Monzino IRCCS, Milan, Italy

^f Department of Clinical Sciences and Community Health, University of Milan, Milan, Italy

^g Department of Biomedical, Surgical and Dental Sciences, University of Milan, Milan, Italy

ARTICLE INFO

Keywords:

Coronary arteries
Geometric deep learning
FFR

ABSTRACT

Coronary artery disease involves the narrowing of coronary vessels due to atherosclerosis and is currently the leading cause of death worldwide. The gold standard for its diagnosis is the fractional flow reserve (FFR) examination, which measures the trans-stenotic pressure ratio during maximal vasodilation. However, the invasiveness and cost of this procedure have prompted the development of computer-based virtual FFR (vFFR) computation, which simulates coronary flow using computational fluid dynamics (CFD) techniques. Geometric deep learning algorithms have recently shown the capability to learn features on meshes, including applications in cardiovascular research. In this work, we aim to conduct a comprehensive empirical analysis of different backends for predicting vFFR fields in coronary arteries, serving as surrogates for CFD simulations. We evaluate six different backends and compare their performance in learning hemodynamics on meshes using CFD solutions as ground truth. This study is divided into two main parts: i) First, we use a dataset of 1,500 synthetic bifurcations of the left coronary artery. Each model is trained to predict various pressure-related fields, from which the vFFR field is reconstructed. We compare the models' performance when different learning variables are used during training. ii) Second, we use a dataset of 427 patient-specific CFD simulations from a previous study by our group. Here, we repeat the experiments conducted on the synthetic dataset, focusing on the learning variable that yielded the best performance in the synthetic dataset. Most backends achieved very good performance on the synthetic dataset, particularly when learning the pressure drop over the manifold. For other network output variables (e.g., pressure and the vFFR field), transformer-based backends outperformed all other architectures. When trained on patient-specific data, transformer-based architectures were the only ones to achieve strong performance, both in terms of average per-point error and in accurately predicting vFFR in stenotic lesions. Our findings indicate that various geometric deep learning backends can serve as effective CFD surrogates for problems involving simple geometries. However, for tasks involving datasets with complex and heterogeneous topologies, transformer-based networks are the optimal choice. Additionally, pressure drop emerged as the optimal network output for learning pressure-related fields.

1. Introduction

Coronary artery disease (CAD) is the leading cause of death worldwide [1]. It consists of the buildup of fibro-lipidic and calcific plaque in the walls of the coronary arteries, leading to narrowing (i.e., stenosis) of the vessel lumen and reduced blood flow. Ultimately, the limited

supply of oxygen can result in myocardial infarction and death [2]. In recent years, fractional flow reserve (FFR), defined as the ratio between the pressure downstream of a stenosis and the pressure at the coronary ostium, has emerged as the gold standard to assess the functional severity of a stenosis [3]. FFR is invasively measured

* Corresponding author.

E-mail address: guido.nannini@polimi.it (G. Nannini).

<https://doi.org/10.1016/j.combiomed.2025.110477>

Received 23 February 2025; Received in revised form 14 May 2025; Accepted 27 May 2025

Available online 18 June 2025

0010-4825/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

in-vivo with a pressure probe during maximal hyperemia induced by vasodilator administration. Stenting is recommended when FFR < 0.8, and medical treatment otherwise [4]. Several clinical trials have shown that invasive FFR is a more robust criterion for identifying potential culprit lesions compared to stenosis degree alone [5–7]. However, its use in clinical practice remains limited due to the invasiveness and cost of the procedure [5,8,9].

Non-invasive computation of FFR requires knowledge of the pressure field within the coronary volume and along the vessel walls. Computational fluid dynamics (CFD) enables *in-silico* simulation of blood flow within the lumen of coronary arteries, typically reconstructed from computed tomography (CT). CFD results can be used to obtain a non-invasive virtual measurement of FFR, known as vFFR [10–12]. Although the primary advantage of vFFR is its ability to overcome the invasiveness of FFR measurement, a full CFD simulation of coronary flow over a single cardiac cycle can be highly computationally demanding (>24 h) and requires expert knowledge throughout the process [13,14]. Different strategies to speed up the computation of vFFR have been proposed in the literature, such as using reduced order modeling [15–17] and hybrid 3D-1D models [18]. While being computationally less expensive, such models fail to capture the complexity of 3D flows, potentially introducing errors in estimating the pressure and velocity field [19]. Additionally, parameters such as time and space discretization, as well as boundary conditions (BCs), are subject to inter-operator variability, introducing uncertainty in the solution of the CFD simulation and making the overall process challenging to standardize [20,21]. In a previous study by our group [22], we proposed a CFD pipeline that automatically tunes the BCs while computing the solution, partially solving the inter-operator bias. Furthermore, we demonstrated that vFFR can be accurately computed by simulating coronary flow in a steady state with BCs averaged over the cardiac cycle, yielding good diagnostic accuracy ($\geq 90\%$), comparable to that of pulsatile simulations with time-varying BCs. Steady-state CFD preserves the 3D flow development while significantly reducing computational costs, although it still requires 30 to 60 min per patient.

Data-driven models allow an additional step forward to accelerate the simulation process. In cardiovascular modeling, deep learning has been used to automate different aspects of the CFD pipeline, such as image segmentation and 3D model reconstruction [23–27]. However, solving the physical problem (e.g., the Navier–Stokes equations), which represents the most computationally demanding step, is still typically reliant on numerical algorithms. Hence, (partial) replacement of this step by data-driven deep learning methods is an active research domain.

Machine learning models can learn to generalize hemodynamics from a training dataset to unseen data based on geometric (e.g., the anatomy) and physical (e.g., the BCs) features. One approach to this involves parametrizing the lumen centerline or the mesh of the arterial wall in a 1D or 2D Euclidean domain and using fully connected networks, such as multilayer perceptrons (MLPs) or convolutional neural networks (CNNs), to estimate scalar or vector fields based on shape descriptors [28,29].

Alternatively, models can be developed that operate directly on the original artery geometry without requiring a mapping to a lower-dimensional Euclidean domain. Such models typically fall in the realm of geometric deep learning (GDL), i.e., deep learning models for data in non-Euclidean domains, such as graphs, meshes, and point clouds [30, 31]. For example, Rygiel et al. [32] proposed a variant to PointNet++ for predicting the pressure drop on a dataset of synthetic anatomies of both left and right coronary arteries. Wang et al. [33] trained a point cloud network (PCN) to infer the pressure and velocity field in the carotid artery and in the aorta; Zhang et al. [34] used a physics-informed variant of the PointNet to predict the 3D velocity field in the abdominal aorta; Suk et al. [35] used a multiscale E(3)-steerable graph neural network (GNN) for estimating the velocity field in synthetic coronary arteries. Suk et al. [36] presented a variant to the geometric

algebra transformer (GATr) model [37], specifically for large biomedical meshes (LaB-GATr), achieving very good wall shear stress (WSS) and velocity estimation in large synthetic coronary arteries models. Pegolotti et al. [38] proposed a data-driven reduced order model based on a GNN that learns 1D quantities along the centerline of the vascular tree. Likewise, Gharleghi et al. [39] proposed a method to map the original surface to a fixed 2D Euclidean domain and applying CNN multiscale-neural network for WSS inference on the surface of the bifurcation of the left coronary artery. The main drawback of such techniques is that parametrization may not be feasible for complex shapes (e.g., in the presence of pathology), limiting their applicability. While demonstrating the potential of GDL techniques for predicting hemodynamic quantities on and in meshes, most studies are conducted on synthetic or simplified geometries due to the commonly limited availability of large patient-specific datasets. This makes it difficult to identify common patterns and answer questions about the right data representation and the most appropriate model architecture and size.

In this work, we address these questions for the problem of estimating hemodynamic fields on the surface of real-life patients' coronary artery 3D anatomies. Specifically, we focus on the problem of learning the scalar pressure field on a surface from which we compute the corresponding FFR field. We perform an extensive empirical comparison of architectures that are different in nature and either well-established or increasingly used in biomedical applications and beyond. These are the multi-layer perceptron (MLP); PointNet++ [40]; DiffusionNet [41]; DeltaConv [42]; and two transformer models, one exploiting geometric algebra representations (i.e., LaB-GATr) and the other, referred to as *vanilla* transformer, operating under linear algebra (i.e., LaB-VaTr) [36]. We train and validate these models using results from steady CFD simulations on two datasets: a synthetic set of 1500 anatomies representing the left coronary artery bifurcation and a patient-specific set of 427 cases validated against invasive FFR measurements. We train our models using both geometric and physical input features, representing the 3D anatomy and the BCs of the corresponding CFD simulation. Moreover, we compare the diagnostic accuracy of the methods.

The main contributions of this paper are the following: (i) We demonstrate the feasibility of leveraging GDL to predict vFFR scalar fields in real-world patient cases; (ii) We systematically compare the performance of various architectures and model sizes for inferring scalar fields on the surface of coronary arteries, finding the superior performance of transformer models in this task; (iii) We train our models to learn various pressure-related variables, demonstrating that predicting the pressure drop is optimal for deriving both the pressure and vFFR fields.

2. Methods

The aim of this study was to systematically compare different models for learning physics-based pressure-related fields on meshes of coronary arteries. Fig. 1 provides a schematic representation of our systematic approach in two different data sets. We first performed experiments using a synthetically generated dataset of geometries representing the left coronary bifurcation. Since our primary interest is in predicting FFR, we trained each model to learn one of three related scalar fields: (i) pressure; (ii) pressure drop, and (iii) the vFFR field. Based on the results obtained in the synthetic dataset, we then selected one output variable among these three for further experiments in patient-specific anatomies. For a fair comparison between models, all models were fed the same vertex-specific input features. Under the assumption of laminar and developed flow, the pressure field (and consequently the pressure drop and the FFR) can be considered constant on the cross-section of the vessel, thus fully represented by the value assumed at the wall. Thus, we consider the pressure to be constant across the vessel cross-section, allowing us to train our models exclusively on the surface mesh while excluding internal nodes of the volumetric mesh. This approach significantly reduces the computational cost both in training and inference.

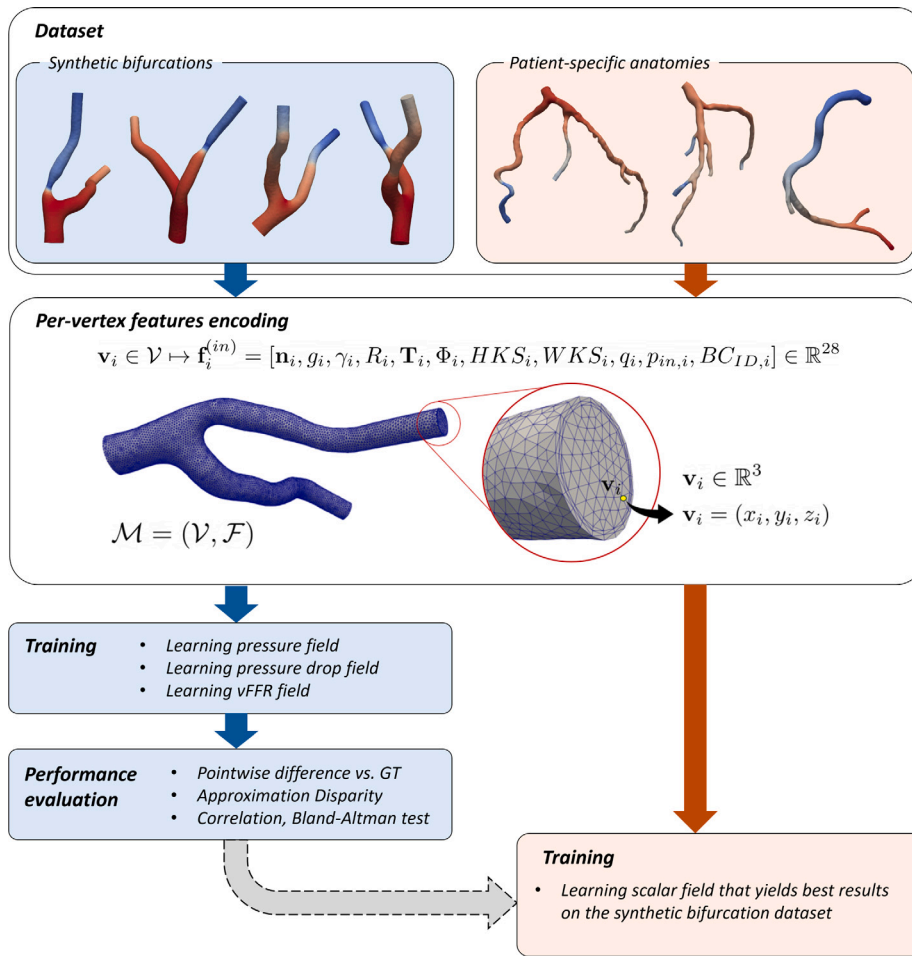


Fig. 1. Representation of the workflow adopted for the training of all models. First, models are trained on the synthetic dataset to learn one of three output variables; after evaluation, they are retrained on the patient-specific dataset using the best-performing variable, with per-vertex feature encoding applied consistently across both datasets. Blue boxes and arrows refer to the pipeline applied to synthetic datasets; orange ones refer to patient-specific data.

2.1. Data

2.1.1. Synthetic dataset

We generated a dataset of 1500 anatomies representing the bifurcation of the left coronary artery into the left anterior descending and circumflex arteries. To create this, we used a bifurcation model reconstructed from a healthy subject, capturing approximately 1 cm upstream and 4 cm downstream of the bifurcation. Randomly positioned stenoses (ranging from 1 to 3 per model) were introduced with severity levels representing caliber reductions between 20% and 60%, including both eccentric and concentric stenoses. Each synthetic model was discretized into a tetrahedral volumetric mesh, then a CFD simulation was conducted in SimVascular [43], setting boundary conditions randomly sampled within physiological ranges to achieve an inflow $Q \sim \mathcal{U}(100, 300)$ ml/min and an aortic pressure $p_{ao} \sim \mathcal{U}(90, 110)$ mmHg. Flow was assumed to be laminar and incompressible, with density $\rho = 1.05$ g/cm³ and dynamic viscosity $\mu = 4$ cP.

2.1.2. Patient-specific dataset

To test the models on real anatomies, we used data from patient-specific CFD simulations generated in a previous study by our group [22]. Briefly, 105 geometries of the left ($n = 76$) and right ($n = 29$) coronary arteries were reconstructed from CT scans acquired at Centro Cardiologico Monzino (Milan, Italy) for patients scheduled for clinically indicated invasive coronary angiography due to suspected CAD. Invasive FFR data was partially available for the cohort at exam. Segmentations were performed in 3DSlicer [44], including the coronary

artery volume downstream of the ostium to where the vessel diameter decreased below 1.5 mm, following established guidelines [45]. The reconstructed anatomies were then discretized into tetrahedral elements with a characteristic size of 0.25 mm using the TetGen algorithm within SimVascular [43]. For each mesh, 3 to 5 steady CFD simulations were run, slightly varying the BCs estimated from baseline clinical measurements (i.e., the average hyperemic inflow \bar{Q}_{hyp} and the outlet resistances R_{out}), matching the patient's aortic pressure. Validation vs. invasive FFR measurements yielded excellent agreement, ensuring the reliability of the results. For further details on the CFD simulations, see Ref. [22]. A total of 427 CFD simulations were conducted, with the results serving as ground truth for training the networks. The present study was performed in accordance with recommendations of the local Ethics Committee, with written informed consent from all subjects, in accordance with the Declaration of Helsinki.

2.2. Network architectures

In this work, we considered six different deep learning architectures, namely a simple, fully connected network (MLP), a point cloud network (PointNet++), two mesh neural network, leveraging approximated local convolutions in the neighborhood of each node (DiffusionNet, Delta-Conv), and two transformer models (LaB-VaTr, LaB-GATr). For each model, we defined a small (S) version and a large (L) version to mitigate any bias related to model size: each S-model had approximately 1M parameters, while each L-model contained 3M parameters. An exception to this is LaB-GATr, which only had 230k and 425k parameters for the

Table 1

Key characteristics of the model variants used in this study. Depth represents the number of building blocks in each network (e.g., layers for MLP, hierarchies for PointNet++). Channels denote the number of latent channels. Downsampling factors (NA = Not Applicable) specify the compression applied to the mesh prior to input into the network.

Model	Variant	Depth	Channels	Downsampling factors
MLP	S variant	6	512	NA
	L variant	12	512	NA
PointNet++ [40]	S variant	4	256	0.05; 0.50; 0.75; 0.90
	L variant	4	512	0.05; 0.50; 0.75; 0.90
DiffusionNet [41]	S variant	3	256	NA
	L variant	6	256	NA
DeltaConv [42]	S variant	8	128	0.10
	L variant	8	256	0.10
LaB-VaTr [36]	S variant	12	NA	0.05
	L variant	24	NA	0.05
LaB-GATr [36]	S variant	8	8	0.05
	L variant	16	8	0.05

Table 2

Input features computed per vertex of the mesh, corresponding definitions, and dimensionality (i.e., number of channels c).

Vertex feature	Definition	c
\mathbf{n}_i	Outward surface normal	3
s_i	Shortest geodesic to inlet	1
γ_i	Curvilinear abscissa	1
R_i	Radius of largest inscribed sphere	1
\mathbf{T}_i	Tangent to centerline	3
Φ_i	Laplacian eigenvectors	6
HKS_i	Heat kernel signature	5
WKS_i	Wave kernel signature	5
q_i	Blood inflow-rate	1
$p_{m,i}$	Pressure at coronary ostium	1
$BC_{TD,i}$	Node type in CFD analysis	1
Total input features		28

S and L versions, respectively, -justified by its task-specific $SO(3)$ equivariant layers and the runtime overhead. A brief description of each architecture, highlighting its main features, is provided below, and the characteristics of the models trained in this study are reported in Table 1. All networks were implemented in Python, using the PyTorch [46] and PyTorch Geometric [47] libraries.

2.2.1. Multi-layer Perceptron (MLP)

We included as a baseline model a point-wise MLP, which consisted of fully connected layers with batch normalization [48]. We used as activation function the Leaky ReLU [49], with a negative slope of 0.2. Throughout our work, we denote MLPs with h_θ .

2.2.2. PointNet++

PointNet++ [40] is a popular point cloud network that applies a per-point MLP followed by max aggregation over the k -nn neighborhood $B(i)$ of the i th vertex of the mesh. The message-passing layers are defined by the following equation:

$$\mathbf{x}_i^{(l+1)} = \max_{j \in B(i)} h_\theta(\mathbf{x}_j^{(l)}, \mathbf{r}_{i-j}) \quad (1)$$

where $\mathbf{x}_i^{(l)}$ are the features of the i th vertex in layer l , h_θ is an MLP of arbitrary depth, and \mathbf{r}_{i-j} represents the relative position between vertices. PointNet++ uses sampling and grouping operations to hierarchically subsample mesh vertices in the contracting pathway and perform interpolation in the expanding pathway. To handle the large scale of the mesh, we utilized a network with four hierarchical levels, employing cumulative sampling factors of 0.05, 0.50, 0.75, and 0.90.

2.2.3. DiffusionNet

DiffusionNet [41] approximates local convolutions to learn through the learned feature diffusion process by leveraging the eigendecomposition of the Laplacian operator and applying spatial gradients. The network consists of multiple blocks, where an MLP is applied to each point to represent pointwise features, learned diffusion (for spatial communication), and spatial gradient features (to enable directional filters). The diffusion layer $d_t(\mathbf{x})$ is evaluated using the precomputed eigenvectors $\Phi := [\phi_i]$ and eigenvalues λ_i of the Laplace–Beltrami operator (denoted as Δ):

$$h_t(\mathbf{x}) = \Phi_\Delta \begin{bmatrix} e^{-\lambda_0 t} \\ e^{-\lambda_1 t} \\ \dots \end{bmatrix} \odot (\Phi_\Delta^T M \mathbf{x})$$

where M denotes the mass matrix, \odot the elementwise product and t is the learned diffusion time. This closed-form solution to the heat equation in eigenvector basis is known as spectral acceleration [41]. Spatial gradient features are obtained by learning a $D \times D$ matrix A (with D equal to the number of features) to compute the spatial gradient $\text{grad}_A(\mathbf{x})$. Features propagation through DiffusionNet blocks is then defined by:

$$\mathbf{x}_i^{(l+1)} = h_\theta \left(d_t(\mathbf{x}_i^{(l)}), \text{grad}_A(\mathbf{x}_i^{(l)}), \mathbf{x}_i^{(l)} \right) \quad (2)$$

2.2.4. DeltaConv

DeltaConv [42] exploits anisotropic and intrinsic convolutional layers to learn on surfaces. DeltaConv utilizes two streams for feature propagation: a scalar feature stream (\mathbf{x}) and a vector feature stream (\mathbf{u}). These streams exchange information through geometric operators that convert scalar features to vector features and vice versa. Message passing in DeltaConv is defined as follows:

$$\mathbf{u}_i^{(l+1)} = h_{\theta_0} \left(\mathbf{u}_i^{(l)}, \text{grad } \mathbf{x}_i^{(l)}, \Delta \mathbf{u}_i^{(l)} \right) \quad (3a)$$

$$\mathbf{x}_i^{(l+1)} = h_{\theta_1} \left(\mathbf{x}_i^{(l)}, \text{div } \mathbf{u}_i^{(l+1)}, \text{curl } \mathbf{u}_i^{(l+1)}, \|\mathbf{u}_i^{(l+1)}\| \right) + \max_{j \in B(i)} h_{\theta_2}(\mathbf{x}_j^{(l)}) \quad (3b)$$

Similar to PointNet++, an MLP is applied point-wisely, followed by maximum aggregation over the neighborhood $B(i)$ of each vertex in the mesh. To account for the large scale of the data, the mesh was downsampled using farthest point sampling before being input into the network. In our implementation, we set a downsampling ratio of 0.1 and defined $B(i)$ to include 30 neighbors.

2.2.5. LaB-GATr and LaB-VaTr

Large biomedical geometric algebra transformers (LaB-GATr) [36] are scalable self-attention models. LaB-GATr applies compression via cross-attention from the source resolution to a coarse token set and learned interpolation back to the original resolution. Across the coarse token set, feature vectors are updated through multi-head attention blocks:

$$a_t = \text{Softmax} \left(\frac{q_t(\mathbf{x}_i^{(l)}) k_t(\mathbf{x}_i^{(l)})^T}{\kappa} \right) v_t(\mathbf{x}_i^{(l)})$$

$$A = \mathbf{x}_i^{(l)} + \text{Linear} \left(\text{Concat}_t a_t \right)$$

$$\mathbf{x}_i^{(l+1)} = A + h_\theta(A)$$

where queries q , keys k and values v are linear layers, $\kappa \in \mathbb{R}$ is a scaling factor and attention heads are indexed by t . LaB-GATr operates in the projective geometric (Clifford) algebra $G(3,0,1)$ which requires expressing $\mathbf{x}_i^{(l)}$ as so-called multivectors. For further details, refer to [37]. The advantage of the attention framework is its global context aggregation after only one block, which facilitates learning on large meshes. To make training the model computationally feasible, we applied a compression factor of 0.05 to downsample the input mesh.

We also included in our benchmark LaB-VaTr, a variant of the LaB-GATr architecture that retains the same layout but replaces geometric algebra with *vanilla* fully connected layers in linear algebra.

Table 3

Synthetic dataset - Performance metrics for different models predicting pressure. Test loss is computed on the predicted pressure. Difference and approximation disparity are computed per point on the whole reconstructed vFFR field. Correlation coefficient and bias are computed in stenotic points. Training cost is expressed as seconds per epoch. For readability, standard deviations and limits of agreement are provided as supplementary material.

Network output	Pressure					
Model	Test loss (Pa $\times 10^{-2}$) \downarrow	Per-point Diff. \downarrow	Approx. Disp. \downarrow	ρ \uparrow	Bias \downarrow	Train. Cost (s/epoch) \downarrow
MLP (S)	9.90	-8.10E-03	8.20E-03	0.658	-0.0094	12
MLP (L)	9.83	-1.24E-02	1.25E-02	0.637	-0.0171	17
PointNet++ (S)	9.25	-1.02E-02	1.03E-02	0.519	-0.0112	16
PointNet++ (L)	8.16	-9.05E-03	9.22E-03	0.471	-0.0093	18
DiffusionNet (S)	5.46	-7.10E-04	1.08E-03	0.973	-0.0008	28
DiffusionNet (L)	5.29	-4.60E-03	4.73E-03	0.698	-0.0049	48
LaB-VaTr (S)	12.98	-1.02E-02	1.06E-02	0.378	-0.0105	35
LaB-VaTr (L)	11.89	-1.44E-02	1.48E-02	0.253	-0.0153	55
LaB-GATr (S)	4.54	-8.47E-04	1.23E-03	0.973	-0.0010	244
LaB-GATr (L)	3.37	-8.83E-04	1.12E-03	0.976	-0.0010	419

2.3. Equivariance

All networks considered in this study exhibit permutation-equivariance. This means that if the input points are permuted (through a permutation P), the output at each stage of the network is permuted in the same way. Denote networks as f_θ . This property can be expressed by the following condition:

$$(f_\theta \circ P)(\cdot) \equiv (P \circ f_\theta)(\cdot), \quad \forall P \in S_n,$$

where \circ denotes composition, and S_n represents the symmetric group of permutations on n elements. Message passing and self-attention mechanisms are examples of permutation-equivariant building blocks commonly used in geometric deep-learning architectures.

In addition to permutation equivariance, the LaB-GATr is SE(3)-equivariant, meaning that if rotation and translation are applied to the input mesh, through a group action ρ , the network will rotate and translate its output in the same manner, preserving the relationship between the mesh and its transformed state:

$$(f_\theta \circ \rho)(\cdot) \equiv (\rho \circ f_\theta)(\cdot), \quad \forall \rho \in \text{SE}(3),$$

Furthermore, DeltaConv is invariant to, i.e., unaffected by, any orientation-preserving isometry such as roto-translation.

2.4. Experiments

First, we trained each model on the synthetic dataset described in Section 2.1.1. The dataset was divided into 1100 training cases, 200 validation cases, and 200 test cases. All models were trained for 500 epochs with a batch size of 2 on NVIDIA L40 (48 GB) GPUs, using the Adam optimizer with an initial learning rate of $3 \cdot 10^{-4}$ and an exponential decay rate of $\gamma_{decay} = 0.987$. We performed three training runs using different target variables, i.e., pressure, pressure drop, and FFR (see Section 2.4.2). For the patient-specific dataset, we only trained models using pressure drop as a target variable. To improve the stability of the training process, all target variables y were standardized using the formula $y_{std,i} = (y_i - \mu_y)/\sigma_y$, where μ_y and σ_y represent the mean and standard deviation of y across the training dataset. All models were trained using the mean squared error (MSE) as the loss function.

2.4.1. Input features

Input features and their definitions are summarized in Table 2. Let $\mathcal{M} = (\mathcal{V}, \mathcal{F})$ represent surface mesh of vertices \mathcal{V} and faces \mathcal{F} . We defined the input features by assigning each vertex $\mathbf{v} \in \mathcal{V}$ a set of c_{in} features, $f^{in} : \mathcal{V} \rightarrow \mathbb{R}^{c_{in}}$. These features capture both local and global morphological properties of the mesh, as well as physical parameters describing the BCs of the CFD simulation. For each anatomy in the dataset, the following steps were carried out to construct the set of input features. For each mesh vertex, we computed the outward surface normal and appended its components as features. Next, the

vessel centerline was reconstructed using the Vascular Modeling Toolkit (VMTK) library [50]. Geometric properties calculated along the centerline (e.g., the radius of the maximum inscribed sphere) were then projected onto the surface by matching each surface point to the nearest point on the centerline. The radius, the curvilinear abscissa, and the projection of the tangent to the centerline onto the surface, approximating the parallel transport field, were appended as input features. The tangent direction to the centerline has previously been used as an input feature for GNNs trained to learn hemodynamic fields, as it serves as an indicator of flow direction through the fluid volume [38]. Next, we appended the shortest geodesic distance from each vertex of the mesh to the inlet surface, which we computed as described in [23]. We included three local shape descriptors derived from the spectral decomposition of the Laplace–Beltrami operator: the Laplacian eigenvectors [51]; the heat kernel signature (HKS) [52], which characterizes heat diffusion on a surface; and the wave kernel signature (WKS) [53], based on the Schrödinger wave equation. These descriptors provide a pointwise characterization of each vertex in relation to the entire mesh. We computed the discretized signature matrices using the Laplace eigenfunctions up to the 5th order and appended the results to the input features. Finally, we added constant scalar fields over \mathcal{V} to represent the global parameters of the CFD simulation BCs, including blood inflow and inlet pressure. We remark that, while the pressure at the inlet surface (i.e., the coronary ostium) is generally a result of the CFD analysis, thus not known *a priori*, it can be safely approximated with the aortic pressure, which is typically known for patients. We also assigned a constant value from the set $[0, 1, 2]$ to each vertex to indicate its corresponding node type in the CFD analysis: [wall, inlet, outlet], respectively. In summary, we associated each vertex of the mesh \mathbf{v}_i to the vector of features $\mathbf{f}_i^{(in)}$:

$$\mathbf{v}_i \in \mathcal{V} \mapsto \mathbf{f}_i^{(in)} = [\mathbf{n}_i, g_i, \gamma_i, R_i, \mathbf{T}_i, \Phi_i, HKS_i, WKS_i, q_i, p_{in,i}, BC_{ID,i}] \in \mathbb{R}^{28}$$

2.4.2. Network output

We trained our models to predict scalar hemodynamic quantities on the domain \mathcal{V} , derived from steady blood flow through the coronary arteries. In experiments with the synthetic dataset, we conducted three groups of training, each targeting a different output variable: (i) pressure p_i , (ii) pressure drop - $\Delta p_i = p_{in} - p_i$ -, and (iii) the vFFR field - $\text{FFR}_i = p_i/p_{in}$. In experiments using patient-specific anatomies, we selected the pressure drop field as the output variable, as it achieved the best evaluation metrics in the synthetic dataset experiments.

2.5. Evaluation metrics

For each model, we reconstructed the corresponding vFFR field from the network output and evaluated its performance by comparing the predicted vFFR field \hat{y}_i to the ground truth y_i at each point. To quantify the similarity between the two fields, we used the per-point difference,

defined as $\hat{y}_i - y_i$, and the approximation disparity [54], defined as follows:

$$\text{Approx. Disp.} := \sum_{i=1}^N (\hat{y}_i - y_i)^2 / \sum_{i=1}^N \hat{y}_i^2$$

We also conducted a per-lesion analysis, comparing the predicted vFFR values in stenotic regions (specifically, the minimum vFFR value within each stenosis) against the ground truth for each model. This evaluation included correlation analysis, confusion matrix analysis, and Bland-Altman analysis. All statistical tests were conducted in Python using the SciPy library [55].

2.6. Ablation studies

We conducted ablation studies to assess the impact of input features on prediction accuracy. Specifically, we trained a single model using seven different combinations of input features and compared the resulting performance metrics. These experiments were conducted using the PointNet++ model, which offers a good trade-off between training time and performance and is a commonly used baseline in GDL tasks.

3. Results

In this section, we present the results of our benchmark. First, we show the performance of each model on the synthetic dataset. Next, we present the results obtained on the patient-specific data. Finally, we provide insights from the ablation study, focusing on the influence of input features on training performance.

3.1. Synthetic dataset

Table 3 lists the performance of each model variant trained to predict the pressure field. During inference on the test set, all models tended to underestimate the per-point values. DiffusionNet and LaB-GATr were the only models to yield a per-point difference between the predicted and CFD-computed vFFR below 10^{-3} . DiffusionNet and LaB-GATr also stood out in terms of approximation disparity. The boxplots for each metric are shown in Fig. 3, while the inferred vFFR fields are illustrated in Fig. 2. Spearman's correlation and Bland-Altman analysis, conducted considering vFFR in stenotic regions, confirmed the superior performance of DiffusionNet and LaB-GATr.

Similarly, Table 4 summarizes the results of training with the standardized pressure drop field as the network output. In inference, the per-point difference and approximation disparity were comparable across all models, averaging approximately $\sim 4.11\text{E}-3$ and $\sim 5.1\text{E}-3$, respectively (Fig. 4). The reconstructed vFFR map derived from the inferred pressure drop field was qualitatively the most accurate (Fig. 2). Notably, the pressure drop was also the only output variable for which we were able to stably train the DeltaConv with good performance. All models demonstrated strong (>0.4) to very strong (>0.7) correlation between the inferred and CFD-based vFFR. The bias was consistent across models, with LaB-GATr exhibiting the narrowest 95% limit of agreement band.

Finally, Table 5 illustrates the results achieved in the experiments using the vFFR field directly as the output variable. In inference, LaB-GATr outperformed all the tested model, yielding a per-point difference of $5.35\text{E}-05$ and an approximation disparity of $5.40\text{E}-04$ (Fig. 5). In particular, LaB-GATr was the only model to predict a vFFR field that qualitatively resembles the ground truth (Fig. 2). The MLP demonstrated comparable test metrics to the LaB-GATr. However, in the statistical analysis of stenotic regions, the MLP showed a lower correlation coefficient and a bias that was an order of magnitude larger.

3.2. Patient-specific dataset

In this section, we present the results of training our models on a dataset of patient-specific CFD analyses. Overall, all models demonstrated reduced performance during inference compared to the results on synthetic anatomies. Among the models, LaB-GATr consistently showed the best generalization capability, particularly in its 'small' variant. Most models correctly predicted the pressure gradient direction from the inlet to the outlet; however, some failed to capture sharp pressure drops in stenotic regions. Fig. 6 shows for each model (rows) the reconstructed vFFR field for 6 cases (columns) of the test set. We observed moderate agreement between the inferred vFFR and the ground truth for Patients 2, 3, 4, and 6. For Patient 5, only the S-variant of LaB-GATr accurately predicted the vFFR gradient in the circumflex artery, while for Patient 1, none of the models successfully captured the vFFR drop in the stenosis of the left anterior descending artery. In general, all models showed greater discrepancies for meshes with numerous bifurcations in sub-branches. Conversely, for cases involving a single main vessel (e.g., right coronary arteries), the models demonstrated better agreement with the ground truth. Tables 6 and 7 report the performance metrics for each model. All models achieved a validation loss of ~ 10 to 15 Pa. In front of a significantly higher computational cost, LaB-GATr was the model that again yielded best results. DiffusionNet (S-variant) yielded the highest correlation coefficient (0.72), and together with LaB-GATr (S-variant) was the only model to achieve an accuracy above 0.7; however it poorly performed in terms of precision and exhibited a larger bias compared to LaB-GATr.

3.2.1. Comparison vs. invasive FFR measurements

Invasive FFR data was available in 5 (out of 14) cases of the test set, corresponding overall to 7 comparable branches (either the LCA, LCX or RCA). For these cases, we compared the clinical data vs. the value of vFFR reconstructed from the inferred pressure-drop field, using the model yielding best performances (i.e., the LaB-GATr). An overview of the comparison is reported in Table 8. We obtained a Spearman's correlation coefficient $r = 0.52$, $p = 0.505$, and a bias of $+0.035$, with ML overestimating invasive data. Accuracy resulted 0.71 using a cut-off of 0.80.

3.3. Impact of different input features

Table 9 summarizes the results of the ablation study conducted on input features, using PointNet++ as the backend and trained on the synthetic dataset of bifurcations. Our findings highlight that specific features, such as surface normals (\mathbf{n}_i), geodesics (g_i), and flow rate (q_i), are critical for achieving strong model performance, while features derived from Laplacian eigenvectors provide only modest improvements. Including the inlet pressure ($p_{in,i}$) as an input feature significantly enhances the predictive capability of the model and correctly sets the inferred pressure range. Local radius (R_i) benefits the generalization capability of the model as well. Interestingly, the model yielded the best results when using either all the features from this study or a limited subset of those features (Ablation 3), which included few local geometric and physical descriptors.

4. Discussion

In this work, we have presented a systematic benchmark between different GDL models to predict scalar fields derived from the solution of steady CFD analysis on both synthetic and real-life patient anatomies of coronary vessels. While the 3D mesh of the artery can be easily reconstructed from CT, the measurement of functional indices used for the assessment of stenosis severity, such as FFR, requires invasive examination or time-consuming CFD simulations, which are both economically expensive solutions. We conditioned the tested models on geometric features that can be computed directly from the mesh and

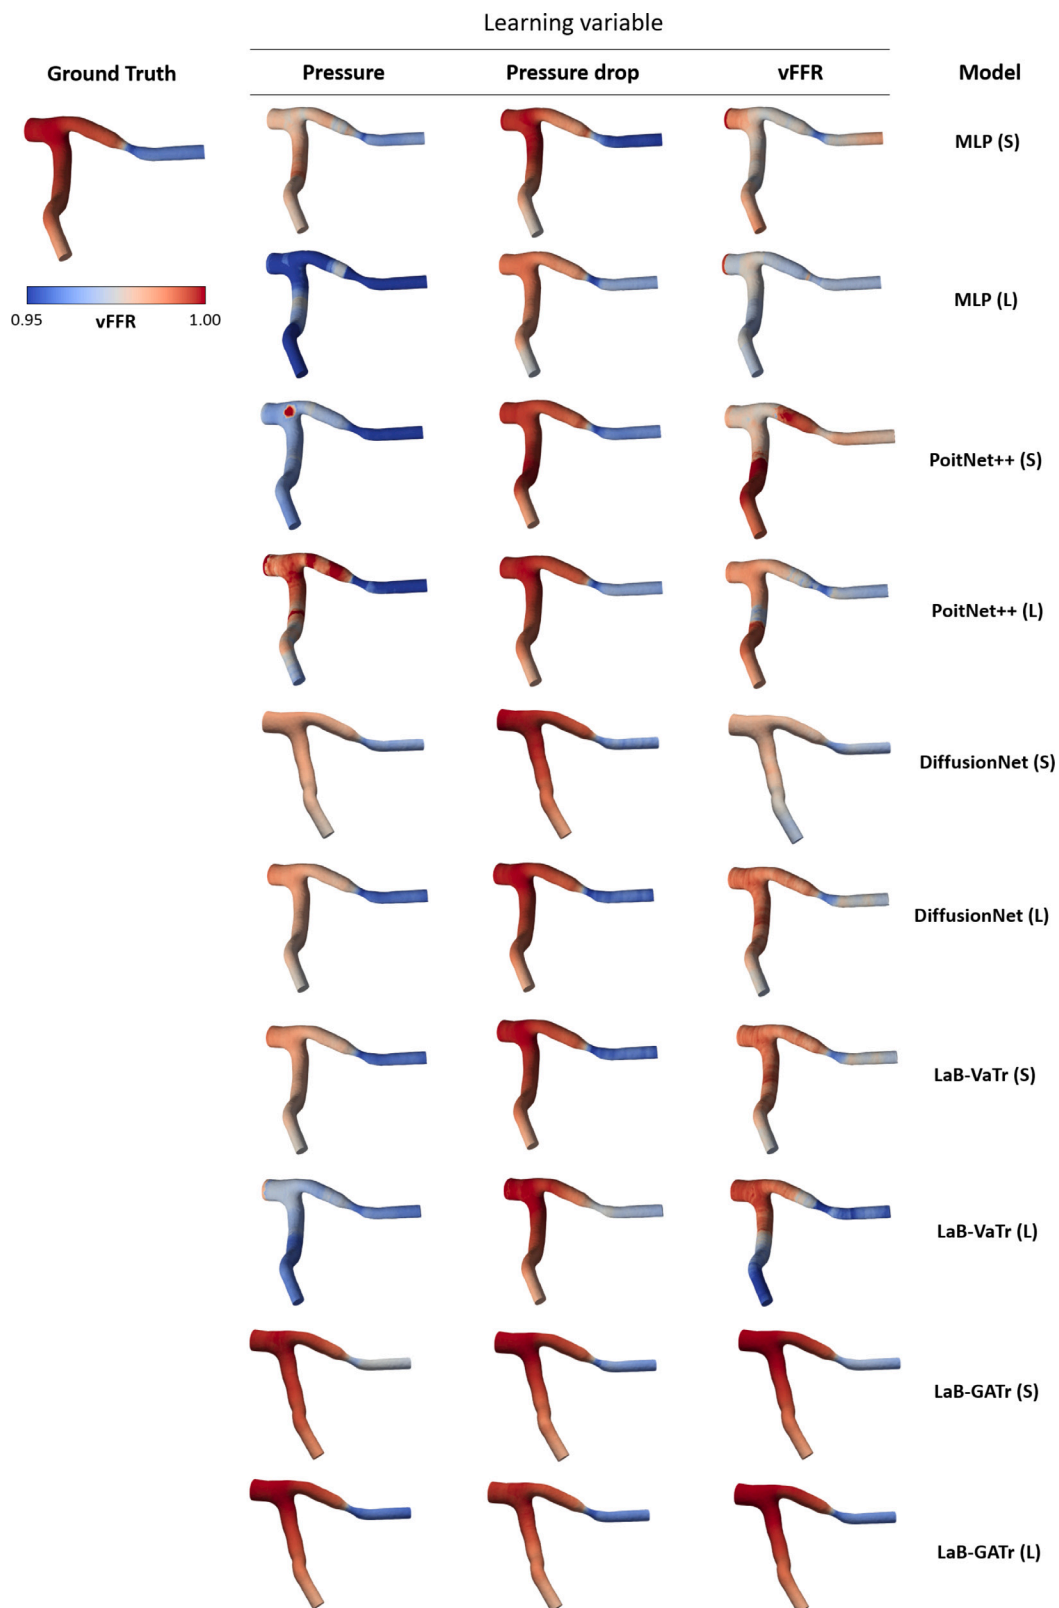


Fig. 2. Maps of vFFR computed from CFD (ground truth) on the synthetic bifurcation dataset, and inferred by each model, using different learning variables for one case of the test set.

Table 4

Synthetic dataset - Performance metrics for different models predicting pressure drop. Test loss is computed on the predicted pressure drop. Difference and approximation disparity are computed per point on the whole reconstructed vFFR field. Correlation coefficient and bias are computed in stenotic points. Training cost is expressed as seconds per epoch. For readability, standard deviations and limits of agreement are provided as supplementary material.

Network output	Pressure drop					
Model	Test loss (Pa $\times 10^{-2}$) \downarrow	Per-point Diff. \downarrow	Approx. Disp. \downarrow	ρ \uparrow	Bias \downarrow	Train. Cost (s/epoch) \downarrow
MLP (S)	7.50	3.23E-03	4.69E-03	0.940	0.0021	12
MLP (L)	9.60	4.15E-03	5.25E-03	0.967	0.0021	17
PointNet++ (S)	8.10	3.93E-03	5.81E-03	0.976	0.0022	16
PointNet++ (L)	7.08	3.98E-03	4.84E-03	0.977	0.0021	18
DiffusionNet (S)	6.84	3.54E-03	5.06E-03	0.971	0.0022	28
DiffusionNet (L)	4.69	3.91E-03	5.14E-03	0.986	0.0024	48
LaB-VaTr (S)	8.58	4.13E-03	4.72E-03	0.935	0.0026	35
LaB-VaTr (L)	11.58	4.59E-03	5.32E-03	0.791	0.0033	55
LaB-GATr (S)	6.34	4.39E-03	4.97E-03	0.979	0.0033	244
LaB-GATr (L)	3.78	4.19E-03	5.12E-03	0.967	0.0030	419
DeltaConv (S)	11.23	4.47E-03	5.28E-03	0.562	0.0021	17
DeltaConv (L)	11.56	4.84E-03	5.42E-03	0.651	0.0026	30

Table 5

Synthetic dataset - Performance metrics for different models predicting vFFR. Test loss, difference and approximation disparity are computed per point on the whole reconstructed vFFR field. Correlation coefficient and bias are computed in stenotic points. Training cost is expressed as seconds per epoch. For readability, standard deviations and limits of agreement are provided as supplementary material.

Network output	vFFR					
Model	Test loss ($\times 10^{-3}$) \downarrow	Per-point Diff. \downarrow	Approx. Disp. \downarrow	ρ \uparrow	Bias \downarrow	Train. Cost (s/epoch) \downarrow
MLP (S)	3.63	3.89E-04	3.66E-03	0.287	0.0016	12
MLP (L)	2.66	-5.76E-05	2.67E-03	0.650	0.0012	17
PointNet++ (S)	4.69	1.08E-03	5.59E-03	0.472	0.0015	16
PointNet++ (L)	2.94	-3.21E-04	3.62E-03	0.743	-0.0004	18
DiffusionNet (S)	4.65	8.21E-04	4.67E-03	0.630	0.0009	28
DiffusionNet (L)	3.23	1.19E-03	3.25E-03	0.726	0.0011	48
LaB-VaTr (S)	2.76	-5.78E-04	2.79E-03	0.616	-0.0016	35
LaB-VaTr (L)	3.28	2.02E-04	3.32E-03	0.643	-0.0017	55
LaB-GATr (S)	0.53	2.28E-04	5.40E-04	0.970	0.0004	244
LaB-GATr (L)	0.31	5.35E-05	3.13E-04	0.988	-0.0002	419

Table 6

Patient-specific dataset - Performance metrics for various models. Test loss, training cost, per-point difference, and approximation disparity are computed pointwisely on the whole mesh. The test loss is computed on the predicted pressure drop and expressed in Pascal. **Bold** font indicates the best result for each metric.

Model	Test Loss (Pa) \downarrow	Per-point Diff. \downarrow	Approx. Disp. \downarrow	Train. Cost (s/epoch) \downarrow
MLP (S)	13.6 \pm 7.6	-0.11 \pm 0.12	0.172 \pm 0.15	23
MLP (L)	13.0 \pm 7.0	-0.10 \pm 0.11	0.169 \pm 0.15	24
DiffusionNet (S)	15.3 \pm 8.3	-0.06 \pm 0.13	0.149 \pm 0.20	44
DiffusionNet (L)	14.7 \pm 8.9	-0.09 \pm 0.14	0.181 \pm 0.19	83
PointNet++ (S)	15.1 \pm 8.9	-0.09 \pm 0.14	0.181 \pm 0.20	17
PointNet++ (L)	14.0 \pm 8.4	-0.09 \pm 0.13	0.170 \pm 0.19	18
LaB-VaTr (S)	14.5 \pm 7.5	-0.10 \pm 0.12	0.180 \pm 0.19	42
LaB-VaTr (L)	13.5 \pm 6.9	-0.11 \pm 0.11	0.165 \pm 0.14	54
LaB-GATr (S)	11.4 \pm 6.5	-0.03 \pm 0.10	0.123 \pm 0.12	104
LaB-GATr (L)	12.0 \pm 8.2	-0.04 \pm 0.13	0.116 \pm 0.14	154

Table 7

Patient-specific dataset - Performance metrics for various models. Bias [LoA=Limits of Agreement], correlation coefficient (ρ), precision (Prec.), recall (Rec.), accuracy (Acc.), and limits of agreement (LoA) are calculated on vFFR predictions in stenotic regions. **Bold** font indicates the best result for each metric.

Model	ρ \uparrow	Bias \downarrow	LoA	Prec. \uparrow	Rec. \uparrow	Acc. \uparrow
MLP (S)	0.58	-0.077	[-0.30; 0.15]	0.40	0.91	0.58
MLP (L)	0.66	-0.082	[-0.29; 0.13]	0.41	1.00	0.58
DiffusionNet (S)	0.72	-0.054	[-0.27; 0.17]	0.50	0.91	0.71
DiffusionNet (L)	0.68	-0.060	[-0.29; 0.18]	0.46	0.91	0.66
PointNet++ (S)	0.53	-0.051	[-0.31; 0.20]	0.46	1.00	0.66
PointNet++ (L)	0.67	-0.059	[-0.29; 0.17]	0.46	1.00	0.66
LaB-VaTr (S)	0.68	-0.080	[-0.31; 0.14]	0.42	1.00	0.61
LaB-VaTr (L)	0.66	-0.087	[-0.32; 0.14]	0.41	1.00	0.58
LaB-GATr (S)	0.69	0.004	[-0.20; 0.14]	0.67	0.55	0.79
LaB-GATr (L)	0.69	-0.033	[-0.21; 0.22]	0.53	0.73	0.68

physical quantities (e.g., flow rate, pressure) that can be directly or indirectly derived from clinical measurements, usually taken during patient screening. We first evaluated each model on a dataset of CFD simulations run on 1500 synthetic anatomies, inferring different outputs from which the vFFR field was then reconstructed. Then, we tested each model on a dataset of 427 CFD simulations run on patient-specific anatomies, inferring the pressure drop, that is, the output variable that yielded the best performances on the synthetic dataset. Pre-computing the input features to inform each model required approximately 1 min on average, while during inference, all models took less than 30 s to produce the output field at hand, limiting the overall cost of the ML-based estimation of vFFR to under 2 min.

4.1. Impact of network output

In our experiments on the synthetic dataset, we observed that all models generally had the best performance, both qualitative (Fig. 2) and quantitative (Table 4) when inferring the pressure drop field. We

Table 8

Comparison of stenosis degree, invasive FFR, and ML-predicted FFR for each coronary artery across patients in the test set.


ID	Stenosis degree (%)			FFR (Invasive)			FFR (LaB-GATr)		
	LAD	LCX	RCA	LAD	LCX	RCA	LAD	LCX	RCA
Patient #1	75	20	95		0.90		0.865	0.889	
Patient #2	10	10	60	0.77	0.78	0.79	0.845	0.927	0.798
Patient #3	75	40	0	0.77	0.78	0.79	0.845	0.927	0.798
Patient #4	10	40	90	0.90		0.50	0.923	0.821	
Patient #5	65	70	30	0.84	0.94	0.90	0.871	0.910	
Patient #6	0	0	0						

Table 9

Comparison of different feature combinations (indicated as *Ablation n°*) for predicting the pressure drop field. Features are indicated with symbols introduced in Table 2. Pointwise difference is reported as evaluation metric. **Bold** values indicate the two best results achieved.

	n_i	g_i	γ_i	R_i	Φ_i	HKS_i	WKS_i	T_i	$BC_{1D,i}$	q_i	$p_{in,i}$	Pointwise Diff. ($\text{Pa} \times 10^{-2}$) ↓
<i>Ablation 1</i>	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	10.56
<i>Ablation 2</i>	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	7.92
<i>Ablation 3</i>	✓	✓	✗	✓	✗	✗	✗	✗	✗	✓	✓	6.96
<i>Ablation 4</i>	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	8.10
<i>Ablation 5</i>	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓	✗	7.50
<i>Ablation 6</i>	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓	8.34
<i>Ablation 7</i>	✗	✗	✗	✗	✓	✓	✓	✓	✗	✓	✓	10.68
<i>All features</i>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	7.08

Output variable: Pressure



Output variable: Pressure drop

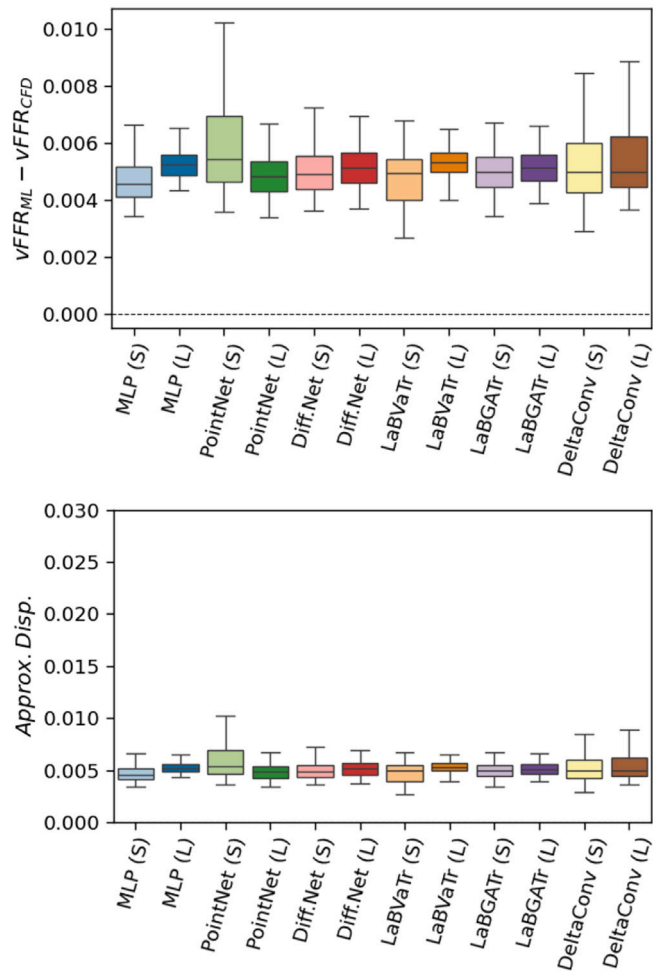


Fig. 3. Boxplots showing the performance of each model on the test set, learning pressure field, in terms of point-wise differences (top row) and approximation disparities (bottom row) between vFFR predictions and CFD-derived values.

Fig. 4. Boxplots showing the performance of each model on the test set, learning pressure drop field, in terms of point-wise differences (top row) and approximation disparities (bottom row) between vFFR predictions and CFD-derived values.




Fig. 5. Boxplots showing the performance of each model on the test set, learning vFFR field, in terms of point-wise differences (top row) and approximation disparities (bottom row) between vFFR predictions and CFD-derived values.

suggest that this is due to the fact that, while the pressure field can range in very heterogeneous intervals across arteries depending both on the pressure drop along the vessel and the pressure at the inlet (i.e., patient's systemic pressure), the pressure drop field is generally more homogeneous across patients, as it always starts from zero at the ostium and is more correlated with geometric features, such as vessel radius and length, especially under Poiseuille's flow hypothesis. Although FFR remains the primary parameter of interest in coronary artery analysis, we propose a strategy centered on learning the full hemodynamic pressure field — particularly the pressure drop — as it offers a more stable and effective basis for FFR prediction. When training the models on the pressure drop, we found a low common bias of $\sim 4E-03$ on the derived vFFR field, while training on the pressure produced for all the models a negative bias (as visible in Fig. 4). This could be due to a common underestimation of the output variable for both cases, which in the case of inference on the pressure field results in a lower distal pressure (P_d), and thus a lower $vFFR_{ML} - vFFR_{CFD}$, as it is defined as $(P_{d,ML} - P_{d,CFD})/P_{in}$. On the other hand, in the case of the pressure drop, underestimating the output leads to systematically overestimating the $vFFR_{ML}$, defined as $(P_{in} - \Delta P_{ML})/P_{in}$. Inferring vFFR directly, produced no systematic bias, but yielded a significantly lower performance, as shown in Fig. 5 by the wider percentile ranges.

Training on the pressure drop resulted, in general, in more stable training runs; in particular, this is the only network output for which

DeltaConv training was stable, achieving good performance metrics. We trained DeltaConv on the other network variable as well, however the performance achieved were not as good as training on the pressure drop, using the architecture implementation as presented in Ref. [42]. Given that the network was successfully trained for one output variable, we expect that with proper adjustment to the backend implementation it is feasible to train it on the other learning variable as well, however this goes beyond the purpose of this work, still hinting a possible development.

4.2. Prediction of vFFR in real-life patients

When inferring the pressure drop onto patient specific data, we tested both the capability to correctly infer pressure on the whole mesh and the diagnostic accuracy in predicting the vFFR in stenotic regions. In general, all models predicted a smooth vFFR decrease going from the inlet to outlets, but LaB-GATr was the only one that accurately permitted to reconstruct the vFFR field, capturing sharper decreases due to stenosis. On average, LaB-GATr yielded an error in the order of 1 to 10 Pa on the predicted pressure drop field (and thus on pressure). When comparing the vFFR values obtained from the FFR field reconstructed using the inferred pressure drop to invasive clinical data, we observed similar accuracy and correlation levels as those obtained when comparing against CFD. Although the significance of these results is limited by the small sample size ($n = 7$), they support the applicability of GDL techniques for FFR prediction. Consistent with our previous experiments (Fig. 4), we observed an average bias of +0.035, with the ML model tending to overestimate invasive FFR.

Previous works have proved the capability of analogous model to learn fluid dynamic fields in vessel-like shaped domains. Rygiel et al. [32] obtained excellent agreement between CFD-based vFFR and machine learning prediction. Similarly, Suk et al. [35,56] proved the capability of the model to learn wall shear stress and velocity field. Nevertheless, in all these studies, experiments were conducted on synthetically generated anatomies, which simplify the task compared to real-life geometries that exhibit significant heterogeneity and complexity.

Two recent studies achieved patient-specific hemodynamic predictions using deep learning. Yang et al. [57] proposed a novel attention-based [58] variant of the popular U-Net architecture [59] to compute vFFR in patient-specific coronary artery anatomies. They trained the model on synthetically generated vessel segments and inferred the vFFR on stenotic segments of patient coronary arteries, thus overcoming the limitation of poor availability of patient-specific data. Pegolotti et al. [38] achieved good performance in predicting flow rate and pressure in patient-specific anatomies by training their network on 1D flow solutions. While both studies successfully predicted hemodynamics in patient-specific anatomies, they introduced significant simplifications, such as reducing the volume of interest to small vessel segments and relying on reduced-order 1D solutions of CFD analyses.

Our benchmark, along with results reported by other authors in the literature, supports the notion that inferring scalar fields on simple topologies (e.g., 2D manifolds representing small portions of anatomy, synthetic data, or 1D manifolds) is a task manageable by mesh neural networks, in general. However, when the topology becomes more complex and exhibits high variability within the dataset, as in the case of patient-specific data, transformer-based architectures emerge as the most suitable option for addressing the task effectively, both due to the efficient global context aggregation and the straight-forward processing of shapes. Specifically, we attribute the higher performance of LaB-GATr to two main factors. First, unlike any other model included in the study, it implements a receptive field mechanism based on a geometric algebra formulation, allowing it to account for long-distance points in the mesh. Second, its learning strategy involves aggressive downsampling of the mesh, followed by large-scale interpolation to reconstruct predictions at the dropped nodes. Pressure-related fields are

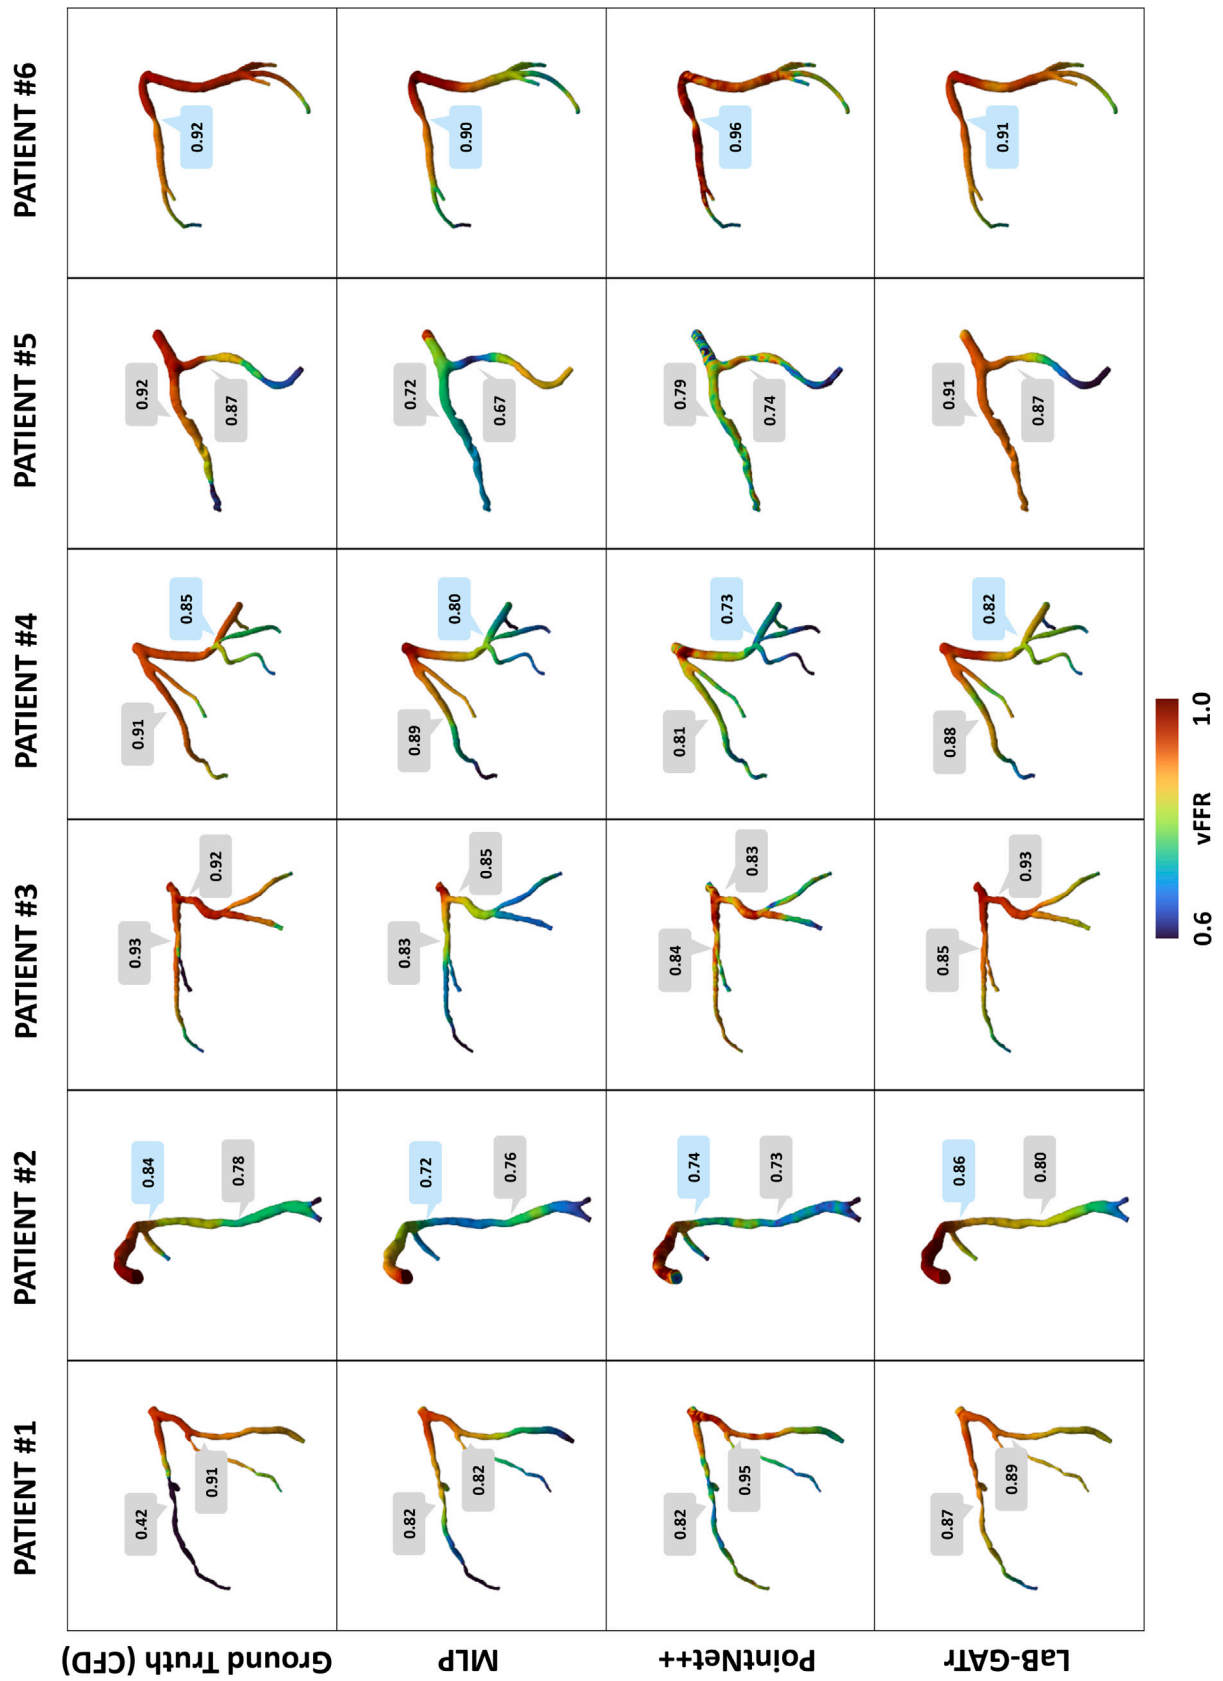


Fig. 6. Maps of vFFR computed from CFD (ground truth) on the patient' specific dataset, and inferred by three representative models (MLP (S), PointNet++ (S) and LaB-GATr (S)) on 6 test set cases, using pressure drop as learning variables. The vFFR value is indicated in correspondence with lesion sites.

generally smooth over a manifold; additionally, deep neural networks tend to learn low-frequency outputs more accurately due to spectral bias. These aspects make LaB-GATr the most suitable architecture for the task at hand.

4.3. Impact of different input features

Our ablation study, conducted to identify the crucial parameters for informing the model and achieving optimal performance, highlighted the importance of certain inputs in granting the model generalization capability. Specifically, geometric parameters such as surface normals and geodesic distance to artery landmarks (intrinsic orientation) proved essential for accurate predictions. When these parameters (*Ablation 7*) were removed from the input features, the evaluation metric (i.e., the vFFR pointwise difference) worsened by 2% to 59% compared to any other tested combination in which they were included. Similarly, the local vessel radius (*Ablation 2* vs. *Ablation 1*) and the pressure at the inlet (*Ablation 3* vs. *Ablation 2*) consistently improved model performance. Providing inlet pressure can be interpreted as setting a Dirichlet condition at the inlet, effectively constraining a subset of the mesh vertices to assume specific values. As a result, the model's predictions were generally more accurate near the inlet nodes and deteriorated as the distance from the ostium increased (as visible in Fig. 2 for the MLP (L) model using vFFR as the learning variable). Local descriptors derived from the spectral decomposition of the Laplacian matrix enhanced the model's performance but, when used as the sole descriptors of the anatomy, produced the poorest results (*Ablation 7*). Overall, we found that a small subset of input features (*Ablation 3*) is sufficient to learn the target scalar field, achieving comparable results to those obtained with the full feature set considered in this study.

For the ablation study, we used PointNet++ as the backend, as it had the lowest training cost in the benchmark and is widely adopted for GDL tasks. Given its efficiency and versatility, PointNet++ could be employed as a preliminary step in future experiments to perform feature selection and identify optimal input combinations.

4.4. Model invariance with respect to mesh transformations

Another relevant aspect to consider is the invariance of the model's output with respect to transformations applied to the mesh. In our experiments, we used models that exhibit permutation invariance, and in the case of LaB-GATr, equivariance with respect to rigid transformations. Mesh density is another factor that may vary across different input models and can affect the final output. In general, all vectorized neural fields demonstrate robustness to mesh resampling, as they learn a continuous field rather than a direct mapping to the mesh itself, as reported Refs. [54,56]. Mesh scaling, on the other hand, would affect the prediction of the model as the spatial dimensions of the 3D model are encoded through the mesh points.

4.5. Limitations

The presented work is not exempt from limitations. Firstly, systematically comparing different backends is challenging due to their varying implementations. We selected the number of network parameters as the primary criterion and mitigated bias from specific model sizes by defining two versions for each model. However, this approach could not be applied to LaB-GATr, as its parameter space is inherently constrained compared to typical GDL models. Moreover, training LaB-GATr with a parameter count comparable to the models tested in this work (~1M to 4M parameters) is infeasible due to its high computational cost. Although the results achieved are promising, mesh neural networks still need refinement to match 3D CFD performance and exceed reduced-order models in resolution and accuracy. Another limitation is the size of the dataset of patient-specific anatomies. While a dataset containing 427 samples might be sufficiently large for other

tasks, the significant variability, both in the topology of the manifolds and in the scalar field we aim to infer, renders this dataset insufficient to represent all the cases required for robust model generalization. Finally, we cannot definitively conclude that the pressure drop, which yielded the best results as the network's output variable in experiments on the synthetic dataset, exhibits the same behavior when applied to real-life patient anatomies; even though evidence in the literature [32, 54] suggests that pressure drop is indeed an optimal learning variable for this type of task.

5. Conclusion

We have presented a benchmark of different model backends, comparing their generalization capabilities in learning scalar fields derived from the solutions of CFD analyses on the surface meshes of coronary arteries. Our results indicate that inferring the pressure drop field on topologically simple anatomies (i.e., a bifurcation) is a task that can be successfully handled by various commonly used geometric deep learning models. Conversely, when predicting pressure drop on real-life patient geometries, only LaB-GATr demonstrated robust generalization capabilities, achieving the best evaluation metrics and diagnostic performance on stenotic segments. To achieve high performance, the model must be provided with essential local geometric and physical descriptors. Our findings indicate that surface normals, geodesics, vessel radius, inflow, and inlet pressure are crucial for this task.

In conclusion, various geometric deep learning models have demonstrated potential as CFD surrogates for FFR estimation when applied to problems involving relatively simple geometries (e.g., a straight vessel or a bifurcation). However, when considering patient-specific anatomies, we only find transformer-based backends (i.e., LaB-GATr) capable of learning the solutions of steady-state CFD analyses. This capability positions them as effective surrogates for accelerating the diagnostic process of CAD by directly predicting hemodynamic parameters.

CRediT authorship contribution statement

Guido Nannini: Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Conceptualization. **Julian Suk:** Writing – original draft, Validation, Methodology, Investigation, Formal analysis. **Patryk Rygiel:** Validation, Methodology, Investigation. **Simone Saitta:** Writing – review & editing, Validation, Supervision, Resources, Methodology, Investigation. **Luca Mariani:** Investigation, Data curation. **Riccardo Maranga:** Data curation. **Andrea Baggiano:** Data curation. **Gianluca Pontone:** Resources, Investigation, Formal analysis, Data curation. **Jelmer M. Wolterink:** Writing – review & editing, Validation, Supervision, Methodology, Investigation, Conceptualization. **Alberto Redaelli:** Writing – review & editing, Supervision, Conceptualization.

Data and codes availability

Data and code can be requested from the corresponding author upon reasonable request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

Acknowledgments

PNRR-POC-2022-12376500 Financed by the Ministry of Health funds PNRR M6C2I2.1 and by the European Union – NextGenerationEU

This work was partly funded by the National Plan for NRRP Complementary Investments (PNC), established with the decree-law 6 May 2021, n. 59, converted by (law n. 101 of 2021) in the call for the funding of research initiatives for technologies and innovative trajectories in the health and care sectors (Directorial Decree n. 931 of 06-06-2022) - project n. PNC0000003 - AdvaNced Technologies for Human-centrEd Medicine (project acronym: ANTHEM). This work reflects only the authors' views and opinions, neither the Ministry for University and Research nor the European Commission can be considered responsible for them. This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101080947 (VASCUL-AID). The authors would like to acknowledge Ruben Wiersma, PhD for his implementation and for guidance on the use of the DeltaConv architecture, which has been adopted in this work (<https://github.com/rubenwiersma/deltaconv>).

References

- [1] K. Okrainec, D.K. Banerjee, M.J. Eisenberg, Coronary artery disease in the developing world, *Am. Heart J.* 148 (1) (2004) 7–15.
- [2] P. Libby, P. Theroux, Pathophysiology of coronary artery disease, *Circulation* 111 (25) (2005) 3481–3488.
- [3] N.H. Pijls, B. Van Gelder, P. Van der Voort, K. Peels, F.A. Bracke, H.J. Bonnier, M.I. El Gamal, Fractional flow reserve: a useful index to evaluate the influence of an epicardial coronary stenosis on myocardial blood flow, *Circulation* 92 (11) (1995) 3183–3193.
- [4] N.H. Pijls, J.-W.E. Sels, Functional measurement of coronary stenosis, *J. Am. Coll. Cardiol.* 59 (12) (2012) 1045–1057.
- [5] L.X. van Nunen, F.M. Zimmermann, P.A. Tonino, E. Barbato, A. Baumbach, T. Engström, V. Klauss, P.A. MacCarthy, G. Manoharan, K.G. Oldroyd, et al., Fractional flow reserve versus angiography for guidance of PCI in patients with multivessel coronary artery disease (FAME): 5-year follow-up of a randomised controlled trial, *Lancet* 386 (10006) (2015) 1853–1860.
- [6] B.-K. Koo, A. Erglis, J.-H. Doh, D.V. Daniels, S. Jegere, H.-S. Kim, A. Dunning, T. DeFrance, A. Lansky, J. Leipsic, et al., Diagnosis of ischemia-causing coronary stenoses by noninvasive fractional flow reserve computed from coronary computed tomographic angiograms: results from the prospective multicenter DISCOVER-FLOW (Diagnosis of Ischemia-Causing Stenoses Obtained Via Non-invasive Fractional Flow Reserve) study, *J. Am. Coll. Cardiol.* 58 (19) (2011) 1989–1997.
- [7] B.L. Nørgaard, J. Leipsic, S. Gaur, S. Seneviratne, B.S. Ko, H. Ito, J.M. Jensen, L. Mauri, B. De Bruyne, H. Bezerra, et al., Diagnostic performance of noninvasive fractional flow reserve derived from coronary computed tomography angiography in suspected coronary artery disease: the NXT trial (Analysis of Coronary Blood Flow Using CT Angiography: Next Steps), *J. Am. Coll. Cardiol.* 63 (12) (2014) 1145–1155.
- [8] P.A. Tonino, B. De Bruyne, N.H. Pijls, U. Siebert, F. Ikeno, M. van Veer, V. Klauss, G. Manoharan, T. Engström, K.G. Oldroyd, et al., Fractional flow reserve versus angiography for guiding percutaneous coronary intervention, *N. Engl. J. Med.* 360 (3) (2009) 213–224.
- [9] F.M. Zimmermann, A. Ferrara, N.P. Johnson, L.X. Van Nunen, J. Escaned, P. Albertsson, R. Erbel, V. Legrand, H.-C. Gwon, W.S. Remkes, et al., Deferral vs. performance of percutaneous coronary intervention of functionally non-significant coronary stenosis: 15-year follow-up of the DEFER trial, *Eur. Heart J.* 36 (45) (2015) 3182–3188.
- [10] C.A. Taylor, T.A. Fonte, J.K. Min, Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis, *J. Am. Coll. Cardiol.* 61 (22) (2013) 2233–2241.
- [11] A. Coenen, M.M. Lubbers, A. Kurata, A. Kono, A. Dedic, R.G. Chelu, M.L. Dijkshoorn, F.J. Gijsen, M. Ouhlous, R.-J.M. van Geuns, et al., Fractional flow reserve computed from noninvasive CT angiography data: diagnostic performance of an on-site clinician-operated computational fluid dynamics algorithm, *Radiology* 274 (3) (2015) 674–683.
- [12] C.K. Zarins, C.A. Taylor, J.K. Min, Computed fractional flow reserve (FFT CT) derived from coronary CT angiography, *J. Cardiovasc. Transl. Res.* 6 (2013) 708–714.
- [13] A. Candreva, G. De Nisco, M.L. Rizzini, F. D'Ascenzo, G.M. De Ferrari, D. Gallo, U. Morbiducci, C. Chiastra, Current and future applications of computational fluid dynamics in coronary artery disease, *Rev. Cardiovasc. Med.* 23 (11) (2022) 377.
- [14] L. Zhong, J.-M. Zhang, B. Su, R.S. Tan, J.C. Allen, G.S. Kassab, Application of patient-specific computational fluid dynamics in coronary and intra-cardiac flow simulations: Challenges and opportunities, *Front. Physiol.* 9 (2018) 742.
- [15] C. Bezerra, P.A. Lemos, F.A. Pinton, L. Müller, C. Bulant, G.M. Talou, R.A. Feijóo, A.E. Filh Esteves, P. Blanco, Tct-619 comparison of one-dimensional (1d) and three-dimensional (3d) models for the estimation of coronary fractional flow reserve through cardiovascular imaging, *J. Am. Coll. Cardiol.* 72 (13S) (2018) B248.
- [16] E. Boileau, S. Pant, C. Roobottom, I. Sazonov, J. Deng, X. Xie, P. Nithiarasu, Estimating the accuracy of a reduced-order model for the calculation of fractional flow reserve (FFR), *Int. J. Numer. Methods Biomed. Eng.* 34 (1) (2018) e2908.
- [17] L. Itu, P. Sharma, V. Mihalef, A. Kamen, C. Suciu, D. Lomaniciu, A patient-specific reduced-order model for coronary circulation, in: 2012 9th IEEE International Symposium on Biomedical Imaging, ISBI, IEEE, 2012, pp. 832–835.
- [18] N. Grande Gutiérrez, T. Sinno, S.L. Diamond, A 1D–3D hybrid model of patient-specific coronary hemodynamics, *Cardiovasc. Eng. Technol.* (2022) 1–12.
- [19] M.R. Pfaller, J. Pham, A. Verma, L. Pegolotti, N.M. Wilson, D.W. Parker, W. Yang, A.L. Marsden, Automated generation of 0D and 1D reduced-order models of patient-specific blood flow, *Int. J. Numer. Methods Biomed. Eng.* 38 (10) (2022) e3639.
- [20] S. Sankaran, H.J. Kim, G. Choi, C.A. Taylor, Uncertainty quantification in coronary blood flow simulations: impact of geometry, boundary conditions and blood viscosity, *J. Biomech.* 49 (12) (2016) 2540–2547.
- [21] K. Valen-Sendstad, A.W. Bergersen, Y. Shimogonya, L. Goubergrits, J. Bruening, J. Pallares, S. Cito, S. Piskin, K. Pekkan, A.J. Geers, et al., Real-world variability in the prediction of intracranial aneurysm wall shear stress: the 2015 international aneurysm CFD challenge, *Cardiovasc. Eng. Technol.* 9 (2018) 544–564.
- [22] G. Nannini, S. Saitta, L. Mariani, R. Maragna, A. Baggiano, S. Mushtaq, G. Pontone, A. Redaelli, An automated and time-efficient framework for simulation of coronary blood flow under steady and pulsatile conditions, *Comput. Methods Programs Biomed.* 257 (2024) 108415.
- [23] G. Nannini, S. Saitta, A. Baggiano, R. Maragna, S. Mushtaq, G. Pontone, A. Redaelli, A fully automated deep learning approach for coronary artery segmentation and comprehensive characterization, *APL Bioeng.* 8 (1) (2024).
- [24] L. Gu, X.-C. Cai, Fusing 2D and 3D convolutional neural networks for the segmentation of aorta and coronary arteries from CT images, *Artif. Intell. Med.* 121 (2021) 102189.
- [25] R. Gharleghi, D. Adikari, K. Ellenberger, S.-Y. Ooi, C. Ellis, C.-M. Chen, R. Gao, Y. He, R. Hussain, C.-Y. Lee, et al., Automated segmentation of normal and diseased coronary arteries—the asoca challenge, *Comput. Med. Imaging Graph.* 97 (2022) 102049.
- [26] S. Saitta, F. Sturla, A. Caimi, A. Riva, M.C. Palumbo, G. Nano, E. Votta, A.D. Corte, M. Glauber, D. Chiappino, et al., A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography, *J. Digit. Imaging* 35 (2) (2022) 226–239.
- [27] D. Alblas, J. Suk, C. Brune, K.K. Yeung, J.M. Wolterink, SIRE: scale-invariant, rotation-equivariant estimation of artery orientations using graph neural networks, 2023, arXiv preprint arXiv:2311.05400.
- [28] L. Itu, S. Rapaka, T. Passerini, B. Georgescu, C. Schwemmer, M. Schoebinger, T. Flohr, P. Sharma, D. Comaniciu, A machine-learning approach for computation of fractional flow reserve from coronary computed tomography, *J. Appl. Physiol.* 121 (1) (2016) 42–52.
- [29] B. Su, J.-M. Zhang, H. Zou, D. Ghista, T.T. Le, C. Chin, Generating wall shear stress for coronary artery in real-time using neural networks: Feasibility and initial results based on idealized models, *Comput. Biol. Med.* 126 (2020) 104038.
- [30] M.M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, Geometric deep learning: going beyond euclidean data, *IEEE Signal Process. Mag.* 34 (4) (2017) 18–42.
- [31] M.M. Bronstein, J. Bruna, T. Cohen, P. Velicković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021, arXiv preprint arXiv:2104.13478.
- [32] P. Rygiel, P. Pluszka, M. Zieba, T. Konopczynski, CenterlinePointNet++: A new point cloud based architecture for coronary artery pressure drop and vFR estimation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2023, pp. 781–790.
- [33] S. Wang, D. Wu, G. Li, Z. Zhang, W. Xiao, R. Li, A. Qiao, L. Jin, H. Liu, Deep learning-based hemodynamic prediction of carotid artery stenosis before and after surgical treatments, *Front. Physiol.* 13 (2023) 1094743.
- [34] X. Zhang, B. Mao, Y. Che, J. Kang, M. Luo, A. Qiao, Y. Liu, H. Anzai, M. Ohta, Y. Guo, et al., Physics-informed neural networks (PINNs) for 4D hemodynamics prediction: an investigation of optimal framework based on vascular morphology, *Comput. Biol. Med.* 164 (2023) 107287.
- [35] J. Suk, C. Brune, J.M. Wolterink, SE (3) symmetry lets graph neural networks learn arterial velocity estimation from small datasets, in: International Conference on Functional Imaging and Modeling of the Heart, Springer, 2023, pp. 445–454.
- [36] J. Suk, B. Imre, J.M. Wolterink, LaB-GATr: geometric algebra transformers for large biomedical surface and volume meshes, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2024, pp. 185–195.

- [37] J. Brehmer, P. De Haan, S. Behrends, T.S. Cohen, Geometric algebra transformer, *Adv. Neural Inf. Process. Syst.* 36 (2024).
- [38] L. Pegolotti, M.R. Pfaller, N.L. Rubio, K. Ding, R.B. Brufau, E. Darve, A.L. Marsden, Learning reduced-order models for cardiovascular simulations with graph neural networks, *Comput. Biol. Med.* 168 (2024) 107676.
- [39] R. Gharleghi, A. Sowmya, S. Beier, Transient wall shear stress estimation in coronary bifurcations using convolutional neural networks, *Comput. Methods Programs Biomed.* 225 (2022) 107013.
- [40] C.R. Qi, H. Su, K. Mo, L.J. Guibas, Pointnet: Deep learning on point sets for 3d classification and segmentation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 652–660.
- [41] N. Sharp, S. Attaiki, K. Crane, M. Ovsjanikov, Diffusionnet: Discretization agnostic learning on surfaces, *ACM Trans. Graph.* 41 (3) (2022) 1–16.
- [42] R. Wiersma, A. Nasikun, E. Eisemann, K. Hildebrandt, Deltaconv: anisotropic operators for geometric deep learning on point clouds, *ACM Trans. Graph.* 41 (4) (2022) 1–10.
- [43] A. Updegrave, N.M. Wilson, J. Merkow, H. Lan, A.L. Marsden, S.C. Shadden, SimVascular: an open source pipeline for cardiovascular simulation, *Ann. Biomed. Eng.* 45 (2017) 525–541.
- [44] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, et al., 3D Slicer as an image computing platform for the Quantitative Imaging Network, *Magn. Reson. Imaging* 30 (9) (2012) 1323–1341.
- [45] G. Pontone, S. Moharem-Elgamal, P. Maurovich-Horvat, O. Gaemperli, F. Pugliese, D. reviewers, M. Westwood, A. Stefanidis, K.F. Fox, B.A. Popescu, Training in cardiac computed tomography: EACVI certification process, *Eur. Hear. J.- Cardiovasc. Imaging* 19 (2) (2018) 123–126.
- [46] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimselshein, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [47] M. Fey, J.E. Lenssen, Fast graph representation learning with PyTorch geometric, 2019, arXiv preprint arXiv:1903.02428.
- [48] S. Ioffe, Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015, arXiv preprint arXiv:1502.03167.
- [49] B. Xu, Empirical evaluation of rectified activations in convolutional network, 2015, arXiv preprint arXiv:1505.00853.
- [50] M. Piccinelli, A. Veneziani, D.A. Steinman, A. Remuzzi, L. Antiga, A framework for geometric analysis of vascular structures: application to cerebral aneurysms, *IEEE Trans. Med. Imaging* 28 (8) (2009) 1141–1155.
- [51] B. Lévy, Laplace-beltrami eigenfunctions towards an algorithm that “understands” geometry, in: *IEEE International Conference on Shape Modeling and Applications* 2006, SMF’06, IEEE, 2006, p. 13.
- [52] J. Sun, M. Ovsjanikov, L. Guibas, A concise and provably informative multi-scale signature based on heat diffusion, in: *Computer Graphics Forum*, Vol. 28, Wiley Online Library, 2009, pp. 1383–1392.
- [53] M. Aubry, U. Schlickewei, D. Cremers, The wave kernel signature: A quantum mechanical approach to shape analysis, in: *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, IEEE, 2011, pp. 1626–1633.
- [54] J. Suk, G. Nannini, P. Rygiel, C. Brune, G. Pontone, A. Redaelli, J.M. Wolterink, Deep vectorised operators for pulsatile hemodynamics estimation in coronary arteries from a steady-state prior, 2024, arXiv preprint arXiv:2410.11920.
- [55] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al., SciPy 1.0: fundamental algorithms for scientific computing in Python, *Nature Methods* 17 (3) (2020) 261–272.
- [56] J. Suk, P. de Haan, P. Lippe, C. Brune, J.M. Wolterink, Mesh neural networks for SE (3)-equivariant hemodynamics estimation on the artery wall, *Comput. Biol. Med.* 173 (2024) 108328.
- [57] H. Yang, B.K. Nallamothu, C.A. Figueroa, K. Garikipati, Attention-based multi-fidelity machine learning model for fractional flow reserve assessment, *Comput. Methods Appl. Mech. Engrg.* 432 (2024) 117338.
- [58] A. Vaswani, Attention is all you need, *Adv. Neural Inf. Process. Syst.* (2017).
- [59] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18, Springer, 2015, pp. 234–241.