



MOX-Report No. 56/2026

**Physics-Informed Learning of Microvascular Flow Models using
Graph Neural Networks**

Botta, P.; Vitullo, P.; Ventimiglia, T.; Linninger, A.; Zunino, P.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Physics-Informed Learning of Microvascular Flow Models using Graph Neural Networks

Paolo Botta^{*,1}, Piermario Vitullo^{*,1}, Thomas Ventimiglia²,
Andreas Linninger^{2,3} and Paolo Zunino^{1,4}

^{*}The first two authors equally contributed to the work

¹MOX, Department of Mathematics, Politecnico di Milano, Italy

²Department of Bioengineering, University of Illinois at Chicago, IL, USA

³Department of Neurosurgery, University of Illinois at Chicago, IL, USA

⁴CMON Lab, Data Science Unit, Fondazione IRCCS Istituto Nazionale dei Tumori, Italy

Abstract

The simulation of microcirculatory blood flow in realistic vascular architectures poses significant challenges due to the multiscale nature of the problem and the topological complexity of capillary networks. In this work, we propose a novel deep learning-based reduced-order modeling strategy, leveraging Graph Neural Networks (GNNs) trained on synthetic microvascular graphs to approximate hemodynamic quantities on anatomically realistic domains. Our method combines algorithms for synthetic vascular generation with a physics-informed training procedure that integrates graph topological information and local flow dynamics. To ensure the physical reliability of the learned surrogates, we incorporate a physics-informed loss functional derived from the governing equations, allowing enforcement of mass conservation and rheological constraints. The resulting GNN architecture demonstrates robust generalization capabilities across diverse network configurations. The GNN formulation is validated on benchmark problems with linear and nonlinear rheology, showing accurate pressure and velocity field reconstruction with substantial computational gains over full-order solvers. The methodology showcases significant generalization capabilities with respect to vascular complexity, as highlighted by tests on data from the mouse cerebral cortex. This work establishes a new class of graph-based surrogate models for microvascular flow, grounded in physical laws and equipped with inductive biases that mirror mass conservation and rheological models, opening new directions for real-time inference in vascular modeling and biomedical applications.

Key words. Microvascular blood flow; Physics-informed machine learning; Graph neural networks; Reduced order modeling;

MSC codes. 35Q92, 68T07, 65N22, 65Z05

1 Introduction

The simulation of microcirculatory blood flow is an essential component in understanding tissue perfusion, vascular function, and the development of pathological conditions such as ischemia, inflammation, and cancer [27, 37, 8], just to give a few examples. On the microscopic scale, the transport of blood through networks of arterioles, capillaries, and venules is governed by intricate hemodynamic phenomena and highly complex topological structures, which pose significant computational challenges.

High-fidelity models are able to accurately capture these dynamics. However, their applicability to large-scale, anatomically realistic microvascular networks is severely hampered by the associated computational burden and the difficulties in accommodating topological variability [20, 12, 13]. In the case of fully resolved three-dimensional models (including or not the description of red blood

cells) even in relatively small tissue volumes, such as a cortical sample of size 1 mm^3 [12, 10, 35, 34], the resolution required to represent individual vessels can easily lead to systems not yet computable using a direct numerical simulation approach. To overcome this limitation, dimensional model reduction has been adopted with the purpose to reduce the mathematical and computational complexity of blood flow to one-dimensional models defined along the vessel centerline, such as in the *geometric multiscale methods* [31]. This approach can be modulated to a high level of physiological and anatomical detail [2, 22, 28, 39, 21], despite the intrinsic multiscale nature of microvascular flow [18, 24, 25, 38]. State of the art simulations reach today the organ scale level [21, 42]. For an accurate description of the pulsatile regime such simulations still require high performance computing resources, while the computational cost is substantially reduced in the case of the steady regime typical of the microcirculation. Despite the great success of one-dimensional model for the simulation of blood flow, the research on surrogate models for blood flow in complex networks is still active [26, 40], with the aim to preserve predictive accuracy while achieving substantial computational speed-ups. These advancement will extend the ability to simulate pulsatile blood flow beyond the current limits and also enable the use of the simpler microcirculation models in the context of computational workflows involving many-query simulations such as inverse problem formulations, optimization, and uncertainty quantification. Vascular networks admit a natural representation as graphs, with vessels corresponding to edges and junctions to nodes. GNNs are designed to operate directly on such graph-structured inputs, learning approximations of the underlying physics from data [6, 11, 3]. GNNs perform inference through forward evaluation of a trained neural network, therefore reducing computational costs by orders of magnitude and enabling real-time prediction [19]. Moreover, GNNs can be trained in a physics-informed fashion, incorporating domain-specific knowledge (such as mass conservation and nonlinear rheology) directly into the loss functional. This strategy mitigates the typical limitations of purely data-driven models and enhances the ability to generalize to unseen network geometries. In this work, we develop a physics-informed computational learning problem specifically designed to address flow problems on graphs using GNN architectures. We apply this framework to define surrogate models for microcirculatory blood flow in the steady regime, although ongoing work confirms that the approach can be successfully extended to the pulsatile regime [4]. The GNN surrogate proposed here predicts both nodal- and edge-based flow quantities and combines data-driven learning with physics-informed constraints. Remarkably, we show that a model trained exclusively on synthetic capillary networks with generic topological and geometric features is able to generalize robustly to large-scale, anatomically realistic vascular reconstructions. In particular, when tested on mouse cerebral cortex networks that exhibit a substantially higher degree of complexity than training data [36, 20, 12, 13], the GNN maintains high predictive precision and reproduces full-order simulation results with excellent agreement. This highlights the capability of the proposed framework to extrapolate beyond the synthetic training distribution and its potential for deployment in biomedical applications that require efficient and reliable vascular flow predictions. Compared to previous studies on GNNs for blood flow [23, 26, 15], our approach introduces several differences. Existing work, such as the reduced-order models of [23] and their recent extensions that integrate recurrent architectures for temporal learning [15], focuses primarily on large-scale cardiovascular flow within arterial trees and exploits data-driven surrogates of one-dimensional flow dynamics. In contrast, we address the steady microcirculatory regime on the scale of large vascular networks, where flow is governed by mass conservation, diffusion mechanisms, and nonlinear rheology. For extensions to the pulsatile regime, we refer to Section 7. The remainder of this manuscript is organized as follows. Section 2 introduces the general mathematical framework for physics-informed learning on graphs, formalizing the parameter-to-solution map and its approximation via GNNs. Section 3 presents the physical models for microvascular blood flow, encompassing both linear and nonlinear rheology, while Section 4 details the algorithms used to generate synthetic and anatomically realistic vascular networks. Section 5 describes the architecture, loss formulations, and training strategies of the proposed GNN surrogates, highlighting the integration of physical constraints within the learning process. Finally, Section 6 reports the numerical experiments that evaluate the predictive accuracy and generalization of the models in synthetic and image-based microvascular networks, followed by the concluding remarks.

2 Physics informed learning using graph neural networks

In this section, we introduce the methodological framework underlying the proposed approach. We formulate parameterized problems on graphs, providing a general setting for network-based physical systems that couple discrete topological structures with continuous quantities defined on nodes and edges. The associated parameter-to-solution map thus exhibits a mixed discrete–continuous nature. Graph Neural Networks (GNNs) are naturally suited to fit this structure, leveraging graph connectivity to propagate information and learn nonlinear relations between nodal and edge variables. This formulation yields a unified paradigm for the operator approximation on graphs, ideal for the representation of microvascular blood flow in complex configurations.

2.1 A learning problem on graphs

Let $G := (V, E)$ be a finite, *directed graph* encoding the topological connectivity of the network, where $V := \{v_j\}_{j=1}^n$ denotes the set of n vertices in \mathbb{R}^3 and $E := \{e_i\}_{i=1}^m$ is the set of m directed edges that connect pairs of vertices. In particular, each edge $e = (v_s, v_t) \subset \mathbb{R}^3$ sets the orientation from the source node v_s to the target node v_t . All relevant information on the connectivity of a directed graph is encoded in the connectivity (or incidence) matrix $C \in \mathbb{Z}^{m \times n}$ with entries:

$$C_{ij} := \{1, \text{ if } e_i \text{ enters } v_j; \quad -1, \text{ if } e_i \text{ leaves } v_j; \quad 0, \text{ otherwise}\}.$$

We then introduce a vector of weights defined on the vertices of G , $W^V := \{\mathbf{w}_j^V\}_{j=1}^n$, and on the edges, $W^E := \{\mathbf{w}_i^E\}_{i=1}^m$, where $\mathbf{w}_j^V \in \mathbb{R}^{p_V} \forall j = 1, \dots, n$ and $\mathbf{w}_i^E \in \mathbb{R}^{p_E} \forall i = 1, \dots, m$. This leads to the definition of the vector of weights $W := [W^V, W^E]^\top \in \mathbb{R}^{n p_V + m p_E}$. A *weighted directed graph* is the triple $\mathcal{G} := (V, E, W)$. Each graph (directed or weighted) \mathcal{G} can be embedded in \mathbb{R}^3 , proceeding as follows:

- (i) we associate with each node in V a point $\mathbf{x}_i \in \mathbb{R}^d$ in a bijective way, with d denoting the number of spatial dimensions ($d = 2, 3$). Consequently, we introduce the matrix that encodes the coordinates of each vertex, $X \in \mathbb{R}^{n \times 3}$ such that $X(i, :) = \mathbf{x}_i^\top$;
- (ii) we connect two points $\mathbf{x}_v, \mathbf{x}_w$ by a three-dimensional simple arc \mathbf{s}_{vw} if and only if v, w are adjacent, in such a way that different arcs do not share any internal points;
- (iii) we define a length function $\rho : E \rightarrow \mathbb{R}_+$ assigning a metric on each edge $e = (v_s, v_t)$, such that its length is $L = \rho(e)$.

We know that for any weighted directed graph \mathcal{G} , the resulting embedded object in \mathbb{R}^3 is a manifold. In particular, assuming for all edges e the convention of identifying the source node v_s with 0 and the target node v_t with $\rho(e)$, it is possible to define a metric over \mathcal{G} .

This construction is the root of the notion of a *metric graph*. Let $\mathcal{G} = (V, E, W)$ be a weighted directed graph and let us define $\mathcal{E} := \prod_{e \in E} \{e\} \times (0, \rho(e))$. Then the triple $\mathfrak{G} := (V, \mathcal{E}, W)$ is called the (weighted) *metric graph* over \mathcal{G} . A point of \mathfrak{G} is an element of either V or \mathcal{E} . The elements of \mathcal{E} are called metric edges.

We now consider steady-state problems governed by differential operators acting on \mathfrak{G} . Under these assumptions, the most general problem we can consider is of the form

$$\begin{cases} \mathfrak{L}^W u^W = f^W & \text{in } \mathfrak{G}, \\ \mathfrak{B}^W u^W = g^W & \text{on } \partial \mathfrak{G}, \end{cases} \quad (1)$$

where \mathfrak{L}^W denotes an operator on the metric graph, and \mathfrak{B}^W represents boundary or vertex conditions enforcing continuity and conservation properties across the nodes of \mathfrak{G} . A prototypical example of \mathfrak{L}^W is a possibly semilinear elliptic partial differential equation in \mathfrak{G} , where the weights W represent the parameters of the operator \mathfrak{L} on the edges $[0, L_i]$.

However, as the physical models used for microvascular blood flow are rooted in Poiseuille’s law, the problems addressed here will reduce to nonlinear algebraic equations on the weighted directed graph $\mathcal{G} = (V, E, W)$. Let us consider for example the following linear relations:

$$\begin{cases} u_{s(i)}^V - u_{t(i)}^V = R_i^E u_i^E, & i = 1, \dots, m, \\ \sum_{i \in E(v_j)} C_{ij} u_i^E = 0, & j = 1, \dots, n. \end{cases} \quad (2)$$

where C is the incidence matrix, and R_i^E are the edge resistances. Setting $\mathbf{u}_j^V \in \mathbb{R}^{p_V} \forall j = 1, \dots, n$ and $\mathbf{u}_i^E \in \mathbb{R}^{p_E} \forall i = 1, \dots, m$, this problem can be formulated in terms of unknown vector values on the vertices and edges of the graph, defined as follows:

$$\mathbf{U}^V := \{\mathbf{u}_j^V\}_{j=1}^n, \quad \mathbf{U}^E := \{\mathbf{u}_i^E\}_{i=1}^m, \quad \mathbf{U}^W := \begin{bmatrix} \mathbf{U}^V \\ \mathbf{U}^E \end{bmatrix}.$$

As a result, problem (2) becomes the following:

$$L^W \mathbf{U}^W = \mathbf{F}^W, \quad L^W := \begin{bmatrix} C & -R^E \\ 0 & C^\top \end{bmatrix}, \quad \mathbf{F}^W := \begin{bmatrix} \mathbf{F}^V \\ \mathbf{F}^E \end{bmatrix}, \quad (3)$$

where $R_i^E = r(\mathbf{w}_i)$, being $r : \mathbb{R}^{p_V + p_E} \rightarrow \mathbb{R}$ a suitable constitutive function. The right-hand side \mathbf{F}^W encodes the boundary conditions that will be addressed later. The system can be reduced to the vertex unknowns, with the corresponding matrix $C^\top (R^E)^{-1} C$ defining a graph Laplacian, which is a possible finite-dimensional realization of the continuous operator \mathfrak{L}^W . The solution depends on both the weights W (representing the physical parameters) and the (simple or directed) graph G , encoded by the matrix of vertices $X \in \mathbb{R}^{n \times 3}$ and the discrete incidence matrix $C \in \mathbb{Z}^{m \times n}$. We therefore define the *parameter-to-solution map*

$$\mathcal{S} : \mathbb{R}^{n \times d} \times \mathbb{Z}^{m \times n} \times \mathbb{R}^{n p_V + m p_E} \longrightarrow \mathbb{R}^{n p_V + m p_E}, \quad \mathcal{S}(G, W) = \mathbf{U}^W.$$

Given $\widehat{\mathcal{H}}$ a set of candidate functions for this map, $\widehat{\mathcal{S}} \in \widehat{\mathcal{H}}$, given a collection of labeled data $\mathcal{S}(G_k, W_k) = \mathbf{U}_k^W$, for $k = 1, \dots, N$, given a loss function measuring the discrepancy between data and predictions, $l(\widehat{\mathbf{U}}^W, \mathbf{U}^W)$, and a feature mapping of the weight vector $\widetilde{W} := \Phi(W)$, we aim to solve the following supervised learning problem, that consists of the minimization of the empirical error \mathcal{L} (from now on called *empirical total loss* or just *loss*, with little abuse of notation) over the candidate functions:

$$\begin{aligned} \mathcal{L}((\widehat{\mathcal{S}}(G, \widetilde{W}), \mathbf{U}^W); (G_k, W_k), \mathbf{U}_k^W) &= \frac{1}{N} \sum_{k=1}^N l(\widehat{\mathbf{U}}_k^W, \mathbf{U}_k^W), \\ \widehat{\mathcal{S}}^* &= \operatorname{argmin}_{\widehat{\mathcal{S}} \in \widehat{\mathcal{H}}} \mathcal{L}((\widehat{\mathcal{S}}(G, \widetilde{W}), \mathbf{U}^W); (G_k, W_k), \mathbf{U}_k^W). \end{aligned}$$

To solve this learning problem, we adopt a Graph Neural Network architecture, which is ideally suited to approximate a parameter-to-solution map featuring a mixed mathematical structure combining discrete and continuous information. In fact, GNNs [33, 6, 3] process node- and edge-based continuous features through standard feed-forward neural networks, while a message passing mechanism exploits the discrete incidence structure of the underlying graph to propagate information across connected entities.

2.2 Graph Neural Networks architecture

Building upon the formalism introduced in Section 2.1, we recall that the physical domain is a weighted directed graph $\mathcal{G} = (V, E, W)$. For computational purposes, we consider its discrete

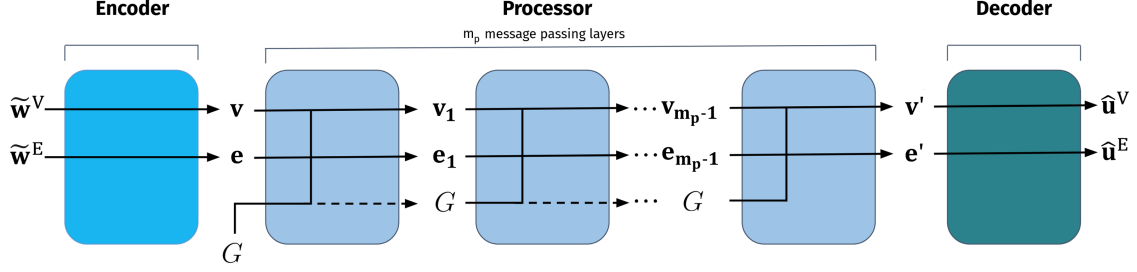


Figure 1: GNN complete architecture, consisting of an encoder, a processor, and a decoder.

counterpart $G = (V, E)$, which retains the same topological connectivity, while the physical and geometric quantities are encoded within the node and edge weights $W = [W^V, W^E]^T$.

Within this setting, we adopt a *supervised learning* approach to approximate the parameter-to-solution map introduced in Eq. (3), acting on weighted directed graphs. Formally, the goal is to learn an approximation $\hat{\mathcal{S}}$ of the exact solution operator \mathcal{S} , namely:

$$\hat{\mathbf{U}}^W = \hat{\mathcal{S}}(G, \tilde{W}) \approx \mathcal{S}(G, W) = \mathbf{U}^W.$$

The adopted GNN architecture follows the standard *Encoder–Processor–Decoder* paradigm, schematically illustrated in Figure 1. The *encoder* (denoted with $\mathcal{E}_v(\tilde{w}^V)$, $\mathcal{E}_e(\tilde{w}^E)$) for the vertices and edges, respectively) maps the physical and geometric input features defined on the nodes and edges into a latent space of fixed dimension. The *processor*, composed of multiple message passing layers (named $\mathcal{P}_v(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E))$ and $\mathcal{P}_e(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E))$), iteratively propagates and updates information throughout the graph according to its incidence structure. We note that all the message passing layers share a fixed hidden dimension l , updating node features $\mathbf{v} : V \rightarrow \mathbb{R}^l$ and edge features $\mathbf{e} : E \rightarrow \mathbb{R}^l$ without modifying their dimension. Finally, *decoders* $\mathcal{D}_v(\mathcal{P}_v(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E)))$ and $\mathcal{D}_e(\mathcal{P}_e(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E)))$ transform the latent representations of the nodes and edges into the predicted physical quantities. More precisely, the composition of these maps can be represented as follows:

$$\hat{\mathbf{u}}^V = \mathcal{D}_v\left(\mathcal{P}_v(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E))\right), \quad \hat{\mathbf{u}}^E = \mathcal{D}_e\left(\mathcal{P}_e(G, \mathcal{E}_v(\tilde{w}^V), \mathcal{E}_e(\tilde{w}^E))\right).$$

The resulting outputs are assembled into the vector-valued approximation of the solution, which provides the learned estimate of the parameter-to-solution map:

$$\hat{\mathcal{S}}(G, \tilde{W}) = \hat{\mathbf{U}}^W := \begin{bmatrix} \hat{\mathbf{U}}^V \\ \hat{\mathbf{U}}^E \end{bmatrix}, \quad \hat{\mathbf{U}}^V := \{\hat{\mathbf{u}}_j^V\}_{j=1}^n, \quad \hat{\mathbf{U}}^E := \{\hat{\mathbf{u}}_i^E\}_{i=1}^m.$$

This architecture, illustrated in Figure 1, enables simultaneous treatment of nodal and edge quantities in the processor and in the decoder, which is essential for predicting coupled variables such as pressure and flow.

2.2.1 Feedforward Neural Networks

Feedforward Neural Networks (FFNNs) constitute the fundamental building blocks of most modern deep learning architectures. They are composed of a sequence of layers that apply affine transformations followed by nonlinear activation functions, enabling the network to approximate complex nonlinear mappings between input and output spaces. Given an input vector $\mathbf{x} \in \mathbb{R}^{d_{\text{in}}}$, a generic FFNN with K layers can be expressed as the composition:

$$\Psi(\mathbf{x}; \theta) = \mathbf{f}_K \circ \mathbf{f}_{K-1} \circ \cdots \circ \mathbf{f}_1(\mathbf{x}), \quad \mathbf{f}_\ell(\mathbf{z}) = \sigma_\ell(\mathcal{W}_\ell \mathbf{z} + \mathbf{b}_\ell),$$

where \mathcal{W}_ℓ and \mathbf{b}_ℓ denote the weight matrix and bias vector of the ℓ -th layer, respectively, σ_ℓ is a nonlinear activation function (e.g., ReLU, tanh), and $\theta = \{\mathcal{W}_\ell, \mathbf{b}_\ell\}_{\ell=1}^K \in \Theta$ represents the set of all trainable parameters of the network.

In the present framework, FFNNs are employed throughout the architecture wherever a nonlinear map Ψ appears. For simplicity of notation, we omit the explicit dependence on the parameters θ in the subsequent sections, writing simply $\Psi(\cdot)$ instead of $\Psi(\cdot; \theta)$, with the understanding that all feedforward mappings implicitly depend on their respective trainable weights and biases.

2.2.2 The Encoder and Decoder modules

Each node $v \in V$ and edge $e \in E$ is associated with feature vectors that represent local physical or geometric information. *Node input features* $\widetilde{W}^V := \{\widetilde{\mathbf{w}}_j^V\}_{j=1}^n$, with $\widetilde{\mathbf{w}}_j^V \in \mathbb{R}^{\tilde{p}_V}$, encode local physical quantities defined at each vertex (e.g., boundary conditions, nodal parameters). *Edge input features* $\widetilde{W}^E := \{\widetilde{\mathbf{w}}_i^E\}_{i=1}^m$, with $\widetilde{\mathbf{w}}_i^E \in \mathbb{R}^{\tilde{p}_E}$, encode the geometric or material properties associated with each edge (e.g., length, diameter, resistance). The specific form of $\widetilde{\mathbf{w}}^V$ and $\widetilde{\mathbf{w}}^E$ is problem-dependent. The collection of these features is $\widetilde{W} := [\widetilde{W}^V, \widetilde{W}^E]^\top := \Phi(W)$ as a feature mapping of the weight vector. We stress that the mapping Φ is a hyperparameter specified *a priori* based on the problem structure. Its choice determines which weights and aspects of the physical system are exposed to the learning architecture. Since the dimensions of the input features generally differ from l , we employ feedforward neural networks to map them to the latent space:

$$\mathcal{E}_v(\widetilde{\mathbf{w}}^V) = \Psi_{\mathcal{E}}^v(\widetilde{\mathbf{w}}^V) \quad \mathcal{E}_e(\widetilde{\mathbf{w}}^E) = \Psi_{\mathcal{E}}^e(\widetilde{\mathbf{w}}^E)$$

The GNN output features corresponding to node- and edge-level quantities are the *node output features* $\mathbf{U}^V := \{\mathbf{u}_j^V\}_{j=1}^n$, with $\mathbf{u}_j^V \in \mathbb{R}^{p'_V}$, representing physical states at the nodes (e.g., pressure). *Edge output features* $\mathbf{U}^E := \{\mathbf{u}_i^E\}_{i=1}^m$, with $\mathbf{u}_i^E \in \mathbb{R}^{p'_E}$, correspond to the quantities defined along the edges (for example, flow rate).

The processor module outputs hidden node and edge features, $\mathbf{v}' = \mathcal{P}_v(G, \mathbf{v}, \mathbf{e})$ and $\mathbf{e}' = \mathcal{P}_e(G, \mathbf{v}, \mathbf{e})$. Since these do not match the output dimensionalities, two separate feedforward neural networks are employed to recover the physical output features:

$$\mathcal{D}_v(\mathbf{v}') = \Psi_{\mathcal{D}}^v(\mathbf{v}'), \quad \mathcal{D}_e(\mathbf{e}') = \Psi_{\mathcal{D}}^e(\mathbf{e}').$$

While the terms *encoder* and *decoder* are traditionally associated with dimensionality reduction and reconstruction, in our context they are used more broadly. The encoder extracts latent features from the input, and the decoder maps latent representations to physically meaningful outputs. Depending on the task, these modules may either increase or decrease the dimension of the features.

2.2.3 The Processor module

The processor module is composed of m_p message passing steps that iteratively update both the node and the edge features. Let H_j and F_j denote the edge and node update functions. Starting from encoded features, the recursive update now reads:

$$\begin{cases} \mathbf{v}_0 = \mathcal{E}_v(\widetilde{\mathbf{w}}^V), & \mathbf{e}_0 = \mathcal{E}_e(\widetilde{\mathbf{w}}^E), \\ \mathbf{e}_j = H_j(G, \mathbf{v}_{j-1}, \mathbf{e}_{j-1}) & j = 1, \dots, m_p, \\ \mathbf{v}_j = F_j(G, \mathbf{v}_{j-1}, \mathbf{e}_j) & j = 1, \dots, m_p, \end{cases}$$

so that the processor outputs

$$\mathcal{P}_v(G, \mathbf{v}, \mathbf{e}) = F_{m_p}(G, \mathbf{v}_{m_p-1}, \mathbf{e}_{m_p}), \quad \mathcal{P}_e(G, \mathbf{v}, \mathbf{e}) = H_{m_p}(G, \mathbf{v}_{m_p-1}, \mathbf{e}_{m_p-1}).$$

This formulation makes explicit that the node and edge states evolve in tandem at each message passing step, by iteratively applying the message passing layers m_p . The parameter m_p is a

hyperparameter of the architecture: a processor with m_p layers enables communication between nodes that are m_p edges apart, thereby extending the receptive field from local to increasingly nonlocal interactions.

The key component of this algorithm is the *message passing layer*, which propagates information along graph edges [33, 6, 3]. In its standard form, the update acts only on node features, while edge features remain unchanged. Specifically, given hidden node features \mathbf{v} and hidden edge features \mathbf{e} , the message passing layer $F = F(G, \mathbf{v}, \mathbf{e})$ takes as input \mathbf{v} , \mathbf{e} and G and outputs a new collection of vertex features $\mathbf{v}' : V \rightarrow \mathbb{R}^l$.

In order to define nodal update function F , let the mappings $\mathbf{v} : V \rightarrow \mathbb{R}^l$ and $\mathbf{e} : E \rightarrow \mathbb{R}^l$ denote hidden node and edge features of dimension l . For a directed edge $(v_s, v_t) \in E$, we set

$$\mathbf{v}_{\text{in}}(v_s, v_t) = \mathbf{v}(v_s), \quad \mathbf{v}_{\text{out}}(v_s, v_t) = \mathbf{v}(v_t).$$

In particular, \mathbf{v}_{in} computes the vertex features of the corresponding source node, while \mathbf{v}_{out} computes the vertex features of the corresponding destination node. The aggregation function for a node $v \in V$ is then defined as

$$\bar{\mathbf{e}}(v) = \sum \mathbf{e}(v_s, v), \quad \forall (v_s, v) \in E, \quad v_s \in V,$$

which collects the contributions of all incoming edges. Given two functions with the same domain $\mathbf{f} : X \rightarrow \mathbb{R}^a$ and $\mathbf{g} : X \rightarrow \mathbb{R}^b$, the concatenation operator \oplus is defined as

$$(\mathbf{f} \oplus \mathbf{g})(x) = [f_1(x), \dots, f_a(x), g_1(x), \dots, g_b(x)]$$

where $x \in X$ is a generic input. The standard message passing update reads

$$F(G, \mathbf{v}, \mathbf{e}) = \Psi_u(\mathbf{v} \oplus \Psi_M(\overline{\mathbf{e} \oplus \mathbf{v}_{\text{in}} \oplus \mathbf{v}_{\text{out}}})) ,$$

which concatenates and aggregates contributions from incoming edges, and Ψ_u and Ψ_M are feedforward neural networks, such that $\Psi_u : \mathbb{R}^{2l} \rightarrow \mathbb{R}^l$ and $\Psi_M : \mathbb{R}^{3l} \rightarrow \mathbb{R}^l$. In GNN architectures, F is considered a transformer that modifies the characteristics of the vertices connected to the graph nodes. A message passing layer, within these architectures, refers to a specific type of graph-forwarding routine, which exploits the local structure of the graph G . This routine allows communication exclusively between neighboring nodes.

In what follows, we will adopt an extended message passing mechanism that updates both node and edge states through local interactions. Specifically, the update rules are:

$$H(G, \mathbf{v}, \mathbf{e}) = \Psi_e(\mathbf{e} \oplus \mathbf{v}_{\text{in}} \oplus \mathbf{v}_{\text{out}}), \quad F(G, \mathbf{v}, \mathbf{e}) = \Psi_u(\mathbf{v} \oplus \Psi_M(\overline{H \oplus \mathbf{v}_{\text{in}} \oplus \mathbf{v}_{\text{out}}})) ,$$

where $\Psi_e : \mathbb{R}^{3l} \rightarrow \mathbb{R}^l$ is an additional feedforward neural network. Note that the nonlinearity of the message passing update arises from the activation functions within the neural networks Ψ_e, Ψ_M, Ψ_u . This formulation makes explicit that both edge and node features evolve at each iteration, a critical property when modeling vascular networks where edges carry hemodynamic variables (e.g., flow, diameter) and nodes represent pressures or bifurcation states. For a graphical representation of this new strategy, see Figure 2.

To improve gradient flow and training stability, we include skip (residual) connections between consecutive message passing layers. This strategy, originally introduced in ResNet architectures [14], has proven effective in GNNs as well [41]. Concretely, the processor module with $k > 0$ skip connections becomes:

$$\begin{cases} \mathbf{v}_0 = \mathcal{E}_v(\tilde{\mathbf{w}}^V), & \mathbf{e}_0 = \mathcal{E}_e(\tilde{\mathbf{w}}^E), \\ \mathbf{e}_j = H_j(G, \mathbf{v}_{j-1}, \mathbf{e}_{j-1}) + \mathbf{e}_{j-k} \mathbb{I}(j \geq k) & j = 1, \dots, m_p, \\ \mathbf{v}_j = F_j(G, \mathbf{v}_{j-1}, \mathbf{e}_j) + \mathbf{v}_{j-k} \mathbb{I}(j \geq k) & j = 1, \dots, m_p, \end{cases}$$

Skip connections are particularly beneficial when stacking a large number of message passing layers ($m_p \gg 1$), as they allow information to propagate across the network without degradation, enabling more expressive but still trainable architectures.

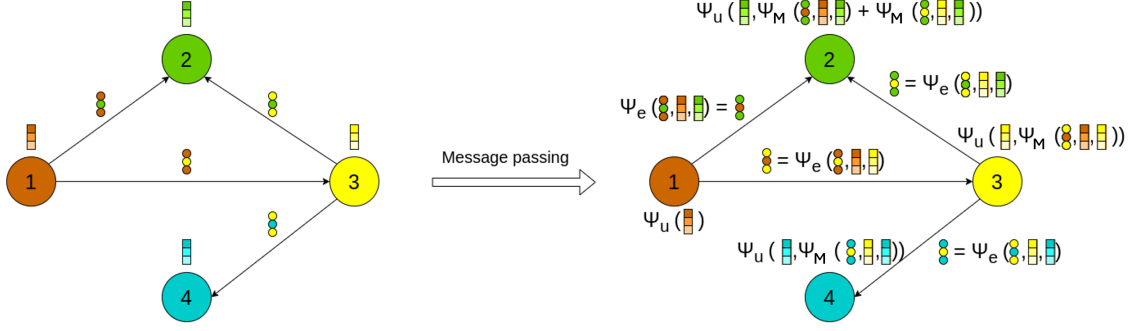


Figure 2: Message propagation and aggregation with edge features update.

Finally, although the vascular network naturally defines a directed graph, reflecting the physical orientation of blood flow, we model the input graph provided to the GNN as undirected. This choice allows the message-passing scheme to propagate information bidirectionally along each vessel segment, enabling both source-to-target and target-to-source interactions at every iteration. Empirically, we observe that this symmetric exchange of information significantly improves convergence of the training algorithm and predictive accuracy.

2.3 Loss functions and training

To train the GNN surrogate models introduced in Section 2.2, we adopt composite loss functions that incorporate both data fidelity, as in the supervised learning setting, and problem-specific physical information within a physics-informed framework. As a consequence, the total loss function used to train the GNN is defined as:

$$\mathcal{L}(\widehat{\mathcal{S}}, \mathbf{U}^W; \theta) := \delta \mathcal{L}_{\text{data}}(\theta) + (1 - \delta) \mathcal{L}_{\text{physics}}(\theta), \quad (4)$$

where $\mathcal{L}_{\text{data}}$ ensures data fidelity and $\mathcal{L}_{\text{physics}}$ is the physics-informed component. The symbol θ denotes the trainable parameters of the network, and $\delta \in [0, 1]$ is a tunable scalar that weights the relative importance of the data-driven and physics-informed components.

Unless otherwise specified, from now on we denote the Euclidean norm of the vector \mathbf{v} by $\|\mathbf{v}\|$. This is also called the L^2 norm of the vector. The data-fidelity loss quantifies the mismatch between the GNN predictions and the full-order model (FOM) ground truth on a dataset of N training graphs:

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{N} \sum_{k=1}^N \left(\frac{1}{n_k} \sum_{j=1}^{p'_V} \alpha_j \|\mathbf{u}_{j,k}^V - \widehat{\mathbf{u}}_{j,k}^V(\theta)\|^2 + \frac{1}{m_k} \sum_{i=1}^{p'_E} \beta_i \|\mathbf{u}_{i,k}^E - \widehat{\mathbf{u}}_{i,k}^E(\theta)\|^2 \right), \quad (5)$$

where $\{\alpha_j\}_{j=1}^{p'_V}$, $\{\beta_i\}_{i=1}^{p'_E}$ are tunable hyperparameters, $\mathbf{u}_{j,k}^V$, $\widehat{\mathbf{u}}_{j,k}^V$ are vectors of FOM outputs and GNN predictions for the output feature of the j -th node and for the k -th graph. Similarly, $\mathbf{u}_{i,k}^E$, $\widehat{\mathbf{u}}_{i,k}^E$ are vectors of FOM outputs and GNN predictions for the i -th edge output feature and for the k -th graph. In particular, the GNN predictions $\widehat{\mathbf{u}}_{j,k}^V$ and $\widehat{\mathbf{u}}_{i,k}^E$ are feedforward mappings explicitly depending on the parameters θ that encode their trainable weights and biases.

The physics-informed component of the loss enforces the governing principles of the graph-based problem by penalizing the residuals of key equations [32]:

$$\mathcal{L}_{\text{physics}}(\theta) = \gamma \mathcal{L}_{\text{constitutive}}(\theta) + (1 - \gamma) \mathcal{L}_{\text{mass}}(\theta), \quad (6)$$

where $\gamma \in [0, 1]$ controls the balance between the constitutive relation and the mass conservation constraint.

The constitutive residual enforces the consistency between the predicted solution and the constitutive law governing the problem on the graph through the term:

$$\mathcal{L}_{\text{constitutive}}(\theta) = \frac{1}{N} \sum_{k=1}^N \|C_k \widehat{\mathbf{U}}_k^V(\theta) - R_k^E \widehat{\mathbf{U}}_k^E(\theta) - \mathbf{F}_k^V\|^2. \quad (7)$$

In the context of microvascular flow, this residual enforces, for example, the Poiseuille constitutive law, ensuring consistency between predicted pressure drops and flow rates along the vessels through the edge resistances computed from predicted features (e.g., diameter, hematocrit).

The mass balance residual enforces the continuity conditions at the junction nodes of the graph, penalizing imbalances of the nodal constraints:

$$\mathcal{L}_{\text{mass}}(\theta) = \frac{1}{N} \sum_{k=1}^N \|C_k^\top \widehat{\mathbf{U}}_k^E(\theta) - \mathbf{F}_k^E\|^2. \quad (8)$$

In the specific case of microvascular flow, this residual reduces to a mass conservation condition that ensures that total inflow and outflow at each bifurcation node are balanced.

3 Microvascular blood flow models

Accurate modeling of blood flow in microvascular networks requires a careful combination of three key components: biologically meaningful vascular geometries (addressed in Section 4), physically consistent hemodynamic models, and computational methods capable of resolving flow across complex networks with high fidelity [27, 37]. In this section, we describe the modeling framework adopted to simulate blood flow through synthetic and anatomically realistic networks, formulated within the abstract setting introduced in Section 2.1. These simulations serve as the basis for the development and evaluation of the surrogate models presented in Section 2. In particular, they are used both to generate training data for Graph Neural Network (GNN) models and to evaluate their predictive performance on out-of-distribution vascular topologies.

3.1 Physical models for microvascular blood flow

The cornerstone of microvascular blood flow models is Poiseuille’s law, which describes the relationship between flow rate, pressure gradient, and hydraulic resistance in laminar flow conditions [27]. Given a vessel of length L and diameter D , the axial velocity of blood v is related to the volumetric flow rate Q by the expression:

$$v = \frac{Q}{\pi \frac{D^2}{4}} = \frac{\Delta P}{\pi R \frac{D^2}{4}}, \quad (9)$$

where ΔP denotes the pressure drop across the vessel and R is the hydraulic resistance. According to Poiseuille’s law, resistance depends on the geometry of the vessel and the viscosity of the blood μ as follows:

$$R = \frac{128\mu L}{\pi D^4}.$$

For a given edge $e \in E$, this expression defines the coefficient R_i^E and the constitutive function $r(\mathbf{w}_i)$ in (2), being $\mathbf{w}_i = \{\mu_i, D_i, L_i\}$. Blood viscosity plays a central role in determining the resistance to flow and may be treated as either constant or variable depending on the physical regime considered. In microvascular domains, where RBC dynamics are non-negligible, the effective viscosity becomes a nonlinear function of both vessel diameter and local hematocrit (i.e. the volumetric concentration of RBCs in the blood [37]). Before detailing the main models, we list the key assumptions underlying our formulation [9, 10, 27, 37]:

(i) Blood flow and hematocrit distribution is stationary;

- (ii) RBC transport is dominated by advection;
- (iii) Hematocrit is assumed to be uniform along the length of each vessel (no axial variations due to RBC production, loss, or plasma leakage);
- (iv) Only Y-junctions with straight branches are considered, and angle dependence is neglected;
- (v) The Fahraeus–Lindqvist and Zweifach–Fung effects are explicitly included to capture key features of microvascular blood flow [9, 10], while other aspects such as leukocyte interactions and diameter irregularities are neglected.

We investigate two distinct modeling frameworks for microvascular blood flow. The first is a linear rheology model, which assumes a constant blood viscosity and neglects the effects of RBCs. This simplification leads to a linear system of equations such as (2) that can be solved efficiently using standard numerical techniques. The second approach is a nonlinear rheology model, which incorporates the dependence of viscosity on hematocrit and accounts for the redistribution of RBC in vascular bifurcations. This added complexity results in a nonlinear system that requires iterative procedures to achieve convergence.

3.1.1 Linear models

The vascular domain is represented as a weighted directed graph $\mathcal{G} = (V, E, W)$ defined before, where the weights on the edges $\mathbf{w}_i^E = \{\mu_i, D_i, L_i\}$ are the diameter D_i and the length L_i of each edge and a uniform blood viscosity $\mu_i = \mu$. The hydraulic state of the network is characterized by the nodal vector \mathbf{U}^V representing the pressure, that we name from now on $\mathbf{P} = \{P_j\}_{j=1}^n$ and an edge vector \mathbf{U}^E representing the flow rates, named $\mathbf{Q} = \{Q_i\}_{i=1}^m$. Assuming laminar flow of a Newtonian fluid and neglecting red blood cells, the pressure drop along each edge follows a Poiseuille-like relation, while the nodal balance of flow enforces mass conservation:

$$\begin{cases} P_{s(i)} - P_{t(i)} = R_i^E Q_i, & i = 1, \dots, m, \\ \sum_{i \in E(v_j)} C_{ij} Q_i = 0, & j = 1, \dots, n, \end{cases} \quad (10)$$

The resistance of each vessel is given by $R_i^E = \frac{128 \mu L_i}{\pi D_i^4}$, with $\mu = 3$ cP denoting the dynamic viscosity of plasma at physiological temperature [27]. Equations (10) describe a linear system reported in (3) that defines the nodal pressures and edge flows on the graph \mathcal{G} . The velocities in the vessels can then be computed using Equation (9). To manage the boundary conditions, we first label the vertices V as interior, inlet and outlet nodes. We set $V_{inlet} := \{v_j\}_{j \in J_{inlet}}$ as J_{inlet} is a multi-index identifying the input nodes. We proceed in a similar way for the set of outlet nodes V_{outlet} . As boundary values, we prescribe Dirichlet conditions, assigning the pressures at the inlet nodes and the outlet nodes named $\mathbf{P}^{in} := \{P_j^{in}\}_{j \in J_{inlet}}$ and $\mathbf{P}^{out} := \{P_j^{out}\}_{j \in J_{outlet}}$, respectively.

We can now assemble a nodal-valued boundary values vector $\bar{\mathbf{P}}$, identifying the weights on the nodes $\mathbf{w}_j^V = \bar{P}_j$, and a diagonal matrices $B^{in}, B^{out} \in \mathbb{Z}^{n \times n}$, encoding the inlet and outlet Dirichlet boundary nodes, respectively, such that:

$$\bar{P}_j = \begin{cases} P_j^{in}, & \forall j \in J_{inlet} \\ P_j^{out}, & \forall j \in J_{outlet} \\ 0, & \text{otherwise} \end{cases} \quad B_{jj}^{in} := \begin{cases} 1, & \forall j \in J_{inlet} \\ 0 & \text{otherwise} \end{cases} \quad B_{jj}^{out} := \begin{cases} 1, & \forall j \in J_{outlet} \\ 0 & \text{otherwise} \end{cases}$$

This entails that we can derive the following algebraic formulation for (10):

$$L^W \mathbf{U}^W = \mathbf{F}^W, \quad L^W := \begin{bmatrix} C & -R^E \\ B^{in} + B^{out} & (I - B^{in} - B^{out})C^T \end{bmatrix}, \quad \mathbf{F}^W := \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{P}} \end{bmatrix}, \quad (11)$$

where the system in (3) is assembled with constant viscosity, leading to a fixed resistance matrix R^E , and is solved using a direct sparse linear solver. Within the linear rheology framework, the solution operator is denoted by \mathcal{S}_L , such that $\mathcal{S}_L(G, W) = \mathbf{U}^W$.

3.1.2 Nonlinear models

Microvascular blood flow exhibits nonlinear rheological behavior due to two main factors: the dependence of blood viscosity on hematocrit and vessel diameter (known as the Fåhræus–Lindqvist effect) and the heterogeneous distribution of RBCs in bifurcations, commonly referred to as plasma skimming (also called the Zweifach-Fung effect) [9, 10]. To capture these effects, we adopt a biphasic nonlinear flow model in which blood is treated as a suspension of plasma and RBCs. Recalling that the hematocrit is the volumetric concentration of RBCs in the blood, we augment the unknowns at the edges including the hematocrit levels. As a result, in this model the unknowns are the pressures on the nodes $\mathbf{U}^V := \mathbf{P} = \{P_j\}_{j=1}^n$, the flow rates, and the hematocrit on the edges, i.e. $\mathbf{Q} = \{Q_i\}_{i=1}^m$, $\mathbf{H} = \{H_i\}_{i=1}^m$, respectively, so that $\mathbf{U}^E := \{[Q_i, H_i]\}_{i=1}^m$. In the nonlinear rheology regime, the governing equations define a solution operator $\mathcal{S}_N(G, W) = \mathbf{U}^W$, which maps the graph and the physical parameters to pressure, flow rate and hematocrit.

Similarly to the linear setting, we introduce the diagonal selector matrices $B^{in}, B^{out} \in \mathbb{Z}^{n \times n}$ to handle boundary nodes. The corresponding boundary vectors \mathbf{P} and \mathbf{H} collect the inlet and outlet pressure and inlet hematocrit values, respectively.

The effective viscosity $\mu(D, H)$ varies with the local discharge hematocrit H and the diameter of the vessel D , incorporating the Fåhræus–Lindqvist effect. Moreover, the hematocrit field evolves under the influence of a kinematic plasma skimming model, ensuring consistent RBC mass conservation across bifurcations [9, 10, 13, 16]. The coupled nonlinear system is solved through a fixed-point iterative scheme alternating between two subproblems: (i) a pressure-flow problem with hematocrit-dependent viscosity and (ii) a convection-like problem for RBC transport governed by pseudo-fluxes. This algorithm is illustrated in the following.

Blood flow with hematocrit-dependent viscosity Given an estimate of the hematocrit values at each edge of the graph $\mathbf{H} = \{H_i\}_{i=1}^m$ and the weights on the vertices $\mathbf{w}_j^V = \{\bar{P}_j, \bar{H}_j\}$ and on the edges $\mathbf{w}_i^E = \{D_i, L_i\}$, we assemble the problem (3) with

$$R_i^E = r(\mathbf{w}_i, L_i) = \frac{128 \mu_i(D_i, H_i) L_i}{\pi D_i^4},$$

where resistances are calculated according to the empirical law of Pries et al. [29],

$$\mu_i(D_i, H_i) = \mu_{\text{plasma}} \left[1 + (\mu_{0.45, i} - 1) \frac{(1 - H_i)^{\gamma_i} - 1}{(1 - 0.45)^{\gamma_i} - 1} \right], \quad (12)$$

where $\mu_{\text{plasma}} = 1$ cP is the plasma viscosity, L_i is the length of vessel i , and the parameters $\mu_{0.45, i}$ and γ_i depend on D_i as:

$$\begin{aligned} \mu_{0.45, i} &= 220e^{-1.3D_i} + 3.2 - 2.44e^{-0.06D_i^{0.645}}, \\ \gamma_i &= (0.8 + e^{-0.075D_i}) \left(-1 + \frac{1}{1 + 10^{-11}D_i^{12}} \right) + \frac{1}{1 + 10^{-11}D_i^{12}}. \end{aligned}$$

This problem is solved using the same matrix formulation as in the linear case, namely (3), with hematocrit-dependent resistances.

Uneven hematocrit and convection formulation In the second stage of this scheme, we update the hematocrit distribution throughout the network while preserving the mass of RBC. The conservation of RBCs in a bifurcation involving one parent vessel (0) and two daughter vessels

(1) and (2) can be expressed as follows (these are local indices, not to be confused with the global numbering of the graph \mathcal{G}):

$$Q_{(0)}H_{(0)} = Q_{(1)}H_{(1)} + Q_{(2)}H_{(2)}. \quad (13)$$

However, due to the strong nonlinear coupling between Q and H —especially through the hematocrit-dependent viscosity law introduced in Eq. (12)—direct application of Eq. (13) across large bifurcating microvascular networks can lead to numerical instability or nonphysical solutions. To mitigate this, we introduce a virtual node potential H^* at each bifurcation node [9]. This reformulation allows plasma skimming to be posed as a linear convection problem over \mathcal{G} , which ensures the uniqueness and stability of the solution [39].

Let us formulate first the model locally, i.e. at the level of one bifurcation, and then extend it to the graph level. The discharge hematocrit $H_{(i)}$ in each downstream daughter branch $i = 1, 2$ is then determined by:

$$H_{(i)} = \lambda_{(i)}H^*, \quad (14)$$

where $\lambda_{(i)}$ is a kinematic plasma skimming coefficient. This coefficient captures the tendency of RBCs to preferentially enter one daughter branch over another, and is defined as a function of the ratio of daughter-to-parent diameters $\lambda_{(i)} = (D_{(i)}/D_{(0)})^{2/M_d}$, where $D_{(i)}$ and $D_{(0)}$ are the diameters of the daughter and parent vessels, respectively, and M_d is the drift parameter. In this work, we set $M_d = 5.25$ based on physiological evidence.

Substituting Eq. (14) into Eq. (13) yields the following,

$$Q_{(0)}H_{(0)} = Q_{(1)}\lambda_{(1)}H^* + Q_{(2)}\lambda_{(2)}H^* = Q_{(1)}^*H^* + Q_{(2)}^*H^*,$$

where the geometry-adjusted pseudo-fluxes $Q_{(i)}^*$ are defined as $Q_{(i)}^* = Q_{(i)}\lambda_{(i)}$. Solving for H^* gives:

$$H^* = \frac{Q_{(0)}^*}{Q_{(1)}^* + Q_{(2)}^*} \bar{H}^*, \quad (15)$$

where \bar{H}^* is the hematocrit potential of the upstream node from which the parent vessel originates.

Equation (15) enables a topologically ordered computation of node hematocrit potentials across the network. At inflow nodes, we prescribe $H^* = H^{in}$ and set $\lambda_i = 1$ for all incident edges. Once all H^* values are computed, the discharge hematocrits H_i along vessels are recovered via Eq. (14).

Let us now discuss how to formulate the previous algorithm at the global level, namely for the entire graph \mathcal{G} . The discrete convection equation for the vector of hematocrit node potentials $H^* \in \mathbb{R}^n$ can be expressed in matrix form:

$$N(\mathbf{Q}^*)\mathbf{H}^* = N^{in}(\mathbf{Q}^*)\bar{\mathbf{H}}. \quad (16)$$

The matrices $N(\mathbf{Q}^*), N^{in}(\mathbf{Q}^*) \in \mathbb{R}^{n \times n}$ are defined as:

$$\begin{aligned} N(\mathbf{Q}^*) &= C^\top \cdot \text{diag}(\mathbf{Q}^*) \llbracket -C \rrbracket - \text{diag}(B^{out} C^\top \mathbf{Q}^*), \\ N^{in}(\mathbf{Q}^*) &= -\text{diag}(B^{in} C^\top \mathbf{Q}^*), \end{aligned}$$

where C is the incidence matrix of the graph and $\llbracket X \rrbracket_{ij} = \max(X_{ij}, 0)$ denotes the positive-part operator.

In vascular confluences (e.g., venular unions), we set $\lambda = 1$ by definition. A more detailed explanation of this discretization can be found in [16] and [39], where it is derived from a finite volume formulation of convection in graphs. Once node potentials are determined, the hematocrit values along the vessel segments are obtained by solving the following system,

$$\mathbf{H} = \Lambda \llbracket -C \rrbracket \mathbf{H}^*, \quad (17)$$

where $\mathbf{H} \in \mathbb{R}^m$ is the vector of hematocrits defined on edges and $\Lambda \in \mathbb{R}^{m \times m}$ is a diagonal matrix with the λ_i coefficients for each edge.

The full nonlinear model is then solved via fixed-point iterations, alternating between the following steps:

- (i) computing the flow field \mathbf{Q} based on hematocrit-dependent viscosity;
- (ii) updating node potentials \mathbf{H}^* via Eq. (16) and recovering edge values \mathbf{H} from Eq. (17).

The iteration is initialized with a uniform hematocrit field (e.g., $H_i = 0.4$) and repeated until both flow and hematocrit fields converge. Convergence is typically achieved within 5–7 iterations, even for large networks. A key computational advantage of this approach is that the convection matrix $N(\mathbf{Q}^*)$ is lower triangular when the nodes are ordered topologically with respect to the flow direction. This guarantees that Eq. (16) has a unique solution, provided that no edge flows vanish. The graph-based convection framework significantly improves numerical stability compared to classical models based on recursive split rules [29, 30].

The linear and nonlinear models described above provide a rigorous computational framework to simulate microvascular blood flow across a wide range of regimes. However, their direct application to large ensembles of vascular graphs is computationally demanding, especially in the nonlinear setting where iterative convergence is required.

4 Mathematical synthesis of vascular networks

The construction of surrogate models for microcirculatory flow requires access to physiologically meaningful vascular graphs that can be efficiently handled in a digital setting. In this Section, we focus on the generation of such networks, presenting two complementary approaches that differ in terms of anatomical fidelity, computational cost, and scalability. Section 4.1 introduces a fast geometry-based procedure that produces large collections of simplified but topologically rich capillary networks, suitable for training data-driven surrogates. Section 4.2, instead, describes an optimization-based strategy for image-based Cerebral Network Synthesis, which produces high-fidelity cerebrovascular architectures that serve as rigorous benchmarks for validation. These two approaches synergistically interact with the generation of GNNs-based surrogate models of blood flow, as we train the GNNs on light vascular graphs obtained through the approach of Section 4.1, while we test them on the anatomically accurate configurations of Section 4.2.

4.1 Fast generation of general-purpose synthetic capillary networks

The core of this algorithm uses Voronoi 3D tessellations and the Dijkstra algorithm to obtain a graph $G^s \in \mathbb{Z}^{m \times n}$ that represents a viable vascular network. The procedure for generating the synthetic networks can be summarized in the following steps: (i) generation of the vascular network in a cubic domain; (ii) removal of all trifurcations; (iii) assignment of diameters; (iv) rescaling the cube dimension to have a coherent Surface-to-Volume (S/V) ratio.

Vascular network generation. Synthetic vascular networks are generated within a unit cubic domain, subdivided into three horizontal layers to introduce heterogeneity throughout depth. In each layer, a point cloud is sampled according to a predefined density function, controlling the spatial variation in vascular density. The aggregate point set is used as input for a three-dimensional Voronoi tessellation, whose edges define the centerlines of the candidate vessel. Unfeasible branches, e.g., isolated or highly tortuous segments, are subsequently removed through filtering procedures. Inlet and outlet planes are then identified geometrically; points in proximity to these planes are projected onto them and connected via straight segments, establishing well-defined boundary inflow and outflow conditions. The resulting vascular graph exhibits structural complexity and spatial variability, while preserving computational feasibility. The number of outlets is about twice as many as the number of inlets.

Removal of trifurcations. Capillary networks can commonly be considered isotropic anastomosing networks, whose vertices typically have up to three connections [7, 5]. To enforce a tree-like topology and eliminate trifurcations, a shortest-path strategy is employed via Dijkstra’s algorithm. Starting from selected inlet nodes, the algorithm traces minimal paths to designated

outlet nodes, assembling a spanning structure that avoids cycles and guarantees flow feasibility. Detected trifurcations are resolved by removing the shortest branch, inserting an auxiliary node on the adjacent vessel, and reconnecting the remaining segments to preserve graph connectivity.

Assignment of diameters. Vessel diameters are assigned following a multi-stage algorithm. Initially, diameters are drawn randomly, following a uniform distribution, to enable a preliminary solution of the linear flow model, which provides reliable orientation for topological sorting. The maximum diameter is then assigned to the input segments and propagated downstream with progressive attenuation at the bifurcation points. For Y-shaped bifurcations, a modified version of Murray’s law is adopted, namely: $r^4 = \sqrt[4]{r_1^4 + r_2^4}$ as proposed in [1], to better represent microvascular scaling laws and mitigate overproduction of vessels of minimal diameter. Throughout, a minimum and maximum diameter threshold is enforced. If r_1 and r_2 were known, r is found applying the formula. On the other hand, if only r is known, r_1 is randomly generated and then r_2 is calculated. In each computation, a check that the diameter is between a minimum and a maximum value is performed.

Rescaling the cube dimension. As a final step, the domain size is rescaled to ensure a physiologically realistic surface-to-volume (S/V) ratio. Specifically, given the total vascular surface $S = \sum_{i \in E} \pi D_i L_i$ the cube volume V is adjusted such that $S/V \approx 7000, \text{m}^{-1}$, in alignment with physiological observations. Figure 3 illustrates a representative output of the synthetic network generator.

4.2 Anatomically accurate cerebral vascular networks

To evaluate the generalization capability of the proposed GNN surrogate, we test its performance on anatomically accurate vascular networks generated through a constrained optimization framework, named the image-based Cerebral Network Synthesis (iCNS) [20, 12], and encoded by $G^a \in \mathbb{Z}^{m \times n}$. These networks are constructed using an image-based extension of the Constrained Constructive Optimization (CCO) methodology, originally proposed for coronary arterial tree synthesis [36, 17].

The CCO paradigm incrementally grows a vascular tree by connecting newly sampled terminal nodes to existing segments of the network, solving at each iteration a constrained minimization problem. The objective is to minimize the total volume of the tree as a function of the position of the bifurcation $x \mapsto V(x)$, while enforcing global hemodynamic consistency in terms of continuity of pressure and flow conservation between bifurcations. Formally, the problem can be cast as a nonlinear optimization problem involving:

- (i) *topological decisions*: the choice of the parent segment to which a new terminal node is attached;
- (ii) *geometric optimization*: the placement of the bifurcation point determining local segment lengths and diameters.

Such a problem is intrinsically NP-hard due to the coupling between discrete (connectivity) and continuous (geometry) variables. We refer the reader to Figure 3 for a visualization of these optimization steps. To ensure anatomical fidelity, the iCNS method is constrained within volumetric regions derived from imaging data, such as micro-CT or mouse brain atlases. These provide anatomical priors that accommodate anatomical realism. Through recursive optimization, the algorithm produces large-scale vascular networks that incorporate arterial, capillary, and venous compartments. A final stage of microvascular closure connects the terminal arterial and venous nodes with tortuous capillary segments, resulting in a fully physiologically connected and stable structure. The resulting graphs reproduce realistic morphometrics, including vessel density, length distribution, diameter hierarchy, and topological depth, and are suitable for the validation of microcirculatory flow models (Figure 3).

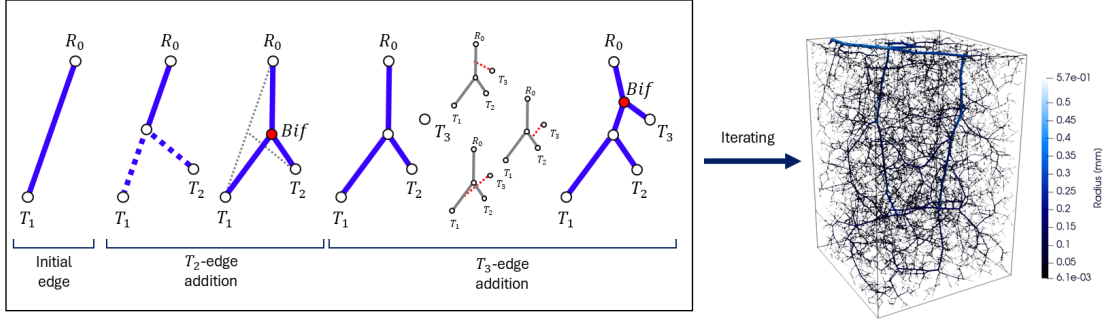


Figure 3: Example of an anatomically guided synthetic arterial tree generated via iCNS, showing the bifurcation optimization process in the CCO framework for each segment addition, adapted from [20, 12].

5 Surrogate blood flow models based on GNNs

We introduce here four GNN configurations designed to approximate the parameter-to-solution map $\mathcal{S}(G, \mathcal{W}) = \mathbf{U}^W = [\mathbf{U}^V, \mathbf{U}^E]^\top$, for the blood flow models introduced in Section 3. Each configuration differs in the complexity of its physical modeling and the composition of its loss function, but all share the same set of input features defined on the nodes and edges of the graph. Table 1 summarizes the features of the four GNN models that are described in what follows.

The GNN models are trained on ensembles of N synthetic vascular graphs $G_k^s = (V_k, E_k)$, $k = 1, \dots, N$, generated using the geometry-based procedure described in Section 4.1. Since these vascular graphs can be created at a low cost and the blood flow models can be efficiently resolved on these setups, this option simplifies the implementation of the supervised learning problem for training the surrogate GNN models.

The input feature vector $\widetilde{W} = [\widetilde{W}^V, \widetilde{W}^E]^\top$ contains geometric and physical quantities. Specifically, the node features vector \widetilde{w}_j^V includes boundary pressures ($P_j^{\text{in}}, P_j^{\text{out}}$) at inlet and outlet vertices together with their extrema ($\max(P^{\text{in}}), \min(P^{\text{out}})$). The edge features \widetilde{w}_i^E include the vessel length L_i and diameter D_i . For the nonlinear rheology model, the hematocrit boundary value H^{in} is also included among the edge inputs. As a result, we have

$$\widetilde{w}_j^V = \{P_j^{\text{in}}, P_j^{\text{out}}, \max(\mathbf{P}^{\text{in}}), \min(\mathbf{P}^{\text{out}})\}, \quad \widetilde{w}_i^E = \{D_i, L_i, H^{\text{in}}\}.$$

The output feature vector $\mathbf{U}^W = [\mathbf{U}^V, \mathbf{U}^E]$ consists of the predicted nodal and edge quantities, which, depending on the model, include pressure P_j , transformed velocity $v_i = \mathcal{T}_v(Q_i)$, and hematocrit H_i . More precisely, from the calculated flow rates, we evaluate the corresponding blood velocities along each segment of the vessel. To mitigate the skewness of the velocity distribution and improve the numerical stability of the learning process, we introduce a *logarithmic velocity transformation* $\mathcal{T}_v : \mathbb{R} \rightarrow \mathbb{R}$, defined as:

$$v := \mathcal{T}_v(Q) = \frac{\text{sign}(Q)}{k_v} \log\left(1 + \frac{4|Q|}{\pi D^2}\right), \quad \mathcal{T}_v^{-1}(v) = \text{sign}(v) \frac{\pi D^2}{4} \left(e^{k_v|v|} - 1\right),$$

where k_v is an empirically chosen scaling constant. The inverse mapping $\mathcal{T}_v^{-1} : \mathbb{R} \rightarrow \mathbb{R}$ is used to recover the physical flow rate values from the transformed representation. Comprehensively, the output features of the nonlinear rheology model are the following:

$$\mathbf{u}_j^V = \{P_j\}, \quad \mathbf{u}_i^E = \{v_i, H_i\},$$

while in the linear models the hematocrit term H_i is not included among the outputs.

Input Features			
Symbol	Description	Type	Models
L_i	Vessel length	Edge feature	1–4
D_i	Vessel diameter	Edge feature	1–4
P_j^{in}, P_j^{out}	Inlet/outlet pressures	Node feature	1–4
$\max(\mathbf{P}^{in}), \min(\mathbf{P}^{out})$	Maximum/Minimum boundary pressure	Node feature	1–4
H^{in}	Hematocrit boundary value	Edge feature	4
Output Features			
P_j	Pressure	Node feature	1–4
v_i	Transformed velocity	Edge feature	2–4
H_i	Hematocrit	Edge feature	4

Table 1: Summary of input and output features for the GNN models.

5.1 GNN model variants

The configuration of the four GNN-based surrogate models is described in the following, with increasing complexity from the simplest to the most complex and physically accurate model.

Model 1: Vanilla GNN (Linear Rheology, Pressure Only) This model approximates only the nodal pressure field $\widehat{\mathbf{U}}^V(\theta^*) = \widehat{\mathbf{P}}_k(\theta^*) = \widehat{\mathcal{S}}_1(G_k^s, \widetilde{W}_k; \theta^*) \approx \mathcal{S}_L(G_k^s, W_k)$, with $\widehat{\mathbf{U}}^V \in \mathbb{R}^{n_k}$, $\widehat{\mathbf{U}}^E = \emptyset$, on each graph G_k^s , using features:

$$\widetilde{\mathbf{w}}_j^V = \{P_j^{in}, P_j^{out}, \max(\mathbf{P}^{in}), \min(\mathbf{P}^{out})\}, j = 1, \dots, n_k, \quad \widetilde{\mathbf{w}}_i^E = \{D_i, L_i\}, i = 1, \dots, m_k.$$

Given labeled data $\mathbf{P}_k = \mathcal{S}_1(G_k^s, W_k) \in \mathbb{R}^{n_k}$ we use a mean-square supervised loss:

$$\mathcal{L}(\widehat{\mathcal{S}}_1, \mathbf{U}^W; \theta) = \mathcal{L}_{\text{pressure}}(\theta) = \frac{1}{N} \sum_{k=1}^N \frac{1}{n_k} \left\| \mathbf{P}_k - \widehat{\mathbf{P}}_k(\theta) \right\|^2.$$

The optimal parameters θ that describe the FFNN in every module of the GNN architecture are determined by solving the following minimization problem (the total number of parameters for this configuration, as well as for the following cases, is reported in Table 2):

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} \mathcal{L}(\widehat{\mathcal{S}}_1, \mathbf{U}^W; \theta).$$

Model 2: Data-Driven GNN (Linear Rheology, Pressure + Velocity) This configuration extends the previous one by jointly predicting nodal pressures and edge velocities

$$[\widehat{\mathbf{U}}^V(\theta), \widehat{\mathbf{U}}^E(\theta)]^\top = [\widehat{\mathbf{P}}_k(\theta), \widehat{\mathbf{v}}_k(\theta)]^\top = \widehat{\mathcal{S}}_2(G_k^s, \widetilde{W}_k; \theta) \approx \mathcal{S}_L(G_k^s, W_k),$$

using the same input features of Model 1. Given labeled data on nodes and edges $[\mathbf{P}_k, \mathbf{v}_k]^\top = \mathcal{S}(G_k, W_k) \in \mathbb{R}^{n_k+m_k}$, we retrieve the minimizer θ^* of the loss that combines the two supervised terms:

$$\mathcal{L}(\widehat{\mathcal{S}}_2, \mathbf{U}^W; \theta) = \gamma_D \mathcal{L}_{\text{pressure}}(\theta) + (1 - \gamma_D) \mathcal{L}_{\text{velocity}}(\theta), \quad \gamma_D \in [0, 1],$$

$$\mathcal{L}_{\text{pressure}}(\theta) = \frac{1}{N} \sum_{k=1}^N \frac{1}{n_k} \left\| \mathbf{P}_k - \widehat{\mathbf{P}}_k(\theta) \right\|^2, \quad \mathcal{L}_{\text{velocity}}(\theta) = \frac{1}{N} \sum_{k=1}^N \frac{1}{m_k} \left\| \mathbf{v}_k - \widehat{\mathbf{v}}_k(\theta) \right\|^2.$$

Model 3: Physics-Informed GNN (Linear Rheology) In this model, we incorporate physics-based regularization to enforce the constitutive and conservation relations of the flow problem. The GNN predicts both nodal pressures and edge velocities in all nodes and edges of the vascular graph,

$$[\widehat{\mathbf{U}}^V(\theta), \widehat{\mathbf{U}}^E(\theta)]^\top = [\widehat{\mathbf{P}}_k(\theta), \widehat{\mathbf{v}}_k(\theta)]^\top = \widehat{\mathcal{S}}_3(G_k^s, \widetilde{W}_k; \theta) \approx \mathcal{S}_L(G_k^s, W_k),$$

and is trained by minimizing a composite loss that combines data fidelity and physical consistency:

$$\mathcal{L}(\widehat{\mathcal{S}}_3, \mathbf{U}^W; \theta) = \delta \mathcal{L}_{\text{data}}(\theta) + (1 - \delta) \mathcal{L}_{\text{physics}}(\theta), \quad \delta \in [0, 1].$$

The supervised term $\mathcal{L}_{\text{data}}$ is identical to Model 2, while the physics-informed component is defined as

$$\mathcal{L}_{\text{physics}}(\theta) = \gamma_P \mathcal{L}_{\text{constitutive}}(\theta) + (1 - \gamma_P) \mathcal{L}_{\text{mass}}(\theta), \quad \gamma_P \in [0, 1].$$

The two physical residuals read:

$$\begin{aligned} \mathcal{L}_{\text{constitutive}}(\theta) &= \frac{1}{C_P} \frac{1}{N} \sum_{k=1}^N \frac{1}{m_k} \left\| C \widehat{\mathbf{P}}_k(\theta) - R^E \mathcal{T}_v^{-1}(\widehat{\mathbf{v}}_k(\theta)) \right\|^2, \\ \mathcal{L}_{\text{mass}}(\theta) &= \frac{1}{C_M} \frac{1}{N} \sum_{k=1}^N \frac{1}{n_k} \left\| (I - B^{\text{in}} - B^{\text{out}}) C^\top \mathcal{T}_v^{-1}(\widehat{\mathbf{v}}_k(\theta)) \right\|^2, \end{aligned}$$

where $C_P = 35 \text{ mmHg}$ and $C_M = 10^6$ are normalization constants. This formulation enforces the Poiseuille constitutive law at the edges of the vascular network and conservation of mass at the junctions of the network.

Model 4: Physics-Informed GNN (Nonlinear Rheology) This most comprehensive configuration extends the previous model by accounting for nonlinear viscosity and hematocrit-dependent effects. The network outputs nodal pressures, edge velocities, and hematocrit levels,

$$[\widehat{\mathbf{P}}_k(\theta), \widehat{\mathbf{v}}_k(\theta), \widehat{\mathbf{H}}_k(\theta)]^\top = \widehat{\mathcal{S}}_4(G_k^s, \widetilde{W}_k; \theta) \approx \mathcal{S}_N(G_k^s, W_k).$$

The total loss function extends that of Model 3 by including hematocrit-dependent terms in both data-based and physics-based components. The data-based contribution reads:

$$\mathcal{L}_{\text{data}}(\theta) = \frac{\gamma_{D_1} + \gamma_{D_2}}{2} \mathcal{L}_{\text{pressure}}(\theta) + \frac{1 - \gamma_{D_1}}{2} \mathcal{L}_{\text{velocity}}(\theta) + \frac{1 - \gamma_{D_2}}{2} \mathcal{L}_{\text{hematocrit}}(\theta),$$

with $\gamma_{D_1}, \gamma_{D_2} \in [-1, 1]$, $\gamma_{D_1} + \gamma_{D_2} \geq 0$, and where

$$\mathcal{L}_{\text{hematocrit}}(\theta) = \frac{1}{N} \sum_{k=1}^N \frac{1}{m_k} \left\| \mathbf{H}_k - \widehat{\mathbf{H}}_k(\theta) \right\|^2.$$

The physics-informed term includes the nonlinear resistance $R^E(\widehat{\mathbf{H}})$ and the hematocrit-weighted conservation residuals:

$$\mathcal{L}_{\text{physics}} = \frac{\gamma_{P_1} + \gamma_{P_2}}{2} \mathcal{L}_{\text{constitutive}} + \frac{1 - \gamma_{P_1}}{2} \mathcal{L}_{\text{mass}}^{(1)} + \frac{1 - \gamma_{P_2}}{2} \mathcal{L}_{\text{mass}}^{(2)},$$

with $\gamma_{P_1}, \gamma_{P_2} \in [-1, 1]$, $\gamma_{P_1} + \gamma_{P_2} \geq 0$, and where

$$\begin{aligned} \mathcal{L}_{\text{constitutive}}(\theta) &= \frac{1}{C_P} \frac{1}{N} \sum_{k=1}^N \frac{1}{m_k} \left\| C \widehat{\mathbf{P}}_k(\theta) - R^E(\widehat{\mathbf{H}}_k(\theta)) \mathcal{T}_v^{-1}(\widehat{\mathbf{v}}_k(\theta)) \right\|^2, \\ \mathcal{L}_{\text{mass}}^{(1)}(\theta) &= \frac{1}{C_M} \frac{1}{N} \sum_{k=1}^N \frac{1}{n_k} \left\| (I - B^{\text{in}} - B^{\text{out}}) C^\top \mathcal{T}_v^{-1}(\widehat{\mathbf{v}}_k(\theta)) \right\|^2, \end{aligned}$$

Model	Rheology	# Param.	Loss hyperparam.	L^2 pres. error	L^2 vel. error
Model 1	Linear	31,144	–	2.83%	–
Model 2	Linear	73,992	$\delta = 1, \gamma_D = 0.4$	5.77%	37.2%
Model 2	Linear	73,992	$\delta = 1, \gamma_D = 0.5$	2.46%	17.7%
Model 2	Linear	73,992	$\delta = 1, \gamma_D = 0.6$	3.04%	21.5%
Model 3	Linear	73,992	$\delta = 0.5, \gamma_D = 0.85, \gamma_P = 0.5$	2.54%	21.7%
Model 3	Linear	73,992	$\delta = 0.5, \gamma_D = 0.9, \gamma_P = 0.5$	2.20%	18.5%
Model 3	Linear	73,992	$\delta = 0.5, \gamma_D = 0.95, \gamma_P = 0.5$	3.10%	28.9%
Model 4	Nonlinear	74,007	$\delta = 0.5, \gamma_{D_1} = \gamma_{D_2} = 0.65, \gamma_{P_1} = \gamma_{P_2} = 0.65$	3.79%	24.7%
Model 4	Nonlinear	74,007	$\delta = 0.5, \gamma_{D_1} = \gamma_{D_2} = 0.75, \gamma_{P_1} = \gamma_{P_2} = 0.75$	2.26%	14.1%
Model 4	Nonlinear	74,007	$\delta = 0.5, \gamma_{D_1} = \gamma_{D_2} = 0.85, \gamma_{P_1} = \gamma_{P_2} = 0.85$	3.97%	33.5%

Table 2: Architecture size, rheology, loss-weight sensitivity, and test errors for the GNN models. Highlighted rows indicate the selected hyperparameter configurations used in the numerical experiments.

$$\mathcal{L}_{\text{mass}}^{(2)}(\theta) = \frac{1}{C_M} \frac{1}{N} \sum_{k=1}^N \frac{1}{n_k} \left\| (I - B^{\text{in}} - B^{\text{out}}) C^{\top} (\hat{\mathbf{H}}_k(\theta) \odot \mathcal{T}_v^{-1}(\hat{\mathbf{v}}_k(\theta))) \right\|^2.$$

Here, the term $\mathcal{L}_{\text{mass}}^{(2)}$ ensures hematocrit-weighted flow conservation, while the nonlinear dependence $R^E(\hat{\mathbf{H}})$ captures the Fåhræus–Lindqvist viscosity effect along each edge.

5.2 Numerical setup

Training and evaluation of GNN models are carried out using synthetic vascular networks generated by the algorithm detailed in Section 3. We generate a dataset of $N = 1200$ graphs with randomized topology and geometry, partitioned as follows: (i) *Training set*: $N_{\text{train}} = 960$ graphs; (ii) *Validation set*: $N_{\text{val}} = 120$ graphs; (iii) *Test set*: $N_{\text{test}} = 120$ graphs.

Each network comprises approximately 300 nodes and 400 edges, with 25 to 40 inlets located at the top surface of the domain, corresponding to about 50–80 outlets. Boundary pressures are randomly sampled within physiological ranges, and all internal quantities are normalized to $[0, 1]$ to ensure numerical stability.

The GNN model employs the *Gaussian Error Linear Unit* (GELU) as activation function and is trained using a processor composed of 30 message passing steps. Each message-passing layer includes 7 hidden layers, with skip connections introduced every 3 layers to enhance gradient propagation and model stability. Given these choices and considering the number of features, the number of trainable parameters is shown in Table 2. The hyperparameters appearing in the loss function have been heuristically chosen to optimize the training process. A sensitivity analysis of such parameters is also reported in the same Table.

Each GNN model is trained using the Adam optimizer, with an initial learning rate of 10^{-3} , decreased by a factor of 10 after 10 consecutive validation epochs without improvement. Training is stopped early if no improvement is observed for 25 epochs. We use a batch size of 1 (graph-level batching) and train for a maximum of 500 epochs. All experiments were performed on an NVIDIA GeForce RTX 4090 supported by an Intel Core i9-14900HX CPU with 64 GB RAM.

A GitHub repository is available (github.com/piermariovitullo/GNN-demo/), providing a plug-and-play demonstration of the GNN surrogate model and offering a minimal yet functional pipeline to reproduce selected simulations. It includes tools for loading vascular graphs, assembling the hydraulic system, solving the linear flow model, and evaluating GNN-based predictions of pressure and velocity.

6 Numerical tests on GNN surrogate models

This section presents a detailed numerical investigation of the GNN-based models to evaluate their ability to accurately reproduce microvascular blood flow in a variety of four configurations and modeling complexities that are detailed in the following. The focus is on predicting the pressure,

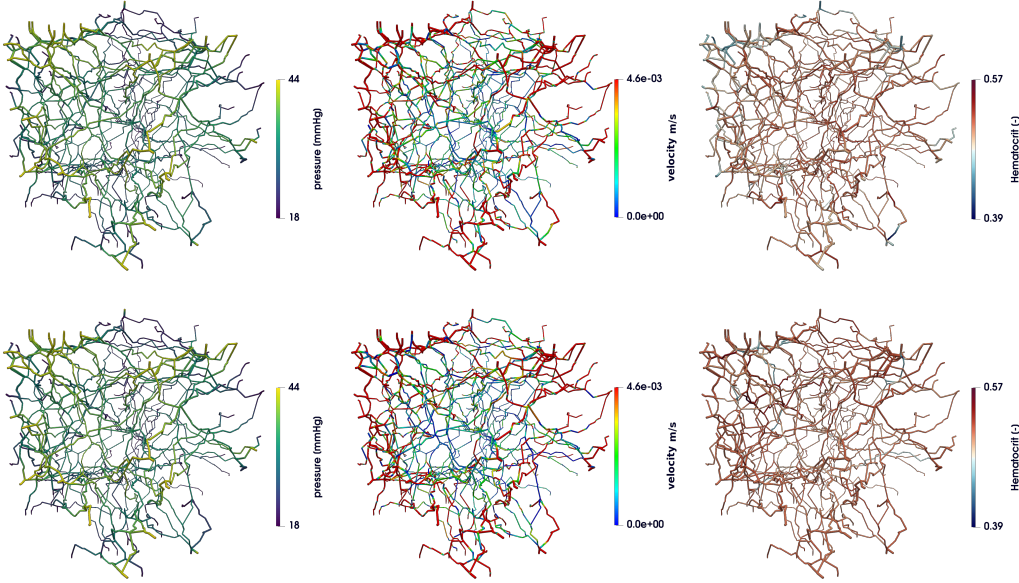


Figure 4: Visual comparison of pressure, velocity and hematocrit solutions (left, center, right panels, respectively) for GNN Model 4 (top), compared with the respective high-fidelity approximations (bottom). In this test case, the L^2 pressure, velocity and hematocrit GNN errors are respectively 2.55%, 16.02% and 4.14%.

velocity and hematocrit distributions under linear and nonlinear rheological laws using the models of Section 5.

6.1 GNN tests on general-purpose capillary networks

We now assess the generalizability of the four GNN models introduced in Section 5 on a set of unseen vascular graphs, focusing on their ability to predict pressure, velocity and hematocrit fields beyond the training distribution. To quantify the accuracy of the model, we compute the mean relative L^p -error, $p = 1, 2$, with respect to the high-fidelity (FOM) solutions for the predicted pressure, velocity and hematocrit fields. The L^2 -error is consistent with the norm used in the loss function, and the L^1 -error is a weaker norm for which lower error levels are expected. For a generic output $\mathbf{U}^W = [\mathbf{U}^V, \mathbf{U}^E]^\top$ and for each testing graph G_k^s , we define

$$\varepsilon_k^* = \frac{\|\mathbf{U}_k^* - \widehat{\mathbf{U}}_k^*\|_{p,*}}{\|\mathbf{U}_k^*\|_{p,*}} \times 100, \quad \text{and} \quad \bar{\varepsilon}^* = \frac{1}{N_{\text{test}}} \sum_{k=1}^{N_{\text{test}}} \varepsilon_k^*, \quad p = 1, 2, \quad * = V, E,$$

where $\widehat{\mathbf{U}}_k^W$ denotes the GNN prediction and $\bar{\varepsilon}^*$ the average error on all the graphs tested.

The corresponding values are summarized in Table 3 showing that the surrogate models comprehensively provide a good approximation. For further insight, the visual comparison of the high-fidelity and surrogate solutions shown in Figure 4 is not particularly informative. A quantitative analysis of the error distribution in the case of Model 4 is reported in Figure 5, where we show the empirical distributions of pressure, velocity, hematocrit and RBC mass balance errors. The pressure error is strongly concentrated around small values and remains relatively well controlled across the network. By contrast, the velocity error displays a broader distribution, with a heavier tail associated with a smaller subset of vessel segments. The systematically larger velocity error reflects the different mathematical roles of pressure and velocity in the graph hemodynamic model. Pressure is a nodal potential governed by an elliptic, graph-Laplacian-like balance law;

Model	Quantity	Training (L^1/L^2)	Validation (L^1/L^2)	Test (L^1/L^2)
Model 1	Pressure	2.11 / 2.95	2.19 / 2.94	2.20 / 2.83
	Velocity	- / -	- / -	- / -
	Hematocrit	- / -	- / -	- / -
Model 2	Pressure	1.93 / 2.46	1.95 / 2.48	1.94 / 2.46
	Velocity	12.5 / 17.0	12.7 / 17.3	12.6 / 17.7
	Hematocrit	- / -	- / -	- / -
Model 3	Pressure	1.66 / 2.21	1.72 / 2.29	1.65 / 2.20
	Velocity	13.8 / 18.0	14.1 / 18.6	13.9 / 18.5
	Hematocrit	- / -	- / -	- / -
Model 4	Pressure	1.65 / 2.26	1.68 / 2.31	1.64 / 2.26
	Velocity	13.5 / 13.3	13.7 / 13.9	13.7 / 14.1
	Hematocrit	2.59 / 3.75	2.60 / 3.76	2.59 / 3.75

Table 3: L^1 and L^2 relative errors (%) w.r.t. the FOM solution for pressure, velocity, and hematocrit on training, validation, and test sets.

therefore, local prediction errors tend to be smoothed by the global network coupling and by the mass-conservation constraint. By contrast, velocity is an edge quantity derived from local pressure drops, resistances, and vessel diameters. Since resistance scales as D^{-4} and velocity scales as Q/D^2 , small errors in pressure differences, flow rates, or local geometric features may be amplified at the edge level, especially in small-diameter vessels. This explains why pressure errors remain concentrated near small values, whereas velocity errors show broader distributions and heavier tails. The hematocrit error is concentrated at small values, while the bifurcation residual is sharply peaked near zero, with only a mild tail. This indicates that the surrogate captures hematocrit partitioning accurately in most vessels and preserves RBC flux conservation at bifurcations. We have performed additional investigations on the dependence on vessel diameter and topological distance from the nearest inlet (not shown). Error stratification data show that pressure error slightly increases in the small-diameter regime, where the solution is more sensitive to local geometric variations and to the stronger diameter dependence of the effective resistance. In addition, the pressure error tends to increase with the topological distance from the nearest inlet. Although the message-passing architecture exchanges information symmetrically across the graph, the target hemodynamic map remains inherently directional, being primarily organized by flow orientation from inlet to outlet. Note that this effect only appears in the general-purpose synthetic capillary networks, characterized by a uniform architecture that does not allow us to discriminate between arterioles, capillaries and venules.

Figure 6 reports the L^2 mean relative pressure and velocity errors across the testing set for the four models (1–4). The evaluation is carried out in an *out-of-sample* regime: the models were trained on vascular networks whose complexity, measured by the number of inlet nodes located on the upper surface of the domain, ranged between 25 and 40 (as indicated by the gray band in Figure 6). The tests were then carried out on 17 additional vascular graphs, characterized by a number of inlets varying from 20 up to 100, with each configuration containing twice as many outlet nodes as inlets. The results highlight three main aspects. All GNN models exhibit stable accuracy far beyond the training window, effectively generalizing to vascular networks whose geometric complexity spans from 20 to 100 inlet nodes. This holds consistently for pressure and velocity fields, confirming the robustness of the learned parameter-to-solution map $\mathcal{S} : (G, W) \mapsto \mathbf{U}^W$ to large variations in graph topology. However, we note that, despite the increasing geometric complexity, the global hemodynamic regime remains coherent across all test graphs. This is confirmed by the near-constant depth of the networks—measured as the average number of nodes from inlets to outlets—whose weak dependence on the inlet count suggests that the flow regime is physically comparable across test configurations (not shown). Finally, the physics-informed models (Model 3 and 4) maintain low residuals for constitutive and mass-balance laws throughout the entire testing range. More precisely, these metrics are defined as $\mathcal{L}_{\text{constitutive}}$ and $\mathcal{L}_{\text{mass}}$ for

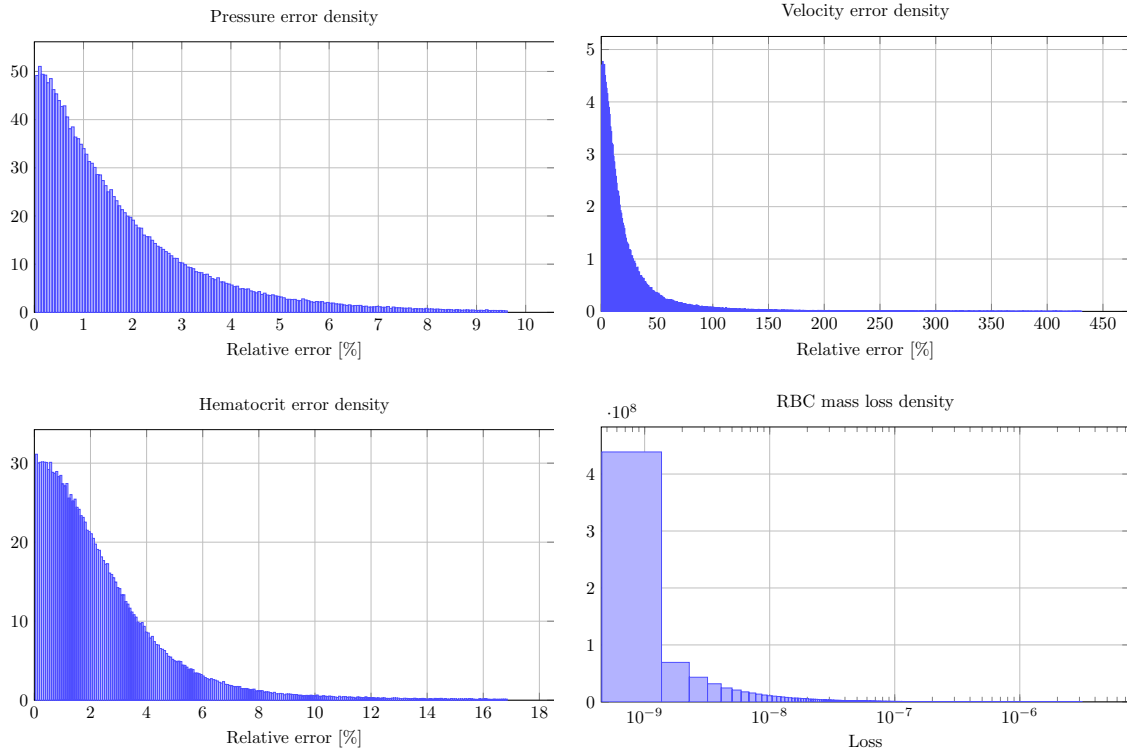


Figure 5: Distributional error analysis (truncated at the 99th percentile for better visualization) on the general-purpose vascular networks. Top row: empirical densities of the relative pressure and velocity errors. Bottom row: empirical density of the relative hematocrit error and of the RBC mass-balance residual at bifurcations.

Model 3 and $\mathcal{L}_{\text{constitutive}}$, $\mathcal{L}_{\text{mass}}^{(1)}$, and $\mathcal{L}_{\text{mass}}^{(2)}$ for Model 4, in Section 5. These residuals remain small in all configurations, as shown in Figure 6 (bottom panel). Overall, these results demonstrate that the proposed GNN architectures not only reproduce local flow quantities with high accuracy but also preserve the global physical structure of the problem, providing reliable predictions across a wide spectrum of vascular geometries.

To evaluate the generalization capability of the proposed GNN surrogate model, we test its performance on a set of anatomically realistic vascular networks generated using the methodology introduced in Section 4.2. In contrast to the synthetic graphs used for training, these cerebral networks are significantly more complex, both topologically and morphologically. In particular, physiological plausibility, hierarchical branching, and image-informed spatial constraints are now enforced.

We consider four synthetic networks representing the cortical circulation of the mouse brain, generated using algorithms developed in [12, 13, 20]. We name these graphs with the following labels: S1.101, S2.102, S3.101, and S4.102, and we refer to Table 4 for a detailed description of their characteristics. These networks faithfully reproduce the statistical morphometric properties observed in two-photon laser scanning microscopy (2PLSM) datasets of mouse somatosensory cortex [5, 12, 20]. Below, we provide a biological overview of their structure and functional roles.

Each vascular network represents an anatomically realistic reconstruction of the murine cortical microcirculation, encompassing interconnected pial arteries, penetrating arterioles, capillaries, and venules generated through the image-based Cerebral Network Synthesis (iCNS) framework [13]. The pial arteries distributed across the cortical surface serve as primary inflow routes, giving rise to 12–14 penetrating arterioles per mm^2 , consistent with morphometric measurements [5]. These vessels descend nearly orthogonally into the cortex, feeding a dense capillary plexus generated by

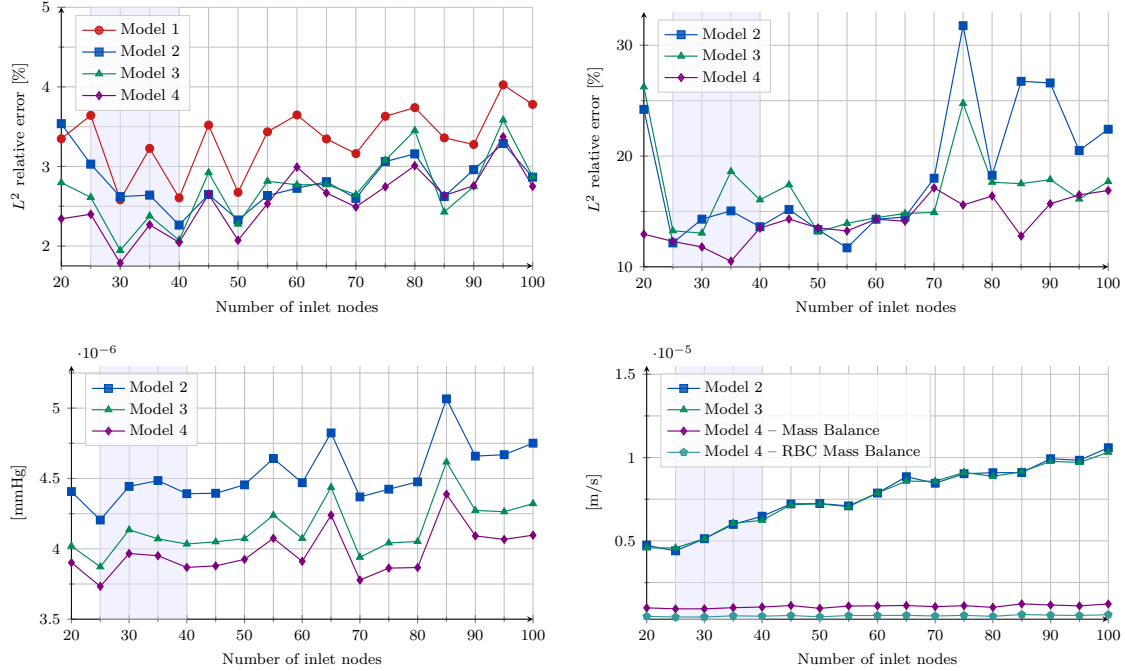


Figure 6: Pressure (top-left) and velocity (top-right) L^2 relative error comparison between the different models. Physics-based residuals of the constitutive law (bottom-left) and the mass balance law (bottom-right). The gray band highlights the interval of number of inlet nodes where the training data are located. These data show the ability of the model to generalize well beyond the vascular complexity of the training dataset.

a microvascular closure algorithm that guarantees topological continuity between the arterioles and venules. The capillary segments display diameter-dependent tortuosity (ranging from 1.1 to 1.8 for $d < 150 \mu\text{m}$), reproducing the geometric variability and flow dispersion observed in vivo. Downstream, post-capillary venules, and ascending veins collect blood and connect to pial veins that drain into the cortical sinuses, with venular densities between 13 and 39 per mm^2 , matching empirical data. Morphometric analysis confirms that segment counts ($\sim 11,500 \text{ mm}^{-3}$), as well as global quantities such as total vascular length, surface area, and volume, fall within the ranges reported for two-photon laser scanning microscopy (2PLSM) datasets of mouse somatosensory cortex [13].

Each network includes one inlet and one outlet node, representing the arterial and venous ends, respectively. To ensure consistency with the training setup, we automatically detect additional inlet and outlet nodes, assigning as inlets all vertices belonging to vessels with diameters greater than $12 \mu\text{m}$ starting from the main arterial node and analogously on the outlet side. We then prescribe the boundary pressures to solve the high-fidelity linear problem (10). These anatomical networks provide physiologically consistent digital analogs of cerebral microcirculation, suitable for evaluating the generalization capabilities of the proposed GNN surrogates on realistic geometries.

6.1.1 Tests on linear blood flow models

Without retraining or fine-tuning the model, we apply GNN Model 3 to predict pressure fields over these domains. The results are summarized in Table 5.

Despite the large complexity gap and the anatomical specificity of the test domains, the surrogate model achieves an L^2 -relative error of less than 10.5% in all cases. Moreover, it produces these predictions in under 30 milliseconds, offering a remarkable speed-up compared to the corresponding full-order solver times, which are on the order of 1 second. These findings demonstrate

Network ID	S1.101	S2.102	S3.101	S4.102
Number of nodes	66,090	88,286	151,626	154,466
Number of segments	65,489	87,545	150,812	153,720
Average diameter [μm]	5.72	5.85	5.67	5.59
Mean segment length [μm]	40.1	39.6	41.0	38.8
Capillary density [segments/ mm^3]	11,300	11,500	11,600	11,400
Total network length [mm]	89.6	113.4	209.8	211.7
Total surface area [mm^2]	5.41	7.12	13.21	13.48
Total volume [mm^3]	0.215	0.287	0.534	0.549

Table 4: Morphometric properties of the four anatomically accurate mouse cortex networks used for the GNN generalization tests. These values are derived from the topological and geometric analysis of the synthesized graphs based on the CCO and microvascular closure algorithms.

Network	1	2	3	4
Pressure error (L^1/L^2)	5.06 / 6.34	8.31 / 10.49	7.64 / 8.80	6.07 / 7.75
Nodes	66090	88286	151626	154466
Solver time	0.43s	0.56s	1.13s	1.32s
GNN time	23.32ms	27.01ms	28.39ms	26.58ms

Table 5: L^1 and L^2 relative pressure error estimates (%) for different anatomical vascular geometries, tested on the linear blood flow model compared with GNN model 3.

that the GNN surrogate, although trained solely on abstract capillary graphs with no anatomical fidelity, is capable of extrapolating to large, biologically realistic domains.

6.1.2 Tests on nonlinear blood flow models

We now assess its ability to generalize to nonlinear microvascular blood flow governed by hematocrit-dependent rheology. The resulting nonlinear system must be addressed using iterative techniques, i.e. the fixed-point iteration strategy described in Section 3.1.2, in which the nonlinear terms are treated using a nested loop approach. Precisely, the hematocrit field is held fixed while solving the linearized pressure and flow problem, and then updated via flow-dependent phase separation laws. Convergence is monitored by tracking changes in hematocrit between successive iterations.

In this test, we deploy GNN Model 4, trained exclusively on synthetic capillary networks with nonlinear rheology, on anatomically realistic cerebral networks. The surrogate is tasked with predicting nodal pressure, edge velocity, and hematocrit distributions without retraining or adjusting. We present the results in Figure 7 and summarize them in Table 6.

The surrogate model tested on nonlinear blood flow models achieves a L^2 -relative error of less than 9.5% in all cases. We observe comparable computational evaluation costs with respect to the linear case, consequently offering a significant speed-up compared to the fixed-point method-based solver, employed for the high-fidelity model.

Figure 7 (bottom row) suggests that the pressure error is not uniformly distributed across the anatomically realistic mouse networks, as the arteriolar and venular segments of the network systematically show errors larger than the average. A more detailed analysis is reported in Figure 8 where vascular labels (Capillaries; Arteries/Arterioles; Veins/Venules) are used in the error distribution plot. The left panel shows that errors are smallest in capillaries and larger in arterial and venous compartments, with the proximal arterial and distal venular sides being the most challenging regions. Consistently, the right panel of Figure 8 shows that the error directly correlates with $\Delta P \cdot Q$, which represents the rate of mechanical energy dissipation along each vessel segment. More precisely, we see that $\Delta P \cdot Q$ sharply increases when it overcomes a threshold ($\log(\Delta P \cdot Q) \approx 12$). This indicates that the surrogate is less accurate in segments undergoing larger energetic loads, typically located near the inflow or outflow side of the network. By comparing these results with those obtained for generic capillary networks, we infer that, in anatomically realistic vascular networks, the error is mainly linked to the vessel hierarchy within the vascular tree. This occurs because the GNN models are attempting to predict microvascular flow under conditions where

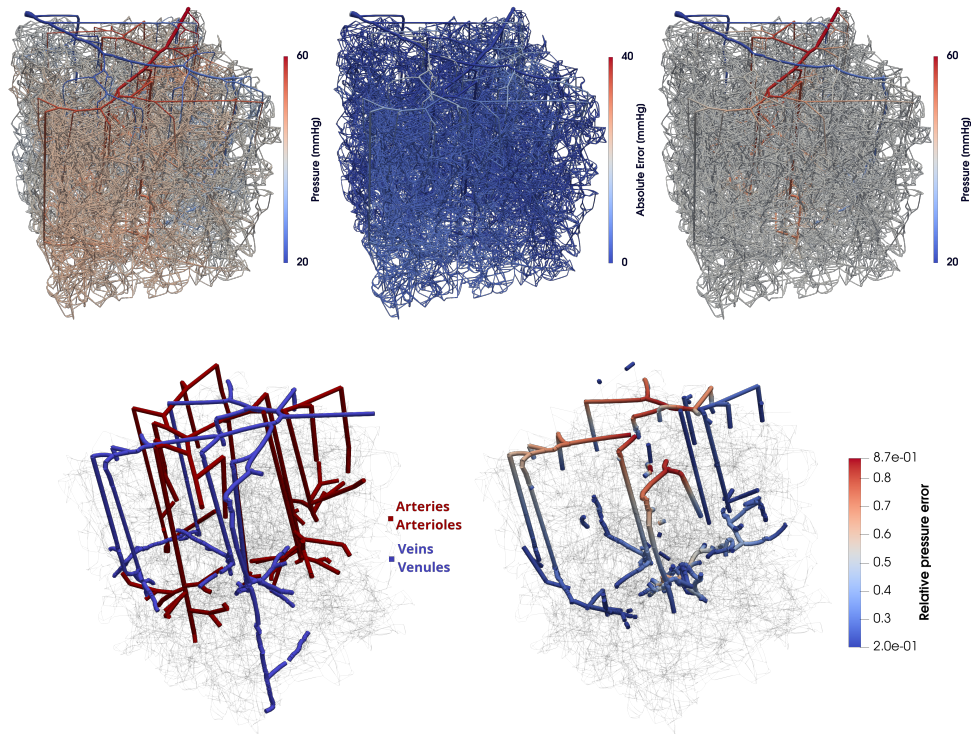


Figure 7: Pressure prediction and spatial distribution of the error for Network S1.101 in the nonlinear blood flow model. Top row: reference pressure field (left), absolute pressure error (center), and GNN-predicted pressure field (right). Bottom row: (left) visualization of the arterial/arteriolar (red) and venous/venular compartments (blue), and spatial distribution of the largest relative pressure errors (right).

this hierarchical information is important, even though they were not trained with it.

6.2 Discussion

The results presented in this section demonstrate the ability of the proposed GNN framework to accurately approximate hemodynamic quantities on complex vascular networks while maintaining strong generalization properties across unseen topologies. Unlike traditional reduced-order or regression-based surrogates, which often rely on fixed connectivity or explicit graph parametrizations, the adopted message-passing formulation enables the model to operate directly on variable graph structures, making it inherently adaptable to different vascular architectures. Incorporating physics-informed loss terms further enhances the robustness of the surrogate model. By enforcing local mass conservation and constitutive relations linking pressure, velocity, and hematocrit, the proposed GNNs achieve not only low data-driven error but also strong physical consistency across a broad range of vascular configurations. This dual data- and physics-driven learning paradigm effectively regularizes the training process, improves the extrapolation capability, and reduces unphysical artifacts in the predictions. More precisely, it delivers accurate predictions for pressure, velocity, and hematocrit in vascular domains with more than 150,000 segments, orders of magnitude more complex than any in the training set, while achieving inference times below 30 milliseconds. In fact, a major advantage of this approach lies in its efficiency. Once trained on a representative but computationally inexpensive set of graphs, the GNN provides near-instantaneous evaluations of the parameter-to-solution map, bypassing the need to repeatedly assemble and solve large nonlinear systems that describe flow and transport in microvascular

Network	1	2	3	4
Pressure error (L^1/L^2)	6.08 / 8.06	6.56 / 9.36	6.09 / 9.14	5.75 / 7.65
Nodes	66090	88286	151626	154466
Solver time	4.29s	7.82s	13.52s	22.53s
GNN time	23.32ms	22.01ms	28.35ms	29.51ms

Table 6: L^1 and L^2 relative pressure error estimates (%) for different complex vascular geometries, tested on the nonlinear blood flow model compared with GNN model 4.

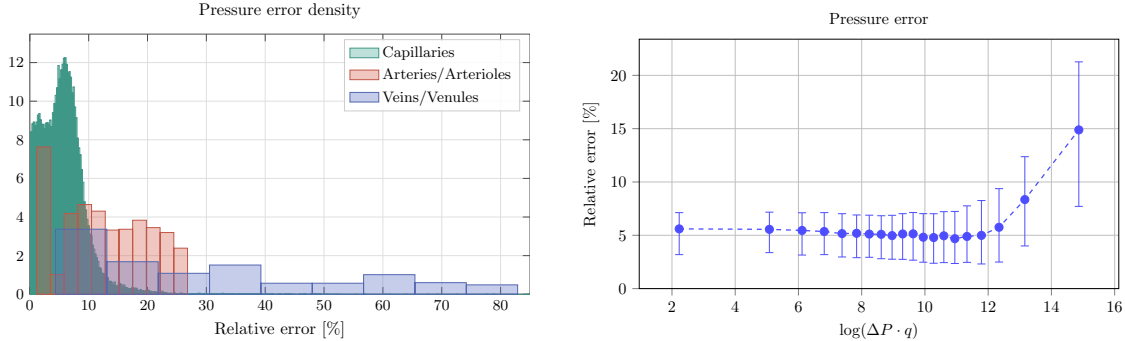


Figure 8: Spatial and structural error analysis for pressure on the anatomically realistic mouse cortical networks. Left Panel: pressure error density stratified by vascular compartment. Right Panel: variation of mean relative pressure error with the logarithm of the local hydraulic power dissipation.

networks. This represents a significant computational gain compared to conventional high-fidelity solvers, whose cost scales unfavorably with network size and stiffness induced by nonlinear rheological effects. Thus, the proposed GNN approach enables the deployment of blood flow models in anatomically accurate geometries such as whole mouse or human brain cortex models [21].

7 Conclusions

This work introduces a physics-informed GNN framework for efficient simulation of microvascular blood flow. The surrogate model is trained on a large dataset of synthetically generated vascular graphs that capture generic capillary features, and is rigorously tested on high-fidelity, anatomically accurate cerebrovascular networks representative of the mouse cortex.

We propose a two-tiered strategy that balances scalability and physiological realism. General-purpose vascular graphs generated via Voronoi tessellations provide a computationally inexpensive means of producing large, diverse training data. Anatomically constrained networks offer a challenging testbed to validate the surrogate’s extrapolation capabilities. These networks include complex topological features, heterogeneous vessel properties, and biological realism that cannot be achieved with heuristic generators. The GNN surrogate, trained solely on the former class of networks, demonstrates strong generalization capabilities when applied to the latter. Our findings emphasize three key strengths of the proposed approach: (i) *Topology-aware modeling*: the graph-based architecture enables flexible learning across diverse vascular structures, supporting application to different organs and pathological conditions; (ii) *Physics-informed learning*: incorporating domain knowledge into the loss function improves prediction accuracy, mass conservation and rheological fidelity, and mitigates overfitting to synthetic patterns; (iii) *Scalable inference*: the trained GNN surrogate enables real-time prediction on large-scale microvascular networks, offering potential for integration into multiscale simulation pipelines and digital twin frameworks. Despite these strengths, current limitations include the restriction to steady-state flow and the assumption of impermeable, non-leaky vessels. Velocity and hematocrit predictions are also more sensitive to

domain shifts, particularly in the presence of strong nonlinearities or unusual branching patterns.

Future work will focus on extending the learning framework introduced here in the context of *capillary* flow (where blood motion is well described by a stationary regime dominated by diffusive mechanisms) to the case of arteriolar flow, where the regime is *pulsatile*. In this case, the governing equations, the physical principles, and the resulting mathematical models differ substantially from those of capillary flow [2, 22]. For example, arterioles are accurately described by one-dimensional hyperbolic systems for pressure, flow, and cross/sectional area. The ongoing extension of the present methodology to the pulsatile setting is the subject of a dedicated study, whose preliminary results developed in [4] are promising. Importantly, it confirms that the transition from capillaries to arterioles introduces substantial conceptual and numerical differences. Primarily, in pulsatile flow, the hierarchical organization of the vascular network becomes a key input feature in the learning architecture. However, despite these differences in the underlying physics, the formulation of the computational learning problem remains conceptually analogous: in both cases, the hemodynamic state is represented on a graph, and a GNN can be designed to approximate the associated parameter-to-solution map. The message-passing mechanism, the encoder–processor–decoder architecture, and the treatment of nodal and edge variables remain essentially unchanged, demonstrating the flexibility and robustness of the GNN framework across different flow regimes.

These observations reinforce the relevance and generality of graph-based surrogate models for hemodynamics, and show that the framework developed in this paper forms a solid foundation for future extensions to more complex vascular regimes.

Declarations

Authorship and Contributorship Statement

All authors have made substantial intellectual contributions to the study conception, execution, and design of the work. All authors have read and approved the final manuscript.

Conflicts of Interest

The authors declare there are no conflicts of interest.

Data & Code Availability Statement

The data that support the findings of this study are openly available at github.com/piermariovitullo/GNN-demo/.

Funding statement

PV and PZ acknowledge support from the MUR PRIN 2022 project 2022WKWZA8 *Immersed methods for multiscale and multiphysics problems* (IMMEDIATE), part of the Next Generation EU program Mission 4, Component 2, CUP D53D23006010006; the European Union’s Euratom research and training programme 2021–2027 under grant agreement No. 101166699 *Radiation safety through biological extended models and digital twins* (TETRIS); PZ acknowledges support from the *Mathematical Challenges in Brain Mechanics* project of the Center for Advanced Study (CAS), The Norwegian Academy of Science and Letters, Oslo, Norway; PB acknowledges support from *Dipartimento di Eccellenza 2023–2027*, Department of Mathematics, Politecnico di Milano. TV and AL gratefully acknowledge financial support through NIH NIA 1R01AG079894. PB, PV, PZ are members of the Gruppo Nazionale per il Calcolo Scientifico (GNCS) of the Istituto Nazionale di Alta Matematica (INdAM).

References

- [1] D. Akita, I. Kunita, M. D. Fricker, S. Kuroda, K. Sato, and T. Nakagaki. Experimental models for Murray’s law. *Journal of Physics D: Applied Physics*, 50(2):024001, 2016.
- [2] Jordi Alastruey, Ashraf W Khir, Koen S Matthys, Patrick Segers, Spencer J Sherwin, Pascal R Verdonck, Kim H Parker, and Joaquim Peiró. Pulse wave propagation in a model human arterial network: assessment of 1-d visco-elastic simulations against in vitro measurements. *Journal of biomechanics*, 44(12):2250–2258, 2011.
- [3] P.W. Battaglia, J.B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V.F. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gülcehre, H.F. Song, A.J. Ballard, J. Gilmer, G.E. Dahl, A. Vaswani, K.R. Allen, C. Nash, V. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M.M. Botvinick, O. Vinyals, Y. Li, and R. Pascanu. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- [4] F. Behrens. Learning pulsatile microvascular blood flow models using graph neural networks. Master’s thesis, Politecnico di Milano, Milan, Italy, 2025.
- [5] P. Blinder, P.S. Tsai, J.P. Kaufhold, P.M. Knutsen, H. Suhl, and D. Kleinfeld. The cortical angiome: an interconnected vascular network with noncolumnar patterns of blood flow. *Nat Neurosci.*, 16(7):889–897, 2013.
- [6] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [7] H.M. Duvernoy, S. Delon, and J.L. Vannson. Cortical blood vessels of the human brain. *Brain Research Bulletin*, 7(5):519–579, 1981.
- [8] C. Alberto Figueroa, Charles A. Taylor, and Alison L. Marsden. *Blood Flow*, pages 1–31. John Wiley & Sons, Ltd, 2017.
- [9] I.G. Gould and A.A. Linninger. Hematocrit distribution and tissue oxygenation in large microcirculatory networks. *Microcirculation*, 22(1):1–18, 2015.
- [10] I.G. Gould, P. Tsai, D. Kleinfeld, and A.A. Linninger. The capillary bed offers the largest hemodynamic resistance to the cortical blood supply. *Journal of Cerebral Blood Flow & Metabolism*, 37(1):52–68, 2017.
- [11] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 2017.
- [12] G. Hartung, S. Badr, S. Mihelic, A. Dunn, X. Cheng, S. Kura, D. A. Boas, D. Kleinfeld, A. Alaraj, and A.A. Linninger. Mathematical synthesis of the cortical circulation for the whole mouse brain—part ii: Microcirculatory closure. *Microcirculation*, 28(5):e12687, 2021.
- [13] G. Hartung, C. Vesel, R. Morley, A. Alaraj, J. Sled, D. Kleinfeld, and A.A. Linninger. Simulations of blood as a suspension predict a depth dependent hematocrit in the circulation throughout the cerebral cortex. *PLoS Computational Biology*, 14(11):e1006549, 2018.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [15] A. Iacovelli, L. Pegolotti, M. Salvador, E. Stoppa, M.D. Santambrogio, and A.L. Marsden. A novel lstm and graph neural networks approach for cardiovascular simulations. In *2024 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*, pages 1–8, 2024.

- [16] M. Jamshidi, T. Ventimiglia, P. Sudres, C. Zhang, F. Lesage, W. Rooney, D. Schwartz, and A.A. Linninger. Impact of stalling events on microcirculatory hemodynamics in the aged brain. *Microcirculation*, 31(3):e12845, 2024.
- [17] R. Karch, F. Neumann, M. Neumann, and W. Schreiner. A three-dimensional model for arterial tree representation, generated by constrained constructive optimization. *Computers in Biology and Medicine*, 29(1):19–38, 1999.
- [18] N.J. Karst, B.D. Storey, and J.B. Geddes. Oscillations and multiple equilibria in microvascular blood flow. *Bulletin of Mathematical Biology*, 77:1377–1400, 2015.
- [19] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, M. Shakir, and P. Battaglia. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- [20] A.A. Linninger, G. Hartung, S. Badr, and R. Morley. Mathematical synthesis of the cortical circulation for the whole mouse brain—part i: Theory and image integration. *Computers in Biology and Medicine*, 110:265–275, 2019.
- [21] A.A. Linninger, T. Ventimiglia, M. Jamshidi, M. Pascal-Suisse, A. Alaraj, F. Lesage, X. Li, D.L. Schwartz, and W.D. Rooney. Vascular synthesis based on hemodynamic efficiency principle recapitulates measured cerebral circulation properties in the human brain. *Journal of Cerebral Blood Flow & Metabolism*, 44(5):801–816, 2024.
- [22] Lucas O Müller, Pablo J Blanco, Sansuke M Watanabe, and Raúl A Feijóo. A high-order local time stepping finite volume solver for one-dimensional blood flow simulations: application to the adan model. *International Journal for Numerical Methods in Biomedical Engineering*, 32(10):e02761, 2016.
- [23] L. Pegolotti, M.R. Pfaller, N.L. Rubio, K. Ding, R. Brugarolas Brufau, E.F. Darve, and A.L. Marsden. Learning reduced-order models for cardiovascular simulations with graph neural networks. *Computers in Biology and Medicine*, 168, 2024.
- [24] R. Penta and A. Gerisch. Investigation of the potential of asymptotic homogenization for elastic composites via a three-dimensional computational study. *Comput. Vis. Sci.*, 17:185–201, 2015.
- [25] M. Peyrounette, Y. Davit, M. Quintard, and S. Lorthois. Multiscale modelling of blood flow in cerebral microcirculation: details at capillary scale control accuracy at the level of the cortex. *PLoS One*, 13(1):e0189474, 2018.
- [26] M.R. Pfaller, L. Pegolotti, J. Pham, N.L. Rubio, and A.L. Marsden. Reduced-order modeling of cardiovascular hemodynamics. In *Biomechanics of the Aorta*, pages 449 – 476. Academic Press, 2024.
- [27] A.S. Popel and P.C. Johnson. Microcirculation and hemorheology. *Annu. Rev. Fluid Mech.*, 37:43–69, 2005.
- [28] L. Possenti, S. di Gregorio, F.M. Gerosa, G. Raimondi, G. Casagrande, M.L. Costantino, and P. Zunino. A computational model for microcirculation including Fahraeus-Lindqvist effect, plasma skimming and fluid exchange with the tissue interstitium. *International Journal for Numerical Methods in Biomedical Engineering*, 35(3), 2019.
- [29] A.R. Pries, D. Neuhaus, and P. Gaetgens. Blood viscosity in tube flow: dependence on diameter and hematocrit. *Am J Physiol*, 263:H1770–H1778, 1992.

- [30] A.R. Pries and T.W. Secomb. Microvascular blood viscosity in vivo and the endothelial surface layer. *American Journal of Physiology-Heart and Circulatory Physiology*, 289(6):H2657–H2664, 2005.
- [31] A. Quarteroni, A. Veneziani, and C. Vergara. Geometric multiscale modeling of the cardiovascular system, between theory and practice. *Computer Methods in Applied Mechanics and Engineering*, 302:193–252, 2016.
- [32] M. Raissi, P. Perdikaris, and G.E. 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.
- [33] F. Scarselli, M. Gori, A.C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- [34] F. Schmid, M.J. Barrett, P. Jenny, and B. Weber. Vascular density and distribution in neocortex. *NeuroImage*, 197:792–805, 2019.
- [35] F. Schmid, P.S. Tsai, D. Kleinfeld, P. Jenny, and B. Weber. Depth-dependent flow and pressure characteristics in cortical microvascular networks. *PLoS Comput. Biol.*, 13(2):e1005392, 2017.
- [36] W. Schreiner, F. Neumann, R. Karch, M. Neumann, S. M. Roedler, and A. End. Shear stress distribution in arterial tree models, generated by constrained constructive optimization. *Journal of Theoretical Biology*, 198(1):27–45, 1999.
- [37] T.W. Secomb. Blood flow in the microcirculation. *Annual Review of Fluid Mechanics*, 49:443–461, 2017.
- [38] Z. Tong, M. Catherall, and S.J. Payne. A multiscale model of cerebral autoregulation. *Med. Eng. Phys.*, 95:51–63, 2021.
- [39] T. Ventimiglia and A.A. Linninger. Mesh-free high-resolution simulation of cerebrocortical oxygen supply with fast fourier preconditioning. *International Journal for Numerical Methods in Biomedical Engineering*, 39(5):e3735, 2023.
- [40] P. Vitullo, A. Colombo, N. R. Franco, A. Manzoni, and P. Zunino. Nonlinear model order reduction for problems with microstructure using mesh informed neural networks. *Finite Elements in Analysis and Design*, 229:104068, 2024.
- [41] K. Xu, M. Zhang, S. Jegelka, and K. Kawaguchi. Optimization of graph neural networks: Implicit acceleration by skip connections and more depth. In *International Conference on Machine Learning (ICML)*, pages 11592–11602, 2021.
- [42] Eduardo G. Zilves, Lucas O. Müller, Gonzalo D. M. Talou, Vladimir Hachinski, J. David Spence, and Pablo J. Blanco. Blood flow in the human cerebral cortex: Large-scale pial vascularization and 1d simulation. *PLoS Computational Biology*, 21(9):1–38, 09 2025.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 52/2026** Bonazzoli M.; Ciaramella G.; Mazzieri I.
On the Unmapped Tent Pitching for the Heterogeneous Wave Equation
- 51/2026** Bellezza P.; Ciaramella G.; Macchini C.; Mazzieri I.; Verani M.
ParaFlow: Parareal Acceleration of Gradient-Flow Minimization
- 53/2026** Dong Z., Jiang Y., Ng M., Ciaramella G., Yin J.
Chebyshev-Filtered Randomized Low-rank Preconditioners for Symmetric Positive Definite Linear Systems
- 55/2026** Beirao da Veiga, L.; Canuto, C.; Nochetto, R.H.; Vacc, G ; Verani, M.
A Virtual Element Method for elliptic problems on trimmed background meshes
- 54/2026** Antonietti, P. F.; Corti, M.; Leimer Saglio, C. B.; Pagani, S.
The lymph 2.0 library: p-adaptive algorithms and parallel assembly strategies for polytopal DG methods
- 50/2026** Donnarumma, A.; Guagliardi, O.; Di Stazio F.; Mazza E.; Tanelli M.; Paganoni A.M.
Modelling Well-Being and Psychological Risk in Doctoral Education: An Integrated Latent Trait Approach
- 49/2026** Bortolotti, T.; Troilo, R.; Casu, F.; Vantini, S.; Menafoglio, A.
Regularized covariance estimation from partially observed interferometric data
- 48/2026** Antonietti, P.F.; Corti, M.; Orlando, G.
Optimized high-order IMEX-RK schemes for degenerate diffusion-reaction problems with application to travelling waves phenomena
- 46/2026** Cancrini, A.; Ciaramella, G.; Antonietti, P.F.
A Scalable Deflated Conjugate Gradient Solver for the Time-Dependent Pseudo-Stress Stokes Problem
- 47/2026** Torri, V.; Barbieri, E.; Cantarutti, A.; Giaquinto, C.; Ieva, F.
Automatic identification of diagnosis from hospital discharge letters via weakly supervised Natural Language Processing