POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Physics-Informed Graph Neural Network Surrogate for Hemodynamic Assessment in Post-Aneurysmal Subarachnoid Hemorrhage Vasospasm

Tesi di Laurea Magistrale in
Biomedical Engineering - Ingegneria Biomedica

Camilla Balzarotti, 10764340

Advisor:
Prof. Alberto Redaelli

Co-advisors:
Elisabetta Criseo

Academic year:
2025-2026

Abstract: Aneurysmal subarachnoid hemorrhage can be followed by vascular complications, including cerebral vasospasm, which narrows intracranial arteries in the Circle of Willis and alters blood supply to downstream territories. Noradrenaline-induced blood pressure augmentation is a standard treatment strategy, yet clinical management remains challenging due to highly heterogeneous hemodynamic responses driven by patient-specific anatomy, collateral capacity, and vasospasm severity. This thesis develops a physics-informed graph neural network surrogate to rapidly predict Circle of Willis hemodynamics under vasospasm and noradrenaline related pressure support regimes, in order to simulate different clinically plausible therapeutic scenarios. The TopCoW Challenge dataset, consisting of 125 cerebral artery reconstructions, was modified to include anatomical variants and controlled vasospasm-induced narrowing. 3D steady-state CFD simulations were performed in SimVascular across three inlet flow regimes mimicking increasing pressure support. Each anatomy was represented as a centerline graph encoding geometric, topological, pathological, and hemodynamic information. Different architectures, training strategies, and feature sets were compared, leading to a fine tuned double decoder model trained with a clinically oriented configuration. Notably, the final model enables the simulation of different scenarios without requiring outlet resistance or inlet flow fraction values, which are not available in clinical practice, while preserving mass conservation. On the internal test set, the final model achieved a pressure relative error of 4.4% and a flow mean absolute error of 5.2%, with prediction times of a few hundredths of a second. External benchmarking showed predictive capability beyond the training domain.

Key-words: vasospasm; Circle of Willis; CFD; graph neural network; physics-informed learning

1. Introduction

1.1. Clinical background

1.1.1 Subarachnoid hemorrhage

Subarachnoid hemorrhage (SAH) is the extravasation of blood into the subarachnoid space, the cerebrospinal fluid–filled compartment between arachnoid and pia mater, which cushions the brain, allows circulation of major

cerebral vessels, and facilitates metabolic exchange [1]. In 80% of cases, the cause of nontraumatic subarachnoid hemorrhage is the rupture of an intracranial aneurysm (aSAH), while the remaining smaller fraction is due to other vascular or systemic causes (e.g., vascular malformations, vasculitis), and a subset remains angiography-negative even after appropriate imaging (naSAH) [2, 3]. The latter cases are further distinguished according to the bleeding pattern. Perimesencephalic naSAH, in which blood is mainly localized around the midbrain, generally has a more benign clinical course and may not require repeat angiography when the initial study is technically adequate. In contrast, non-perimesencephalic or diffuse naSAH is managed more cautiously and commonly requires follow-up vascular imaging, since an initially occult aneurysm or vascular lesion may still be present [4].
From an epidemiologic standpoint, population-based data report overall SAH incidence on the order of single-digit cases per 100,000 person-years in European settings (with variability by region, age, and sex). In one Italian population-based registry (2011–2020), the incidence of first-ever non-traumatic SAH was 6.5 per 100,000 person-years, with most cases classified as aneurysmal [2]. In the US, it affects 21,000 to 33,000 people each year [5].
Despite accounting for only approximately 5% of all strokes [2, 5], subarachnoid hemorrhage disproportionately affects younger patients compared to other cerebrovascular events and is associated with substantial mortality and long-term morbidity. Among patients with aneurysmal subarachnoid hemorrhage who survive the acute event, up to 50% experience persistent neuropsychological deficits [6], resulting in a marked reduction in functional independence and quality of life [7] and one third of survivors require lifelong assistance or care. Furthermore, long-term cognitive impairment has been reported in up to 46% of cases [8], significantly impacting daily activities, employment, and social reintegration. In addition to its individual consequences, subarachnoid hemorrhage imposes a considerable burden on healthcare systems, largely driven by prolonged intensive care and hospitalization [9].

1.1.2 Acute phase presentation and management

Pathophysiologically, following the rupture of an intracranial aneurysm, arterial blood is rapidly released into the subarachnoid space, causing an abrupt increase in intracranial pressure (ICP), that may transiently approach mean arterial pressure. This sudden ICP surge can result in a short period of global cerebral hypoperfusion or ischemia [10]. In parallel, the presence of blood and its breakdown products [11] within the cerebrospinal fluid triggers a cascade of secondary processes, including blood–brain barrier disruption [12], cerebral edema, oxidative stress [13], microvascular dysfunction, and neuroinflammation, collectively referred to as early brain injury (EBI) [10].
Clinically, this causes severe thunderclap headache, often accompanied by nausea or vomiting, neck stiffness, photophobia, and, depending on severity, transient or persistent loss of consciousness or focal deficits.
Because early rebleeding and secondary brain injury can occur soon after rupture, current guidelines emphasize SAH as a neurologic emergency requiring rapid recognition and diagnostic confirmation [5]. Initial neurological assessment relies on standardized clinical grading systems, including the Glasgow Coma Scale (GCS) [14], the World Federation of Neurosurgical Societies (WFNS) scale, and the Hunt–Hess classification. Non-contrast computed tomography (CT) of the head is the first-line imaging modality for rapid confirmation of hemorrhage and identification of its source. Although non-contrast CT demonstrates high sensitivity in the early hours following symptom onset, its diagnostic accuracy progressively decreases over time [15]. Consequently, when CT findings are negative but clinical suspicion for subarachnoid hemorrhage remains high, a lumbar puncture may be performed to detect cerebrospinal fluid abnormalities consistent with hemorrhage. Once SAH is confirmed, acute management is initiated immediately in the intensive care unit to stabilize the patient and avoid neurological complications. CT angiography (CTA) or digital subtraction angiography (DSA) [16, 17] is then performed to identify the ruptured aneurysm and to guide definitive treatment, either by surgical clipping or endovascular coiling, reducing the risk of rebleeding [10, 16].

1.1.3 The role of the Circle of Willis and its variants

The Circle of Willis (CoW) is a polygonal arterial anastomotic network located at the base of the brain that connects the anterior and posterior cerebral circulations. It is formed by the bilateral internal carotid arteries (ICA), from which the middle cerebral arteries (MCA) branch off, the anterior cerebral arteries (ACA) connected by the anterior communicating artery (ACom), and the posterior cerebral arteries (PCA), which are linked to the internal carotid arteries through the posterior communicating arteries (PCom). The posterior cerebral arteries originate from the basilar artery (BA), which provides the main inflow to the posterior circulation. Together, these vessels form a circular configuration surrounding the optic chiasm and the pituitary stalk.

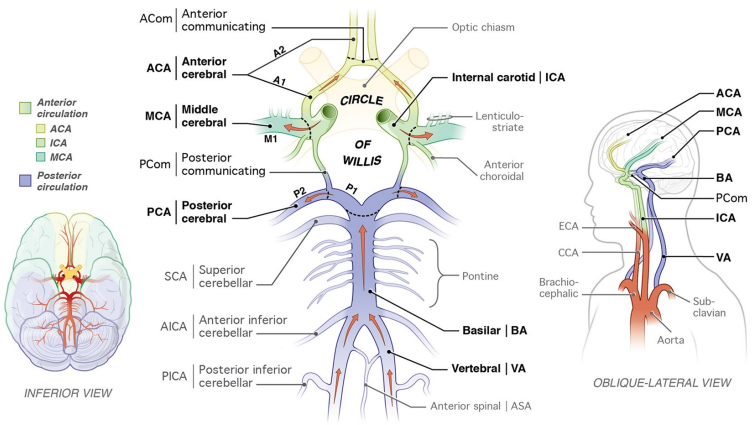

Figure 1: Circle of Willis anatomy and location.

Functionally, the CoW has historically been attributed a compensatory role, allowing, through its communicating arteries, redistribution of blood flow between the left and right hemispheres and between the anterior and posterior circulations when perfusion through one arterial territory is compromised. Additionally, Vrselja et al. (2014) suggested the CoW also acts as a passive pressure-dissipating system, which dampens sharp pressure spikes caused by sudden changes in blood velocity, shielding the fragile microcirculation and blood-brain barrier from hemodynamic stress [18].
However, the anatomical configuration of the CoW shows considerable inter-individual variability, with hypoplasia (narrowing) or aplasia (absence) of segments occurring in 52% to 88% of the population [19], while complete and symmetric configurations are present in only a minority of individuals. These asymmetric configurations create hemodynamic stress by forcing a single artery to supply a larger portion of the brain, which is strongly associated with the formation of aneurysms [20], and directly contribute to rupture risk [21].

1.1.4 The role of vasospasm and delayed cerebral ischemia

Cerebral vasospasm consists of the delayed, sustained narrowing of large and medium-sized intracranial arteries, most commonly involving the proximal vessels of the CoW [22]. Its development arises largely from the degradation of subarachnoid blood, which releases hemoglobin breakdown products that induce nitric oxide scavenging, oxidative stress, endothelial dysfunction, and vascular smooth muscle hypercontractility, leading to prolonged arterial narrowing and compromised perfusion of downstream brain territories. Inflammatory and structural changes of the vascular wall further maintain vasoconstriction [23, 24]. Angiographic or radiographic vasospasm has been reported in up to approximately 70% of patients after aneurysmal SAH, highlighting its high prevalence and clinical relevance [10]. Specifically, the main medical concern associated with vasospasm is its potential contribution to delayed cerebral ischemia (DCI), a major driver of secondary neurological injury, which typically develops between days 3 and 14 after the initial hemorrhage and represents one of the leading causes of secondary neurological deterioration and long-term disability among patients who survive the early insult [10, 16]. It affects approximately 20–30% of patients, despite appropriate early management and aneurysm securement [24].
Formally, DCI is defined by new focal neurological impairment and/or a decrease of at least 2 points on the Glasgow Coma Scale, lasting $\geq$1 hour, occurring after the immediate peri-treatment period, and not attributable to other alternative causes after appropriate clinical assessment and imaging/laboratory evaluation [23, 25], such as rebleeding, hydrocephalus, metabolic disturbances, or infection.
While for decades DCI was attributed primarily to cerebral vasospasm [10], it has more recently been recognized as a multifactorial process [23, 26], influenced also by microcirculatory dysfunction, microthrombosis, thromboinflammation, neuroinflammation, impaired cerebral autoregulation, endothelial and glycocalyx dysfunction, and cortical spreading depolarizations [23, 24]. This shift is supported also by the incomplete correspondence between large-vessel narrowing and delayed ischemic injury: for example, up to 25% of delayed cerebral infarcts have been reported not to correlate with territories of arterial narrowing or to occur in patients without evidence of vasospasm [27]. These contributors are not necessarily all present in every patient, and even when appearing in combination, they are not mutually exclusive, but may interact through coupled and feedforward pathways. Therefore, it is not possible to assign an independent and fixed percentage contribution to each individual mechanism.
Nevertheless, in this complex scenario, vasospasm remains problematic even before it results in a formal DCI endpoint, because neurological symptoms attributed to vessel narrowing indicate a state of threatened perfusion in which downstream territories are likely to be approaching a critical ischemic threshold. This is clinically

interpreted as a warning state of a possibly evolving ischemic condition and commonly motivates intensified monitoring and timely therapeutic interventions aimed at improving cerebral blood flow, such as blood-pressure augmentation or endovascular treatment, such as intra-arterial vasodilators or angioplasty, in selected cases. Furthermore, in a systematic review and meta-analysis, Kumar et al. [28] reported that transcranial-Doppler evidence of vasospasm was highly predictive of DCI, with a pooled sensitivity of 90% and specificity of 71%. Consistently, the 2023 AHA/ASA guidelines [3] explicitly state that in aSAH patients with symptomatic vasospasm, elevating systolic blood pressure may be reasonable to reduce progression and severity of DCI, and that cerebral angioplasty may be reasonable in patients with severe vasospasm. Therefore, modelling vasospasm-induced pressure and flow redistribution is relevant not only for interpreting established DCI, but also for assessing patient-specific hemodynamic vulnerability during symptomatic or evolving vasospasm, where neurological deterioration or perfusion compromise is still potentially reversible and timely treatment may prevent progression toward clinically defined DCI or infarction. In particular, a patient-specific surrogate model can help investigate whether a given vascular anatomy and collateral configuration may compensate for arterial narrowing, or whether the imposed reduction in vessel caliber produces redistribution patterns compatible with increased ischemic risk.

1.1.5 Induced Hypertension

A common first-line treatment to augment cerebral perfusion is induced hypertension (IH), thus deliberately increasing systemic arterial blood pressure using vasopressor agents. The most common is norepinephrine (noradrenaline), a catecholamine neurotransmitter and hormone [3, 29], which binds predominantly to α_1-adrenergic receptors on vascular smooth muscle cells, producing systemic vasoconstriction. This effect occurs predominantly in peripheral vascular beds, such as the skin, gastrointestinal tract, and skeletal muscle, increasing total peripheral resistance (i.e. systemic vascular resistance, SVR), leading to an elevation in mean arterial pressure (MAP) [30]. Although norepinephrine exhibits β_1-adrenergic activity that could theoretically increase heart rate by elevating myocardial contractility, the rise in arterial pressure activates the arterial baroreflex [31], leading to increased parasympathetic tone and inhibition of sympathetic outflow to the heart. As a result, heart rate frequently remains unchanged or may even decrease through reflex bradycardia despite β_1 stimulation. Therefore, in this context, the increase in arterial pressure is mainly driven by the α_1-mediated increase in systemic vascular resistance. The effect of noradrenaline administration on the brain critically depends on the state of cerebral autoregulation. Cerebral blood flow is governed by the balance between the pressure available to drive blood through the cerebral vasculature and the resistance opposed by the vascular bed. This can be summarized by expressing cerebral blood flow as proportional to cerebral perfusion pressure (CPP), defined as the difference between MAP and intracranial pressure (ICP), and inversely proportional to cerebrovascular resistance, according to Equation 1:

$$Q_{cerebral} = CBF = \frac{CPP}{R_{cerebral}} = \frac{MAP - ICP}{R_{cerebral}} \tag{1}$$

from which it can be derived that by increasing MAP, norepinephrine increases cerebral perfusion pressure (CPP), provided intracranial pressure (ICP) does not rise proportionally.
However, the impact of the variation of CPP on cerebral blood flow (CBF) is modulated by cerebral autoregulation, which refers to the intrinsic ability of the cerebral vasculature to maintain relatively constant cerebral blood flow despite changes in cerebral perfusion pressure (CPP) within a physiological range. Consequently, in individuals with intact autoregulation, when CPP increases, cerebral arterioles dynamically constrict to prevent excessive flow; conversely, when CPP decreases, arterioles dilate to preserve perfusion, adjusting to maintain stable flow, as demonstrated by Meng et al. (2024) [32]. In the presence of vasospasm, autoregulation is frequently impaired [33, 34], allowing increases in MAP to translate more directly into higher CPP and improved cerebral blood flow, despite cerebrovascular resistance remaining elevated.

1.1.6 Limitations of pressure-based management

There are several aspects suggesting an arbitrary non-patient specific pressure-based management may fail to capture the complexity of DCI.
First, clinical and imaging evidence highlights the non-deterministic relationship between vasospasm and infarction. Severe arterial narrowing may be observed in the absence of ischemic lesions, while infarction can occur without significant proximal vasospasm [25, 35]. This dissociation suggests that large-vessel diameter reduction alone is insufficient to explain the spatial distribution of delayed cerebral ischemia. Consistently, perfusion imaging studies further support that regional cerebral blood flow deficits do not always spatially correspond to the most severely narrowed arterial segments [35]. Moreover, while vasospasm is frequently a focal phenomenon (affecting specific large arteries), autoregulation impairment is often global, extending the brain areas vulnerable

to ischemia [34].
In addition, the perfusion response to hypertension is heterogeneous, meaning that raising MAP does not produce a uniform increase in cerebral blood flow (CBF) across patients or even across different brain regions within the same patient. This is closely linked to the substantial variability in collateral circulation and cerebrovascular anatomy, which determine the capacity of collateral pathways to redistribute flow across vascular territories.
Further evidence arises from randomized clinical trials. Although it was terminated early, the HIMALAIA study [30, 36] showed that induced hypertension successfully increased mean arterial pressure by approximately 10–11 mmHg during the intervention period, but this physiological effect did not translate into improved functional outcome at three months. Moreover, induced hypertension was associated with a higher rate of serious adverse events, including cardiopulmonary complications [30]. Another retrospective study [37] similarly observed that patients receiving induced hypertension experienced significantly higher rates of specific adverse effects, including myocardial ischemia, electrolyte imbalances and polyuria. Nevertheless, the treatment showed a temporal link to clinical improvement in 48% of patients (such as an improved Glasgow Coma Score or resolution of focal deficits), highlighting the importance of meticulously weighing the risks on a patient-by-patient basis and to carefully assess patient-specific optimal pressure targets.

1.2. Objectives of the work

Taken together, these observations suggest that relying solely on angiographic severity or standardized blood pressure targets may be insufficient to guide optimal management. Instead, patient-specific modeling of vascular geometry, pressure distribution, and collateral configurations may provide insight into the regional distribution of flow and help identify territories and cases in which pressure augmentation is likely to be beneficial, as opposed to those in which structural limitations constrain compensatory redistribution. This represents a relevant clinical gap in the management of aneurysmal subarachnoid hemorrhage, currently relying on empirical blood pressure augmentation without fast individually-tailored hemodynamic quantification. While computational fluid dynamics (CFD) can provide detailed insight into cerebral blood flow patterns, its high computational cost limits its use in time-sensitive clinical decision making. Alternative imaging techniques, such as 4D flow MRI, can provide in vivo information on blood-flow patterns but their use also remains limited by acquisition time, cost, spatial and temporal resolution constraints, and restricted accessibility in routine clinical practice.
The central purpose of the project is therefore to develop and validate a physics-informed graph neural network for near real-time prediction of patient-specific pressure and flow redistribution in the CoW. The intended use of the framework is not to replace clinical monitoring or high-fidelity CFD, but to support rapid assessment of collateral capacity, vasospasm-induced alterations, and the impact of pressure augmentation across anatomical variants. Starting from a patient-specific vascular reconstruction, the operator could define different clinically plausible conditions, such as different pressure-support regimes mimicking noradrenaline administration or even alternative vasospasm locations and severities if recent imaging is not yet available. For each scenario, the model estimates relative pressure and normalized flow redistribution across the vascular graph, allowing comparison of how perfusion patterns change under different anatomical and hemodynamic conditions.
By providing insight into patient-specific perfusion mechanisms, the model supports clinical decision-making by helping identify regions at risk of insufficient perfusion and by enabling timely intervention, ultimately aiming to reduce the consequences of DCI following SAH. Moreover, it may also assist in tailoring noradrenaline administration, avoiding complications associated with the side effects of excessive hypertension.
To the best of our knowledge, it is the first investigation of the ability of machine learning models trained on simulated hemodynamic data for this setting, where patient-specific anatomy, vasospasm-induced narrowing, and pressure-support conditions are jointly considered.

1.3. Current state of the art

Given the clinical significance of the CoW, many studies have investigated its hemodynamics, in order to understand its physiological and pathological phenomena, including aneurysm formation and dynamics, as well as changes in collateral flow rates and pressure in the presence of stenosis or following surgical procedures. Computational fluid dynamics (CFD) models are especially attractive for their ability to simulate blood flow, without the need for invasive instrumentation, while also allowing the estimation of additional parameters (e.g. pressure, wall shear stress), that otherwise could not be continuously recorded by currently available clinical devices [38].
As reported by Straccia et al. [39], even though experimental setups vary, most studies share common modeling assumptions or characteristics. Most of the existing literature assumes rigid vascular walls, even though blood vessels are inherently compliant, passively expanding early in the cardiac cycle and relaxing later in the cardiac cycle. Autoregulation is usually not incorporated. One exception is Moore et al. [40], in which a feedback loop

is used to modify the resistance of the downstream vascular beds, in response to the current flow rate, until the expected flow rate is reached. Outlet boundary conditions are imposed either by deriving Womersley profiles (pulsatile velocity profiles) from patient-data or defined using resistance values, often in conjunction with a capacitance term to represent downstream vascular compliance via RC or RCR conditions [41]. By contrast, constant pressure boundary conditions are generally considered unrealistic since they implicitly assume identical downstream resistance for all outlets. Similarly, prescribing the flow rate as inlet boundary constraint, either in a steady-state or pulsatile setting, is preferred over imposing constant pressure values.

Different model configurations can be adopted based on the aim of the hemodynamics analysis [38]. Reduced-order models, such as 0D and 1D formulations, are commonly employed when only specific aspects of the problem are of interest or when computational efficiency is required (e.g. in applications requiring repeated simulations, such as boundary condition calibration, or in large-scale and microvascular networks). 0D models, also referred to as lumped parameter models, provide rapid estimates of global hemodynamic quantities, such as pressure and flow, but neglect spatial and geometric details of the vasculature [42]. 1D models retain a simplified spatial description, which allows them to account for variations in vessel area and axial propagation of pressure and flow, while remaining computationally efficient [43]. However, the predictive performance of reduced-order models can be compromised by their inability to accurately represent pressure losses at vascular bifurcations, as well as their dependence on simplified mathematical descriptions that are typically based on empirical data or heuristic assumptions, particularly in the presence of pathological conditions such as stenoses or aneurysms. Therefore, this configuration is usually employed when the characterization of global systemic hemodynamics is prioritized over the resolution of complex, local flow topologies, accepting a limited loss of accuracy in exchange for significantly faster simulations (typically a few seconds compared to several hours for full 3D models) [44]. This approach is adopted, for instance, by Huang et al. [45] to evaluate how the anatomical completeness of the CoW influences the risk of cerebral hyperperfusion following carotid artery surgery.

Three-dimensional models, however, represent the highest level of spatial resolution, capturing the full volumetric complexity of vascular structures. Their governing physics are described by the Navier-Stokes equations, together with the continuity equation for an incompressible fluid, which represent the conservation of momentum and mass at every infinitesimal point within the three-dimensional domain. To solve these complex differential equations, the continuous geometry is discretized into a computational mesh or grid. The main advantage of this approach is the ability to extract localized mechanical parameters that lower-dimensional models cannot compute, albeit at the cost of significantly higher computational expense. In the literature, 3D CFD simulations have been widely used to investigate the relationship between vascular geometry, flow patterns, and clinically relevant indicators, such as aneurysm development and rupture risk [46–48]. These studies highlight the potential of high-fidelity models to capture complex hemodynamic phenomena, while also emphasizing the challenges associated with data availability and computational cost.

Despite providing detailed insight into blood flow behavior, CFD models require simulation-related expert knowledge [49] and high computational costs, due to the iterative optimization, as well as the preprocessing, mesh generation, and postprocessing steps [50]. This restricts the use of numerical simulations both for population-scale studies and for exploring multiple physiological or pathological scenarios in an individual patient (e.g. different levels of stenosis) [50], limiting its integration into clinical workflow for real-time or patient-specific applications. To address this challenge, increasing attention has been directed toward deep learning-based surrogate models, capable of approximating hemodynamic quantities, while demanding a fraction of the time and computational cost [51].

Concerning the learning paradigm, early models were mostly purely data-driven, learning a mapping from geometry and boundary conditions to flow variables from simulation data alone. For example, Li et al. [52] proposes a purely data-driven framework designed to rapidly and accurately predict blood flow patterns in cerebral aneurysms before and after the placement of flow-diverting (FD) stents. Although such models can approximate complex physical systems without any physical knowledge involved, they require a massive amount of simulation data [53], making them unsuitable for several application domains, including the medical sector, where data scarcity is still an ongoing challenge. Consequently, increasing attention has been recently given to Physics-Informed Neural Networks (PINN) [53]. These networks are designed to solve supervised learning tasks while strictly adhering to physical laws described by nonlinear partial differential equations (PDEs). Specifically, this is achieved by embedding the governing equations of the system directly into the model's loss function, which combines the data consistency loss (i.e. the physics-unaware difference between the predicted value and the ground truth) with a physics loss, which penalizes any predictions that do not satisfy the underlying PDE at a set of "collocation points" (e.g., boundary points). This approach significantly constrains the space of possible solutions to those that are physically realistic (e.g., conserving mass or momentum), allowing the model to generalize accurately even with very few training examples [53]. This makes them particularly suitable also for hidden problems and hidden-variable estimation because physics constraints let the model infer quantities that are not directly measured from the governing equations. For instance, Sarabian et al. [54] shows how a physics-informed deep learning framework can provide a complete picture of brain hemodynamics by fusing Transcranial Doppler (TCD) ultrasound measurements, which are spatially limited to only a few locations in

the cerebral arteries, with the underlying physics of blood flow embedded in the network.

However, while physics-informed learning improves physical consistency, the choice of the underlying architecture remains a critical factor.
Convolutional Neural Networks (CNNs) [55] have been widely used to learn the dynamics of complex systems across different domains. CNNs operate on data defined on regular grids, applying convolutional filters that capture local spatial patterns within a fixed neighborhood (the receptive field) and propagate this information through successive layers. By stacking multiple convolutional layers, the network can learn hierarchical representations, progressively encoding higher-level features from local interactions. Nonetheless, the regularity of this grid-based formulation limits the applicability of CNNs to problems defined on irregular domains. In particular, vascular systems such as the CoW are inherently sparse and structured as irregular branching networks. As a result, applying CNNs typically requires interpolation or voxelization of the geometry, which can lead to loss of geometric fidelity and reduced ability to capture the underlying topology. This representation mismatch can also hinder generalization to unseen anatomical configurations, especially when structural variability plays a central role in determining the system dynamics. Additionally, a key limitation of CNNs in this context is their weak relational inductive bias. This concept, introduced by Battaglia et al. [56], refers to the set of assumptions a learning algorithm uses, due to architectural constraints that are independent from the data itself, to learn specifically about how entities interact with one another and to prioritize one solution over another. CNNs have a strong bias for locality (i.e. assuming an entity only has a meaningful relationship with its immediate neighbors) and translation invariance (i.e. the same filter is applied everywhere on the grid). While these characteristics make CNNs extremely powerful in many fields (i.e. image classification or segmentation), they also limit their effectiveness in other areas, for example when the understanding of long-range dependencies and of arbitrary relational structures is required [56]. Another non-negligible aspect is that the repetitive application of complex feature extractors can reduce model interpretability, which represents a huge criticality in clinical settings.
To address the limitations imposed by regular grid formulations on complex vascular geometries, geometric deep learning approaches like PointNet++ [57] bypass voxelization entirely by operating directly on raw, unstructured 3D point clouds. PointNet++ learns local geometric descriptors by repeatedly sampling points, grouping nearby points into neighborhoods, and aggregating their features hierarchically. For instance, Suk et al. [58] developed a feed-forward surrogate framework utilizing an equivariant extension of PointNet++ integrated with physics-informed Navier-Stokes regularization to estimate patient-specific 3D volumetric blood velocity fields directly within carotid artery bifurcations. However, because PointNet++ relies on local neighborhoods and hierarchical aggregation, long-range vascular connectivity and changes in network topology are not explicitly encoded. This is a limitation in structures such as the CoW, where locations that are distant in Euclidean space may still be strongly coupled hemodynamically through the vascular network, depending on vessel geometry, collateral pathways, and the presence or absence of communicating arteries.
Transformer-based [59] architectures have also been explored for learning fluid dynamics and flow-field prediction, showing strong capability in capturing long-range dependencies and complex interactions. Their self-attention mechanism enables flexible modeling of relationships between different elements of the system, making them highly expressive. When evaluating how different deep learning backends perform in predicting virtual Fractional Flow Reserve (vFFR), a critical non-invasive diagnostic metric in coronary artery disease, Nannini et al. [60] demonstrated the optimality of transformer-based backends in predicting hemodynamics of complex, heterogeneous topologies of real patient-specific data. However, transformers are structurally agnostic, in the sense that they do not inherently encode the underlying topology of the system. In particular, they do not explicitly assume graph connectivity, spatial adjacency, or vessel topology. While positional or geometric encodings can be introduced to partially recover structural information, the relationships between elements must still be learned from data through attention mechanisms rather than being explicitly defined. In the context of cerebrovascular hemodynamics, and in particular for the CoW, these structural properties are not incidental but fundamental, as connectivity and topology directly determine pressure distribution and flow redistribution across vascular territories. As a result, relying on models that must infer these relationships may introduce additional complexity in learning.
Recent state-of-the-art frameworks have combined transformer backends with operator learning. While traditional neural networks learn mappings between finite-dimensional Euclidean spaces or fixed discrete representations, neural operators learn mappings between function spaces [61, 62]. This formulation makes them well suited for approximating solutions to parametric partial differential equations (PDEs). Notably, Kaczmarek et al.[63] proposed a patient monitoring framework based on the General Neural Operator Transformer (GNOT) to address the inverse problem of reconstructing full 1D velocity, cross-sectional area and pressure pulsatile waveforms across the CoW from sparse data, specifically utilizing blood velocity from a limited set of arteries obtained via transcranial Doppler (TCD) ultrasound. However, neural-operator formulations also introduce structural and computational challenges. First, they typically require large synthetic datasets. For instance, Kaczmarek et al. [63] trained their framework on around 15,000 1D numerical simulations. Furthermore, al-

though a trained neural operator can substantially reduce the cost of each forward evaluation, solving the inverse problem to reconstruct patient-specific states still relies on an outer iterative optimization loop, possibly requiring a significant computational cost and time [64]. Finally, while neural operators are theoretically designed to be discretization-invariant, generalization across irregular geometries and topologically distinct domains remains challenging in practice [65]. This is particularly relevant for the CoW, where patient-specific anatomy, non-uniform meshes, and missing or variant vessels introduce substantial intra- and inter-patient variability.
In contrast, graph neural networks (GNNs) provide a representation that is naturally aligned with the structure of the vascular system [66]. By explicitly encoding vessels as edges and junctions as nodes, GNNs allow information to propagate along physically meaningful connections. This enables the model to incorporate prior knowledge about the system's spatial organization and provides a more suitable inductive bias for modeling flow redistribution in vascular networks such as the CoW. The adoption of GNNs as solvers for mesh-based simulation has been introduced by Pfaff et al. [67], presenting MeshGraphNets, a framework that utilizes an Encode-Process-Decode architecture to learn the dynamics of complex physical systems directly from mesh-based data. The original MeshGraphNet implementation has been adapted by Pegolotti et al. [68] to iteratively predict the pressure and flow rate in cardiovascular blood flow simulations. Specifically, they demonstrated the ability of the GNN-based surrogate model to maintain the efficiency of 1D models while achieving the accuracy of 3D simulations across diverse patient-specific geometries. Botta et al. [66] further advanced the application of GNNs to vascular hemodynamics by incorporating physics-informed learning. By embedding physics laws directly in the loss function, the model indeed achieves physical consistency and geometrical flexibility, making it a viable tool for real-time biomedical applications [66].

1.4. Guidance to the reader

Taken together, Pegolotti et al. [68] and Botta et al. [66] demonstrate the potential of graph-based approaches for modeling vascular hemodynamics, particularly in terms of capturing complex geometries and flow interactions. In this context, the present work builds upon these developments by proposing a physics-informed graph neural network specifically designed for rapid prediction of patient-specific pressure and flow redistribution in the CoW under post-subarachnoid hemorrhage vasospasm-inspired conditions and pressure-support regimes. By combining the structural advantages of graph-based representations with clinically accessible inputs and near real-time inference capabilities, this framework aims to bridge the gap between computational modeling and practical clinical applicability.
A first relevant contribution is the generation of a large patient-derived Circle of Willis hemodynamic dataset. Starting from clinical imaging, the dataset systematically incorporates anatomical variants, vasospasm-induced narrowing, and multiple hemodynamic regimes mimicking the effect of noradrenaline administration. 3D CFD simulations performed in SimVascular are used as ground truth, and the resulting three-dimensional pressure and velocity fields are projected onto centerline-based vascular graphs. This provides a structured representation specifically designed to study patient-specific pressure and flow redistribution in the CoW under clinically motivated pathological and physiological conditions.
A second innovative aspect is the formulation of the surrogate model itself. The vascular domain is represented as a graph, allowing the model to encode vessel connectivity and propagate information along physiologically meaningful pathways. The proposed physics-informed graph neural network predicts relative pressure and normalized flow fraction from node, edge, and global-context features, combining data-driven learning with conservation-based constraints. This is essential because a clinically useful surrogate should not only minimize numerical error, but also produce physiologically plausible flow distributions.
Finally, the tool is designed with clinical deployability in mind. It relies exclusively on clinically available data, including the CoW reconstruction obtained from CT or MRI imaging, geometry-derived vascular parameters and mean arterial pressure, available from standard blood pressure monitoring (see Section 2.3.1 for further insight), while avoiding dependence on quantities such as outlet resistance or inlet flow fraction that would normally require empirical estimation, calibration procedures, or additional flow measurements by TCD Doppler that are usually not timely available in routine clinical practice. Moreover, incorporating physics-informed constraints, including mass conservation at the boundaries, promotes physiologically plausible predictions even when generalizing to unseen anatomical variants or pathological states. This design seeks to facilitate the integration of the surrogate into real-world clinical workflows, rather than remaining limited to research environments.
Overall, the thesis demonstrates the feasibility of using physics-informed graph neural networks for fast, patient-specific estimation of CoW hemodynamics, with potential applications in vasospasm assessment, collateral-capacity evaluation, and pressure-support scenario exploration.

2. Materials and Methods

2.1. Dataset Generation

The development of data-driven surrogate models, and in particular graph neural networks, requires the availability of diverse datasets capturing both anatomical variability and corresponding hemodynamic behavior. However, in the context of cerebrovascular modeling, no publicly available dataset currently provides a comprehensive collection of patient-specific Circle of Willis geometries coupled with 3D computational fluid dynamics (CFD) simulations. Existing studies on the COW typically rely on small cohorts, often limited to a few patient-specific cases [69], which are insufficient to train models capable of generalizing across anatomical variations.

To address this limitation, a novel large-scale dataset was developed, having the TopCoW2024 dataset as the foundational source of vascular geometries [70]. This dataset, born with the similar intent to provide open data with verified annotations for CoW segmentation benchmarking, is composed of both MRI and CTA scanners of patients admitted to the Stroke Center of the University Hospital Zurich (USZ) between 2018 and 2019. As reported by the study, within the segmented CoW anatomies, 13 vessels have been annotated: left and right internal carotid artery (ICA), left and right anterior cerebral artery (ACA), left and right middle cerebral artery (MCA), anterior communicating artery (Acom), left and right posterior communicating artery (Pcom), left and right posterior cerebral artery (PCA), and basilar artery (BA). In rare cases, the anterior part of the CoW can have a third A2 artery arising from the Acom, and we labeled it with class 3rd-A2.

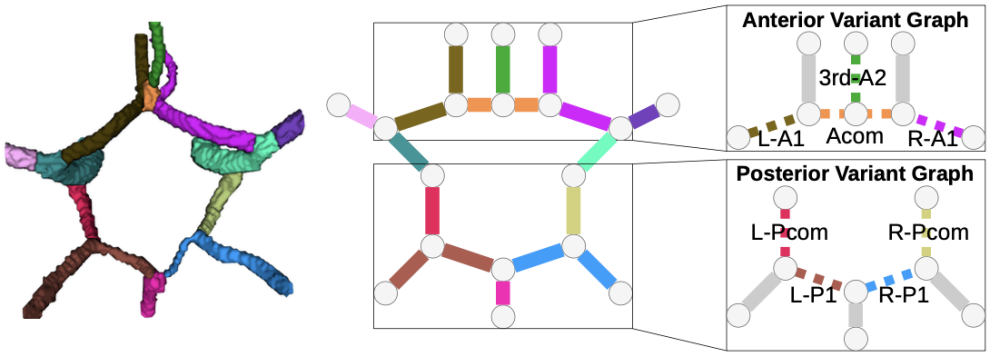

Figure 2: Example of CoW graph annotation, taken from the TopCoW2024 challenge [70]

The 125 patient-specific cases provided as part of the publicly available training set were used as the basis of this study, including both imaging data and anatomical annotations. The advantages of such adoption are several. Firstly, this dataset provides a large and diverse collection of patient-specific Circle of Willis reconstructions derived from clinical imaging, capturing a wide range of anatomical configurations representative of real-world variability. Such diversity is essential for modeling collateral flow redistribution, which is strongly influenced by vascular topology. Moreover, this representation also allows for direct geometric manipulation, enabling controlled modifications to simulate vasospasm severity and explore anatomical variability in a systematic and reproducible manner.

All vascular geometries, originally provided as volumetric segmentations in `.nii.gz` format, were first converted into surface representations (`.stl`). This conversion enables explicit geometric manipulation of the vascular structures, which is not directly feasible in voxel-based representations and guarantees compatibility with downstream processing steps such as mesh generation for CFD simulations. Moreover, for each geometry the corresponding centerline is automatically computed using the "Extract Centerline" module in the 3D Slicer [71] and saved in `.vtp` format.

2.1.1 Inclusion of anatomical variants

The original TopCoW2024 dataset was used as the anatomical starting point of this work, but, having been developed primarily for Circle of Willis segmentation benchmarking, it does not exhaustively cover the full spectrum of anatomical configurations observed in the general population, which can significantly influence collateral flow and hemodynamic response (see Section 1.1.3). Moreover, the visibility and annotation of small vessels may depend on the imaging modality and segmentation quality, especially for segments such as the PComs and ACom. To bridge this gap, the original dataset was curated and modified to systematically comprise anatomically-adherent variants, which are defined as configurations where vessels are absent, hypoplastic, or alter the polygon's connectivity, while diameter asymmetries between sides are treated as continuous geometric properties rather than discrete anatomical classes.

Specifically, six topological categories were included, with frequencies matching population distributions reported in the literature. The complete CoW configuration accounts for 41 subjects, matching the approximately 40% prevalence reported by Kizilgoz et al. [72]. Bilateral PComA absence appears in 25 subjects, reflecting the roughly 28% population frequency, while unilateral PComA absence also appears in 25 subjects, corresponding to approximately 16% prevalence. Fetal-type PCA configurations account for 25 subjects, within the 10% to 30% range reported across studies [21]. A1 absence appears in 6 subjects, matching the 5% to 8% prevalence, and AComA absence in 3 subjects, reflecting the 1% to 3% population frequency. The number of patients per anatomical variants is listed in Table 1. Anatomies from TopCoW were assigned to categories requiring minimal modifications, preserving geometric fidelity while ensuring topological accuracy.

Variant distribution

Variant type	#Patients	Percentage
Complete	41	33%
L-PCom missing	15	12%
R-PCom missing	10	8%
Both PComs missing	25	20%
Hypoplastic L-ACA	3	2.5%
Hypoplastic R-ACA	3	2.5%
A-Com missing	3	2%
L-PCA not connected to BA	10	8%
R-PCA not connected to BA	10	8%
Both PCAs not connected to BA	5	4%

Table 1: Distribution of the anatomical variants across the 125 different patient topologies.

2.1.2 Vasospasm simulation and evaluation

All 125 original patient-specific geometries were retained because, even though vasospasm introduces localized increases in vascular resistance and can therefore amplify redistribution patterns, the governing hemodynamic principles remain the same in vasospasm-induced and healthy vascular networks. In this work, healthy geometries thus refer to the original non-narrowed vascular configurations, used as anatomical references for the corresponding vasospasm-induced cases. Variations in vessel caliber, cross-sectional area, vessel shape, anatomical completeness, and collateral connections already influence the general relationship between vascular geometry, topology, and pressure and flow patterns, which is accentuated by vasospasm. Moreover, as further explained in Section 2.2.2, all CFD simulations were performed under boundary-condition assumptions consistent with impaired autoregulation, so that the healthy geometries did not represent normally autoregulated physiology. This allowed the model to learn general geometry-topology-hemodynamics relationships that remain relevant when pathological narrowing is introduced.

To explicitly account for the impact of pathological conditions on cerebral hemodynamics, vasospasm was introduced in a subset of the geometries. Specifically, arterial narrowing was applied to 60 of the 125 cases, enabling systematic investigation of the interaction between anatomical structure and pathological alterations. Since vasospasm occurrence after SAH is not known to be preferentially associated with one specific Circle of Willis variant, the vasospasm-induced cases were distributed so as to maintain the anatomical-variant proportions of the original cohort and, as far as possible, of the population, as reported in Figure 5, thus ensuring coverage of vasospasm patterns across different anatomical variants, locations and severity levels. To limit selection bias, the cases to be modified within each anatomical-variant group were selected randomly. Furthermore, including both healthy and vasospasm-affected geometries allows the model to learn not only the baseline relationship between vascular anatomy and hemodynamics, but also how flow redistribution changes under pathological conditions, which are central to the intended clinical application of the surrogate.

Vasospasm was simulated by locally modifying the vessel geometry using Meshmixer (Autodesk Inc.), applying controlled reductions in vessel diameter to selected arterial segments. Affected vessels were chosen considering that vasospasm is usually a multi-vessel phenomenon and to represent the macrovascular territories most commonly considered in angiographic assessment, including vessels of both the anterior and posterior circulation, belonging to the internal carotid, middle cerebral, anterior cerebral, basilar, and posterior cerebral arterial territories [73]. The specific segments modified and the degree of narrowing in each patient are reported in the

Appendix [6].
To quantitatively assess the severity of the induced vasospasm, an automated algorithm was developed. The method takes as input the paired healthy and modified surface meshes, together with the corresponding vessel centerline. Radii are sampled along the centerline by computing the implicit distance to the vessel surface for both configurations. A radius preservation ratio is then calculated at each point as the ratio between the radius of the modified geometry and the radius of the corresponding healthy geometry, and a stenosis mask is derived to identify regions of significant narrowing. Three ranges of vessel shrinking (none, mild, moderate, or severe) were included, as reported in Figure 3, based on the distribution of the caliber reduction along the vessel, using percentile-based criteria. This percentile-based approach reduces sensitivity to isolated local fluctuations and allows the classification to reflect the global extent of narrowing rather than a single-point diameter reduction. While severity thresholds are not fully standardized and directly comparable [74], the adopted labels were chosen according to the most common criterion to ensure a consistent and objective classification of the severity of the vasospasm in all geometries. For the same reason, no unique reference distribution of the frequency of each severity class can be unambiguously derived from the literature. Therefore, the dataset was intentionally enriched in moderate and severe narrowing patterns relative to mild cases, as these configurations are expected to induce more pronounced hemodynamic alterations, making them both more clinically relevant and more informative for evaluating the surrogate model under hemodynamically challenging conditions.

Severity	Diameter preservation [%]
Mild	$\geq 75\%$
Moderate	$55\% \leq x < 75\%$
Severe	$\leq 55\%$

Figure 3: Definition of vasospasm narrowing.

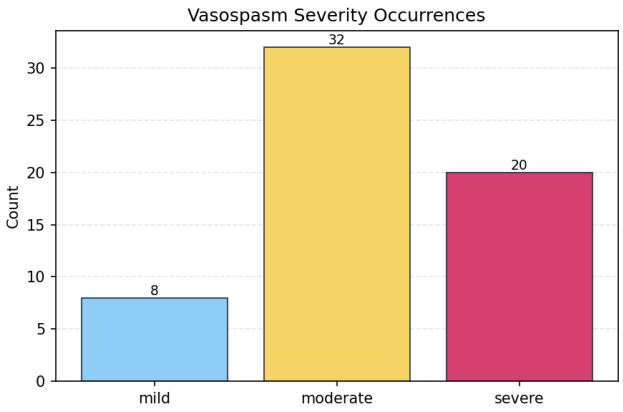

Figure 4: Distribution of vasospasm severity among 60 patients.

2.2. Computational Fluid Dynamics

Building upon this dataset, a large-scale CFD simulation pipeline was developed to generate corresponding hemodynamic data under different physiological and pathological conditions, which would constitute the ground truth for the AI-surrogate model.

2.2.1 Theoretical background

The mathematical foundation of any CFD solver is represented by the Navier-Stokes equations, which can be interpreted as a translation of Newton's Second Law of Motion to fluid dynamics. In fact, the total force acting on a fluid element, having an infinitesimal volume dV, can be written as (2):

$$F = \rho \, dV \frac{D\mathbf{v}}{Dt} \tag{2}$$

or alternatively, per unit volume:

$$\frac{F}{dV} = \rho \frac{D\mathbf{v}}{Dt} = \rho \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} \right) \tag{3}$$

where the term $\frac{\partial \mathbf{v}}{\partial t}$ represents the change of velocity in time, while $(\mathbf{v}\cdot\nabla)\mathbf{v}$ represents the convective acceleration, accounting for changes in velocity experienced by a fluid particle as it moves through spatial velocity gradients. Considering that the forces acting on a fluid can be distinguished in body forces, accounting for the effect of gravity, and surface forces, namely the divergence of the stress tensor, derived from the conversion of the traction vector to force per unit of volume, Equation 3 can be expressed as the Cauchy Momentum Equation

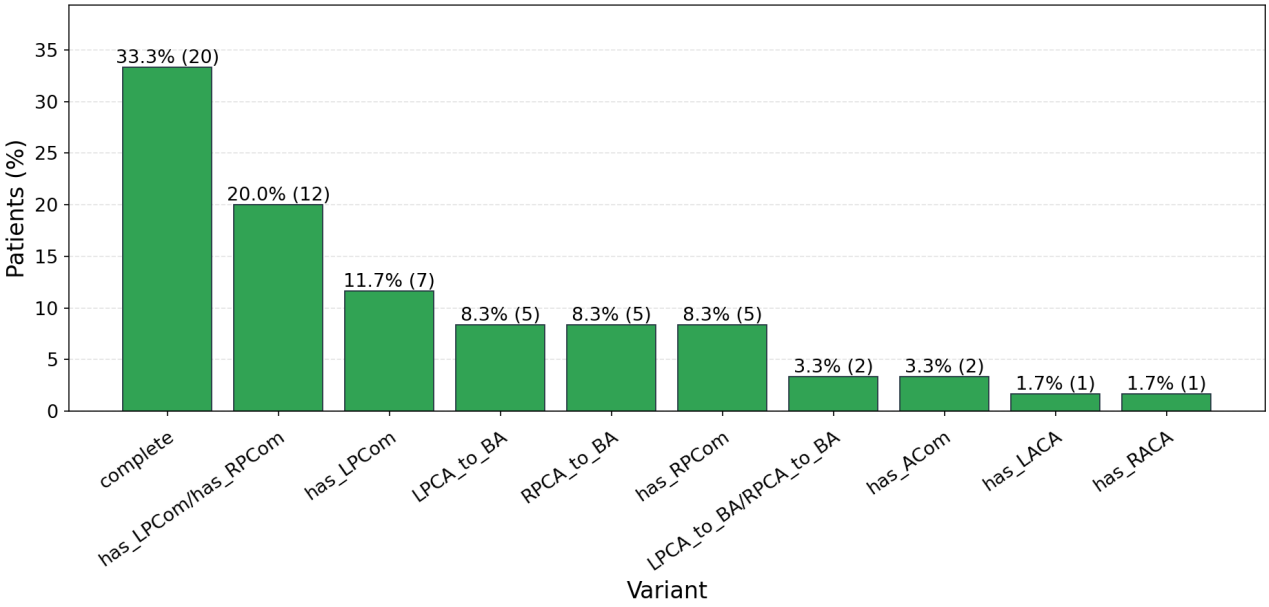

Figure 5: Distribution of anatomical variants among the 60 geometries selected for vasospasm induction. Bars report the percentage of vasospasm-induced patients belonging to each variant category, with the corresponding number of cases reported above each bar. Variant labels indicate the anatomical segment or connection that is absent or altered with respect to a complete Circle of Willis configuration: for example, *has_LPcom* denotes absence of the left posterior communicating artery, *LPCA_to_BA* denotes absence of the connection between the left posterior cerebral artery and the basilar artery, and combined labels indicate simultaneous alterations. The label *complete* identifies anatomies without the considered topological alterations.

(4):

$$\rho\left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v}\right) = \nabla\cdot\boldsymbol{\tau} + \rho\mathbf{g} \tag{4}$$

which gets coupled with the Conservation of Mass Equation (5), embodying the relationship between the change of density over time and the divergence of the velocity:

$$\frac{1}{\rho}\frac{\partial \rho}{\partial t} = -\nabla\cdot(\mathbf{v}) \tag{5}$$

As it is, this system, consisting of 4 equations in 10 unknowns, remains open, unless specific hypotheses concerning the properties of the fluid under analysis are provided. Specifically, vascular blood is usually modeled as a Newtonian fluid [39], meaning the viscous (deviatoric) stress tensor depends linearly on the rate of strain tensor, with the constant of proportionality given by the dynamic viscosity μ.
Thus, the viscous (deviatoric) contribution to the total stress (6), is substituted by the Constitutive Equation for Viscous Stress (7).

$$\boldsymbol{\tau} = -p\mathbf{I} + \boldsymbol{\sigma} \tag{6}$$

$$\boldsymbol{\sigma} = \mu\left(\nabla\mathbf{v} + (\nabla\mathbf{v})^T\right) \tag{7}$$

where ∇v describes how the velocity changes in space as a result of stretching, rotation and shearing, whereas the sum $\nabla\mathbf{v} + (\nabla\mathbf{v})^T$ allows to only keep the deformation component of the change of velocity in space, representing the rate of strain. Equation 4 can be therefore rewritten as:

$$\rho\left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v}\right) = \nabla\cdot\left(-p\mathbf{I} + \mu\left(\nabla\mathbf{v} + (\nabla\mathbf{v})^T\right)\right) + \rho\mathbf{g} \tag{8}$$

$$\rho\left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v}\right) = -\nabla p + \nabla\cdot\mu\left(\nabla\mathbf{v} + (\nabla\mathbf{v})^T\right) + \rho\mathbf{g} \tag{9}$$

The additional assumption of blood as an incompressible fluid, meaning the density ρ and viscosity μ do not change over time, implies:

$$\frac{\partial \rho}{\partial t} = 0 \tag{10}$$

which simplifies Equation 5 to:

$$\nabla \cdot (\mathbf{v}) = 0 \tag{11}$$

Considering $\nabla \cdot \nabla \mathbf{v} = \nabla^2 \mathbf{v}$ and $\nabla \cdot \nabla \mathbf{v}^T = \nabla(\nabla \cdot \mathbf{v}) = 0$, the viscous stress tensor becomes:

$$\nabla \cdot \mu \left(\nabla \mathbf{v} + (\nabla \mathbf{v})^T \right) = \mu \nabla^2 \mathbf{v} \tag{12}$$

This allows to obtain the Final Equation for Newtonian Incompressible Fluid:

$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{v} \right) = -\nabla p + \mu \nabla^2 \mathbf{v} + \rho \mathbf{g} \tag{13}$$

2.2.2 SimVascular Simulations

CFD simulations were conducted using SimVascular, a fully open-source software package, originally developed for cardiovascular modeling, allowing the solution of the Navier-Stokes equations in an arbitrary domain, usually confined to the volumetric vascular model.

Set-up

Prior to mesh generation, a series of preparatory steps were applied, employing VTK (Visualization Toolkit) [75], a publicly available code library for processing, analyzing and visualizing 3D data, and VMTK (Vascular Modeling Toolkit) [76], a framework built on top of VTK specializing in vascular modeling. First, endpoints are extracted from centerlines, in order to define the location of flow boundaries (inlets and outlets) and to provide the local vessel direction, enabling the identification of cross-sectional planes orthogonal to the centerline. Based on this information, the vascular surface was clipped to generate well-posed inlet and outlet faces, perpendicular to the local vessel axis, and converted to `vtp`. Crucially, this procedure allows physically meaningful boundary conditions to be imposed, as it ensures that velocity and pressure are defined correctly, to improve numerical stability of the CFD solutions and to guarantee consistency across the computation domain.
Since SimVascular operates on volumetric data, the original clipped surfaces were capped at all inlets and outlets in order to create a closed domain suitable for volumetric discretization. The resulting boundary faces were then identified and classified as wall, inlet, or outlet based on their geometric properties and spatial correspondence with predefined anatomical landmarks. Then, they were converted into a tetrahedral CFD mesh using SimVascular TetGen interface [77], with a target global edge size of 0.4 mm as the primary discretization parameter, allowing for a uniform and sufficiently detailed representation of the vascular geometry while maintaining a computational cost compatible with repeated simulation runs. To improve the numerical resolution of near-wall flow gradients, an internal boundary layer was generated along the vessel wall using four prismatic layers, with a thickness defined as 50% of the local edge size and a layer growth ratio of 0.7.
For each vascular geometry, three hemodynamic regimes were simulated with progressively increasing inlet flow rates, defined as (1) patient-specific baseline flow (Q_{base}), (2) low norepinephrine-induced flow increase ($Q_{\text{NE}}^{\text{low}}$), and (3) high norepinephrine-induced flow increase ($Q_{\text{NE}}^{\text{high}}$), mimicking the hemodynamic effect of noradrenaline administration. To introduce controlled variability, the total inflow was randomly sampled from three physiological intervals: 10000–12000 mm^3/s, 12000-13500 mm^3/s and 13500-15000 mm^3/s. The prescribed inflow was then distributed among the inlets according to anatomical considerations. Specifically, 25% of the total flow was assigned to the basilar artery (BA), while the remaining 75% was distributed between the left and right internal carotid arteries (ICAs), which is further split proportionally to their respective cross-sectional areas, ensuring consistency with physiological flow partitioning.
Regarding outlet boundary conditions, a Neumann resistance-based (R) Windkessel model was adopted to account for the effect of distal vasculature not explicitly included in the computational domain. In patient-specific CFD studies, outlet resistance values can be calibrated when additional hemodynamic measurements are available, for example by matching simulated outlet flow rates or vessel velocities to phase-contrast MRI, 4D-flow MRI, transcranial Doppler, or other clinical flow measurements through iterative forward simulations. However, such patient-specific measurements were not available for the imaging dataset used in this work, which was originally developed for segmentation and anatomical reconstruction tasks rather than for hemodynamic calibration. For this reason, resistance values were assigned based on literature data, while introducing controlled variability to enhance the representativeness of the dataset. In particular, the study by Alastruey et al. [78] provides quantitative estimates of cerebral microvascular resistance for an idealized geometry, from which a baseline value of 5.97 $Pa \cdot s/mm^3$ for the middle cerebral artery (MCA) is taken as reference. To account for inter-subject variability, resistance values were randomly perturbed within a $\pm 20\%$ interval. The corresponding resistances for the anterior cerebral artery (ACA) and posterior cerebral artery (PCA) were then scaled according to the territorial proportions reported by Vikström et al. [79], namely 34:59:78 for MCA:ACA:PCA. These proportions reflect the inverse relationship between vascular resistance and the volume of perfused brain tissue,

and have been shown to remain stable across individuals, including in hemodynamically altered conditions. Importantly, resistance values were kept constant across all hemodynamic regimes for a given patient, and between the healthy and vasospasm configurations. While in healthy subjects cerebral autoregulation typically leads to an increase in distal resistance in response to an elevation in the mean arterial pressure, thereby maintaining approximately constant cerebral blood flow, this mechanism is known to be impaired in vasospastic conditions (see Section 1.1.4). Consequently, maintaining fixed outlet resistances allows the simulations to capture the expected increase in flow associated with elevated pressure in the absence of effective autoregulation, reflecting the choice of explicitly modeling vasospasm dynamics and their impact.
Geometries for which mesh quality was insufficient, mainly due to the number of non-manifold edges, were excluded from the study. During this stage, some fundamental parameters were also set. Blood was modeled as an incompressible Newtonian fluid, with constant density of 1060 kg/m^3 and viscosity of $4 \cdot 10^{-3} kg/(m \cdot s)$. Vessel walls were assumed rigid, neglecting fluid–structure interaction effects. Simulations were performed under steady-state conditions, over 100 time steps with a time step size of $5 \cdot 10^{-4} s$. A maximum of 10 nonlinear iterations per time step was imposed to ensure convergence. Backflow stabilization was applied with a coefficient of 0.2 to improve numerical stability at the outlets.

Post Processing

The output of the simulations, mapped by the post-solver back onto the mesh to produce interpretable physical fields, were inspected using ParaView [80]. Simulations exhibiting non-physiological pressure or velocity values, or for which the solver failed to converge (residual $\geq 10^{-2}$), were systematically excluded from the final dataset, as they could introduce spurious patterns and degrade the model's ability to learn meaningful hemodynamic relationships, ultimately compromising generalization performance.
To characterize the CFD-derived pressure fields, two complementary analyses were performed. First, kernel density estimates were used to visualize the distribution of absolute centerline pressure across the three hemodynamic regimes. Distributions were computed separately for all geometries, healthy geometries, and vasospasm-induced geometries. To avoid over-representing geometries with a larger number of centerline nodes, each geometry was assigned equal total statistical weight. Second, the spatial pressure variation within each geometry was quantified using a robust pressure range, separating healthy and vasospasm-induced geometries to highlight their distinct hemodynamic behaviour. Since several geometries contained multiple inlets, a single inlet-to-outlet pressure drop could not be defined unambiguously. Therefore, pressure variation was described by comparing the 95th percentile and the 5th percentile of the centerline pressure values, whose difference was used as a measure of within-geometry pressure heterogeneity.
Velocity was analyzed directly from the 3D CFD solutions at the inlet and outlet cross-sectional cap surfaces. These surfaces corresponded to the anatomical segments used as model boundaries: the basilar artery (BA) and the right and left internal carotid arteries (R-ICA and L-ICA) at the inlets, and the right and left middle cerebral arteries (R-MCA and L-MCA), anterior cerebral arteries (R-ACA and L-ACA), and posterior cerebral arteries (R-PCA and L-PCA) at the outlets. The vector magnitude was computed from the Cartesian velocity components and converted from mm/s to cm/s. Considering TCD assesses the average of peak or near-peak velocity during a cardiac cycle and that CFD steady-state simulations already represent time-averaged hemodynamic conditions, the area-weighted 95th percentile of nodal velocity magnitude was considered the most appropriate quantity for comparison with Doppler-oriented reference values. Specifically, the 95th percentile was preferred to the absolute maximum to reduce sensitivity to isolated numerical peaks and local mesh-refinement effects, while the area weighting was applied to account for non-uniform surface discretization.

2.3. AI Model Development

2.3.1 Graph Generation

To represent the vascular network in a form suitable for graph-based learning, each anatomical configuration (uniquely defined by *patient id*, *anatomy* number indicator, *shrinked* flag) is modeled as a graph constructed from the centerline polydata.
Each node is defined as one sampled point of the 3D VTK centerline polydata. Near-duplicate points are removed by quantizing coordinates with a tolerance of $\geq 10^{-4}$ and collapsing points falling in the same bin into one node.
An edge is the connectivity relation between two consecutive points along a polyline cell of the centerline file.
If multiple centerline components exist and are treated as one anatomy (i.e. 3D Slicer was not able to generate a unique centerline path for a single anatomy), all points and edges are concatenated into a single global index space, then stitching edges are added from the leaf nodes of each appended component to the nearest node in the main component.
A trimming step is performed to remove spurious terminal centerline portions that lie outside the valid vessel

region after the clipping procedure.
To increase spatial resolution consistency and stability in message passing, the graph is then resampled by iterative shortest-edge collapse until the number of nodes approximately matched the total centerline length divided by the prescribed spacing of 0.75 mm.
Node and edge features described in the following paragraphs are selected based on prior knowledge about the problem.

Node features

The cross-sectional area A_i is formally defined from the local vessel radius stored at each centerline node. Under the hypothesis that the local vessel cross section can be approximated as circular, if r_i denotes the radius associated with node i, then the cross-sectional area at that node is defined as $A_i = \pi r_i^2$. The inclusion of this feature is motivated by the fundamental role of vessel caliber in determining hemodynamic behaviour. According to Poiseuille law (Equation 14), under idealized conditions, the pressure drop ΔP across a cylindrical vessel is calculated as:

$$\Delta P = RQ = \frac{8\mu L}{\pi r^4} Q \tag{14}$$

in which the fourth power inverse proportionality between the resistance R of the vasculature segment and its radius r suggests small geometric variations can lead to significant changes in flow and pressure distribution. While Poiseuille condition assumptions are not strictly satisfied globally in the Circle of Willis (e.g. fully developed velocity profile and straight cylindrical pipe), the strong non-linear dependence of vascular resistance on vessel radius (and therefore cross sectional area) still provides useful local physical intuition of hemodynamic behaviour.
Node types are determined based on graph topology according to four mutually-exclusive categories, encoded in a one-hot vector α_i: *inlets*, *outlets*, *junctions* and *branch* nodes. Specifically, leaf nodes are associated with a degree, defined as the number of incident undirected graph edges, of 1 and discriminated based on the distance to precomputed landmark coordinates. Junctions are then identified as nodes with degree strictly greater than 2. Any node that is neither inlet, nor outlet, nor junction is assigned the default type branch. Distinguishing between branch and junction nodes allows to explicate the different geometric properties of these regions [68]. In particular, the cross-sectional area defined along the centerline may exhibit abrupt variations within junctions, as slices can locally intersect multiple vessels simultaneously.
Each node is also assigned a tangent vector $t_i \in R^3$ describing the local orientation of the vascular centerline. Following VTMK implementation [76], the tangent is estimated from neighboring centerline points using finite differences along each polyline: forward difference at the first point, backward difference at the last point, and centered difference at interior points. The resulting vectors are normalized to unit norm. For points shared by multiple centerline segments (i.e. at bifurcations), tangent contributions are averaged and renormalized, yielding a single representative local direction per node. The node feature set also includes the fraction of the total inlet flow Q_i^{in}, assigned to each inlet node according to the same literature proportions applied for the CFD simulations (0.25 for the BA, 0.75 divided into R-ICA and L-ICA proportionally to their area), with zero value on all non-inlet nodes. This attribute acts as a normalized indicator of inflow distribution across multiple vascular entrances, thereby allowing the model to distinguish not only where the inflow boundaries are, but also their relative hemodynamic importance.
A resistance value R_i^{out} is also assigned to each node to encode the boundary conditions. This value corresponds to the outlet resistance prescribed in the CFD simulations for terminal nodes, and is set to zero elsewhere.
Additionally, two area-derived features were considered. Firstly, the area reduction ρ_i^A quantifies the local loss of lumen area relative to the corresponding healthy anatomy: for healthy geometries it is simply set to 1 everywhere, while for vasospasm-induced geometries, each node is matched to the closest node in the corresponding healthy graph, and the feature is computed as the ratio between the current area and the matched healthy area. This utility gives the model an explicit descriptor of vasospasm location and severity, highlighting where the pathological geometry differs from its healthy counterpart. Secondly, the area gradient g_i^A captures local geometric variations within the vessels themselves, by computing the relative difference between the current node area and the previous node area. Describing how abruptly the vessel caliber changes locally allows to detect geometric transitions, whether physiological or pathological, which may influence pressure losses and flow rerouting.

Node features

Node feature	Definition	Dimensionality
A_i	Cross-sectional area	1
α_i	Node type	4
t_i	Tangent to centerline	3
Q_i^{in}	Fraction of the inlet flow rate	1
R_i^{out}	Outlet resistance	1
ρ_i^A	Area reduction	1
g_i^A	Area gradient	1

Table 2: Node features, their definition and their dimensionality after one-hot encoding.

Edge features

Various geometry-derived features are incorporated to convey knowledge about the connectivity of the graph. However, due to the structure of the centerline representation, the resulting graph exhibits a predominantly chain-like topology, with most nodes connecting to only two neighbors. As highlighted by Pegolotti et al. [68], this limits the efficiency of information propagation across the graph, particularly between boundary nodes and interior nodes. In a graph-like network, this implies many iterations (encoded in message passing layers) are required for the boundary condition information driving hemodynamic behavior to effectively influence internal nodes. To alleviate this limitation, auxiliary edges are introduced, along with physical edges, to directly connect each interior node n_i to the boundary n_j that is closest in terms of graph distance, thus minimizing the shortest path length between n_i and n_j (computed via the Dijkstra's algorithm). This choice is motivated by the observation that hemodynamic influence is indeed primarily governed by local interactions and progressively attenuates with distance along the vascular network. These artificial connections therefore contribute to improving information transmission throughout the network while preserving its intrinsic structure. For the same reason, all edges are treated as bidirectional, enabling information exchange in both directions along the vasculature.
The attribute $d_{ij} \in R^3$ is defined as the normalized vector distance between node i and node j, encoding the edge direction and orientation in the 3-dimensional space. The scalar z_{ij} represents the length of the shortest path between node i and j, capturing the distance magnitude between linked nodes. Edge types, encoded as a one-hot vector β_{ij} are distinguished according to their different structural role in the graph: edges connecting branch nodes, edges incident to at least a junction node, and edges linking a node to its nearest inlet or to its nearest outlet.

Edge features

Edge feature	Definition	Dimensionality
d_{ij}	Normalized vector distance between node i and j	3
z_{ij}	Shortest path length between node i and j	1
β_{ij}	Edge type	4

Table 3: Edge features, their definition and their dimensionality after one-hot encoding.

Global context

The global context represents the global state shared by the whole vascular graph, concatenating a set of patient-level, anatomy-level, and regime-level descriptors.
The Mean Inlet Pressure (MIP) is computed as the average of the pressure at the inlets of the Circle of Willis (BA, L-ICA, R-ICA) in the current hemodynamic regime. This pressure is generally expected to be approximately equal to the systemic Mean Arterial Pressure (MAP), measured in supine position, since mean pressure varies only slightly along large proximal conduit arteries.
The anatomical variant mask v_g is meant to summarize the large-scale topology of the patient's vascular anatomy, independently of the local node-level graph geometry. Its purpose is to encode which major branches are present or absent with respect to a reference complete anatomy, possibly allowing combinatorial variations,

thereby informing the model of broader anatomical differences that may condition its predictions.
The degree of vasospasm severity, denoted as s_v, is one-hot encoded according to the criteria established in Section 2.1.2 (i.e. 0/mild/moderate/severe).
Additional features indicate if noradrenaline has been administered, defined according to whether a total inlet flow is currently at least 10% higher than the baseline, adopting the binary flag n_i, and how much the MIP has increased compared to the benchmark after drug administration, represented by the scalar ΔMIP.

Global context features

Global context feature	Definition	Dimensionality
MIP	Mean pressure at the inlets	1
v_g	Anatomical variant mask	7
s_v	Vasospasm severity	4
n_i	Noradrenaline administration binary flag	1
ΔMIP	MIP increase	1

Table 4: Global context features, their definition and their dimensionality after one-hot encoding.

Ground truth

The ground truth definition is a key step in the modeling pipeline, as it determines both the quantities to be predicted and the way in which the underlying hemodynamic behavior is represented. Specifically, in this work, the pressure variation with respect to the MIP and the fraction of the total input flow are selected to characterize the hemodynamic state of the vascular network.
For each graph node, a local cross-section is extracted from the CFD solution by intersecting the volumetric mesh with a plane orthogonal to the node tangent. Since pressure is approximately uniform across a vessel cross-section in the Circle of Willis, and its gradient is primarily axial (along the vessel), not radial, the nodal absolute pressure value is obtained as the mean of pressure values belonging to the slice, and then normalized into a relative pressure by subtracting the MIP and converting from Pascals to mmHg.
The flow is defined as the surface integral of the normal component of the velocity field through the cross-section (Equation 15).

$$Q_i = \int_{S_i} (v \cdot n_i) dS \tag{15}$$

This quantity is approximated, after triangulating the slice, as shown in Equation 16:

$$Q_i \approx \sum_k (\overline{v}_k \cdot n_i) A_k \tag{16}$$

where, for each triangle k, A_k is its area, $\overline{v}_k \cdot n_i$ the velocity averaged over the triangle vertices, and n_i is the unit slice normal. The flow target is then expressed as a normalized flow fraction by dividing the absolute nodal flow magnitude by the total inlet flow prescribed in the simulation.
The choice of flow and pressure as prediction targets is motivated by their fundamental role in describing vascular hemodynamics. Together, they provide a complete characterization of the system, where pressure represents the driving force of the circulation and flow quantifies the resulting transport of blood through the network. Their relationship is governed by the underlying physical laws of fluid dynamics (see Section 2.2.1), with pressure gradients determining flow distribution in interaction with vascular resistance and geometry. From a clinical perspective, both quantities are directly linked to tissue perfusion and the risk of ischemia. In particular, alterations in pressure and flow patterns are central to the pathophysiology of conditions such as cerebral vasospasm, where vessel narrowing modifies resistance and leads to complex redistribution of blood supply across the Circle of Willis.

2.3.2 GNN Architecture

Physics Informed Graph Neural Networks [56, 81], also denoted as PiGNNs, are deep learning architectures specifically designed for graph-structured data, formally defined as $G = (V, E)$, with V and E indicating the nodes and edges, respectively. These models operate through a message-passing mechanism [82], in which node and edge representations are iteratively updated by aggregating and transforming information from their local neighborhood, enabling the network to capture local and complex global relational dependencies within

the graph. Specifically, the architecture is composed of three main components, described below: encoder, processor, and decoder. A representation of such network is shown in Figure 6.

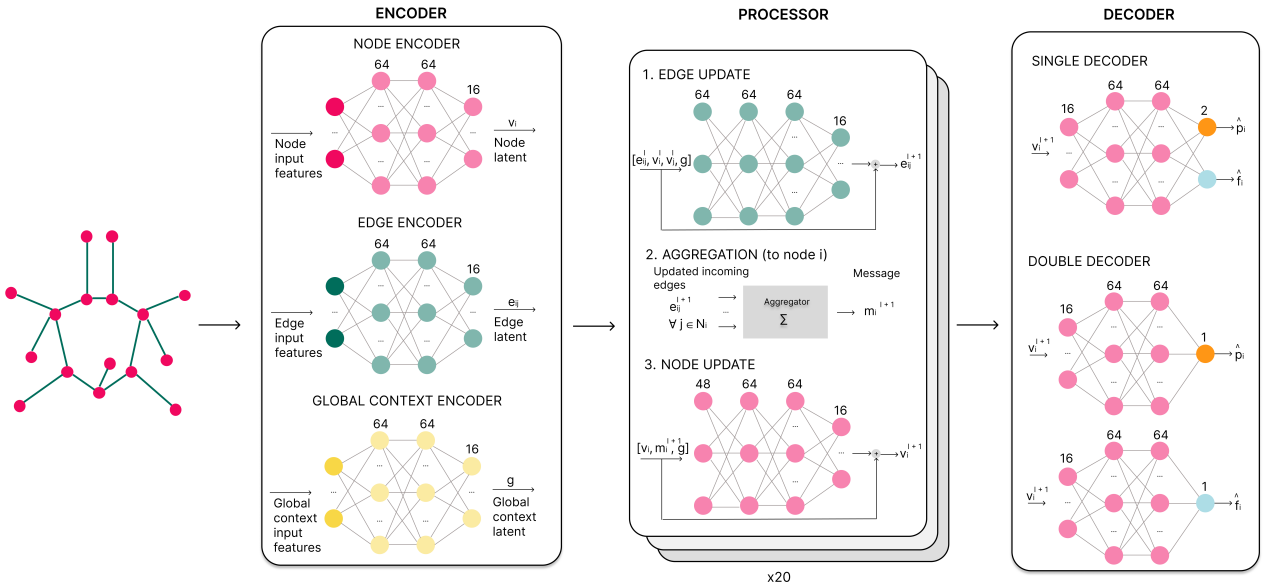

Figure 6: Schematic representation of the PiGNN architecture.

The input vascular network is represented as a graph with node, edge, and global-context features. Separate encoder networks map node, edge, and global inputs into latent representations. The processor then performs iterative message passing, consisting of edge updates, aggregation of incoming edge messages at each node, and node updates, repeated for 20 message-passing steps. Finally, the decoder maps the updated node latent representations to the target variables. Two decoder configurations are considered: a single shared decoder predicting both relative pressure and flow fraction, and a double-decoder architecture with separate output heads for pressure and flow.

Encoder

The encoder module computes embeddings for nodes, edges, and global context, projecting the raw input features into a latent space of fixed dimension $l \in N$, which generally differs from that of the original input space. This is achieved through separate encoder Multi-Layer Perceptrons (MLPs):

$$\begin{cases} v = \Psi_E^v(u) \\ e = \Psi_E^e(\xi) \\ g = \Psi_E^g(\gamma) \end{cases} \tag{17}$$

in which u and ξ denote the raw node and edge features, respectively, while γ represents the global context vector associated with each dataset sample (see Section 2.3.1).

Processor

The processor is composed of m message-passing layers M_k, iteratively updating the latent descriptions of nodes and edges according to a forwarding routine based on the structure of the graph. A single message-passing layer indeed allows each node to exchange information only with its immediate neighbors, locally. By stacking multiple layers, information progressively propagates across the graph, enabling each node to integrate signals originating from nodes that are up to m edges away, thereby increasing the receptive field. This is particularly important in the present context, due to the global coupling characterizing hemodynamic behavior, in which changes in boundary conditions and geometry in one branch affect pressure and flow patterns in distant regions of the network. The number of message-passing layers therefore constitutes a key hyperparameter, as it controls the extent of information propagation across the graph and directly impacts the model's ability to capture long-range dependencies. Based on a validation set evaluation, the value is set to 20, providing the best trade-off between predictive performance and training stability. Residual connections are implemented, so that each layer effectively learns a correction to the current representation rather than redefining it entirely. This stabilizes training and facilitates deeper propagation of information across the vascular networks.
Each message passing layer k can be divided into 3 steps. First, every directed edge $(i, j) \in E$ is updated

from the current edge embedding e_{ij}^k, the embeddings of the incident source v_i^k and destination nodes v_j^k, and the encoded global context g, by concatenating them and then applying a function Ψ_e^k (an MLP) according to Equation 18:

$$\widetilde{e}_{ij}^{k+1} = \Psi_e^k([e_{ij}^k, v_i^k, v_j^k, g]) \tag{18}$$

to which the residual connection is then applied:

$$e_{ij}^{k+1} = e_{ij}^k + \widetilde{e}_{ij}^{k+1} \tag{19}$$

Once all edges have been updated, each edge transmits its latent vector as a message to its destination node. From a physical perspective, nodes represent localized vascular entities along the centerline, each associated with geometric and hemodynamic attributes, while edges encode the connectivity of the vascular network, capturing both anatomical adjacency and the pathways through which flow and pressure variations propagate across the system. Within this framework, the aggregation of incoming messages at a node represents the cumulative effect of neighboring segments on its local hemodynamic state, enabling the model to approximate the underlying physical interactions governing blood flow. For each node $j \in V$, incoming messages are gathered using a permutation-invariant aggregator, as graph neighborhoods have no inherent ordering and therefore the update should not depend on how neighbors are indexed. In this project, the final node-wise message is defined according to Equation 20:

$$\overline{e}_j^{k+1} = \rho_{e\to v}(\{e_{ij} | (i \to j) \in E\}) = \sum_{i | e_{ij} \in E} e_{ij}^{k+1} \tag{20}$$

Finally, node states are updated using its previous state v_j^k, aggregated incoming messages, and global context, by applying a function Ψ_v^{k+1} according to Equation 21:

$$\widetilde{v}_j^{k+1} = \Psi_v^{k+1}([v_j^k, \overline{e}_j^{k+1}, g]) \tag{21}$$

from which the final representation is yielded by means of residual connections:

$$v_j^{k+1} = v_j^k + \widetilde{v}_j^{k+1} \tag{22}$$

Decoder

The output of the processor, namely the node embeddings produced by the final message passing layer, is given to the decoder, which maps the learned latent node embeddings back into the nodal output space of physically interpretable quantities, via a node-wise function Ψ_{out}^v, according to Equation 23:

$$\widehat{y}_i = \Psi_{out}^v(v_i^m) \tag{23}$$

2.3.3 Training strategies

All components of the architecture are trained jointly in an end-to-end fashion, allowing the model to learn task-specific representations directly from the data. Several strategies are adopted to improve generalization and prevent overfitting. Training was performed on an NVIDIA A30 GPU (24 GB VRAM), enabling efficient processing of large-scale graph data and accelerated model optimization. Training and validation curves were logged on `wandb`, for further systematic analysis and comparison.

Data splitting

Patients are split into training, validation, and test sets (60%, 20%, 20% respectively) at the patient level, rather than at the sample level, meaning that all rows associated with a given patient, including different anatomies, vasospasm conditions, and simulation settings, are assigned exclusively to a single subset. This prevents data leakage, which would otherwise occur if similar geometries or hemodynamic patterns from the same patient were present in both training and testing sets, leading to overly optimistic performance estimates. Patient-level splitting instead forces the model to predict on entirely unseen patients, which is a much more realistic measure of generalization. Furthermore, the split procedure aims to preserve anatomical-variant coverage across all subsets. This avoids biases due to rare configurations being present only in one split, which could cause performance estimates to become unstable and misleading.

Feature normalization

Inputs are standardized using statistics computed exclusively on the training split. For each selected node, edge, global-context, and target feature, the mean μ and standard deviation σ are estimated over all finite training values, and samples are normalized by z-score transformation (i.e. subtracting μ and dividing by σ). The only

exception is the pressure target, whose distribution is strongly skewed and heavy-tailed. In this case, robust scaling (i.e. subtracting the training-set median and dividing by the interquartile range) is adopted, reducing sensitivity to outliers and improving numerical stability during training. The same statistics are then used to normalize validation and test samples, ensuring consistency while preventing information leakage.

Network parametrization

The function Ψ in the encoder, processor, and decoder blocks are implemented as multilayer perceptrons [83], composed of linear layers, Layer Normalization, and LeakyReLU activations [84], followed by a final linear output layer. Nonlinear activation functions are essential to capture the complex non-additive relationships between vascular geometry, boundary conditions, and hemodynamic targets. Without nonlinearities, stacking several linear layers would instead still be equivalent to a single linear transformation, precluding the model from learning the multifaceted dependencies, thresholds, and interactions inherent to the problem.

Regularization

Several regularization techniques are employed to improve training stability and generalization. Early stopping is used by monitoring the validation metric and interrupting training when no improvement is observed for a predefined number of epochs (patience), restoring the best-performing model checkpoint. Weight decay is applied through Adam optimizer, acting as an L_2-type penalty on the model weights and helps discourage overly large parameters. Additionally, a cosine-annealing learning-rate scheduling is adopted, gradually reducing the learning rate over training epochs, which helps convergence and mitigates noisy updates in later stages.

Multi-task learning

The model is trained to jointly predict pressure and flow, leveraging the fact that both quantities are governed by the same underlying physical principles and depend on shared geometric and topological features of the vascular network. Two architectural configurations are explored: (i) a single-decoder model with two output neurons, jointly predicting pressure and flow, and (ii) a double-decoder model, in which two separate decoding branches are used, each specialized for one target. Two training strategies are investigated on both layouts. In the first approach, models are trained from scratch using random initialization (Xavier initialization). In the second approach, a pretraining and fine-tuning strategy is adopted. Specifically, a model is first trained to predict flow only, which represents a relatively more stable and easier learning task, as flow distribution across different hemodynamic regimes is more directly constrained by mass conservation and network topology. However, flow prediction is still governed by the same underlying geometric and topological information used for pressure inference. The learned weights are then used to initialize the full multi-task model, which is subsequently fine-tuned to jointly predict both flow and pressure using a reduced learning rate.

2.3.4 Loss Function

The loss function minimized during training is a combination of a supervised data-fitting term and a physics-informed constraint term, balanced by a weighting coefficient. For each batch of B graphs, the total loss is defined as:

$$\mathcal{L} = \frac{1}{B}\sum_{b=1}^{B}\left[\delta\,\mathcal{L}_{data}^{(b)} + (1-\delta)\,\mathcal{L}_{physics}^{(b)}\right] \tag{24}$$

with $\delta = 0.5$, assigning equal importance to the data-driven and physics-based contributions.
The data component enforces agreement between predictions and CFD-derived ground truth at the node level. For each graph b, the data fidelity term is:

$$\mathcal{L}_{data}^{(b)} = \frac{1}{N_b}\sum_{i=1}^{N_b}\alpha_i\left(w_p\,l_{Huber}(\widehat{p}_i, p_i) + w_f\,(\widehat{q}_i - q_i)^2\right) \tag{25}$$

in which N_b is the number of nodes in graph b, $\widehat{p}_i$ and p_i denote the predicted and reference relative pressure, and $\widehat{q}_i$ and q_i indicate the predicted and reference flow fraction at node i. The coefficient α_i is a node-wise weighting factor that assigns higher importance to boundary nodes (inlets and outlets) by setting $\alpha_i = 10$ for boundary nodes and $\alpha_i = 1$ otherwise. This choice reflects the fact that boundary conditions strongly influence the global hemodynamic solution, and inaccuracies at these locations may propagate throughout the network [68]. The weights w_p and w_f (both set to 1 in this work) allow balancing the contribution of pressure and flow errors, tuning the importance of the target based on their relative difficulty in prediction. The pressure error is modeled using the Huber loss:

$$l_{Huber}(\widehat{p}_i, p_i) = \begin{cases} \frac{1}{2}(\hat{p}_i - p_i)^2, & |\hat{p}_i - p_i| \le \Delta_H, \\ \Delta_H\left(|\hat{p}_i - p_i| - \frac{1}{2}\Delta_H\right), & |\hat{p}_i - p_i| > \Delta_H, \end{cases} \tag{26}$$

in which $\Delta_H = 1$ defines the transition threshold between quadratic and linear regimes. For small errors, the loss behaves as a mean squared error, ensuring smooth optimization, while for large deviations it becomes linear, reducing sensitivity to outliers. This is particularly suitable for pressure prediction, where localized large errors may occur due to numerical or geometric variability.
The physics term encourages consistency with mass conservation by enforcing that predicted flow fractions satisfy global balance constraints. For each graph b, it is computed as:

$$\mathcal{L}_{physics}^{(b)} = \lambda_{physics} \left[\left(q_{tot} - \sum_{i \in \mathcal{O}_b} \widehat{q}_i \right)^2 + \gamma_{inlet} \left(q_{tot} - \sum_{i \in \mathcal{I}_b} \widehat{q}_i \right)^2 \right] \tag{27}$$

with $\mathcal{O}_b$ and $\mathcal{I}_b$ denoting the sets of outlet and inlet nodes, respectively, and q_{tot} indicating the total flow fraction entering (and, by conservation, leaving) the geometry. This variable is generally equal to 1, except for the variant with both PComs absent, where two separate graphs are defined and the flow is redistributed between them as described in Section 2.2. The coefficient $\lambda_{physics}$ controls the overall contribution of the physics-based term (set to 1 in this work), while γ_{inlet} regulates the enforcement of inlet consistency, set to 1 to explicitly enforce mass conservation at both inlet and outlet boundaries, 0 otherwise.
Combining Equations 24, 25, 26, 27, the final loss formulation can be obtained as:

$$\begin{aligned} \mathcal{L} = \frac{1}{B} \sum_{b=1}^{B} \Bigg[& \delta \frac{1}{N_b} \sum_{i=1}^{N_b} \alpha_i \left(w_p l_{Huber}(\hat{p}_i, p_i) + w_f (\hat{q}_i - q_i)^2 \right) \\ & + (1-\delta) \lambda_{physics} \left(\left(q_{tot} - \sum_{i \in \mathcal{O}_b} \hat{q}_i \right)^2 + \gamma_{inlet} \left(q_{tot} - \sum_{i \in \mathcal{I}_b} \hat{q}_i \right)^2 \right) \Bigg]. \end{aligned} \tag{28}$$

guiding the model toward physically plausible predictions while retaining flexibility to learn from data, effectively combining data-driven learning with domain-specific constraints.

2.3.5 Evaluation Metrics

The performance of the proposed model is assessed using quantitative metrics designed to evaluate the agreement between predicted and CFD-derived hemodynamic quantities, capturing both local accuracy and global consistency. All metrics are first computed on a per-graph basis and subsequently averaged across graphs, ensuring that larger graphs do not disproportionately influence the overall evaluation.
The Pressure Relative Error (PRE) is defined, for graph g with N_g nodes, as:

$$\mathrm{PRE}^{(g)} = \sqrt{\frac{\sum_{i=1}^{N_g} \left(p_{i,g}^{\mathrm{abs}} - \hat{p}_{i,g}^{\mathrm{abs}} \right)^2}{\sum_{i=1}^{N_g} (p_{i,g}^{\mathrm{abs}})^2}} = \sqrt{\frac{\sum_{i=1}^{N_g} \left(p_{i,g}^{\mathrm{rel}} - \hat{p}_{i,g}^{\mathrm{rel}} \right)^2}{\sum_{i=1}^{N_g} \left(\mathrm{MIP}_g + p_{i,g}^{\mathrm{rel}} \right)^2}} \tag{29}$$

which, accounting for all the graphs in the test set, becomes:

$$\overline{\mathrm{PRE}}_{\text{graph-wise}} = \frac{1}{G} \sum_{g=1}^{G} \mathrm{PRE}^{(g)} \tag{30}$$

This formulation corresponds to a relative root-mean-square error normalized by the global pressure scale of each graph. The numerator quantifies the prediction error, while the denominator rescales it using the overall magnitude of the pressure field, reconstructed from the relative pressure through the addition of the MIP term. This normalization ensures that errors are comparable across cases with different pressure levels and prevents artificial inflation in regions where relative pressure values are close to zero. Moreover, the metric reflects the physical significance of the error, weighting it according to the underlying pressure scale rather than treating all deviations uniformly.
The flow prediction is evaluated using the Mean Absolute Error (MAE), defined, for graph g with N_g nodes, as:

$$\mathrm{MAE}_{\mathrm{flow}}^{(g)} = \frac{1}{N_g} \sum_{i=1}^{N_g} |\hat{q}_i - q_i| \,, \tag{31}$$

with $\hat{q}_i$ and q_i indicating the predicted and ground-truth flow fraction at node i, respectively. The graph-wise average is then computed as:

$$\mathrm{MAE}_{\mathrm{flow}} = \frac{1}{G} \sum_{g=1}^{G} \mathrm{MAE}_{\mathrm{flow}}^{(g)} \tag{32}$$

The use of an absolute error for flow is motivated by the fact that flow is expressed as a normalized fraction, making deviations directly interpretable. In this context, relative error measures would be less stable and harder to interpret, particularly for small flow values, where even minor absolute discrepancies could lead to inflated relative errors [66].
Finally, while some previous works restrict evaluation to a subset of node types [68], all nodes, including inlets, outlets, and junctions, are here considered in the metric computation. This choice provides a more comprehensive assessment of model performance and is particularly relevant in the present application, where boundary and junction regions play a critical role in determining the overall hemodynamic behavior.
Overall, the selected metrics ensure a balanced evaluation of predictive accuracy, capturing both point-wise discrepancies and physically meaningful global behavior.
In addition to prediction accuracy, the physical consistency of the model is evaluated through a mass conservation metric. For each graph g, the set of outlet nodes is identified as $\mathcal{O}g = i \mid$ node i is an outlet. The predicted outlet flow fractions are then summed as:

$$Q_{out}^{(g)} = \sum_{i \in \mathcal{O}g} \hat{q}_i \tag{33}$$

Ideally, due to conservation of mass, this quantity should be equal to the total input fraction Q_{in}. The signed residual is therefore defined as:

$$r_{out}^{(g)} = Q_{in}^{(g)} - Q_{out}^{(g)} \tag{34}$$

and the corresponding absolute conservation error is given by:

$$e_{\text{out}}^{(g)} = \left| Q_{in}^{(g)} - \sum_{i \in \mathcal{O}_g} \hat{q}_i \right|. \tag{35}$$

A graph is considered to violate mass conservation if $e_{\text{out}}^{(g)} > 0.05$, a tolerance threshold adopted to account for numerical and modeling uncertainties. The percentage of violating graphs is then reported as an additional evaluation metric. Similarly, $e_{\text{in}}^{(g)}$ is defined for inlet mass conservation.

2.3.6 Ablation study

To investigate the contribution of different input features to model performance, an ablation study was performed. In addition, this analysis was designed to evaluate whether input quantities that are not directly available in routine clinical practice could be removed without compromising model accuracy or physical consistency, with the aim of improving the applicability of the model within a medical workflow.
Starting from a predefined baseline architecture and training strategy, for every ablation setting, selected features are systematically removed by excluding the corresponding entries from the node, edge, or global-context feature vectors. Each ablated model is then retrained using the same train/validation/test split as the full-feature model. Performance is evaluated using the same metrics, namely pressure relative error (PRE), flow mean absolute error (MAE), and mass-conservation residuals, as defined in Section 2.3.5. Specifically, the percentage improvement of ablation experiment a relative to the full-feature baseline, for metric m, is computed as the relative reduction of the metric value with respect to the baseline:

$$\text{Improvement}_m^{(a)}(\%) = 100 \cdot \frac{M_m^{(\text{full})} - M_m^{(a)}}{M_m^{(\text{full})}} \tag{36}$$

The input variables included in the analysis range from geometric quantities and boundary-condition descriptors encoded in the node and edge feature vectors (such as resistance, cross-sectional area, area gradient, area reduction, centerline tangent and inlet flow fraction), clinical information from the global context vector (e.g. MIP, pressure increase, noradrenaline assumption) and categorical graph-structural labels (e.g. node-type and edge-type encodings).
Both individual and grouped ablations are considered to assess possible redundancy or interaction effects between correlated inputs. In this context, a feature interaction indicates that the simultaneous removal of multiple features produces a different effect than expected from the removal of each feature independently, suggesting that the corresponding inputs may encode overlapping or complementary information. In particular, combined ablations were included when features were physically related, potentially redundant, or jointly relevant to clinical applicability. For example, resistance and inlet flow fraction are both boundary-condition-related quantities not routinely available in clinical practice; distance and relative position both encode spatial edge information; area gradient and area reduction describe local caliber changes; pressure increase and noradrenaline administration together represent the imposed hemodynamic regime.

2.3.7 Benchmarking

To further assess the robustness of the final model, an external benchmark was performed on four patient-specific vascular geometries obtained from the clinical cohort of Policlinico di Milano. These cases were reconstructed from computed tomography angiography (CTA) images using a semi-automatic segmentation workflow. Differently from the geometries used for model training, the external reconstructions included a more extended vascular domain, comprising the Circle of Willis together with upstream carotid and vertebral arteries and additional distal branches. An example of the anatomical difference between the training domain and the external benchmark geometries is shown in Figure 7.
For each external case, centerlines were extracted using 3D Slicer and used to generate the corresponding graph representation. The hemodynamic reference solution was obtained from patient-specific CFD simulations. Inlet flow rates were estimated from available clinical information, including height, weight, sex, heart rate, and systemic arterial pressure measurements. Outlet resistance values were then calibrated so that the simulated velocities matched the velocities measured in the middle cerebral arteries by transcranial Doppler ultrasound. This calibration procedure provided patient-specific boundary conditions for the CFD simulations used as benchmark ground truth.

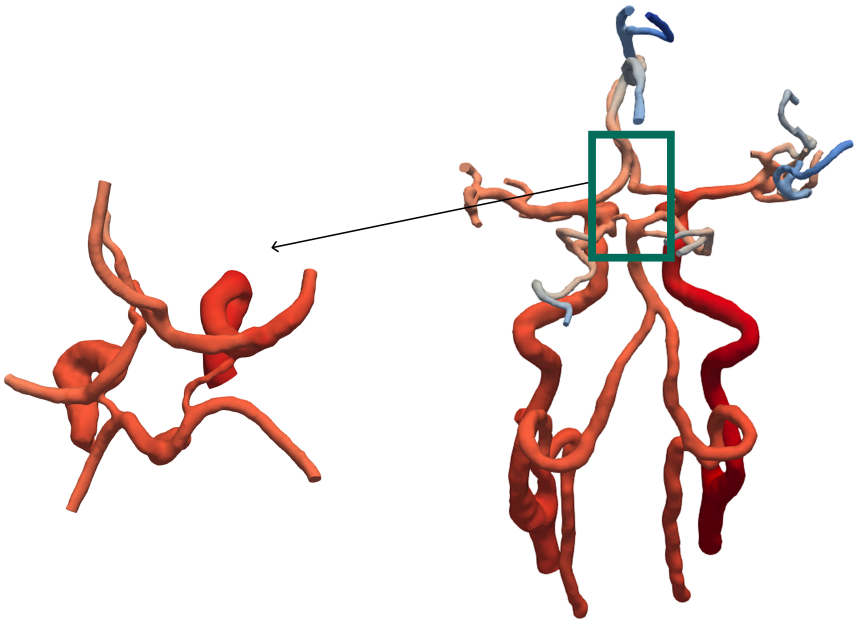

Figure 7: Example of the anatomical difference between the original training domain (left) and the external clinical benchmark geometries (right). The highlighted region indicates the approximate Circle of Willis domain represented in the training dataset, while the external reconstruction includes additional upstream carotid and vertebral arteries and further distal branches.

3. Results

3.1. CFD analysis

The CFD simulations were analyzed in terms of the primary hemodynamic quantities directly obtained from SimVascular, namely pressure and velocity fields, in order to characterize the simulated hemodynamic dataset across the three hemodynamic regimes and to verify the expected differences between healthy and vasospasm-induced configurations. The pressure and velocity fields were subsequently processed as described in Section 2.3.1 to obtain the graph-level targets used for model training.
The distribution of absolute pressure values across hemodynamic regimes is reported in Figure 8, where the three hemodynamic regimes correspond to progressively raised inlet-flow regimes used to mimic increasing pressure-support conditions. Considering both all geometries together and healthy and vasospasm-induced geometries separately, the pressure distributions gradually shift toward higher values from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, as reported in the Appendix [6], Table 10 and Table 11. Beyond this global pressure shift, the robust within-geometry pressure range was analyzed to evaluate spatial pressure heterogeneity within each vascular network. As shown in

Figure 9, vasospasm-induced geometries exhibit broader distributions and generally higher values, as numerically detailed in the Appendix [6], Table 12. Overall, the CFD simulations show a regime-dependent increase in absolute pressure and a larger within-geometry pressure variation in vasospasm-induced configurations.

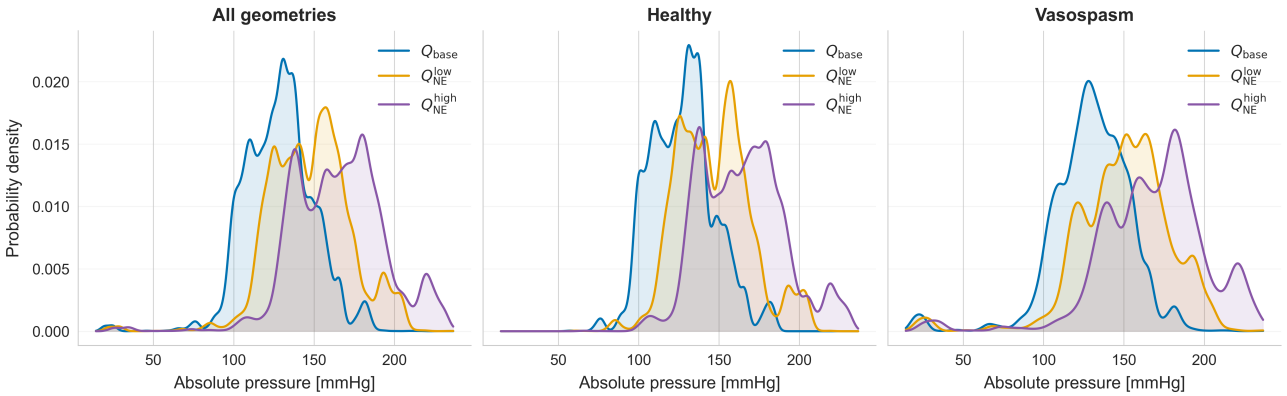

Figure 8: Geometry-balanced distributions of absolute centerline pressure across the three hemodynamic regimes, corresponding to progressively increased inlet-flow regimes. Kernel density estimates are shown for all geometries, healthy geometries, and vasospasm-induced geometries separately. Each geometry was assigned equal total statistical weight, so that geometries containing a larger number of centerline nodes did not disproportionately influence the distributions. The distributions show a progressive shift toward higher absolute pressure values from Q_{base} to $Q_{\text{NE}}^{\text{high}}$.

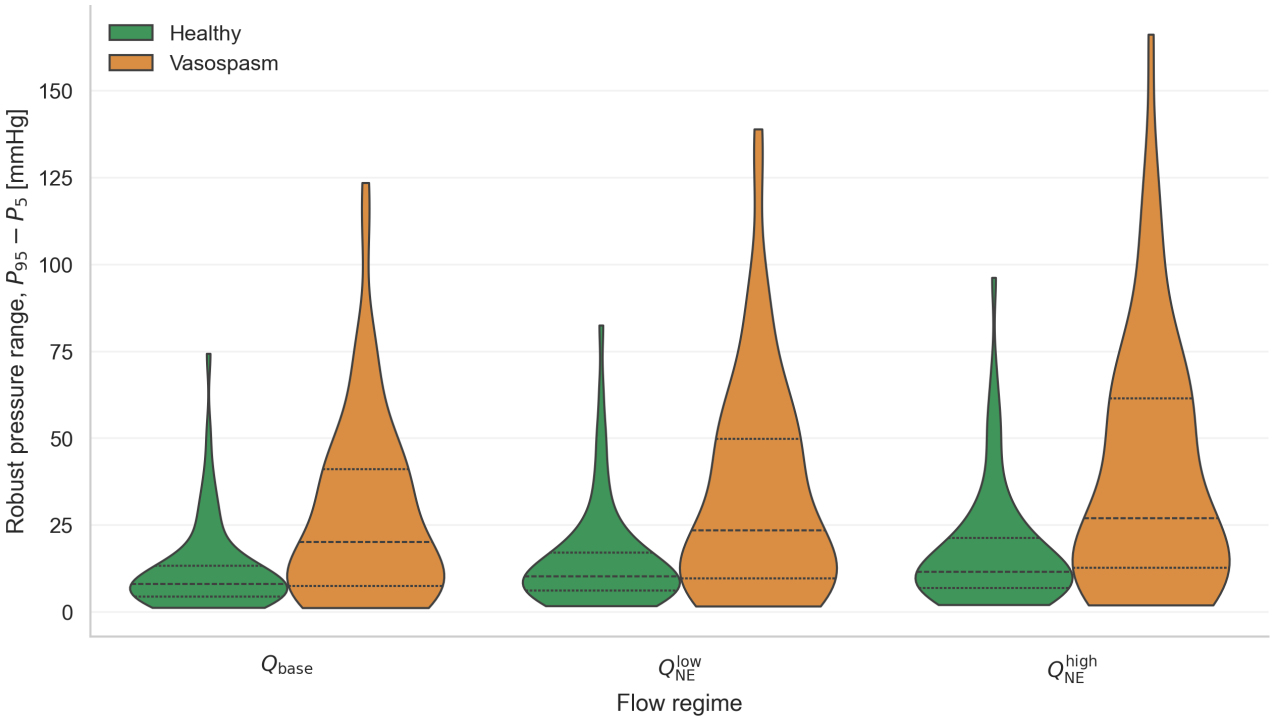

Figure 9: Distribution of the robust within-geometry pressure range for healthy and vasospasm-induced geometries across the three hemodynamic regimes. The robust pressure range was computed as the difference between the 95th and 5th percentiles of centerline pressure values within each geometry to capture the pressure variation across most of the vascular network. Internal dashed lines indicate the quartiles of each distribution. Vasospasm-induced geometries show broader distributions and higher pressure ranges than the corresponding healthy geometries, particularly in $Q_{\text{NE}}^{\text{high}}$.

Regarding velocity, the inlet and outlet velocity distributions are reported in Appendix [6], Table 13, while the corresponding inlet and outlet distributions are shown in Appendix [6], Figures 18–21. As described in Section 2.2.2, velocity was analyzed using the area-weighted 95th-percentile velocity magnitude at the inlet and outlet

boundary sections.
As illustrated by Figures 18 and 19, inlet near-peak velocities are mostly concentrated within a comparable range across vessels and hemodynamic regimes, with the main density generally remaining below approximately 100 cm/s and only a small number of cases show higher-velocity tails, particularly in vasospasm-induced configurations. When compared to healthy geometries, vasospasm-induced configurations display a shift towards higher values and broader tails. This behavior is evident for the R-ICA and L-ICA, whereas it is more moderate for the BA, whose distributions remain largely overlapping between healthy and vasospasm-induced cases. The regime-wise comparison reported in Figure 19 reveals that, for both healthy and vasospasm-induced configurations, inlet velocity distributions shift slightly toward higher values from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, although without abrupt or disproportionate changes in their distributions.
The same general dynamics are even more evident at the outlets, as observed from Figures 20 and 21. In healthy geometries, velocity distributions are relatively compact, particularly for the MCA and ACA outlets. In vasospasm-induced configurations, the distributions become broader and show more pronounced high-velocity tails, especially in the R-MCA, L-MCA, R-ACA, and L-ACA. The PCA outlets show more overlapping distributions between healthy and vasospasm-induced cases instead. The regime-wise comparison in Figure 21 also exhibits a moderate shift toward higher outlet velocities from Q_{base} to $Q_{\text{NE}}^{\text{high}}$. This trend is clearly visible in healthy geometries and is also present in vasospasm-induced cases, although in the latter the broader velocity distribution associated with vessel narrowing partly masks the regime-dependent shift.
Overall, both inlet and outlet analyses show that the simulated hemodynamic regimes produce a controlled increase in near-peak velocities, while vasospasm-induced narrowing increases velocity variability and introduces higher-velocity tails, especially in anterior-circulation outlet branches.

3.2. Comparison of decoder architectures and training strategies

As explained in Section 2.3.3, two multi-task decoder layouts were considered: a single-decoder architecture, in which pressure and flow are predicted through the same output branch, and a double-decoder architecture, in which the two targets are predicted through separate decoder branches. Each architecture was trained both from scratch and through fine tuning from a flow-only model. This comparison was designed to evaluate two modelling choices: the decoder layout used for multitask pressure-flow prediction, and the training strategy used to initialize the multitask model.
Since model capacity, latent dimensionality, and depth of the message passing directly affect the ability of the model to capture both local and long-range vascular interactions, these architectural and training parameters were selected through a tuning process, in which multiple combinations were empirically evaluated to identify the setup yielding the best validation performance. The chosen hyperparameters, reported in Table 5, were then adopted across all experiments to ensure a fair comparison between the different model configurations.
The MLP encoder, processor, and decoder modules have a hidden dimension of 64 and a latent space of dimension 16, while the processor is composed of 20 message-passing layers. Training is performed with a batch size of 8 for up to 500 epochs, using an initial learning rate of 10^{-3} and a reduced rate of 10^{-4} during fine-tuning. Regularization is introduced through weight decay (10^{-5}), learning rate scheduling, and early stopping with a patience of 20 epochs.

Hyperparameters

Category	Hyperparameter	Value
MLP	Hidden dimension	64
MLP	Latent dimension	16
MLP	Hidden layers	2
MLP	Message-passing layers	20
Training	Batch size	8
Training	Maximum epochs	500
Training	Weight decay	1e-5
Training	Learning rate (LR)	1e-3
Training	Fine-tuning LR	1e-4
Training	LR decay factor	1e-3
Training	Minimum LR	1e-6
Training	Early-stopping patience	20

Table 5: Summary of the hyperparameters adopted for the model, including architectural design choices and training settings.

Table 6 reports the performance of the different model configurations. The flow-only model achieves a flow MAE of $3.5 \cdot 10^{-2}$, but does not provide pressure predictions. Among the multitask models trained from scratch, the single- and double-decoder architectures achieve comparable PRE values ($4.4 \cdot 10^{-2}$) and similar flow MAE ($5.1 \cdot 10^{-2}$ and $5.2 \cdot 10^{-2}$ respectively). However, both configurations exhibit high outlet mass-conservation violation rates, equal to 36.7% for the single decoder and 61.7% for the double decoder.
Given the high outlet violation rates, fine-tuning was evaluated to assess whether initialization from a flow-specialized representation affected the performance and physical consistency of the multitask architectures. The fine-tuned single-decoder model attains a PRE of $4.5 \cdot 10^{-2}$ and reduces the outlet violation rate to 5.83%, although its flow MAE increases to $6.7 \cdot 10^{-2}$. The fine-tuned double-decoder model reaches comparable pressure and flow accuracy compared to the trained-from-scratch models, with a PRE of $4.7 \cdot 10^{-2}$ and a flow MAE of $5.2 \cdot 10^{-2}$, while eliminating outlet mass-conservation violations.
However, all configurations exhibit poor performance in enforcing inlet mass conservation, with high violation rates across both single- and double-decoder architectures, regardless of the training strategy.

Model comparison

Category	Training	PRE (10^{-2}/%)	Flow MAE (10^{-2}/%)	%Outlet Violations
Flow only	From scratch	N/A	3.5	0.0
Single decoder	From scratch	4.4	5.1	36.7
Double decoder	From scratch	4.4	5.2	61.7
Single decoder	Fine-tuned	4.5	6.7	5.83
Double decoder	**Fine-tuned**	**4.7**	**5.2**	**0.0**

Table 6: Comparison of the different model configurations in terms of pressure relative error (PRE), flow mean absolute error (MAE), and mass conservation violations at the outlets, used to identify the architecture giving the best balance among multitask prediction and outlet conservation (highlighted in bold). PRE and flow MAE values are reported in units of 10^{-2}, or, equivalently, on a percentage scale.

Overall, the double-decoder fine-tuned model yields the best balance between pressure accuracy, flow prediction, and physical consistency, and is therefore selected as the baseline architecture for subsequent analysis. Nevertheless, one important limitation of this configuration is that it still relied on the full input feature set, including

outlet resistance and inlet flow fraction. These quantities are commonly available in simulation-based settings, but they are not directly accessible in a routine clinical workflow and would therefore need to be prescribed or estimated a priori, undermining the reliability and translational applicability of the framework.

3.3. Ablation study

Given the motivation described above, an ablation study was performed, not only to assess the contribution of different input feature groups, but also to evaluate whether clinically unavailable inputs could be removed while preserving predictive accuracy and physical consistency. Each ablation consists of removing one feature or group of features from the full input representation and retraining the model, in order to evaluate how the missing information affects predictive accuracy and physical consistency.

Figure 10 reports the percentage variation in pressure relative error (PRE) and flow mean absolute error (flow MAE) across the three hemodynamic regimes for the different ablation experiments. Positive values indicate an improvement with respect to the full-feature baseline, whereas negative values indicate a degradation in performance. Figure 11 illustrates the corresponding outlet mass-conservation residuals, expressed both as mean absolute residual and as percentage of violating samples, according to the definition introduced in Section 2.3.5.

The effect of feature ablation on pressure prediction is strongly dependent on the three simulated hemodynamic regimes with progressively increasing inlet flow rates, in order to mimic noradrenaline-related pressure support. In Q_{base}, representing a physiological patient-specific baseline flow regime, most ablations lead to a degradation in PRE, with the largest negative variations observed when removing variant encoding, resistance together with area, cross-sectional area, MIP, and pressure-related global-context information. As flow starts to increase in $Q_{\text{NE}}^{\text{low}}$, the magnitude of these degradations is generally reduced, while some ablations begin to show small performance gains. In $Q_{\text{NE}}^{\text{high}}$, corresponding to the highest flow regime, all ablation settings lead to a favorable variation in PRE relative to the full-feature model evaluated in $Q_{\text{NE}}^{\text{high}}$. The largest improvement is observed when removing edge type, reaching approximately 7%, followed by vasospasm severity, node type, inlet flow fraction, and resistance together with inlet flow fraction, all producing reductions in error between approximately 4% and 5%. In particular, the simultaneous removal of resistance and inlet flow fraction produces only a limited degradation in Q_{base} , equal to 0.59%, while improving PRE relative to the regime-matched full-feature baseline in $Q_{\text{NE}}^{\text{low}}$ and $Q_{\text{NE}}^{\text{high}}$.

Flow prediction shows a more consistent behavior across ablation settings. Differently from PRE, flow MAE improves for all feature removals and across all hemodynamic regimes, suggesting that the information required to infer flow redistribution is distributed across multiple partially redundant input features. The largest improvements are observed when removing distance together with relative position, with variations of approximately 20% across Q_{base} and $Q_{\text{NE}}^{\text{low}}$ and exceeding 20% in $Q_{\text{NE}}^{\text{high}}$. Other feature groups also lead to substantial positive variations, generally above 10%, including pressure increase together with noradrenaline intake, area gradient together with area reduction, edge type, node type, variant encoding, and resistance together with inlet flow fraction. The smallest gains are observed when removing resistance together with area and, to a lesser extent, tangent and MIP. Furthermore, the improvement in flow MAE is generally more pronounced in $Q_{\text{NE}}^{\text{high}}$.

Outlet mass conservation is also affected by the selected input feature set. The full-feature model presents the lowest outlet violation rate, with no violating samples under the adopted threshold (see Section 2.3.5) and a mean outlet residual of approximately 0.013. Among the ablated models, the simultaneous removal of resistance and inlet flow fraction preserves outlet conservation close to the full-feature baseline, with a low mean residual and a violation rate of approximately 0.84%. Similarly low violation rates are observed when removing resistance together with area, tangent, and vasospasm severity. In contrast, removing MIP, inlet flow fraction, edge type, and area gradient together with area reduction leads to higher outlet conservation failure rates, reaching approximately 22%, 19%, 15%, and 13%, respectively.

In addition to differences in mean residuals and violation rates, several ablation settings exhibit substantial variability in outlet mass-conservation residuals across test graphs, as indicated by the larger standard deviation bars in Figure 11. This heterogeneity is particularly evident for the removal of MIP, inlet flow fraction, edge type, area gradient together with area reduction, and pressure increase together with noradrenaline intake. In contrast, the full-feature model and the ablation removing resistance together with inlet flow fraction show lower residual dispersion.

Interestingly, some ablation settings show improved PRE and/or flow MAE while simultaneously increasing outlet mass-conservation violations, suggesting that pointwise predictive accuracy and outlet physical consistency capture complementary aspects of model performance. Consequently, ablations that improved error metrics but degraded outlet conservation were not selected as final candidates, since the final feature set was required to preserve both numerical accuracy and physiologically plausible flow redistribution.

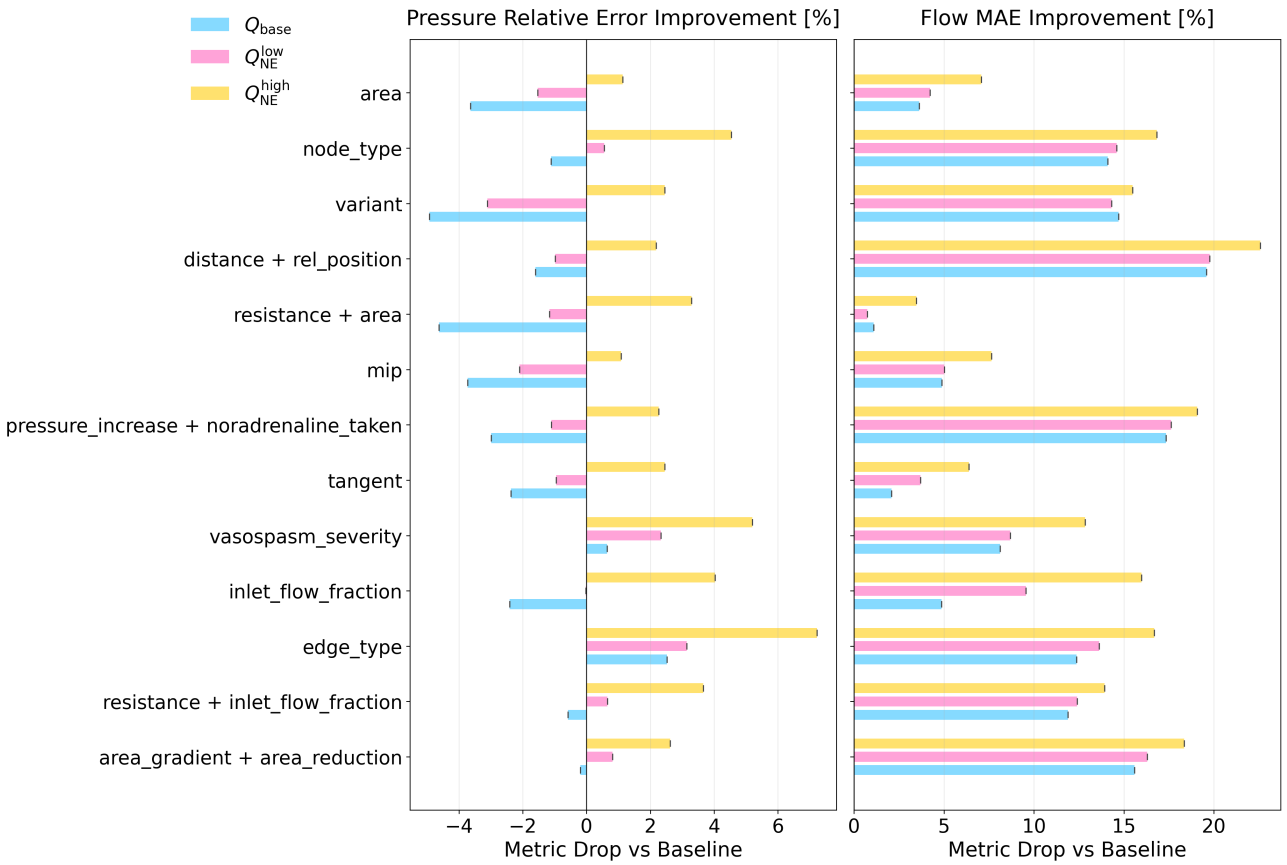

Figure 10: Regime-wise improvement in pressure relative error (PRE) and flow mean absolute error (MAE) resulting from the ablation of different feature groups, expressed as percentage change with respect to the baseline (full-feature) model.

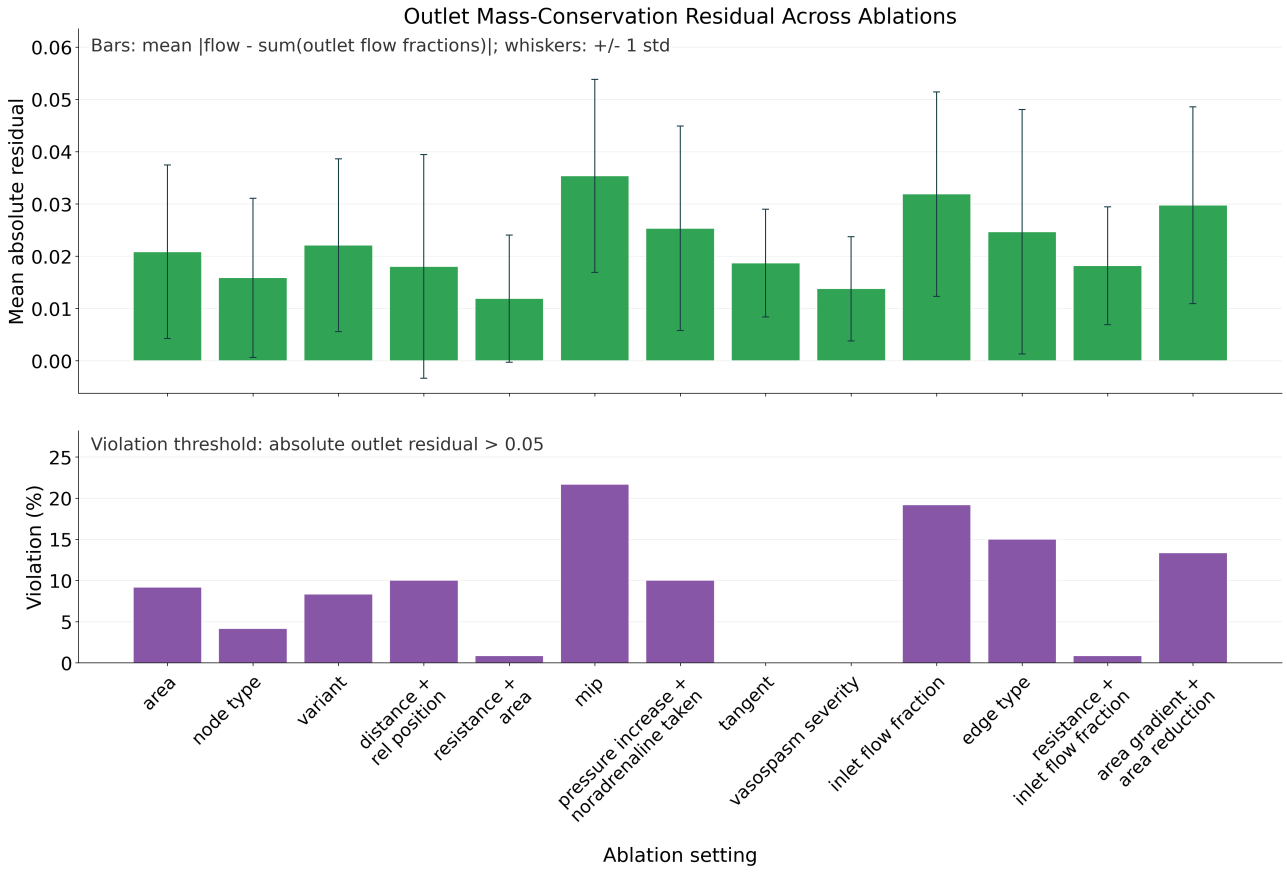

Figure 11: Outlet mass-conservation residuals across ablation settings. The upper panel reports the mean absolute outlet residual, defined as $|1 - \sum Q_{\text{out}}|$, with whiskers indicating the standard deviation across test graphs. The lower panel reports the percentage of test graphs violating outlet mass conservation, where a violation is defined as $|1 - \sum Q_{\text{out}}| > 0.05$.

3.4. Final model definition

The architectural comparison described in Section 3.2 and the ablation analysis reported in Section 3.3 guided the definition of the final model architecture and input feature set. The final configuration consists of a fine-tuned double-decoder PiGNN trained without resistance and inlet flow fraction as input features. This choice provides a balanced trade-off between clinical applicability and predictive performance: it improves flow MAE across all hemodynamic regimes, preserves PRE close to the full-feature model, and maintains outlet mass conservation close to the baseline, while removing input quantities that are not commonly available in clinical practice. A complete list of the input features retained in the final model is reported in the Appendix [6], Table 14. Importantly, the reduced-feature model already represents a substantial step toward a clinically deployable framework, since it avoids requiring prescribed outlet resistances or inlet flow split fractions at inference time. Nevertheless, a secondary residual limitation was still observed at the inlet level, where mass-conservation violations remained frequent. Although this issue was less critical than the dependence on clinically unavailable inputs, the possibility of further improving the physical consistency of the selected model was explored by introducing an inlet mass-conservation term in the loss formulation. Since the inlet flow distribution was no longer provided directly through the inlet flow fraction feature, this term did not duplicate information already available to the model, but acted as a non-redundant physics-based constraint on the predicted inlet flows. In this way, the loss modification served as an additional refinement to further optimize the clinically oriented reduced-feature model.
A representative prediction example is shown in Figure 12, selected as a test case with error close to the median test performance.

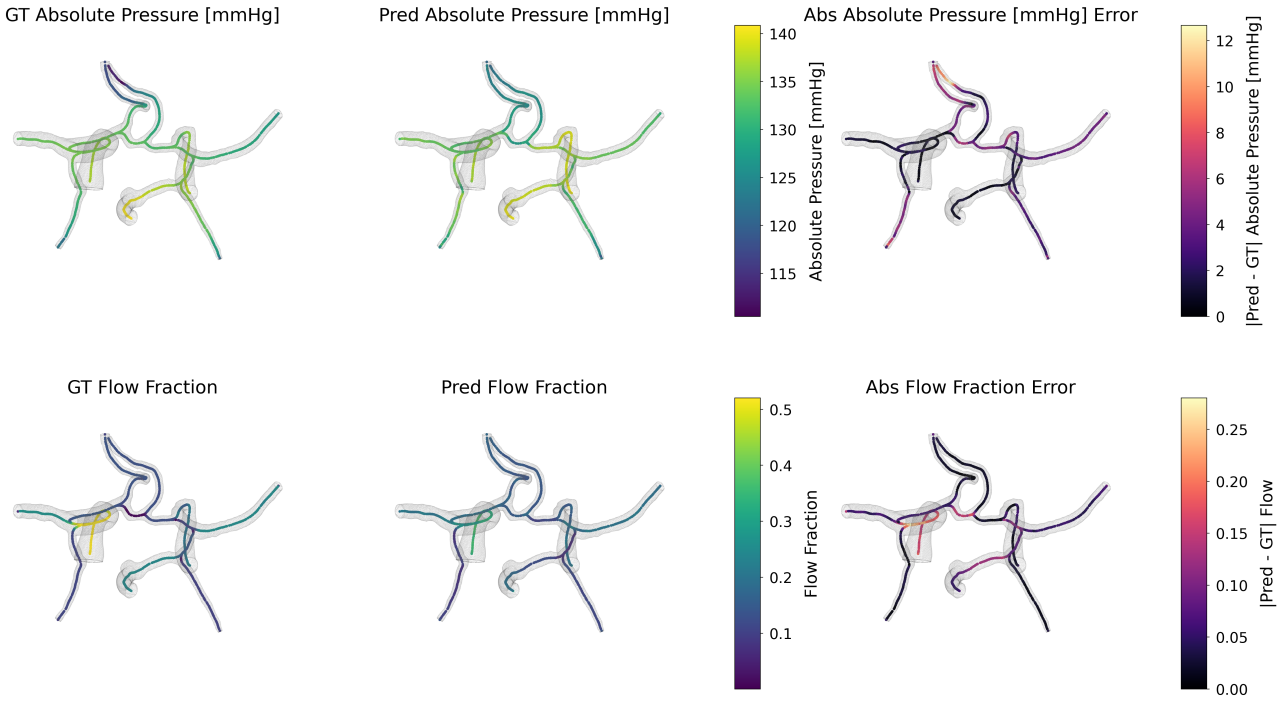

Figure 12: Representative example of final-model predictions on a test graph with error close to the median test performance. The first row shows ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error. The second row shows ground-truth flow fraction, predicted flow fraction, and absolute flow fraction error. The vascular surface is shown in grey for anatomical reference, while graph values are mapped onto the corresponding centerline representation. Local discrepancies are highlighted in the error maps. Absolute pressure errors remain low over most of the central vascular network, but localized higher discrepancies are visible in the ACAs, where the error reaches values on the order of $10 - 12$ mmHg. flow fraction errors are also spatially localized, with the largest discrepancy observed near the R-ICA, reaching approximately 0.20. Additional errors can be seen around the communicating or branching regions, with values roughly between 0.10 and 0.15.

The final model achieved a global PRE of $4.4 \cdot 10^{-2}$ and a flow MAE of $5.2 \cdot 10^{-2}$ on the test set. As reported in Table 7, the performance remains stable across the three hemodynamic regimes, corresponding to the baseline inflow regime and two progressively increased pressure-support regimes, indicating consistent predictive performance across the different hemodynamic regimes, with PRE increasing slightly from $4.2 \cdot 10^{-2}$ in Q_{base} to $4.6 \cdot 10^{-2}$ in $Q_{\text{NE}}^{\text{high}}$, and flow MAE increasing from $5.1 \cdot 10^{-2}$ to $5.3 \cdot 10^{-2}$. The inference time required for a sample prediction is $3.5 \cdot 10^{-2} s$ with a standard deviation of $3.9 \cdot 10^{-2} s$.

Regime-wise performance

Evaluation set	PRE (10^{-2}/%)	Flow MAE (10^{-2}/%)
Q_{base}	4.23	5.14
$Q_{\text{NE}}^{\text{low}}$	4.43	5.21
$Q_{\text{NE}}^{\text{high}}$	4.62	5.34
Global average	**4.43**	**5.23**

Table 7: Regime-wise and global performance of the final model in terms of pressure relative error (PRE) and flow mean absolute error (MAE), reported in units of 10^{-2}, or, equivalently, on a percentage scale. Metrics are computed graph-wise and averaged over the corresponding test graphs.

Compared with the ablation-selected model, the final model preserved comparable pressure accuracy while introducing only a small absolute increase in flow MAE and a minor reduction in outlet mass-conservation performance. Specifically, flow MAE increased from $4.7 \cdot 10^{-2}$ to $5.2 \cdot 10^{-2}$, and outlet violation percentage increased

from 0.84% to 6.67%. However, this limited deterioration was accompanied by a substantial improvement in inlet mass conservation, with inlet violations decreasing from 100% in the reduced-feature ablation to 1.67% in the final model. Figure 13 compares the signed inlet mass-conservation residuals for the full-feature baseline, the reduced-feature model without resistance and inlet flow fraction, and the final model with the explicit inlet conservation loss. The full-feature baseline and the reduced-feature model show residuals largely shifted toward positive values, with most samples concentrated between approximately 0.2 and 0.5, indicating a systematic mismatch between the predicted total inlet flow fraction and the expected value. In contrast, the final model shows residuals tightly concentrated around zero and a much narrower spreading interval, approximately from -0.05 to 0.08, leading to significantly smaller absolute residuals.

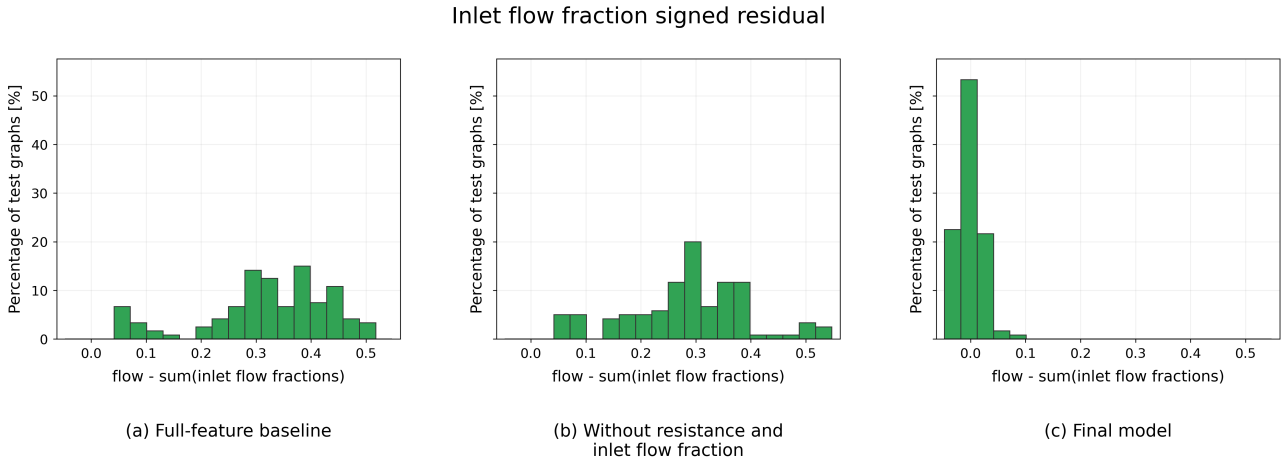

Figure 13: Comparison of signed inlet flow fraction residuals for the full-feature baseline, the reduced-feature ablation without resistance and inlet flow fraction, and the final model with explicit inlet mass-conservation loss. The x-axis reports the signed residual, computed as the difference between the predicted total flow and the sum of the predicted inlet flow fractions, while the y-axis reports the percentage of test graphs falling within each residual interval. Values close to zero indicate better inlet mass conservation. The full-feature baseline and the reduced-feature ablation show broad residual distributions shifted toward positive values, indicating systematic inlet imbalance. In contrast, the final model produces residuals tightly concentrated around zero, showing that the explicit inlet mass-conservation loss substantially improves inlet flow consistency.

Therefore, the final model was selected because it provided a substantially more balanced physical behavior across both inlet and outlet boundaries, while maintaining pointwise pressure and flow errors close to those of the ablation-selected model.

Figure 14 reports the final model performance stratified by anatomical variant. The number of test graphs differs across variants, reflecting the original dataset distribution, which was designed to reproduce the relative occurrence of anatomical variants in the population. Since the train/validation/test split was performed at the patient level while preserving this distribution as much as possible, less frequent variants are also represented by fewer samples in the test set.

The complete configuration shows the lowest pressure relative error, with a mean PRE of approximately $1.6{\cdot}10^{-2}$. Higher pressure errors are observed for several non-complete variants, particularly when the ACom, the PCAs and the left PCom are missing, which also show larger inter-graph variability. Flow MAE follows a similar but less pronounced pattern: the variants associated with higher pressure errors are generally also among those with higher flow errors, although the variation across variants is smaller for flow than for pressure. Across all variants, flow MAE indeed remains within a relatively narrow range, approximately between $4.0 \cdot 10^{-2}$ and $7.0 \cdot 10^{-2}$.

Taken together, the results therefore show that anatomical configuration affects both targets, but the effect is more pronounced for pressure prediction than for flow prediction.

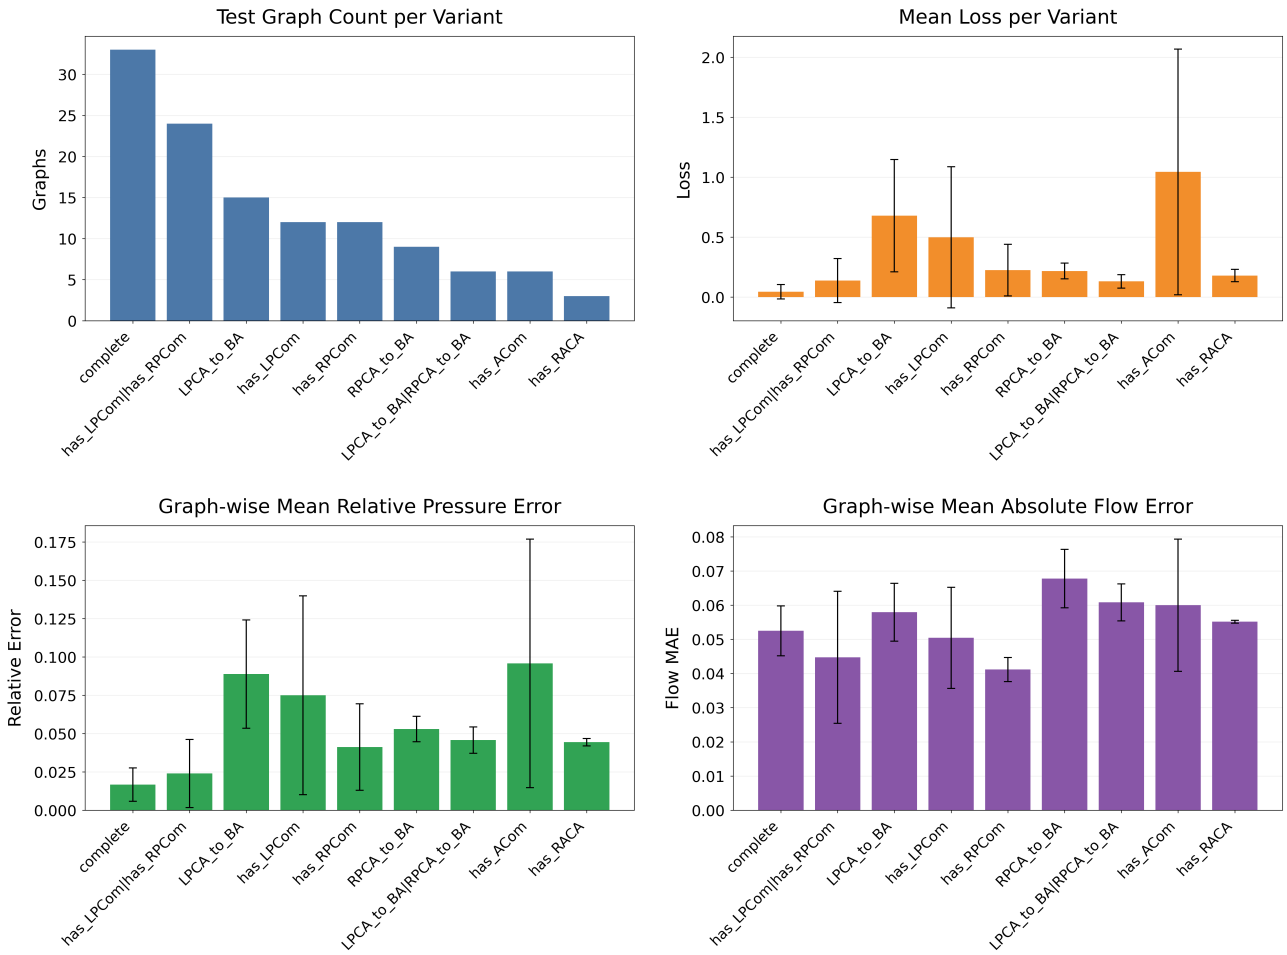

Figure 14: Variant-wise performance of the final model on the test set. The upper-left panel reports the number of test graphs for each anatomical variant. The name of the variant denotes the flag set to 0 (absent) in the global context vector. The upper-right panel shows the mean test loss per variant. The lower panels report graph-wise pressure relative error (PRE) and flow mean absolute error (MAE), respectively. Bars indicate the mean across test graphs belonging to each variant, while whiskers indicate one standard deviation.

Figure 15 reports node-wise pressure and flow error distributions stratified by node type and hemodynamic regime. For pressure prediction, normal, junction, and inlet nodes present comparable distributions, with mean PRE values remaining below approximately $5.0 \cdot 10^{-2}$ in all hemodynamic regimes. Outlet nodes are instead associated with higher errors, characterized by larger mean values and broader distributions than the other node categories. A mild increase in pressure error is observed across hemodynamic regimes for most node types, while the overall node-type pattern remains consistent. Flow errors display a stronger dependence on node type. Across all hemodynamic regimes, representing progressively increasing flow regimes, the largest absolute flow errors occur at inlet nodes, for which the inlet flow fraction is no longer prescribed as input, followed by junction nodes, whereas normal and outlet nodes present lower mean errors. Inlet and junction nodes also exhibit broader distributions, indicating greater inter-case variability in the prediction of flow redistribution at boundary and branching locations.

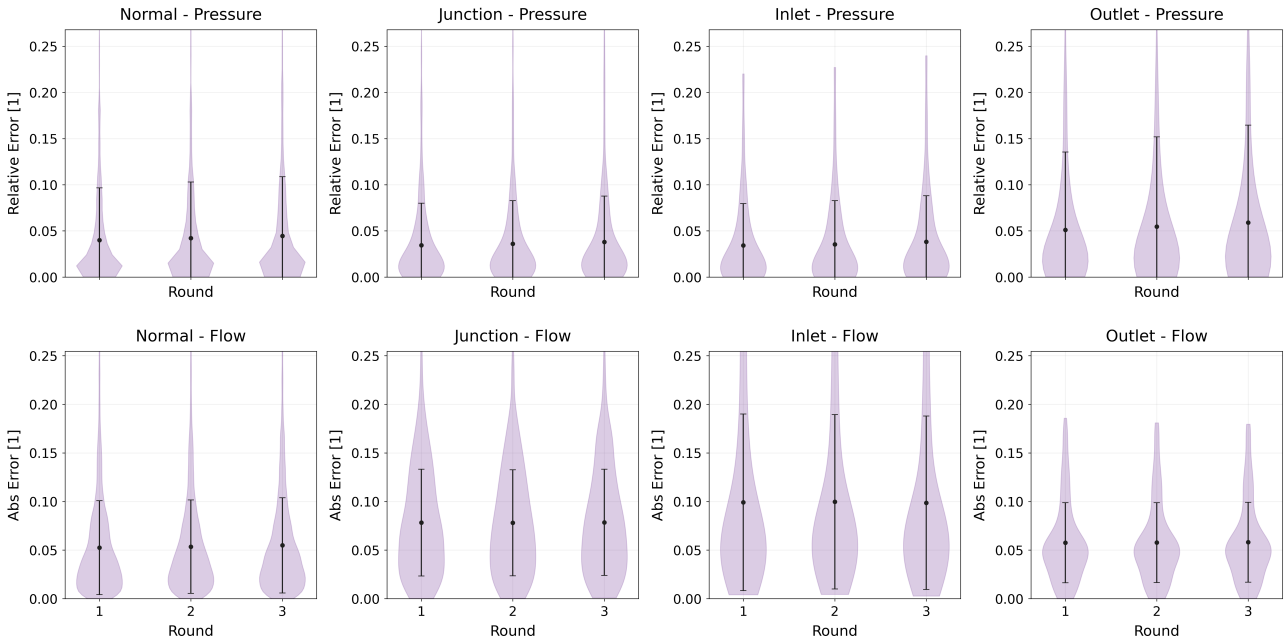

Figure 15: Node-type and regime-wise error distributions for the final model. The upper row reports node-wise pressure relative error for normal, junction, inlet, and outlet nodes across hemodynamic regimes, while the lower row reports the corresponding absolute flow fraction error. Violin plots show the distribution of node-wise errors, while markers and whiskers indicate the mean and standard deviation for each hemodynamic regime.

The effect of vasospasm-induced narrowing is assessed by comparing the errors of matched healthy and vasospasm-induced geometries, both quantitatively through the corresponding error variation and visually through paired scatter plots, reported in Figure 16. This analysis is particularly important because global test metrics are computed over the entire test set and therefore do not show whether the model performs equally well on healthy and vasospasm-induced geometries. By comparing each vasospasm-induced geometry with its corresponding healthy geometry, the underlying patient-specific anatomy and graph topology are kept fixed, while the main difference is the imposed reduction in vessel caliber. This makes it possible to evaluate whether vasospasm introduces an additional prediction error with respect to the same baseline geometry. However, since vasospasm-induced geometries represent the most clinically relevant application domain of the proposed tool, it is relevant to verify that good average performances are not hiding systematic failure on vasospastic cases.

The scatter plots should be read as a paired robustness analysis: each point compares the model error before and after narrowing for the same patient geometry, and the vertical distance from the identity line quantifies the additional prediction error introduced by vasospasm.

For each pair, the error variation is defined as the error in the vasospasm-induced geometry minus the error in the corresponding healthy geometry; therefore, positive values indicate worse performance after vessel narrowing. Across all matched pairs ($n = 30$), PRE increases by an average of $5.28 \cdot 10^{-2}$ after narrowing, whereas flow MAE increases by $9.56 \cdot 10^{-3}$. The mean increase in PRE is therefore approximately five times larger than the corresponding increase in flow MAE.

The increase in PRE depends on narrowing severity. Mild cases ($n = 3$) exhibit an average PRE increase of $4.28 \cdot 10^{-2}$, moderate cases ($n = 21$) of $4.54 \cdot 10^{-2}$, and severe cases ($n = 6$) of $8.37 \cdot 10^{-2}$. Flow MAE shows smaller variations, with a slight average decrease for mild cases ($-1.91 \cdot 10^{-3}$), and increases of $6.87 \cdot 10^{-3}$ and $2.47 \cdot 10^{-2}$ for moderate and severe cases, respectively. This trend is also visible in the scatter plots: most points lie above the identity line, indicating higher errors in the vasospasm-induced geometries than in the corresponding healthy configurations. The separation from the identity line is more pronounced for PRE, especially in moderate and severe cases, where several vasospastic geometries show PRE values above $1.0 \cdot 10^{-1}$. For flow MAE, the points are generally closer to the identity line and remain mostly within the range $3.0 \cdot 10^{-2}$–$8.0 \cdot 10^{-2}$.

In summary, the results show that the model remains applicable to vasospasm-like configurations, although pressure prediction becomes more sensitive as narrowing severity increases.

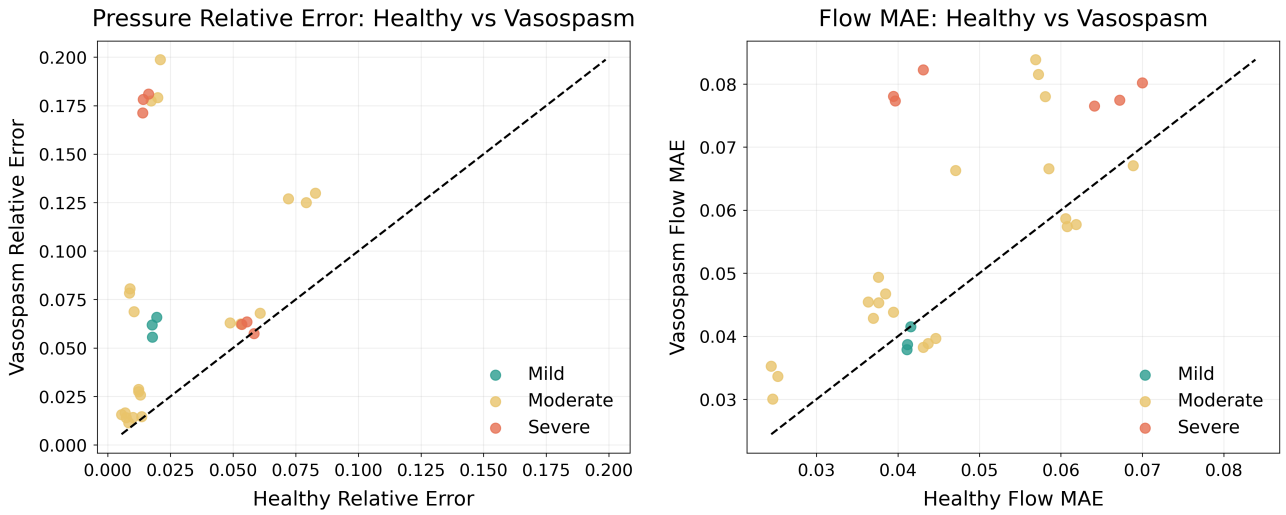

Figure 16: Comparison of final-model performance between healthy and vasospasm-induced geometries for the same patients, pressure relative error (PRE) (left) and flow mean absolute error (MAE) (right) are considered. Each point corresponds to one matched patient configuration, with the healthy error shown on the x-axis and the corresponding vasospasm-induced-geometry error shown on the y-axis. Points are colored according to vasospasm severity. The dashed line indicates equal performance between healthy and vasospasm-induced samples; points above the line correspond to higher error in the vasospasm-induced configuration. This paired comparison assesses whether model performance remains robust on vasospasm-induced geometries, which represent the main pathological application setting of the proposed tool.

3.5. Benchmarking

The final model was also evaluated on the four external benchmark geometries as described in Section 2.3.7. This analysis was performed to assess the ability of the model to generalize to a broader vascular domain than the one used during training, since the benchmark geometries include additional upstream and distal vessels. The reported PRE and flow MAE values are computed as averages over the four benchmark geometries.
On these external cases, PRE increased from $4.4 \cdot 10^{-2}$ on the original test set to $5.1 \cdot 10^{-2}$, while flow MAE rose from $5.2 \cdot 10^{-2}$ to $8.2 \cdot 10^{-2}$. Outlet mass conservation also deteriorated substantially, with the absolute outlet mass-conservation residuals ranging approximately between 1.4 and 2.6.
A representative benchmark prediction is shown in Figure 17, selected as a benchmark case with error close to the median benchmark performance. The remaining three benchmark cases are reported in the Appendix [6].

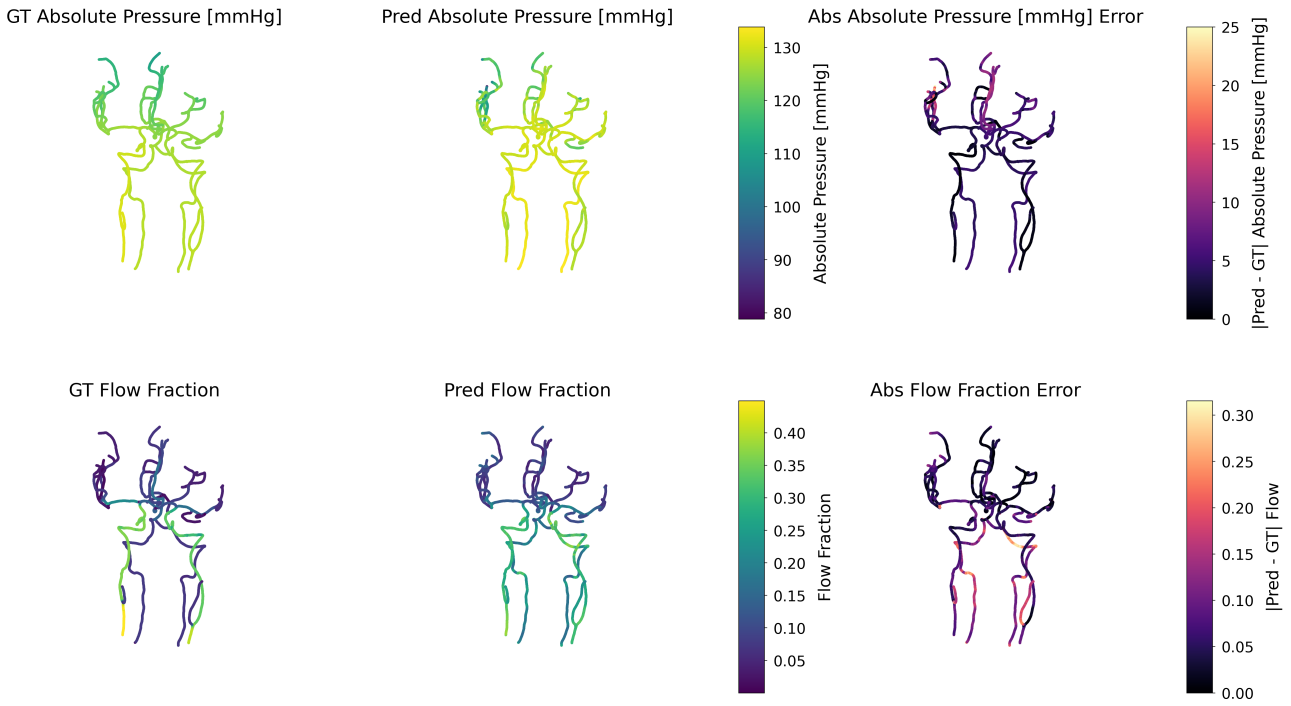

Figure 17: Representative prediction on an external benchmark geometry from Policlinico di Milano. The first row shows ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error. The second row shows ground-truth flow fraction, predicted flow fraction, and absolute flow fraction error. The error maps show that the largest discrepancies are spatially located in branches not represented in the original training geometries, while predictive ability is maintained in the central vascular region. Absolute pressure errors locally increase up to approximately $20 - 25$ mmHg in distal vessels branching off the R-MCA and R-ACA. flow fraction errors are more pronounced in the extended upstream vertical and carotid vessels, reaching a flow MAE of around 0.25-0.30.

4. Discussion

The present work developed and evaluated a physics-informed graph neural network surrogate for predicting pressure and flow redistribution in patient-specific Circle of Willis geometries. Starting from a CFD-based dataset including anatomical variants, vasospasm-induced narrowing, and hemodynamic regimes mimicking noradrenaline-related pressure support, different model architectures and input-feature sets were evaluated to define the final model formulation. An external benchmark was also performed to assess the generalization capability of the framework. The following discussion interprets the methodological and physiological implications of these findings, with particular attention to the results of the CFD, the target formulation, the distinction between pointwise prediction accuracy and physical consistency, the different behavior of pressure and flow targets, and the role of graph-based surrogate models in clinically motivated hemodynamic scenario assessment.

4.1. CFD analysis

The analysis of the CFD results is important because the simulated pressure and velocity fields constitute the ground truth from which the graph-level learning targets were derived. Before evaluating the surrogate model, it is therefore necessary to verify that the CFD dataset contains the expected hemodynamic trends.
Pressure behavior confirms that the simulated dataset captures two distinct hemodynamic effects. First, the progressive shift of absolute pressure distributions from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, shown in Figure 8, reflects the imposed pressure-support conditions used to mimic the effect of noradrenaline administration [30]. This behavior is expected, since $Q_{\text{NE}}^{\text{low}}$ and $Q_{\text{NE}}^{\text{high}}$ were designed to represent higher inlet-pressure conditions compared with the baseline Q_{base}. While these variations should not be considered as a direct prediction of the pressure increase caused by noradrenaline in vivo, they are consistent with the incremental induced hypertension protocol described by Veldeman et al. [85]. Second, the larger robust pressure ranges observed in vasospasm-induced geometries indicate that vessel narrowing increases spatial pressure heterogeneity. This is consistent with the additional pressure losses expected across narrowed segments, where local increases in resistance can create

stronger pressure gradients [86]. The fact that this effect becomes more pronounced especially in $Q_{\mathrm{NE}}^{\mathrm{high}}$ suggests that the hemodynamic impact of vasospasm is amplified under higher-flow or higher-pressure-support conditions. Overall, these analyses show that the CFD dataset captures both the imposed global pressure increase across hemodynamic regimes and the local pressure redistribution induced by vasospasm-like narrowing. This provides a physically meaningful pressure ground truth for the subsequent graph-based learning task.
The CFD velocity analysis further supports the physiological consistency of the simulated dataset. The near-peak velocity distributions obtained at inlet and outlet boundary sections remain within ranges compatible with values reported in transcranial Doppler and patient-specific CFD studies of the Circle of Willis [48]. In particular, healthy configurations show compact distributions in the major intracranial arteries, while vasospasm-induced configurations exhibit broader distributions and higher-velocity tails, especially in the MCA and ACA outlets. This behavior is expected, since arterial narrowing reduces the local lumen area and can therefore increase velocity for a given flow rate. This same principle underlies the clinical use of transcranial Doppler as a non-invasive tool for vasospasm monitoring, where elevated flow velocities in basal cerebral arteries are interpreted as a marker of vessel narrowing [22, 87]. While the values obtained are comparable with the TCD benchmarks, these results should not be interpreted as a direct reproduction of transcranial Doppler measurements, which depend on insonation location, angle, cardiac-cycle averaging, and vessel-specific acquisition protocols. The progressive increase in velocity from Q_{base} to $Q_{\mathrm{NE}}^{\mathrm{high}}$ is also consistent with the imposed hemodynamic design of the simulations, especially considering autoregulatory compensation is explicitly not modeled. $Q_{\mathrm{NE}}^{\mathrm{low}}$ and $Q_{\mathrm{NE}}^{\mathrm{high}}$ were defined to reproduce higher pressure-support and inlet-flow conditions; therefore, for comparable vessel caliber, an increase in flow rate necessarily produces higher velocities. In narrowed vessels, this effect may be further amplified by the reduced cross-sectional area. Importantly, the observed shift across hemodynamic regimes is moderate rather than abrupt, suggesting that the simulated noradrenaline-like conditions increase the hemodynamic load without producing non-physiological velocity distributions.

4.2. Interpretation of the target formulation

In surrogate modelling, the way in which the output targets are defined is not only a technical detail, but directly affects the physical information that the network is encouraged to learn. In this work, the targets were chosen to represent the two quantities most relevant for Circle of Willis hemodynamics for vasospasm management: pressure redistribution and flow redistribution.
Jointly predicting pressure and flow enables the model to capture two complementary aspects of the vascular response. Pressure describes the driving potential and the losses occurring along the vascular network, whereas flow describes how blood is redistributed across branches and collateral pathways. Their simultaneous prediction is therefore relevant for assessing collateral capacity and understanding how vascular anatomy, boundary conditions, and pathological narrowing interact to determine the global hemodynamic state. At the same time, pressure and flow are not expected to behave identically as learning targets, since they have different physical sensitivities and different levels of dependence on local versus global information.
A central design choice of this work is the prediction of relative hemodynamic quantities rather than absolute pressure and flow values. This formulation is motivated by both physical and practical considerations.
Pressure is predicted relative to the mean inlet pressure (MIP), since, for incompressible flow, the governing equations are driven by pressure gradients rather than by an arbitrary absolute pressure offset. Subtracting the MIP therefore preserves the physically relevant pressure differences while improving numerical stability and reducing sensitivity to differences in the absolute pressure scale across simulations. This choice follows the same general rationale adopted by Nannini et al. [60], who used pressure-drop fields as network targets to improve learning stability in the context of coronary hemodynamics. Although the anatomical district and clinical application are different, the underlying motivation is similar: learning pressure differences is generally more physically meaningful than learning an arbitrary absolute pressure level. Differently from approaches that predict absolute pressure values, such as Pegolotti et al. [68] and Botta et al. [66], the relative-pressure formulation used in this work focuses the model on spatial pressure variations, which are more directly linked to vascular resistance, stenosis-induced losses, and downstream pressure redistribution. Similarly, because the clinical objective of this work is to assess flow redistribution across the vascular graph, flow rate was selected as a more directly interpretable output than velocity, predicted for instance by Botta et al. [66] for cerebral microvasculature simulations. In the present context, the main interest is in fact not the local speed of blood within a vessel section, but how the available inflow is partitioned among vascular branches and collateral pathways under different anatomical and pathological conditions. However, differently from approaches that predict absolute volumetric flow rate, such as Pegolotti et al. [68], here flow is expressed as a fraction of the total inlet flow. This choice has both modelling and interpretative motivations. Absolute flow values depend on the total inflow prescribed at the boundaries, which varies across patients, hemodynamic regimes, and physiological conditions, and is not generally available in routine clinical practice. In contrast, the redistribution of flow within the Circle of Willis is more directly related to vascular geometry, topology, collateral pathways, and boundary

conditions. Expressing flow as a normalized fraction therefore provides a scale-invariant description of flow redistribution, facilitating comparison across different inlet regimes and reducing target variability associated only with changes in total inflow magnitude. From a physiological and clinical perspective, flow fraction provides a more immediate description of relative flow redistribution. Since the Circle of Willis acts as a collateral pathway between vascular territories, the distribution of the available inflow across its branches is relevant for assessing whether flow is redirected through collateral routes or reduced toward specific downstream regions. In this sense, the relevant information is not only the absolute amount of flow passing through a vessel, but also the share of the available cerebral inflow assigned to that branch or territory. The same absolute flow value may have different implications depending on the total inflow of the system, whereas a reduced flow fraction in a major downstream branch can more directly indicate a reduced relative contribution to the corresponding vascular territory.
This formulation also shapes the interpretation of the results. Flow prediction is expected to be more robust across hemodynamic regimes and feature perturbations because the fractional target removes the effect of absolute inflow magnitude and emphasizes redistribution patterns. Pressure prediction, instead, remains more sensitive to local vessel caliber, pressure drops, and boundary-condition effects, since relative pressure still reflects spatial gradients and losses along the vascular paths. This distinction provides a useful reading key for the following analyses. For example, the relatively stable flow errors observed across ablation settings and hemodynamic regimes, shown in Figure 10, can be interpreted in light of the normalized flow-fraction target, which emphasizes network-level redistribution. Conversely, the stronger variability of PRE, particularly under vessel narrowing or higher-flow regimes (Figure 16 and Figure 10 respectively), reflects the higher sensitivity of pressure to local resistance changes and pressure losses. Therefore, the different behavior of pressure and flow observed in the model comparison, ablation study, and final-model evaluation should be interpreted considering these target definitions.

4.3. Choice of the baseline architecture

As described in Sections 2.3.3 and 3.2, this work explicitly compared different decoder layouts and training strategies in order to identify a reference configuration that provides a suitable balance between accuracy, physical consistency, and training stability. The evaluation was motivated by the design choices adopted in related graph-based hemodynamic surrogate models. For example, Pegolotti et al. [68] used a single-decoder architecture trained from scratch, whereas Botta et al. [66] investigated fine-tuning to improve model performance while still relying on a single decoder for the predicted quantities. In addition, Sarabian et al. [54] suggested that using a pretrained network rather than starting from random initialization could reduce computational cost and improve model performance. Based on these considerations, both single- and double-decoder layouts and training either from scratch or through fine-tuning from a flow-only model were assessed.
The model comparison highlights that pointwise prediction accuracy and physical consistency are not necessarily aligned. Although the multitask models trained from scratch achieve competitive pressure relative error and flow MAE, they still exhibit substantial outlet mass-conservation violations. This points to the consideration that minimizing local prediction errors alone is not sufficient to guarantee that the predicted flow field satisfies global physics constraints. This behavior may be related to the difficulty of the multitask learning problem. When trained from a random initialization, the model must simultaneously learn a representation of the vascular graph, the pressure field, the flow redistribution, and the implicit physical constraints linking them. As a result, it can achieve low average errors while still producing flow predictions whose global sum is inconsistent at the outlet level.
Fine tuning instead leverages the simpler task of flow prediction to guide the learning of a more complex one by initializing the model from a more informative representation. Despite not predicting pressure, the only-flow model achieves the lowest flow MAE and satisfies outlet mass conservation in all test cases. This indicates that the flow-only training stage allows the shared graph representation to capture flow redistribution patterns more effectively. This ability is enhanced by the design decision of predicting flow as a fractional redistribution, more strongly constrained by vascular geometry and topology than by the absolute flow scale. However, it is important to consider that different flow-only models can achieve very similar standalone flow errors while still learning different internal representations and optimizer states. Fine-tuning then starts from these different states, and because downstream optimization is nonlinear and couples pressure prediction with physics constraints, small upstream differences can be amplified into noticeable differences in final pressure and flow performance. For this reason, the flow-only checkpoint should be treated not as a neutral pretraining artifact, but as a major conditioning factor for convergence, stability, and generalization of the final model.
Fine-tuning is particularly beneficial for the double-decoder architecture: although pressure relative error and flow MAE remain comparable to the corresponding model trained from scratch, the outlet violation rate decreases from 61.7% to 0%. Thus, the main advantage of fine-tuning is not a reduction in pointwise accuracy, but rather in the physical consistency of the predicted flow field.

The different behavior of the fine-tuned single- and double-decoder frameworks may be explained by task interference. In the single-decoder architecture, pressure and flow are predicted through the same final mapping, which may force the two targets to share a representation even at the output level, but, since the two outputs have different numerical distributions and different physical sensitivities, this can lead to competition between tasks. In contrast, the double-decoder architecture shares the encoder and processor while allowing pressure and flow to use separate output mappings. This distinction may help preserve the flow representation learned during pretraining while adapting the model to pressure prediction. This interpretation is supported by the fine-tuned single-decoder model, which substantially reduces outlet violations but shows a higher flow MAE than the other multitask configurations. This suggests that the single shared decoder may improve global consistency at the expense of local flow accuracy, whereas the double-decoder architecture achieves both comparable flow MAE and perfect outlet conservation.
For this reason, the fine-tuned double-decoder model is selected as the reference architecture. This choice is not based on optimizing a single metric, but on achieving the best compromise across pressure accuracy, flow prediction, and outlet mass conservation. In particular, it is the only multitask configuration that maintains comparable pressure and flow errors while eliminating outlet mass-conservation violations. More generally, these results show that physical consistency should be evaluated independently from pointwise prediction accuracy.
However, the framework still relies on boundary-condition quantities, namely outlet resistance and inlet flow fraction, which, despite meaningful in the context of CFD simulations and widely used by similar works on machine learning applied to CFD [60, 68], in this context are not readily available in routine clinical practice. Their use would therefore reduce the applicability of the model, since predictions would depend on quantities that would have to be assumed or estimated rather than directly obtained from patient-specific data. Indeed, previous work on clinical prediction models has highlighted that the inclusion of variables that are difficult to measure in routine care can restrict clinical applicability [88]. Moreover, studies on clinical decision support systems suggest that clinician comfort and trust may be reduced when predictions rely on imputed or otherwise not directly observed patient information [89, 90].

4.4. Ablation study

As explained in Section 2.3.6, the ablation study was conducted with a twofold objective: to determine whether a clinically more usable input configuration could be obtained without sacrificing performance, and to disentangle the contribution of different groups of input features to predictive accuracy and physical consistency. This analysis is particularly relevant because the input representation contains several quantities that encode related aspects of the same underlying hemodynamics, such as area, resistance, local geometric descriptors, and boundary-condition information. For example, vessel caliber and resistance are both related to the opposition to flow, while inlet flow fraction and resistance both encode boundary-condition information. Similarly, area reduction and area gradient describe local changes in vessel caliber, and distance-based edge features encode related aspects of spatial vascular connectivity. For this reason, selected features were removed not only individually, but also in physically meaningful combinations. Grouped ablations make it possible to assess whether the information lost by removing one feature can be compensated by correlated inputs, or whether the simultaneous removal of related features exposes a stronger dependence on that specific type of information. Therefore, evaluating features in isolation and in combination provides insight into which information is essential, which can be recovered from other inputs, and which may introduce unnecessary complexity.
However, it is important to note that ablation results should not be interpreted as direct measures of the intrinsic physical relevance of each feature. In particular, a gain in performance following feature removal does not necessarily imply that the removed feature is uninformative. Instead, such behavior may underline redundancy or multicollinearity within the feature set, whereby the model is able to compensate by recovering similar information from correlated inputs. In this case, removing certain features can reduce noise, limit overfitting, or simplify the learning problem, thereby improving generalization, even when the excluded variable remains physically meaningful.

4.4.1 Flow performance

As can be seen from Figure 10, a first relevant observation from the ablation study is that flow MAE improves after most feature removals, including the exclusion of several geometric, categorical, and global-context descriptors, highlighting that flow redistribution can be inferred from multiple partially redundant sources of information. The different effect of removing area alone compared with removing area together with resistance also supports this consideration: when area is excluded, resistance remains available as a correlated geometric and hemodynamic proxy, whereas removing both features reduces this compensatory information. This behavior is consistent with the formulation of the flow target adopted in this work (see Section 4.2). As a result, removing selected features does not necessarily eliminate essential information, but may instead reduce input

redundancy and simplify an otherwise over-parameterized learning problem. Therefore, the improvement observed after removing variables such as variant encodings, node or edge type information, and pressure-related global context should not be interpreted as evidence that these variables are physically irrelevant. Rather, for the flow task, their contribution may already be partially encoded by the remaining geometric and topological features, or their inclusion may introduce additional complexity that is not beneficial for generalization.

4.4.2 Pressure performance

The ablation analysis also shows that pressure prediction is more sensitive and more regime-dependent than flow prediction. While flow MAE improves rather consistently after most feature removals, PRE exhibits a more heterogeneous behavior, with several ablations degrading performance in Q_{base} and $Q_{\text{NE}}^{\text{low}}$ but improving it in $Q_{\text{NE}}^{\text{high}}$. This reinforces the hypothesis that pressure estimation depends more strongly on the specific input representation and on the underlying hemodynamic regime. The pattern is in line with the physical nature of the pressure field, which is more directly linked to local vessel caliber, boundary conditions, and systemic pressure level, provided via the global context (see Section 4.2). In contrast, flow fraction is more strongly constrained by vascular topology and by the redistribution of the total inflow, making it more robust to changes in the input feature set.

4.4.3 Behavior across different hemodynamic regimes

The ablation findings in Figure 10 reveal a clear regime-dependent effect of feature reduction.
In absolute terms, model performance progressively worsens from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, while remaining within the range of comparable surrogate-model results. This trend is expected, since $Q_{\text{NE}}^{\text{low}}$ and $Q_{\text{NE}}^{\text{high}}$ correspond to higher-flow hemodynamic regimes. Under these conditions, pressure gradients and flow redistribution become more pronounced, particularly in the presence of vessel narrowing, bifurcations, and variant collateral pathways, making higher-flow regimes intrinsically more challenging to predict.
However, the relative effect of removing selected input features is more favorable under the noradrenaline-induced flow-increase regimes. The positive PRE variations observed in $Q_{\text{NE}}^{\text{high}}$ should also be interpreted in light of the relative nature of the metric: since PRE normalizes the prediction error by the magnitude of the true absolute pressure field, the same error has a smaller contribution when the pressure scale is higher. This is not a limitation, but rather the intended behavior of the metric, as it enables fair comparison across simulations with different pressure levels. Nevertheless, the simultaneous improvement observed in flow MAE provides evidence that the regime-wise trend is not solely a consequence of pressure normalization, but also reflects stable model behavior under the noradrenaline-induced regime.
This phenomenon may further support the observation that the initial input set contained several correlated features. For example, inlet flow fraction, resistance, area, node type, global MIP, pressure-support indicators, topology, and vessel caliber all encode relatively overlapping information about flow redistribution. When many correlated features are provided simultaneously, the model may learn feature-specific associations that reduce the average training error but are not equally informative across all hemodynamic regimes. In the full-feature model, the network could therefore rely on correlations that are more robust for the baseline regime but do not necessarily transfer optimally to the higher-flow regimes.
Conversely, a simplified input representation might help generalize better, reducing the sensitivity to noise and redundant interactions. The higher impact under the noradrenaline-induced flow-increase regimes could thus be explained partly due to their sharpest regime shift with respect to the baseline, and partially due to their larger errors, which leave more room for visible improvement after feature reduction with respect to lower-flow regimes, whose performance may be already close to its best achievable level. Consequently, removing redundant features can produce a more favorable relative gain under higher-flow conditions, while the absolute prediction task remains more challenging.
The fact that feature removal provides preserved or improved performance relative to the full-feature baseline under the noradrenaline-induced flow-increase regimes is particularly relevant from a clinical perspective, because, by simulation design, they correspond to conditions in which noradrenaline has been administered, leading to a higher-pressure hemodynamic regime. Overall, the results confirm that the reduced feature representations remain effective, in scale-normalized terms, under the higher-pressure conditions representative of pharmacological pressure support in SAH patients.

4.4.4 Physical plausibility

The outlet conservation analysis reported in Figure 11 further confirms that pointwise predictive accuracy and physical consistency are distinct properties of the model. Several ablations improve flow MAE while simultaneously increasing the fraction of outlet mass-conservation violations compared with the full-feature

model, as observed when removing inlet flow fraction, MIP, edge type or area gradient alongside area reduction. Therefore, enhancements in flow MAE do not necessarily imply that the predicted flow distribution satisfies global conservation constraints. The different behavior observed across ablations also indicates that input features can contribute differently to local flow accuracy and to global conservation: a feature may be non-essential for reducing pointwise MAE while still helping the model preserve physically consistent flow sums. This matches the results observed in the architecture comparison, where models with competitive PRE and flow MAE could still exhibit substantial outlet mass-conservation violations, highlighting how standard prediction errors are insufficient to assess physical plausibility. The conservation residual provides complementary information to standard accuracy metrics, capturing whether the predicted flow fractions are globally coherent across the graph. Finally, the magnitude and variability of the outlet residuals also reveal that conservation behavior is not uniform across test cases. Some ablation settings show large standard deviations, indicating that the effect of feature removal may be graph-dependent. This may reflect differences in anatomical configuration, outlet distribution, or hemodynamic regime, which can make some cases more sensitive to the removal of specific feature groups.

4.4.5 Selection of the final feature set

Given the results of the ablation study, the final feature set, summarized in Table 14 of the Appendix [6], was selected by considering both predictive performance and clinical deployability. In this respect, resistance and inlet flow fraction represent two particularly relevant inputs. Resistance is a derived hemodynamic quantity, strongly related to vessel geometry and therefore partially encoded by the cross-sectional area and graph structure already provided to the model. Moreover, it is not generally available in patient-specific clinical practice and would therefore require estimation within literature-based ranges or empirical assumptions, potentially introducing additional uncertainty and limiting robustness. Similarly, inlet flow fraction represents boundary-condition information that is generally not expected to be available a priori in the intended clinical setting. Removing these quantities therefore reduces potential sources of systematic bias associated with non-directly measurable or assumption-dependent inputs, while preserving physics-consistent predictions. This choice is also consistent with limitations and open directions identified in previous hemodynamic surrogate-modeling studies. Pegolotti et al. [68] explicitly identified the elimination of dependence on prescribed boundary-condition parameters as an important advantage and future direction for surrogate approaches, while Sarabian et al. [54] highlighted the uncertainty associated with unknown patient-specific boundary conditions as a major limitation of purely physics-based computational models, especially when additional invasive measurements would be required for calibration.

The ablation results also support the exclusion of these two quantities. Integrated analysis across statistical performance (Figure 10) and physics compliance (Figure 11) indeed reveals that removing resistance and inlet flow fraction simultaneously provides a favorable compromise across multiple objectives: (1) preservation of pressure estimation close to the full-feature baseline, (2) improvement of flow prediction across regimes, approximately in the range of 10-15%, and (3) maintenance of low outlet mass conservation violations (1% vs. 0% baseline), which is instead compromised in other ablation settings showing comparable flow MAE and PRE. This suggests that the remaining clinically accessible geometric, topological, and global-context features contain sufficient information to recover the relevant flow redistribution patterns without explicitly providing resistance or inlet flow fractions as inputs.

In summary, the reduced feature set was not selected solely to maximize pointwise accuracy metrics in isolation, but because it provides a clinically more realistic input representation while preserving predictive accuracy and outlet physical consistency. While other features, such as vasospasm severity, showed comparable or greater negative importance in individual ablation experiments, they can be derived from imaging and patient records and are therefore more compatible with the intended clinical workflow. This formulation is indeed more adherent to the initial and fundamental design requirement of operating only with directly and reliably obtainable inputs for clinical deployment.

4.5. Final model

The final model, defined in Section 3.4, consists of a fine-tuned double-decoder PiGNN trained with the reduced clinically oriented feature set, excluding resistance and inlet flow fraction as input features. It represents the outcome of the architectural comparison and the ablation study, aiming at balancing predictive accuracy, physical consistency, and clinical applicability. Crucially, this configuration represents a substantial step toward a clinically usable surrogate model, as it preserves the hemodynamic predictive performance of the baseline framework while removing input quantities that would otherwise need to be prescribed or estimated rather than directly obtained in a clinical workflow.

The introduction of an inlet mass-conservation term also proved to be an effective refinement of this clinically

oriented model, enabling it to learn an inlet flow distribution that is both data-driven and constrained by conservation, rather than relying on prescribed inlet fractions. Inlet conservation is therefore not only evaluated after prediction but also explicitly enforced during training. Since violations of the inlet constraint directly contribute to the optimization objective, physically inconsistent solutions become costly during training and the learning process is shifted toward predictions that both minimize pointwise pressure and flow errors and satisfy physical consistency at the boundaries. This is reflected in the decrease of inlet mass-conservation violations from nearly all test graphs in the previous formulations to 1.67% in the final model, as shown in Figure 13.

From a broader perspective, the result also highlights the role of the loss function as more than a numerical training criterion: it defines which types of errors the model is incentivized to avoid, actively changing the learning problem and trajectory. This is consistent with previous observations in graph-based and physics-informed surrogate modelling. For example, Pegolotti et al. [68] identified the incorporation of mass-conservation constraints into the graph neural network structure as a relevant future direction, while Lannelongue et al. [91] emphasized the importance of embedding physics-based laws and constraints into the learning process. Suk et al. [92] also cited the addition of a balance term for inlet and outlet flows in the training loss as a future extension to improve consistency with physics.

Beyond global performance and boundary consistency, the final model was further analyzed with respect to anatomical variants, node types and vasospasm-induced narrowing, in order to identify where prediction is most reliable and which configurations remain more challenging.

Starting from the anatomical-variant analysis, Figure 14 highlights that anatomical configuration influences both pressure and flow prediction, but not to the same extent. The variants associated with larger pressure errors are generally also among those showing higher flow errors, revealing that certain anatomical configurations are intrinsically more challenging for the model. However, the variability across variants is more pronounced for PRE than for flow MAE, supporting the observation that pressure prediction is more sensitive to anatomical configuration, changes in vessel caliber, pathway connectivity, and local resistance distribution. Flow fraction, in contrast, remains more constrained by mass conservation and by the network topology, which may explain the narrower range of flow errors across variants. The interpretation of variant-wise differences should nevertheless account for the dataset design. Since anatomical variants are not equally represented in the dataset, the estimated mean error for rare variants is more sensitive to individual cases, as also indicated by the larger standard deviations observed in some groups. This supports the interpretation that variant labels alone do not fully determine prediction difficulty; local geometry, vasospasm distribution, and boundary conditions likely contribute to the observed inter-graph variability.

The node-type analysis shown in Figure 15 provides complementary information by identifying the vascular locations where the final model is most accurate and where prediction remains more challenging. The largest errors at the outlets observed for pressure may be related to the fact that outlet pressure reflects the cumulative effect of upstream pressure losses along the vascular path and is therefore sensitive to local geometry, vessel narrowing, and boundary-condition effects. Small errors accumulated along the graph may consequently become more visible at terminal locations. For flow prediction, the larger errors observed at inlet nodes suggest that these locations remain challenging independently of whether inlet flow fraction is explicitly provided as an input feature and despite the inclusion of an inlet mass-conservation term in the loss function. In the final formulation, this component constrains the total predicted inflow, but does not explicitly prescribe how this total should be partitioned across the individual inlet nodes. As a result, the model can strongly reduce global inlet mass-conservation violations while still producing local discrepancies in the predicted flow assigned to each inlet. Therefore, inlet nodes can remain a challenging location for flow prediction even when the global inlet conservation constraint is satisfied. Junction nodes are also expected to be challenging because they correspond to branching regions where flow is redistributed between multiple downstream paths, possibly causing small inaccuracies in the learned splitting behavior to lead to larger local flow errors. An additional source of difficulty at junctions is related to the projection of the three-dimensional CFD solution onto graph-level quantities. Near bifurcations, the vessel geometry changes rapidly, cross-sections may be less clearly defined, and the local flow field can include asymmetric profiles, secondary motion, or flow separation. As a result, the pressure or flow value assigned to a single graph node may be more sensitive to the exact location and orientation of the extraction plane, increasing both the intrinsic complexity of the target and the uncertainty associated with its graph-based representation. In contrast, internal vessel nodes show lower errors, particularly for flow. This is likely because these nodes belong to continuous vessel segments where hemodynamic quantities vary more smoothly and are less affected by branching or boundary effects. The long-tailed distributions observed in Figure 15 for some node types demonstrate that the model performs well for most nodes but still encounters localized difficult cases, possibly corresponding to complex anatomical configurations, severe narrowing, or regions where the graph representation only partially captures the underlying three-dimensional hemodynamics. Therefore, node-type analysis complements global metrics by identifying spatial locations where future model improvements should be focused.

The matched healthy-versus-vasospastic comparison in Figure 16 shows that vessel narrowing affects pressure predictions more strongly than flow predictions. This is consistent with the hemodynamic role of stenosis: a

reduction in lumen diameter increases local resistance and can produce increased blood-flow velocity and larger pressure drops across the stenotic segment. As a result, the pressure field may develop sharper local gradients around stenotic segments, making node-wise relative pressure prediction harder for the model to capture. The higher robustness of flow to the transition from healthy to pathological geometries could be motivated by the target formulation used in this work, which prioritizes vascular connectivity constraints and conservation requirements. This is clinically relevant because vasospasm-induced geometries represent the main pathological application setting of the proposed tool, and the paired comparison indicates that predictive performance is largely preserved after narrowing, especially for flow. At the same time, the increased pressure error in severe vasospasm-like configurations reveals that predictions in these cases should be interpreted with greater caution and motivates future enrichment of the dataset with additional severe stenosis configurations.

4.6. Benchmarking

Benchmarking on external geometries is an important step to assess whether a surrogate model retains predictive capability beyond the anatomical and topological distribution used during training. This aspect has been identified as a relevant direction in recent physics-informed graph-learning studies. Suk et al. [58] defined as an interesting future development the possibility of assessing the model's ability to generalize to substantially different vascular topologies, while Lannelongue et al. [91] framed benchmarking not only as a robustness test, but also as a way to assess the potential transferability of learned models toward future clinical extensions. Against this background, the external benchmark should be interpreted as an out-of-distribution robustness test rather than as a direct extension of the internal test set.

The Policlinico di Milano geometries include a substantially larger vascular domain, with upstream carotid and vertebral arteries and additional distal branches that were not represented in the training data. Therefore, the moderate degradation observed in PRE and flow MAE is expected, since the model is asked to extrapolate beyond the anatomical and topological distribution on which it was trained. Nevertheless, both pressure and flow errors remain comparable to values reported in the literature [66, 68].

Interestingly, as detailed in Section 3.5, pressure prediction appears more robust than flow prediction on these external cases. This may be related to the smoother nature of the pressure field and to the target formulation adopted in this work, which may help preserve generalization across larger domains. In contrast, flow prediction is more directly affected, as flow fraction must be redistributed across a larger number of branches and outlets, many of which are not present in the original Circle of Willis training domain. This is also visible in the prediction maps, as reported in Figure 17, where the largest flow errors are concentrated in peripheral or additional branches. The deterioration of outlet mass conservation should be interpreted in the same context. Although the pointwise flow MAE remains within a reasonable range, the global outlet residual accumulates errors over all terminal branches. When the number of outlets increases substantially, even small local flow-fraction errors can add up to a large deviation in the total outlet sum. Therefore, the high outlet residuals do not necessarily contradict the moderate pointwise flow error, but rather reveal that the conservation behavior learned on the original outlet topology does not extrapolate reliably to graphs with a much larger and more complex terminal structure.

Taken together, the benchmark suggests that the final model retains partial robustness on external patient-specific geometries, while showing sensitivity to changes in graph topology and outlet multiplicity. This result is encouraging because it shows that the model does not completely fail outside the training distribution, but it also highlights the need for future training or fine-tuning on larger vascular domains before applying the surrogate to extended clinical geometries.

4.7. Clinical relevance and impact of the work

The clinical relevance of the proposed framework is motivated by the dynamic nature of post-subarachnoid hemorrhage cerebrovascular complications. In SAH patients, the hemodynamic state can evolve over time, and clinicians may need to assess whether different cerebral territories receive sufficient perfusion under changing vascular and pressure-support conditions. Current clinical measurements provide important but incomplete information. Transcranial Doppler ultrasound is non-invasive, bedside-compatible, and suitable for repeated monitoring, but it provides velocity information only at accessible arterial locations and remains dependent on acoustic window quality and operator expertise. Conversely, phase-contrast and 4D-flow MRI can provide richer flow information, but their longer acquisition time and limited practicality in acute or repeated monitoring settings restrict their routine use for dynamic patient-specific assessment [63]. Patient-specific CFD can overcome some of these limitations by providing detailed pressure and flow fields, but its clinical applicability is limited by the complexity of the full simulation workflow. In practice, CFD requires vascular segmentation and reconstruction, mesh generation, boundary-condition assignment, numerical solution, and post-processing.

In addition, calibration of Windkessel parameters may require multiple forward simulations to match available measurements [63], further increasing the time and expertise required. These constraints make conventional CFD difficult to integrate into rapid clinical decision-making or repeated scenario exploration.
Several machine-learning approaches have been proposed for cerebrovascular image analysis, including segmentation from MRI or CT data, vessel classification, and aneurysm-related risk prediction [91]. Graph-based surrogate models for cardiovascular hemodynamics have already demonstrated the potential of GNNs to approximate hemodynamic quantities from simulation data [66, 68].
Within this context, to the best of our knowledge, this work is the first to develop and analyze a graph-based surrogate model specifically targeting Circle of Willis pressure and flow redistribution under post-SAH vasospasm-inspired conditions and noradrenaline-related pressure-support regimes, addressing a specific and unique clinical and physiological setting. In addition, the final formulation is explicitly designed only around clinically realistic inputs, which can be obtained or derived from routine patient-specific imaging and contextual information, rather than relying on parameters that would require additional calibration or invasive measurements. From a translational perspective, this is important because the model is not intended merely as a faster numerical solver, but as a step toward rapid and accessible patient-specific hemodynamic assessment. At the current stage, this use still assumes that the patient-specific geometry and vascular graph have already been extracted from clinical images, since automatic end-to-end segmentation and reconstruction are not yet integrated into the pipeline, as further discussed in Section 4.8. However, once the vascular graph has been constructed, the trained model can provide near-instantaneous estimates of relative pressure and flow redistribution.
The clinically relevant advantage is not limited to the speed of a single prediction, but also includes the possibility of evaluating multiple patient-specific scenarios. Because inference is extremely fast, the model could be used to explore numerous anatomical or physiological hypotheses without requiring a new CFD simulation for each configuration. For instance, this could allow rapid comparison of different pressure-support settings. Since the model includes pressure-related global-context variables, once clinically defined pressure targets or noradrenaline-supported regimes are specified, the surrogate could estimate how pressure and flow redistribution are expected to change under each condition. This would not replace clinical judgment or pharmacological dose selection, but could help assess whether additional pressure support is likely to provide a meaningful hemodynamic benefit or whether its effect may be limited for a given patient-specific vascular configuration. The same framework could also be used to explore plausible disease-progression trajectories. Different vasospasm locations and severity levels could be imposed on the same patient-specific Circle of Willis geometry to evaluate how worsening arterial narrowing may alter pressure patterns and flow redistribution, thus investigating whether collateral compensation remains sufficient under more severe configurations, or whether specific vascular territories become more vulnerable to reduced supply. A further potential use involves situations in which a baseline or pre-vasospasm patient-specific vascular reconstruction is available, but updated angiographic information is unavailable, incomplete, or not repeated. In this scenario, clinically plausible narrowing hypotheses could be defined from the available evidence, such as neurological symptoms, suspected hypoperfusion in a specific vascular territory, transcranial Doppler findings, or previous imaging. The surrogate could then be used to rapidly evaluate whether these hypothesized narrowing patterns produce pressure and flow-redistribution changes consistent with the observed clinical presentation. The clinical value would therefore not lie only in identifying a possible narrowing location or autonomously diagnose the location of vasospasm, but in estimating its patient-specific hemodynamic relevance: whether collateral pathways remain sufficient, which vascular territories become most vulnerable, and whether simulated pressure-support regimes are expected to improve flow redistribution or have limited benefit. In this sense, the model could support patient-specific hemodynamic interpretation and comparison of plausible vascular scenarios when the current vascular configuration is uncertain, while waiting for or complementing updated imaging. Beyond vasospasm management, the same framework could be extended to pre-surgical planning scenarios, in which the tolerance of a vascular territory to altered or temporarily interrupted inflow needs to be estimated. This situation can occur during intracranial aneurysm surgery, where temporary arterial occlusion may be used to facilitate aneurysm dissection and clipping, but may also expose downstream tissue to ischemic risk [93]. In this context, a rapid patient-specific estimate of collateral redistribution could help explore whether the Circle of Willis and communicating pathways are likely to compensate for temporary flow reduction. A related application concerns patients with ischemic symptoms not fully explained by a major-vessel occlusion. Since the Circle of Willis provides collateral flow between the anterior and posterior circulations, anatomical variants or hypoplastic communicating arteries may influence the ability to maintain perfusion when one vascular pathway is compromised [94]. The proposed graph-based framework could therefore support the investigation of whether subtle anatomical variants or limited collateral pathways contribute to regional flow redistribution deficits.
Finally, a secondary but essential contribution of this work is the creation of a comprehensive Circle of Willis hemodynamics dataset, comprising around 624 steady-state 3D CFD simulations generated using SimVascular across 125 unique patient-derived anatomies extracted from clinical imaging, with systematic coverage of anatomical variants according to population-representative frequencies based on anatomical studies. In addition to baseline conditions, vasospasm pathology is applied at three severity grades (mild, moderate, severe) to 60

original geometries, representing the clinical spectrum of post-SAH arterial narrowing. The simulations also mimic physiological variations in total cerebral blood flow associated with noradrenaline administration. The relevance of this dataset is further supported by recent work highlighting the scarcity of large-scale, high-quality training datasets as a major limitation for machine-learning models in neurovascular applications [91]. To the best of our knowledge, this dataset is among the most comprehensive CFD datasets specifically focused on CoW hemodynamics under anatomical variation and vasospasm-inspired conditions and will be made publicly available to accelerate future research in deep-learning based cerebrovascular modeling.

4.8. Limitations

Despite the encouraging performance of the proposed surrogate model, the present study is not exempt from limitations related to dataset composition, ground-truth generation, model representation, and clinical deployment.

One first limitation concerns the variability of performance across anatomical variants. As explained in Section 4.5, the variant-wise analysis showed that some configurations, particularly those involving communicating arteries or altered posterior circulation pathways, are associated, especially for pressure estimation, with higher errors and larger inter-graph variability. At the same time, the dataset was intentionally generated to reflect the relative occurrence of anatomical variants in the population; therefore, less frequent variants are represented by fewer samples. As a consequence, the statistical confidence of variant-specific performance estimates is lower for rare configurations, and the model may be less exposed during training to uncommon anatomical patterns. Future work should therefore include targeted enrichment of rare or hemodynamically challenging variants, while preserving a realistic population-inspired distribution for global evaluation.

Another relevant methodological caveat concerns the representation and prediction of vasospasm-induced cases. In this work, vasospasm was synthetically imposed, not obtained from real vasospasm scans, which may introduce a bias toward the imposed narrowing patterns used during dataset generation. A related remaining challenge is the behavior of pressure predictions in the presence of stenosis. Although overall performance remains comparable to literature benchmarks, severe or highly localized narrowing may create pressure patterns that are less well captured by the learned representation. Future improvements could therefore include real patient-derived vasospasm geometries, when available, to validate and enrich the synthetic cases, additional local descriptors of stenosis severity, increased representation of severe vasospasm cases in the training set, or loss terms that place greater emphasis on regions with high pressure gradients.

An additional important limitation concerns the definition of the ground truth itself, since the model is trained on CFD-derived quantities that are projected from the three-dimensional solution onto the graph representation, thus inheriting the assumptions and uncertainties of the numerical simulations used as ground truth. These include uncertainties related to vascular geometry reconstruction, boundary-condition assignment, numerical discretization, blood rheology assumptions, and the steady-state formulation. Any systematic error in the simulated pressure or flow fields may be transferred to the learned surrogate, since the network is optimized to reproduce the CFD solution rather than direct in-vivo measurements. Therefore, the accuracy of the surrogate cannot exceed the reliability of the computational ground truth on which it is trained. A further source of uncertainty arises from the projection of the 3D CFD solution onto graph-level targets. As described in Section 2.3.1, pressure and flow quantities are extracted on cross-sections obtained by cutting the vascular geometry with planes perpendicular to the local centerline tangent. While this provides a consistent way to associate 3D CFD fields with graph nodes, the procedure can become less robust in regions of high curvature or near bifurcations. In such locations, the centerline tangent may vary rapidly, adjacent cross-sections may become non-parallel or partially overlapping, and a plane normal to the centerline may not correspond to a well-defined anatomical vessel section. Near junctions, the slicing plane may also intersect complex branching regions, making the integrated pressure or flow value more sensitive to the exact centerline position and orientation. Consequently, part of the prediction error, especially in highly curved or bifurcating regions, may reflect uncertainty in the definition of the projected ground truth rather than only limitations of the learning model. Additionally, although the final model removes resistance and inlet flow fraction from the input set, it still depends on the availability and reliability of global clinical context variables. In practice, uncertainty in these inputs may affect prediction quality.

Moreover, the graph representation necessarily reduces the vascular domain to centerline-based quantities. As a result, the model does not predict fully three-dimensional hemodynamic fields such as wall shear stress, secondary flow structures, or cross-sectional velocity profiles. This is acceptable for the present objective, which focuses on pressure and flow redistribution, but it limits the use of the model for applications requiring detailed near-wall or volumetric flow information. Similarly, although the graph representation encodes bifurcation topology, it does not explicitly resolve three-dimensional flow structures at bifurcations, such as local separation, recirculation, or asymmetric velocity profiles. This caveat is partly mitigated by the fact that the training targets are derived from 3D CFD simulations and then projected onto graph-level quantities. Nevertheless, any

information not represented in the projected targets cannot be recovered by the model. A natural extension of the work would therefore be to operate directly on surface or volumetric meshes using a 3D GNN. This could allow prediction of richer hemodynamic fields, including wall shear stress or spatially resolved velocity. However, such a formulation would require substantially larger datasets, higher memory consumption [95, 96], and more complex preprocessing, whereas the present framework is designed for rapid prediction of clinically relevant reduced quantities. Moreover, since simulations are steady-state, the model does not capture pulsatile effects, pressure wave propagation, or time-resolved flow dynamics.
Finally, patient-specific CFD workflows are widely recognized as time-consuming, complex, and labor-intensive [97]. While the current work integrates graph construction and inference is extremely fast, it does not yet include an automatic end-to-end pipeline from clinical images. Segmentation, reconstruction of the Circle of Willis, and centerline extraction, including point coordinates, connectivity, radius, and tangent information, are assumed to be available as a basis for computing the geometric features used as model input. Therefore, an important future extension toward clinical deployment would be the integration of automated image-processing tools and a graphical interface suitable for medical use. Furthermore, the model has not yet been validated prospectively on independent clinical cohorts or linked to clinical outcomes. Its current role should thus be interpreted as a hemodynamic surrogate rather than as a validated decision-support tool.

5. Conclusion

This thesis developed and evaluated a physics-informed graph neural network surrogate for the prediction of patient-specific hemodynamics in the Circle of Willis. The work was motivated by the need for rapid assessment of pressure and flow redistribution in post-subarachnoid hemorrhage conditions, where vasospasm, anatomical variability, and pressure-support strategies can strongly influence cerebral perfusion. To this aim, a dataset of patient-derived Circle of Willis geometries was generated, incorporating anatomical variants, vasospasm-induced narrowing, and multiple hemodynamic regimes mimicking noradrenaline-related pressure support. 3D CFD simulations performed in SimVascular were used to obtain the reference pressure and flow fields, which were then projected onto centerline-based graph representations. Fine-tuning from a flow-only model and adopting a double-decoder architecture provided the best compromise between pressure accuracy, flow prediction, and outlet mass conservation. This result highlights the importance of jointly considering predictive performance and physical consistency, since models with comparable pointwise errors may still violate global conservation constraints. The feature set was chosen based on an ablation analysis and clinical availability considerations, preserving accuracy and physics consistency while relying only on quantities more realistically derived from patient-specific imaging and clinical context. The analysis of anatomical variants, node types, and matched healthy-versus-vasospastic geometries revealed that pressure prediction is more sensitive than flow prediction to anatomical configuration, vessel narrowing, and local graph structure.
Overall, this work demonstrates the feasibility of using physics-informed graph neural networks as rapid surrogates for Circle of Willis hemodynamics under clinically relevant post-SAH conditions. The framework does not replace high-fidelity CFD, but provides a computationally efficient and physically constrained approximation that could support repeated patient-specific scenario exploration, including vasospasm assessment, collateral-capacity evaluation, and pressure-support analysis. Future work should focus on automatic image-to-graph pipeline integration, enrichment of rare anatomical variants and severe vasospasm cases, extension toward richer three-dimensional or time-resolved outputs, and prospective validation on independent clinical cohorts before clinical deployment.

References

[1] David G McLone. The subarachnoid space: a review. *Pediatric Neurosurgery*, 6(3):113–130, 1980.

[2] Chiara Ragaglini, Matteo Foschi, Federico De Santis, Anna Laura Molliconi, Francesco Conversi, Enrico Colangeli, Raffaele Ornello, and Simona Sacco. Epidemiology and treatment of atraumatic subarachnoid hemorrhage over 10 years in a population-based registry. *European Stroke Journal*, 9(1):200–208, 2024.

[3] Brian L Hoh, Nerissa U Ko, Sepideh Amin-Hanjani, Sherry Hsiang-Yi Chou, Salvador Cruz-Flores, Neha S Dangayach, Colin P Derdeyn, Rose Du, Daniel Hänggi, Steven W Hetts, et al. 2023 guideline for the management of patients with aneurysmal subarachnoid hemorrhage: a guideline from the american heart association/american stroke association. *Stroke*, 54(7):e314–e370, 2023.

[4] Luís Guilherme Bastos Silva Aguiar Coelho, José Manuel Dias Costa, and Elsa Irene Peixoto Azevedo Silva. Non-aneurysmal spontaneous subarachnoid hemorrhage: perimesencephalic versus non-perimesencephalic. *Revista Brasileira de Terapia Intensiva*, 28:141–146, 2016.

[5] Jose I Suarez, Robert W Tarr, and Warren R Selman. Aneurysmal subarachnoid hemorrhage. *New England Journal of Medicine*, 354(4):387–396, 2006.

[6] Maree L Hackett and Craig S Anderson. Health outcomes 1 year after subarachnoid hemorrhage: an international population-based study. *Neurology*, 55(5):658–662, 2000.

[7] Jeannette W Hop, Gabriël JE Rinkel, Ale Algra, and Jan Van Gijn. Case-fatality rates and functional outcome after subarachnoid hemorrhage: a systematic review. *Stroke*, 28(3):660–664, 1997.

[8] SA Mayer, KT Kreiter, D Copeland, GL Bernardini, JE Bates, S Peery, J Claassen, YE Du, and ES Connolly Jr. Global and domain-specific cognitive impairment and outcome after subarachnoid hemorrhage. *Neurology*, 59(11):1750–1758, 2002.

[9] Sami Ridwan, Horst Urbach, Susanne Greschus, Johanna von Hagen, Jonas Esche, and Azize Boström. Health care costs of spontaneous aneurysmal subarachnoid hemorrhage for rehabilitation, home care, and in-hospital treatment for the first year. *World neurosurgery*, 97:495–500, 2017.

[10] Hendrik Stragier, Hans Vandersmissen, Sofie Ordies, Steven Thiessen, Dieter Mesotten, Dieter Peuskens, and Hugo Ten Cate. Pathophysiological mechanisms underlying early brain injury and delayed cerebral ischemia in the aftermath of aneurysmal subarachnoid hemorrhage: a comprehensive analysis. *Frontiers in Neurology*, 16:1587091, 2025.

[11] Robert P Ostrowski, ART Colohan, and John H Zhang. Molecular mechanisms for early brain injury after subarachnoid hemorrhage. *Handbook of Neurochemistry and Molecular Neurobiology*, pages 235–253, 2007.

[12] T Peeyush Kumar, Devin W McBride, Pramod K Dash, Kanako Matsumura, Alba Rubi, and Spiros L Blackburn. Endothelial cell dysfunction and injury in subarachnoid hemorrhage. *Molecular neurobiology*, 56(3):1992–2006, 2019.

[13] Matyáš Jelínek, Michal Jurajda, and K Duris. The role of oxidative stress in early brain injury after subarachnoid hemorrhage. *Oxidative Medicine and Cellular Longevity*, 2020(1):8877116, 2020.

[14] Kiyoshi Takagi, Akira Tamura, Tadayoshi Nakagomi, Hitoshi Nakayama, Osamu Gotoh, Kensuke Kawai, Mamoru Taneda, Nobuyuki Yasui, Hiromu Hadeishi, and Keiji Sano. How should a subarachnoid hemorrhage grading scale be determined? a combinatorial approach based solely on the glasgow coma scale. *Journal of neurosurgery*, 90(4):680–687, 1999.

[15] Nicole M Dubosh, M Fernanda Bellolio, Alejandro A Rabinstein, and Jonathan A Edlow. Sensitivity of early brain computed tomography to exclude aneurysmal subarachnoid hemorrhage: a systematic review and meta-analysis. *Stroke*, 47(3):750–755, 2016.

[16] Michael T Lawton and G Edward Vates. Subarachnoid hemorrhage. *New England Journal of Medicine*, 377(3):257–266, 2017.

[17] Stanlies D'Souza. Aneurysmal subarachnoid hemorrhage. *Journal of neurosurgical anesthesiology*, 27(3):222–240, 2015.

[18] Zvonimir Vrselja, Hrvoje Brkic, Stefan Mrdenovic, Radivoje Radic, and Goran Curic. Function of circle of willis. *Journal of Cerebral Blood Flow & Metabolism*, 34(4):578–584, 2014.

[19] Sumaiya Iqbal. A comprehensive study of the anatomical variations of the circle of willis in adult human brains. *Journal of clinical and diagnostic research: JCDR*, 7(11):2423, 2013.

[20] Lu Feng, He-Jiao Mao, Ding-Ding Zhang, Yi-Cheng Zhu, and Fei Han. Anatomical variations in the circle of willis and the formation and rupture of intracranial aneurysms: A systematic review and meta-analysis. *Frontiers in Neurology*, 13:1098950, 2023.

[21] Nebojša N Stojanović, Aleksandar Kostić, Radisav Mitić, Luka Berilažić, and Miša Radisavljević. Association between circle of willis configuration and rupture of cerebral aneurysms. *Medicina*, 55(7):338, 2019.

[22] J Max Findlay, Joshua Nisar, and Tim Darsaut. Cerebral vasospasm: a review. *Canadian Journal of Neurological Sciences*, 43(1):15–32, 2016.

[23] Brandon Foreman. The pathophysiology of delayed cerebral ischemia. *Journal of Clinical Neurophysiology*, 33(3):174–182, 2016.

[24] Katharina M Busl, Elisa Gouvea Bogossian, Jan Claassen, Raimund Helbok, Jose Javier Provencio, Chiara Robba, Mervyn DI Vergouwen, Stefan Wolf, Eliza R Zanier, and Giuseppe Citerio. Beyond the bleed: complications after aneurysmal subarachnoid hemorrhage. pathophysiology, clinical implications, and management strategies: a review. *Critical Care*, 29(1):414, 2025.

[25] Mervyn DI Vergouwen, Marinus Vermeulen, Jan van Gijn, Gabriel JE Rinkel, Eelco F Wijdicks, J Paul Muizelaar, A David Mendelow, Seppo Juvela, Howard Yonas, Karel G Terbrugge, et al. Definition of delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage as an outcome event in clinical trials and observational studies: proposal of a multidisciplinary research group. *Stroke*, 41(10):2391–2395, 2010.

[26] Devin W McBride, Spiros L Blackburn, Kumar T Peeyush, Kanako Matsumura, and John H Zhang. The role of thromboinflammation in delayed cerebral ischemia after subarachnoid hemorrhage. *Frontiers in Neurology*, 8:555, 2017.

[27] Amr Abdulazim, Marina Heilig, Gabriel Rinkel, and Nima Etminan. Diagnosis of delayed cerebral ischemia in patients with aneurysmal subarachnoid hemorrhage and triggers for intervention. *Neurocritical care*, 39(2):311–319, 2023.

[28] Gyanendra Kumar, Reza Bavarsad Shahripour, and Mark R Harrigan. Vasospasm on transcranial doppler is predictive of delayed cerebral ischemia in aneurysmal subarachnoid hemorrhage: a systematic review and meta-analysis. *Journal of neurosurgery*, 124(5):1257–1264, 2016.

[29] Andrea Cattaneo, Christoph Wipplinger, Caroline Geske, Florian Semmler, Tamara M Wipplinger, Christoph J Griessenauer, Judith Weiland, Alexandra Beez, Ralf-Ingo Ernestus, Thomas Westermaier, et al. Investigating the relationship between high-dose norepinephrine administration and the incidence of delayed cerebral infarction in patients with aneurysmal subarachnoid hemorrhage: a single-center retrospective evaluation. *PloS One*, 18(3):e0283180, 2023.

[30] Celine S Gathier, Walter M van den Bergh, Mathieu van der Jagt, Bon H Verweij, Jan Willem Dankbaar, Marcella C Müller, Annemarie W Oldenbeuving, Gabriel JE Rinkel, and Arjen JC Slooter. Induced hypertension for delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage: a randomized clinical trial. *Stroke*, 49(1):76–83, 2018.

[31] Eduardo E Benarroch. The arterial baroreflex: functional organization and involvement in neurologic disease. *Neurology*, 71(21):1733–1738, 2008.

[32] Lingzhong Meng, Yanhua Sun, Xu Zhao, Mads Rasmussen, Yazan Al-Tarshan, Deyi M Meng, Ziyue Liu, David C Adams, and David L McDonagh. Noradrenaline-induced changes in cerebral blood flow in health, traumatic brain injury and critical illness: a systematic review with meta-analysis. *Anaesthesia*, 79(9):978–991, 2024.

[33] Gabriela A Santos, Nils Petersen, Amir A Zamani, Rose Du, Sarah LaRose, Andrew Monk, Farzaneh A Sorond, and Can Ozan Tan. Pathophysiologic differences in cerebral autoregulation after subarachnoid hemorrhage. *Neurology*, 86(21):1950–1956, 2016.

[34] Karol P Budohoski, Marek Czosnyka, Peter J Kirkpatrick, Peter Smielewski, Luzius A Steiner, and John D Pickard. Clinical relevance of cerebral autoregulation following subarachnoid haemorrhage. *Nature Reviews Neurology*, 9(3):152–163, 2013.

[35] Jan W Dankbaar, Mienke Rijsdijk, Irene C Van Der Schaaf, Birgitta K Velthuis, Marieke JH Wermer, and Gabriel JE Rinkel. Relationship between vasospasm, cerebral perfusion, and delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage. *Neuroradiology*, 51(12):813–819, 2009.

[36] CS Gathier, WM van den Bergh, AJC Slooter, and HIMALAIA-Study Group. Himalaia (hypertension induction in the management of aneurysmal subarachnoid haemorrhage with secondary ischaemia): a randomized single-blind controlled trial of induced hypertension vs. no induced hypertension in the treatment of delayed cerebral ischemia after subarachnoid hemorrhage. *International Journal of Stroke*, 9(3):375–380, 2014.

[37] Timothy P Webber, Benjamin J Young, Marianne J Chapman, and Mark E Finnis. Complications of induced hypertension for delayed cerebral ischaemia secondary to aneurysmal subarachnoid haemorrhage. *Anaesthesia and Intensive Care*, 53(5):315–323, 2025.

[38] Haipeng Liu, Defeng Wang, Xinyi Leng, Dingchang Zheng, Fei Chen, Lawrence Ka Sing Wong, Lin Shi, and Thomas Wai Hong Leung. State-of-the-art computational models of circle of willis with physiological applications: a review. *Ieee Access*, 8:156261–156273, 2020.

[39] Angela Straccia. *Computational fluid dynamics modeling of hemodynamics in the Circle of Willis during vasospasm and in the left ventricle in the presence of a left ventricular assist device*. University of Washington, 2024.

[40] SM Moore, KT Moorhead, JG Chase, Tim David, and J Fink. One-dimensional and three-dimensional models of cerebrovascular flow. 2005.

[41] Leopold Grinberg, Elizabeth Cheever, Tomer Anor, Joseph R Madsen, and GE Karniadakis. Modeling blood flow circulation in intracranial arterial networks: a comparative 3d/1d simulation study. *Annals of biomedical engineering*, 39(1):297–309, 2011.

[42] Chi Zhang, Shuyu Li, Fang Pu, Yubo Fan, and Deyu Li. The effect of anatomic variations of circle of willis on cerebral blood distribution during posture change from supination to standing: a model study. *Bio-medical materials and engineering*, 24(6):2371–2380, 2014.

[43] John Fillenwarth. *Improved One-Dimensional Model of Blood Flow in the Human Arterial Network, with Application to Cerebral Flow*. PhD thesis, University of Delaware, 2016.

[44] Ilija D Šutalo, Anh V Bui, Shakil Ahmed, Kurt Liffman, and Richard Manasseh. Modeling of flow through the circle of willis and cerebral vasculature to assess the effects of changes in the peripheral small cerebral vasculature on the inflows. *Engineering Applications of Computational Fluid Mechanics*, 8(4):609–622, 2014.

[45] Suqin Huang, Bao Li, Jincheng Liu, Liyuan Zhang, Hao Sun, Huanmei Guo, Yanping Zhang, Fuyou Liang, Yanjun Gong, and Youjun Liu. Quantitative evaluation of the effect of circle of willis structures on cerebral hyperperfusion: A multi-scale model analysis. *Biocybernetics and Biomedical Engineering*, 44(4):782–793, 2024.

[46] Masato Shiba, Fujimaro Ishida, Kazuhiro Furukawa, Hiroshi Tanemura, Masanori Tsuji, Shinichi Shimosaka, and Hidenori Suzuki. Relationships of morphologic parameters and hemodynamic parameters determined by computational fluid dynamics analysis with the severity of subarachnoid hemorrhage. *Journal of Neuroendovascular Therapy*, pages oa–2016, 2017.

[47] Martin Sandve Alnæs, Jørgen Isaksen, Kent-André Mardal, Bertil Romner, Michael K Morgan, and Tor Ingebrigtsen. Computation of hemodynamics in the circle of willis. *Stroke*, 38(9):2500–2505, 2007.

[48] Angela Straccia, Fanette Chassagne, David I Bass, Guilherme Barros, Daniel F Leotta, Florence Sheehan, Deepak Sharma, Michael R Levitt, and Alberto Aliseda. A novel patient-specific computational fluid dynamics study of the activation of primary collateral pathways in the circle of willis during vasospasm. *Journal of Biomechanical Engineering*, 145(4):041008, 2023.

[49] Noah Maul, Katharina Zinn, Fabian Wagner, Mareike Thies, Maximilian Rohleder, Laura Pfaff, Markus Kowarschik, Annette Birkhold, and Andreas Maier. Transient hemodynamics prediction using an efficient octree-based deep learning model. In *International Conference on Information Processing in Medical Imaging*, pages 183–194. Springer, 2023.

[50] Amirtahà Taebi. Deep learning for computational hemodynamics: A brief review of recent advances. *Fluids*, 7(6):197, 2022.

[51] Dhruvil Panchigar, Kunal Kar, Shashank Shukla, Rhea Mary Mathew, Utkarsh Chadha, and Senthil Kumaran Selvaraj. Machine learning-based cfd simulations: a review, models, open threats, and future tactics. *Neural Computing and Applications*, 34(24):21677–21700, 2022.

[52] Gaoyang Li, Xiaorui Song, Haoran Wang, Siwei Liu, Jiayuan Ji, Yuting Guo, Aike Qiao, Youjun Liu, and Xuezheng Wang. Prediction of cerebral aneurysm hemodynamics with porous-medium models of flow-diverting stents via deep learning. *Frontiers in Physiology*, 12:733444, 2021.

[53] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[54] Mohammad Sarabian, Hessam Babaee, and Kaveh Laksari. Physics-informed neural networks for brain hemodynamic predictions using medical imaging. *IEEE transactions on medical imaging*, 41(9):2285–2303, 2022.

[55] Keiron O'shea and Ryan Nash. An introduction to convolutional neural networks. *arXiv preprint arXiv:1511.08458*, 2015.

[56] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

[57] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

[58] Julian Suk, Dieuwertje Alblas, Barbara Hutten, Albert Wiegman, Christoph Brune, Pim Van Ooij, and Jelmer M Wolterink. Physics-informed graph neural networks for flow field estimation in carotid arteries. *Medical Image Analysis*, page 103974, 2026.

[59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[60] Guido Nannini, Julian Suk, Patryk Rygiel, Simone Saitta, Luca Mariani, Riccardo Maranga, Andrea Baggiano, Gianluca Pontone, Jelmer M Wolterink, and Alberto Redaelli. Learning hemodynamic scalar fields on coronary artery meshes: A benchmark of geometric deep learning models. *Computers in biology and medicine*, 195:110477, 2025.

[61] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

[62] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

[63] Wojciech Kaczmarek, Jakub Magdziarz Ibrahim-El-Nur, Magdalena Łoś, and Tomasz Roleder. Reconstructing cerebral hemodynamics from sparse data using neural operator transformers. *Computers in Biology and Medicine*, 195:110492, 2025.

[64] Adam J Thorpe, Stepan Tretiakov, Dibakar Roy Sarkar, Krishna Kumar, and Ufuk Topcu. Learning generalizable neural operators for inverse problems. *arXiv preprint arXiv:2512.18120*, 2025.

[65] Hrishikesh Viswanath, Md Ashiqur Rahman, Abhijeet Vyas, Andrey Shor, Beatriz Medeiros, Stephanie Hernandez, Suhas Eswarappa Prameela, and Aniket Bera. Neural operator: Is data all you need to model the world? an insight into the paradigm of data-driven scientific ml. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

[66] Paolo Botta, Piermario Vitullo, Thomas Ventimiglia, Andreas Linninger, and Paolo Zunino. Physics-informed learning of microvascular flow models using graph neural networks. *arXiv preprint arXiv:2512.10792*, 2025.

[67] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv e-prints*, pages arXiv–2010, 2020.

[68] Luca Pegolotti, Martin R Pfaller, Natalia L Rubio, Ke Ding, Rita Brugarolas Brufau, Eric Darve, and Alison L Marsden. Learning reduced-order models for cardiovascular simulations with graph neural networks. *Computers in biology and medicine*, 168:107676, 2024.

[69] Xin Liu, Zhifan Gao, Huahua Xiong, Dhanjoo Ghista, Lijie Ren, Heye Zhang, Wanqing Wu, Wenhua Huang, and William Kongto Hau. Three-dimensional hemodynamics analysis of the circle of willis in the patient-specific nonintegral arterial structures. *Biomechanics and modeling in mechanobiology*, 15(6):1439–1456, 2016.

[70] Kaiyuan Yang, Fabio Musio, Yihui Ma, Norman Juchler, Johannes C Paetzold, Rami Al-Maskari, Luciano Höher, Hongwei Bran Li, Ibrahim Ethem Hamamci, Anjany Sekuboyina, et al. Benchmarking the cow with the topcow challenge: Topology-aware anatomical segmentation of the circle of willis for cta and mra. *ArXiv*, pages arXiv–2312, 2025.

[71] Steve Pieper, Michael Halle, and Ron Kikinis. 3d slicer. In *2004 2nd IEEE international symposium on biomedical imaging: nano to macro (IEEE Cat No. 04EX821)*, pages 632–635. IEEE, 2004.

[72] Volkan Kızılgöz, Mecit Kantarcı, and Şevket Kahraman. Evaluation of circle of willis variants using magnetic resonance angiography. *Scientific Reports*, 12(1):17611, 2022.

[73] Axel Neulen, Svenja Kunzelmann, Michael Kosterhon, Tobias Pantel, Maximilian Stein, Manfred Berres, Florian Ringel, Marc A Brockmann, Carolin Brockmann, and Sven R Kantelhardt. Automated grading of cerebral vasospasm to standardize computed tomography angiography examinations after subarachnoid hemorrhage. *Frontiers in neurology*, 11:13, 2020.

[74] Helena Merkel, Dirk Lindner, Khaled Gaber, Svitlana Ziganshyna, Jennifer Jentzsch, Simone Mucha, Thilo Gerhards, Sabine Sari, Annika Stock, Felicitas Vothel, et al. Standardized classification of cerebral vasospasm after subarachnoid hemorrhage by digital subtraction angiography. *Journal of clinical medicine*, 11(7):2011, 2022.

[75] Berk Geveci, Will Schroeder, A Brown, and G Wilson. Vtk. *The architecture of open source applications*, 1:387–402, 2012.

[76] LADS Antiga and David A Steinman. Vmtk: vascular modeling toolkit. *VMTK, San Francisco, CA, accessed Apr*, 27:2015, 2006.

[77] Si Hang. Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw*, 41(2):11, 2015.

[78] JPKH Alastruey, KH Parker, J Peiró, SM Byrd, and SJ Sherwin. Modelling the circle of willis to assess the effects of anatomical variations and occlusions on cerebral flows. *Journal of biomechanics*, 40(8):1794–1805, 2007.

[79] Axel Vikström, Petter Holmlund, Madelene Holmgren, Anders Wåhlin, Laleh Zarrinkoob, Jan Malm, and Anders Eklund. Establishing the distribution of cerebrovascular resistance using computational fluid dynamics and 4d flow mri. *Scientific Reports*, 14(1):14585, 2024.

[80] James Ahrens, Berk Geveci, and Charles Law. Paraview: An end-user tool for large data visualization. *The visualization handbook*, 717(8), 2005.

[81] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.

[82] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. Pmlr, 2017.

[83] Marius-Constantin Popescu, Valentina E Balas, Liliana Perescu-Popescu, and Nikos Mastorakis. Multilayer perceptron and neural networks. *WSEAS transactions on circuits and systems*, 8(7):579–588, 2009.

[84] Jin Xu, Zishan Li, Bowen Du, Miaomiao Zhang, and Jing Liu. Reluplex made more practical: Leaky relu. In *2020 IEEE Symposium on Computers and communications (ISCC)*, pages 1–7. IEEE, 2020.

[85] Michael Veldeman, Miriam Weiss, Walid Albanna, Omid Nikoubashman, Henna Schulze-Steinen, Hans Clusmann, Anke Hoellig, and Gerrit Alexander Schubert. Incremental versus immediate induction of hypertension in the treatment of delayed cerebral ischemia after subarachnoid hemorrhage. *Neurocritical Care*, 36(3):702–714, 2022.

[86] Jean F Soustiel, Eli Levy, Roni Bibi, Sergei Lukaschuk, and Dan Manor. Hemodynamic consequences of cerebral vasospasm on perforating arteries: a phantom model study. *Stroke*, 32(3):629–635, 2001.

[87] Alexander L Loomis and Mathew N Chakko. Doppler trans-cranial assessment, protocols, and interpretation. In *StatPearls [Internet]*. StatPearls Publishing, 2023.

[88] Mohammad ZI Chowdhury and Tanvir C Turin. Precision health through prediction modelling: factors to consider before implementing a prediction model in clinical practice. *The Journal of Primary Health Care*, 12(1):3–9, 2020.

[89] Saskia Haitjema, Steven WJ Nijman, Inge Verkouter, John JL Jacobs, Folkert W Asselbergs, Karel GM Moons, Ines Beekers, Thomas PA Debray, and Michiel L Bots. The use of imputation in clinical decision support systems: a cardiovascular risk management pilot vignette study among clinicians. *European Heart Journal-Digital Health*, 5(5):572–581, 2024.

[90] Jessica M Schwartz, Maureen George, Sarah Collins Rossetti, Patricia C Dykes, Simon R Minshall, Eugene Lucas, and Kenrick D Cato. Factors influencing clinician trust in predictive clinical decision support systems for in-hospital deterioration: qualitative descriptive study. *JMIR Human Factors*, 9(2):e33960, 2022.

[91] Vincent Lannelongue, Paul Garnier, Pablo Jeken-Rico, Aurèle Goetz, Philippe Meliga, Yves Chau, and Elie Hachem. Physics constrained graph neural network for real time prediction of intracranial aneurysm hemodynamics. *NPJ Digital Medicine*, 2026.

[92] Julian Suk, Guido Nannini, Patryk Rygiel, Christoph Brune, Gianluca Pontone, Alberto Redaelli, and Jelmer M Wolterink. Deep vectorised operators for pulsatile hemodynamics estimation in coronary arteries from a steady-state prior. *Computer methods and programs in biomedicine*, page 108958, 2025.

[93] Shima Shahjouei, Seyed Mohammad Ghodsi, Morteza Zangeneh Soroush, Saeed Ansari, and Shahab Kamali-Ardakani. Artificial neural network for predicting the safe temporary artery occlusion time in intracranial aneurysmal surgery. *Journal of Clinical Medicine*, 10(7):1464, 2021.

[94] Mohammed Oumer, Mekuriaw Alemayehu, and Abebe Muche. Association between circle of willis and ischemic stroke: a systematic review and meta-analysis. *BMC neuroscience*, 22(1):3, 2021.

[95] Mohammad Amin Nabian, Chang Liu, Rishikesh Ranade, and Sanjay Choudhry. X-meshgraphnet: Scalable multi-scale graph neural networks for physics simulation. *arXiv preprint arXiv:2411.17164*, 2024.

[96] Yawei Li, He Chen, Zhaopeng Cui, Radu Timofte, Marc Pollefeys, Gregory S Chirikjian, and Luc Van Gool. Towards efficient graph convolutional networks for point cloud handling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3752–3762, 2021.

[97] Pan Du, Delin An, Chaoli Wang, and Jian-Xun Wang. Ai-powered automated model construction for patient-specific cfd simulations of aortic flows. *Science Advances*, 11(36):eadw2825, 2025.

6. Appendix

Vasospasm simulation cases

Patient ID	Variant	Laterality	Side	Severity	Vessels to shrink (IDs)
3	ACA missing - Left	Unilateral	L	Severe	10,12
24	ACA missing - Right	Bilateral	Bilateral	Severe	10,12,7
30	ACom absent	Bilateral	Bilateral	Severe	4,5,6,7
59	ACom absent	Unilateral	L	Moderate	6,7
14	Both PCom absent	Bilateral	Bilateral	Mild	4,5,11,6,7,12
31	Both PCom absent	Bilateral	Bilateral	Moderate	4,5,11,6,7,12
45	Both PCom absent	Bilateral	Bilateral	Moderate	4,5,11,6,7,12
60	Both PCom absent	Bilateral	Bilateral	Moderate	4,5,11,6,7,12
72	Both PCom absent	Bilateral	Bilateral	Severe	4,5,11,6,7,12
76	Both PCom absent	Unilateral	R	Mild	4,5,11
80	Both PCom absent	Unilateral	R	Mild	4,5,11

Vasospasm simulation cases – continued

Patient ID	Variant	Laterality	Side	Severity	Vessels to shrink (IDs)
81	Both PCom absent	Unilateral	R	Moderate	4,5,11
87	Both PCom absent	Unilateral	L	Moderate	6,7,12
138	Both PCom absent	Unilateral	L	Moderate	6,7,12
153	Both PCom absent	Unilateral	L	Severe	6,7,12
158	Both PCom absent	Unilateral	L	Severe	6,7,12
1	Complete CoW	Bilateral	Bilateral	Mild	6,7,1,4,5
8	Complete CoW	Bilateral	Bilateral	Moderate	6,7,4,5
19	Complete CoW	Bilateral	Bilateral	Moderate	6,7,1,4,5
27	Complete CoW	Bilateral	Bilateral	Moderate	6,7,4,5
29	Complete CoW	Bilateral	Bilateral	Severe	6,7,10,1,4,5
40	Complete CoW	Bilateral	Bilateral	Severe	6,7,4,5
47	Complete CoW	Bilateral	Bilateral	Severe	6,7,4,5
56	Complete CoW	Unilateral	R	Mild	4,5
61	Complete CoW	Unilateral	L	Mild	6,7,1
66	Complete CoW	Unilateral	R	Mild	4,5
71	Complete CoW	Unilateral	R	Mild	4,5
85	Complete CoW	Unilateral	L	Moderate	6,7,10
131	Complete CoW	Unilateral	R	Moderate	4,5
137	Complete CoW	Unilateral	L	Moderate	6,7,10
141	Complete CoW	Unilateral	L	Moderate	6,7
151	Complete CoW	Unilateral	L	Moderate	6,7,1
155	Complete CoW	Unilateral	R	Moderate	4,5
161	Complete CoW	Unilateral	L	Moderate	12,7
162	Complete CoW	Unilateral	R	Severe	4,5
163	Complete CoW	Unilateral	L	Severe	6,7,10
15	Fetal PCA - Bilateral	Bilateral	Bilateral	Moderate	4,8,2,6,9,3,1
78	Fetal PCA - Bilateral	Unilateral	L	Severe	7,9,3,6,1
12	Fetal PCA - Left	Bilateral	Bilateral	Moderate	6,9,3,4
17	Fetal PCA - Left	Bilateral	Bilateral	Moderate	6,9,3,4
32	Fetal PCA - Left	Unilateral	R	Mild	4,8,2
77	Fetal PCA - Left	Unilateral	L	Severe	6,9,3
148	Fetal PCA - Left	Unilateral	L	Severe	6,9,3
10	Fetal PCA - Right	Bilateral	Bilateral	Moderate	4,8,2,6
41	Fetal PCA - Right	Bilateral	Bilateral	Moderate	4,8,2,6
69	Fetal PCA - Right	Unilateral	R	Mild	4,8,2
79	Fetal PCA - Right	Unilateral	L	Severe	6,9,3
157	Fetal PCA - Right	Unilateral	L	Severe	6,9,3
4	One PCom missing - Left	Bilateral	Bilateral	Moderate	4,5,11,8,2,6
16	One PCom missing - Left	Bilateral	Bilateral	Moderate	4,5,11,8,2,6
22	One PCom missing - Left	Bilateral	Bilateral	Severe	4,5,11,8,2,6
43	One PCom missing - Left	Unilateral	L	Mild	6,7,12,9,3
58	One PCom missing - Left	Unilateral	L	Moderate	6,7,12,9,3

Vasospasm simulation cases – continued

Patient ID	Variant	Laterality	Side	Severity	Vessels to shrink (IDs)
134	One PCom missing - Left	Unilateral	R	Severe	4,5,11,8,2
164	One PCom missing - Left	Unilateral	R	Severe	4,5,11,8,2
18	One PCom missing - Right	Bilateral	Bilateral	Mild	6,7,12,9,3,4
50	One PCom missing - Right	Bilateral	Bilateral	Moderate	6,7,12,9,3,4
64	One PCom missing - Right	Unilateral	R	Moderate	4,5,11,8,2
132	One PCom missing - Right	Unilateral	L	Severe	6,7,12,9,3
136	One PCom missing - Right	Unilateral	L	Severe	6,7,12,9,3

Table 8: Definition of vasospasm simulation cases. For each patient, the table reports the anatomical variant, laterality, side, severity level, and vessel segment IDs selected for narrowing. Vessel IDs refer to the segment-label correspondence reported in Table 9.

Vessel segment ID mapping

Vessel ID	Anatomical vessel label
1	BA
2	R-PCA
3	L-PCA
4	R-ICA
5	R-MCA
6	L-ICA
7	L-MCA
8	R-PCom
9	L-PCom
10	ACom
11	R-ACA
12	L-ACA

Table 9: Correspondence between vessel segment identifiers and anatomical vessel labels used in the vasospasm simulation plan.

Absolute pressure across hemodynamic regimes

Hemodynamic regime	Mean absolute pressure [mmHg]
Q_{base}	128.3 ± 20.3
$Q_{\text{NE}}^{\text{low}}$	147.8 ± 23.6
$Q_{\text{NE}}^{\text{high}}$	166.0 ± 27.0

Table 10: Descriptive statistics of absolute centerline pressure across the three hemodynamic regimes, considering all 203 simulated geometries/conditions. Values are reported as mean ± standard deviation.

Regime-wise increase in absolute pressure

Comparison among regimes	Mean pressure increase [mmHg]
$Q_{\text{base}} \rightarrow Q_{\text{NE}}^{\text{low}}$	19.5 ± 8.7
$Q_{\text{NE}}^{\text{low}} \rightarrow Q_{\text{NE}}^{\text{high}}$	18.1 ± 7.3
$Q_{\text{base}} \rightarrow Q_{\text{NE}}^{\text{high}}$	37.7 ± 11.2

Table 11: Regime-wise increase in mean absolute centerline pressure across all 203 simulated geometries/conditions. For every case, absolute pressure increased between consecutive hemodynamic regimes. Values are reported as mean ± standard deviation.

Robust spatial pressure variation in matched geometries

Hemodynamic regime	Healthy median [IQR] [mmHg]	Vasospasm-induced median [IQR] [mmHg]
Q_{base}	8.0 [4.4–13.3]	20.1 [7.4–41.0]
$Q_{\text{NE}}^{\text{low}}$	10.2 [6.1–17.0]	23.4 [9.6–49.7]
$Q_{\text{NE}}^{\text{high}}$	11.6 [6.9–21.3]	26.9 [12.7–61.4]

Table 12: Robust within-geometry pressure range for matched healthy and vasospasm-induced geometries across the three hemodynamic regimes. The robust pressure range was defined as the difference between the 95th and 5th percentiles of centerline pressure values within each geometry. Values are reported as median and interquartile range (IQR).

Table 13: **Near-peak velocity statistics.** Descriptive statistics of near-peak velocity distributions at inlet and outlet boundary sections. Velocities were computed as the area-weighted 95th percentile of velocity magnitude and are reported in cm/s. Values are shown for each vessel, hemodynamic regime, and geometrical condition as mean ± standard deviation and median with interquartile range (IQR). The mean ± standard deviation describes the average velocity level and the overall variability across cases, and is therefore sensitive to high-velocity tails or isolated extreme values. The median [IQR] provides a more robust description of the central tendency and spread of the distribution, where the IQR represents the range between the 25th and 75th percentiles. Comparing mean and median values helps identify skewed distributions: when the mean is higher than the median, this indicates the presence of right-skewed distributions or high-velocity tails.

Vessel	**Hemod regime**	**Healthy**		**Vasospasm-induced**	
		Mean ± SD	**Median [IQR]**	**Mean ± SD**	**Median [IQR]**
BA	Q_{base}	52.4 ± 27.7	49.1 [40.9–54.9]	58.1 ± 43.6	50.0 [43.3–56.5]
	$Q_{\text{NE}}^{\text{low}}$	60.4 ± 31.2	55.2 [49.0–64.0]	66.4 ± 48.2	57.3 [50.1–66.6]
	$Q_{\text{NE}}^{\text{high}}$	67.1 ± 32.0	62.2 [53.8–71.0]	73.7 ± 50.4	63.3 [56.0–75.4]
R-ICA	Q_{base}	40.0 ± 10.7	37.9 [33.8–47.0]	53.9 ± 21.8	48.2 [39.6–66.9]
	$Q_{\text{NE}}^{\text{low}}$	46.2 ± 12.4	46.2 [36.2–54.1]	61.8 ± 25.2	58.4 [44.5–73.3]
	$Q_{\text{NE}}^{\text{high}}$	51.5 ± 13.2	51.6 [41.0–60.5]	69.0 ± 27.4	63.7 [51.4–82.0]
L-ICA	Q_{base}	40.0 ± 10.7	37.9 [33.7–46.7]	53.7 ± 21.4	48.5 [39.2–67.2]
	$Q_{\text{NE}}^{\text{low}}$	46.2 ± 12.3	46.4 [36.4–54.2]	61.6 ± 24.7	58.3 [44.4–73.5]
	$Q_{\text{NE}}^{\text{high}}$	51.4 ± 13.1	51.6 [41.1–60.3]	68.8 ± 26.9	63.4 [51.0–82.0]
R-MCA	Q_{base}	72.9 ± 26.3	66.5 [54.8–78.2]	123.2 ± 78.2	97.2 [72.6–146.2]
	$Q_{\text{NE}}^{\text{low}}$	83.4 ± 30.1	78.0 [64.0–85.7]	139.8 ± 85.8	113.9 [86.0–171.5]
	$Q_{\text{NE}}^{\text{high}}$	92.3 ± 33.4	85.1 [71.5–95.9]	155.6 ± 96.2	123.2 [94.8–195.7]
L-MCA	Q_{base}	71.1 ± 38.0	60.9 [52.9–72.3]	136.3 ± 87.5	106.8 [67.2–184.7]
	$Q_{\text{NE}}^{\text{low}}$	81.1 ± 41.5	70.8 [61.3–84.3]	155.6 ± 102.0	118.6 [76.1–217.5]
	$Q_{\text{NE}}^{\text{high}}$	90.1 ± 46.7	78.3 [67.7–92.5]	172.1 ± 110.5	132.3 [88.4–237.2]
R-ACA	Q_{base}	71.4 ± 25.5	67.1 [54.4–84.0]	96.0 ± 51.3	85.0 [66.3–104.8]
	$Q_{\text{NE}}^{\text{low}}$	80.9 ± 28.6	75.7 [64.0–92.6]	109.3 ± 60.2	92.7 [75.3–119.2]
	$Q_{\text{NE}}^{\text{high}}$	89.3 ± 31.2	83.7 [70.3–103.9]	121.4 ± 68.0	102.2 [84.6–136.9]
L-ACA	Q_{base}	75.2 ± 21.8	73.2 [59.6–84.2]	106.4 ± 63.0	91.5 [68.5–123.8]
	$Q_{\text{NE}}^{\text{low}}$	85.2 ± 24.3	83.0 [68.8–94.8]	121.1 ± 71.0	104.6 [76.7–149.0]
	$Q_{\text{NE}}^{\text{high}}$	93.5 ± 26.0	93.6 [76.4–103.0]	133.1 ± 76.8	114.8 [85.1–160.5]
R-PCA	Q_{base}	74.5 ± 46.8	61.3 [52.7–80.7]	87.5 ± 53.5	67.8 [57.2–102.5]
	$Q_{\text{NE}}^{\text{low}}$	85.4 ± 53.4	71.4 [59.5–93.8]	99.7 ± 60.4	80.2 [63.5–116.2]
	$Q_{\text{NE}}^{\text{high}}$	94.5 ± 57.3	78.1 [68.2–102.2]	111.0 ± 66.7	89.3 [69.5–127.5]
L-PCA	Q_{base}	80.0 ± 72.0	60.5 [48.5–79.5]	82.9 ± 50.3	68.6 [50.5–96.9]
	$Q_{\text{NE}}^{\text{low}}$	91.5 ± 83.6	68.9 [55.3–90.8]	94.6 ± 56.2	76.8 [61.2–106.8]
	$Q_{\text{NE}}^{\text{high}}$	100.6 ± 88.3	76.2 [61.6–100.2]	104.9 ± 61.6	86.4 [71.5–117.0]

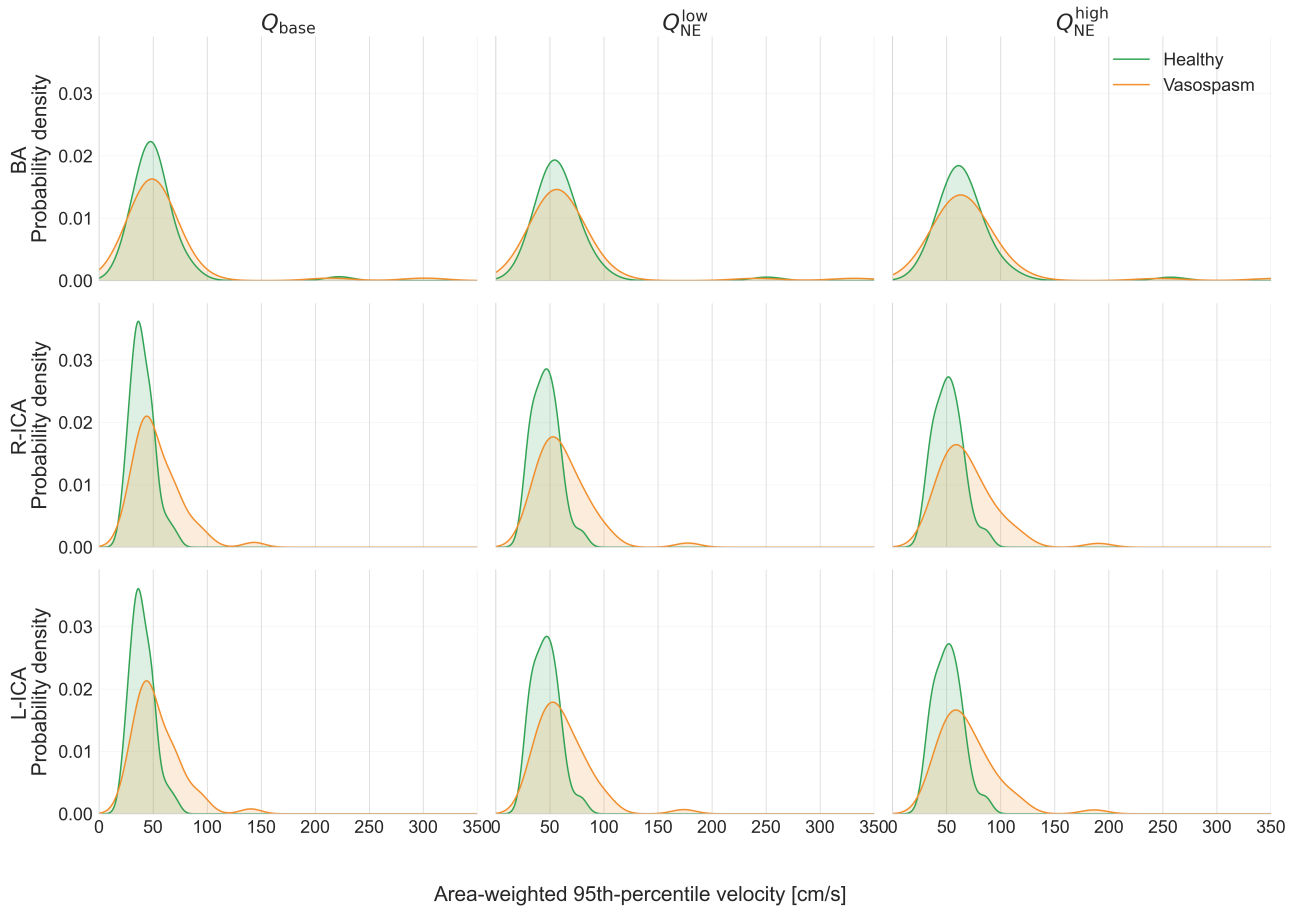

Figure 18: Distribution of inlet near-peak velocities for the BA, R-ICA, and L-ICA across the three hemodynamic regimes. For each one, physiological and vasospasm-induced geometries are compared. The BA distributions are largely overlapping between conditions, whereas the R-ICA and L-ICA show broader distributions and a shift toward higher velocities in vasospasm-induced geometries.

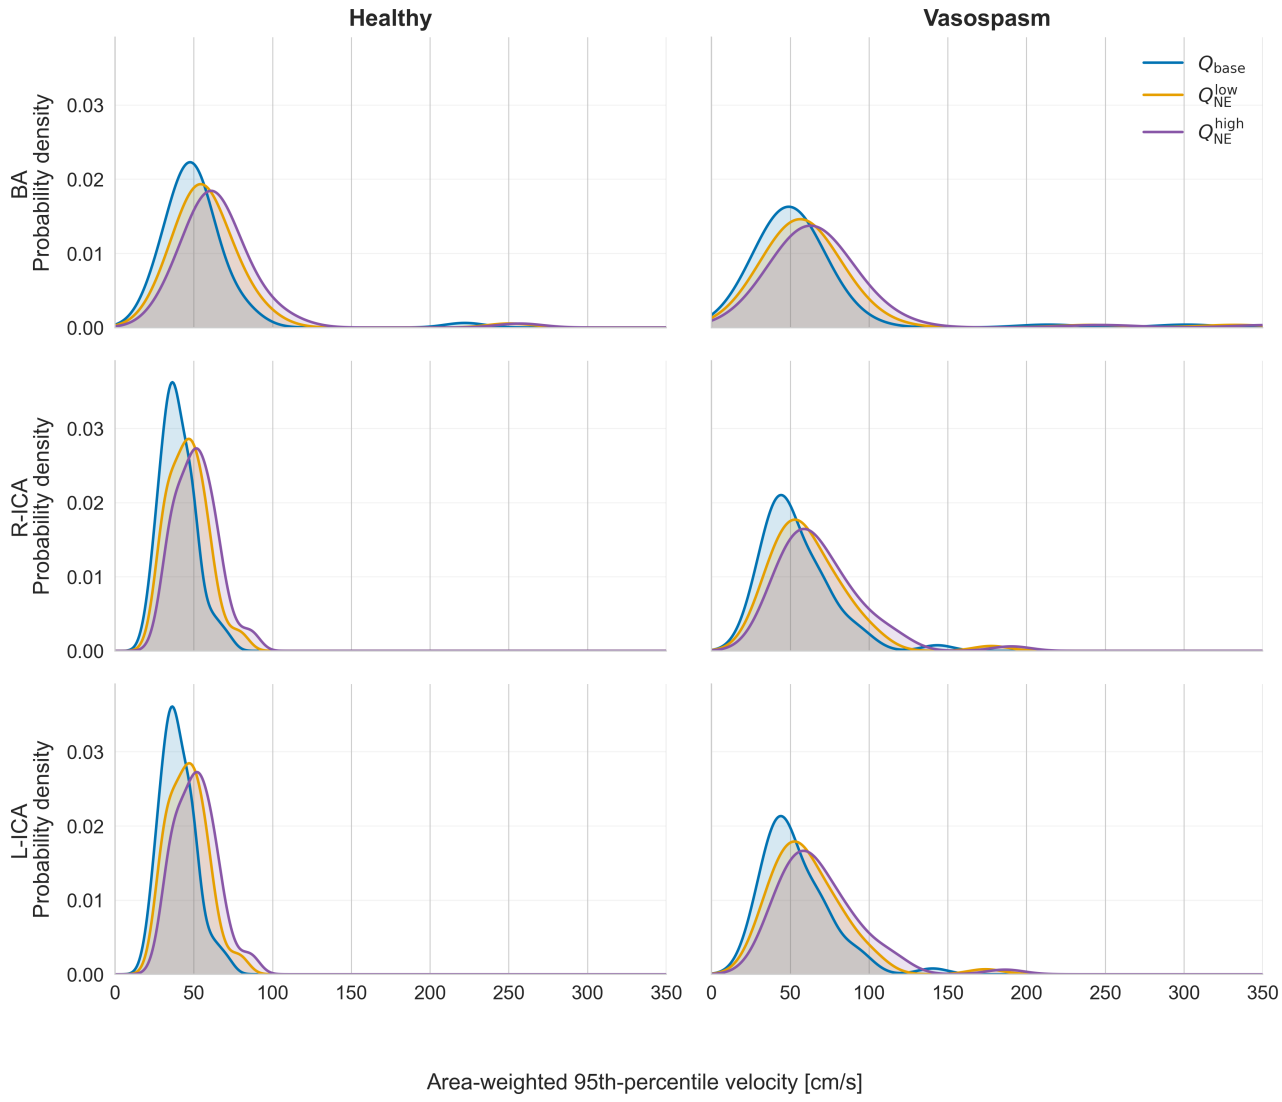

Figure 19: Regime-wise comparison of inlet near-peak velocity distributions for healthy and vasospasm-induced geometries, shown separately for the BA, R-ICA, and L-ICA. In both healthy and vasospasm-induced configurations, the distributions shift moderately toward higher values from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, consistently with the increasing hemodynamic support imposed under the noradrenaline-induced flow-increase regimes.

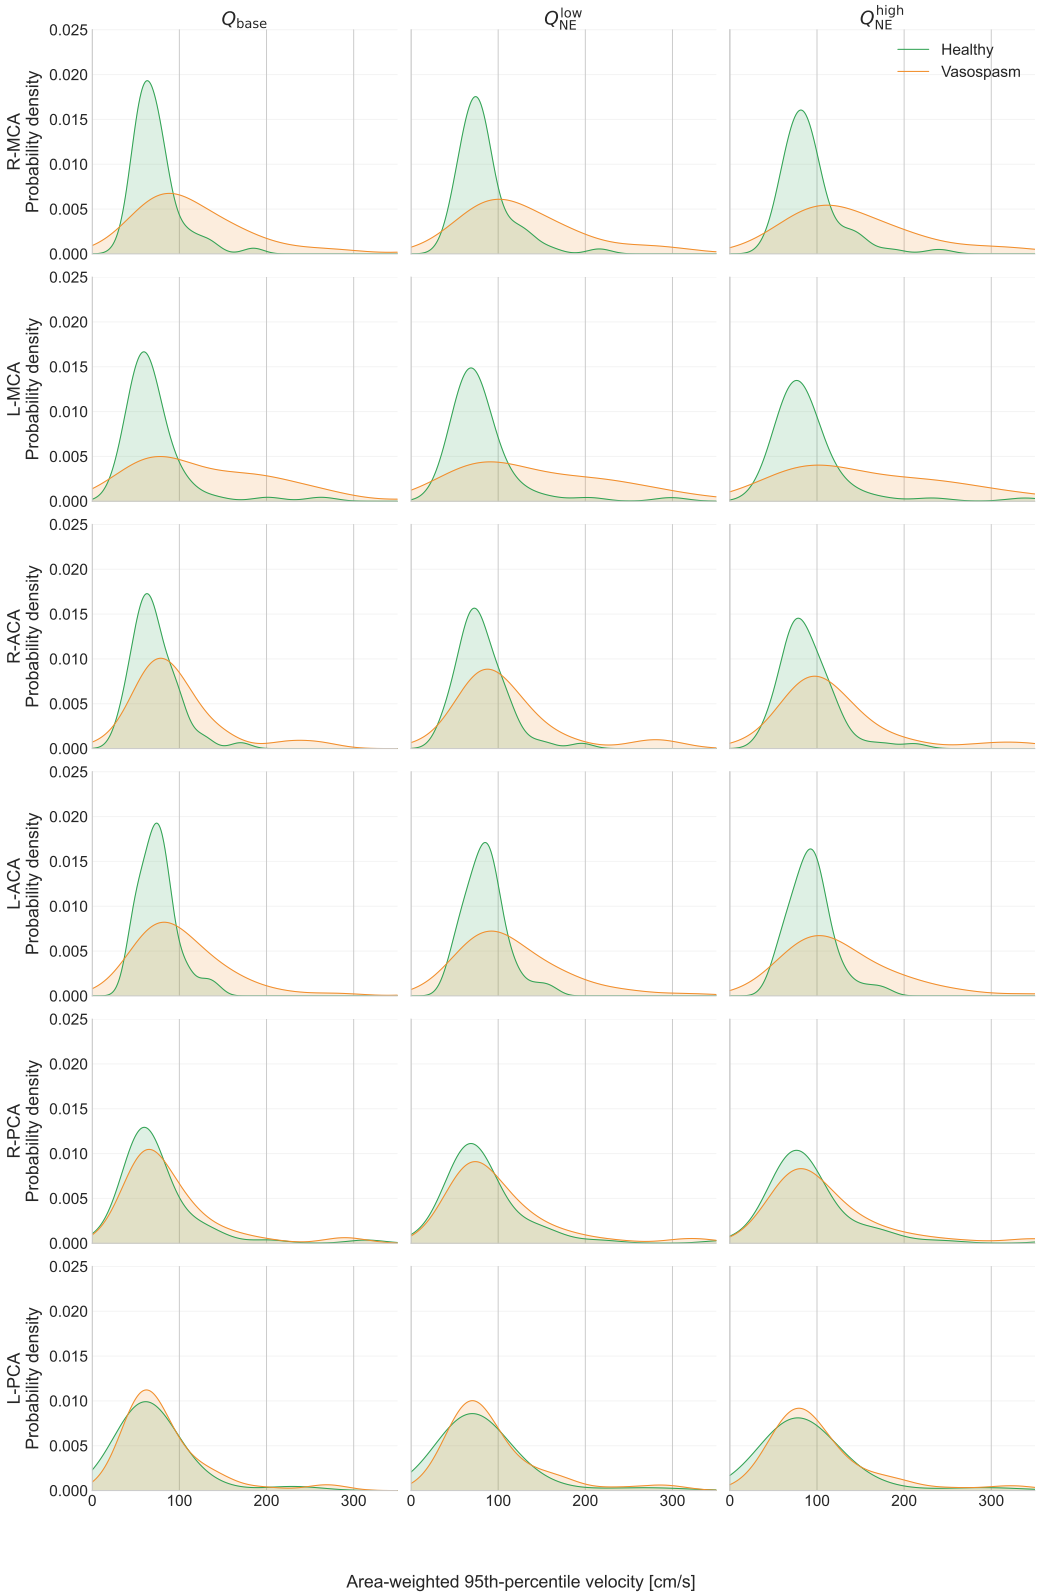

Figure 20: Distribution of outlet near-peak velocities for the R-MCA, L-MCA, R-ACA, L-ACA, R-PCA, and L-PCA across the three hemodynamic regimes. Velocity was quantified as the area-weighted 95th percentile of velocity magnitude on the outlet boundary section. For each one, physiological and vasospasm-induced geometries are compared. Vasospasm-induced configurations show broader distributions and more evident high-velocity tails, particularly in the MCA and ACA outlets, while PCA distributions remain comparatively more overlapping and within a lower velocity range.

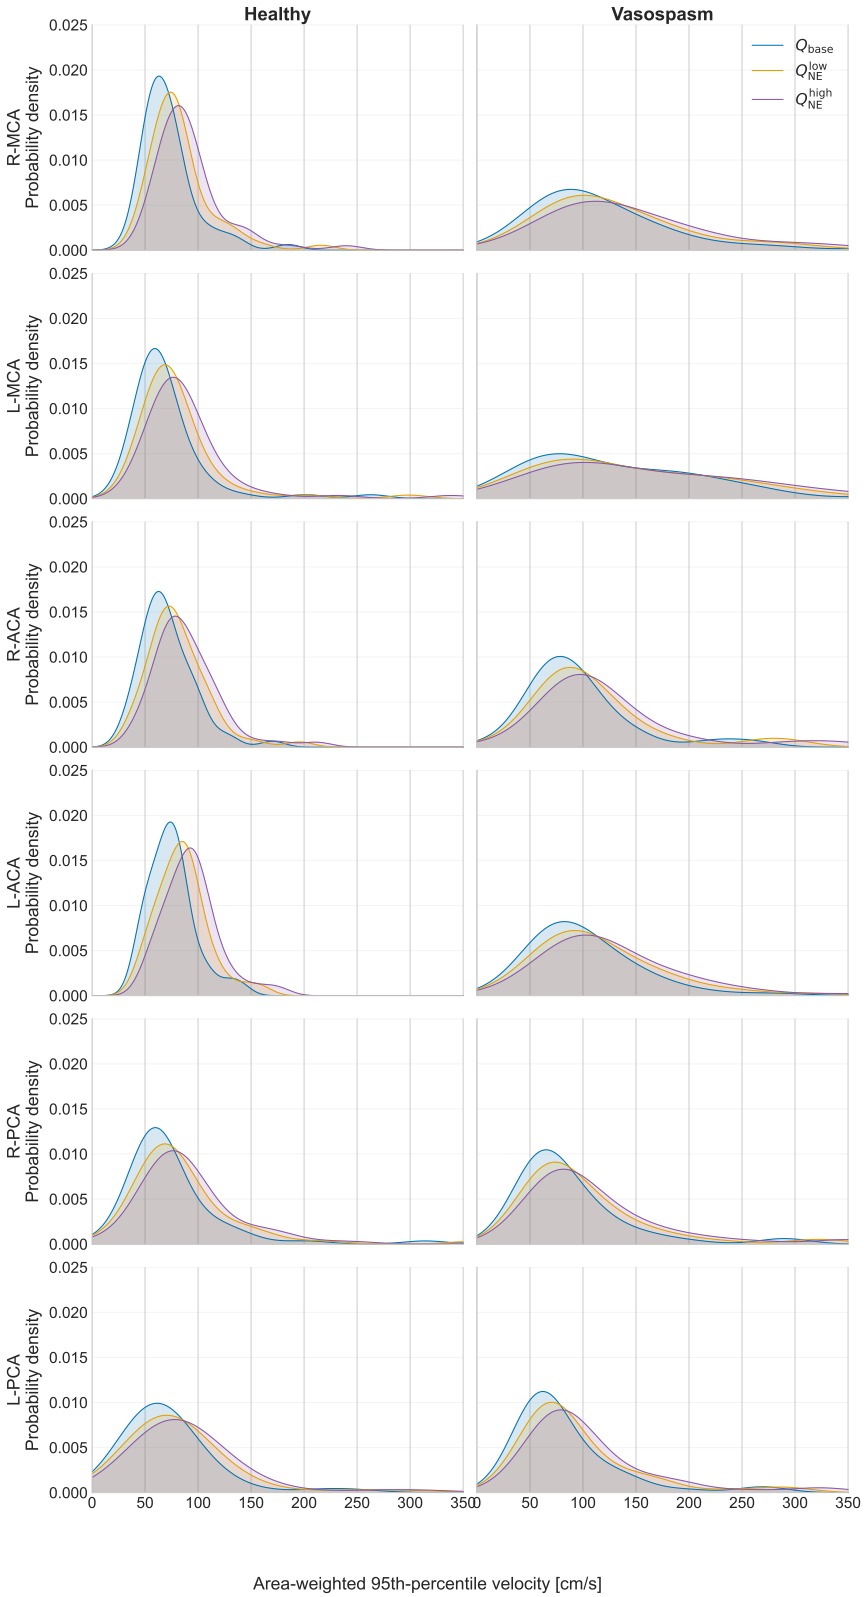

Figure 21: Regime-wise comparison of outlet near-peak velocity distributions for healthy and vasospasm-induced geometries. Velocity was computed as the area-weighted 95th percentile of velocity magnitude at the outlet boundary section. Distributions are shown separately for the R-MCA, L-MCA, R-ACA, L-ACA, R-PCA, and L-PCA. In healthy geometries, the distributions shift moderately toward higher values from Q_{base} to $Q_{\mathrm{NE}}^{\mathrm{high}}$. In vasospasm-induced geometries, the regime-dependent shift is still present, although it is partly masked by the broader variability introduced by arterial narrowing.

Final-model input features

Feature type	Symbol	Feature name	Dimensionality
Node	A_i	Cross-sectional area	1
Node	α_i	Node type	4
Node	t_i	Tangent to centerline	3
Node	ρ_i^A	Area reduction	1
Node	g_i^A	Area gradient	1
Edge	d_{ij}	Normalized vector distance between node i and node j	3
Edge	z_{ij}	Shortest path length between node i and node j	1
Edge	β_{ij}	Edge type	4
Global	MIP	Mean pressure at the inlets	1
Global	v_g	Anatomical variant mask	7
Global	s_v	Vasospasm severity	4
Global	n_i	Noradrenaline administration flag	1
Global	ΔMIP	MIP increase	1

Table 14: Input features included in the final model. The clinically unavailable inlet flow fraction and outlet resistance features are excluded from the final input representation.

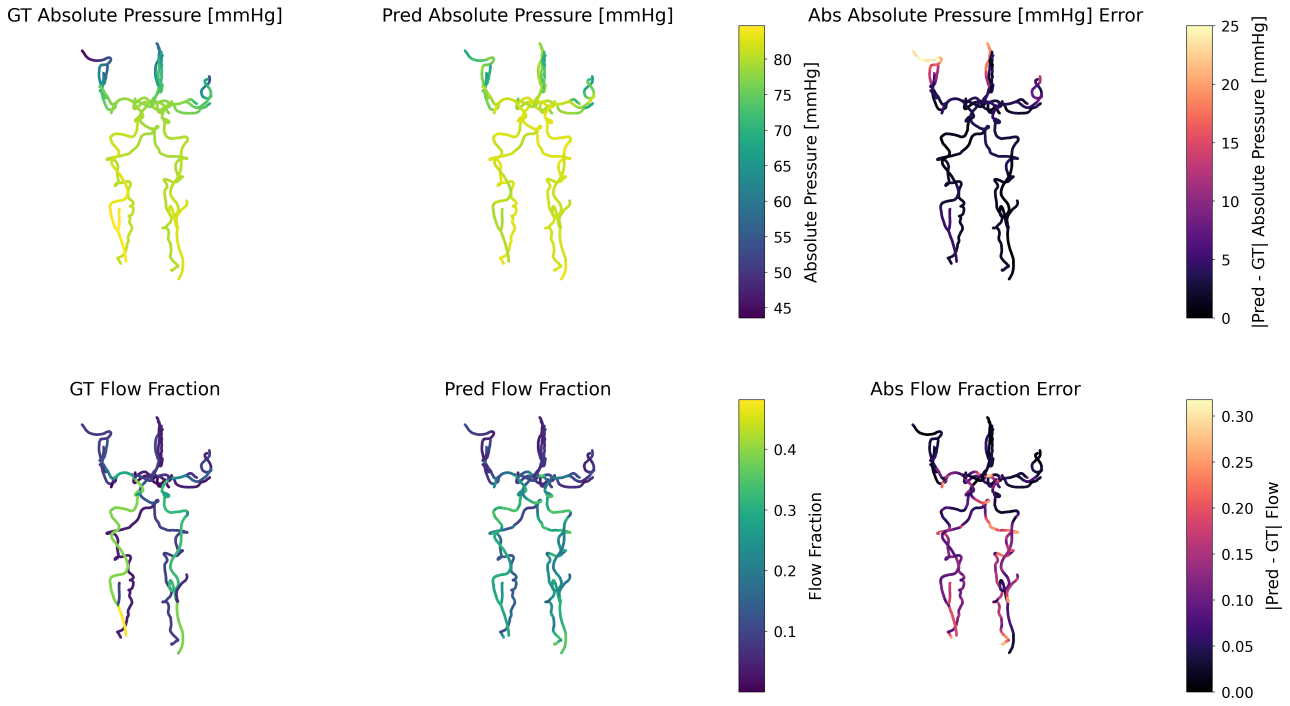

Figure 22: Additional external benchmark prediction from Policlinico di Milano, case 2. The first row reports ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error, while the second row reports ground-truth flow fraction, predicted flow fraction, and absolute flow-fraction error. In this case, pressure errors remain limited in the central Circle of Willis region, while localized discrepancies up to approximately $20 - 25$ mmHg are visible in distal vessels branching from the R-ACA and R-MCA. Flow-fraction errors around $0.25 - 0.30$ can be seen in the upstream vertebral and carotid vessels, while the central vascular region shows comparatively lower flow MAE.

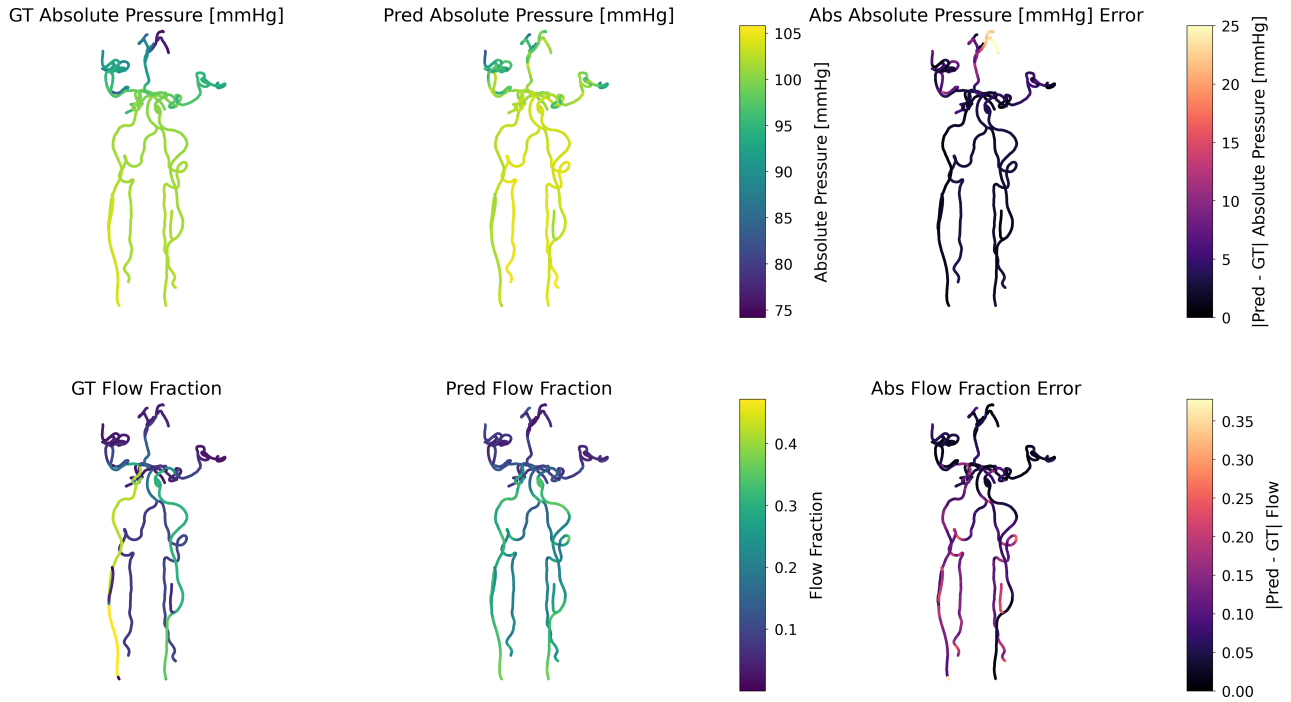

Figure 23: Additional external benchmark prediction from Policlinico di Milano, case 3. The first row reports ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error, while the second row reports ground-truth flow fraction, predicted flow fraction, and absolute flow-fraction error. Pressure discrepancies are mainly localized in upper distal branches, with local errors approaching $20 - 25$ mmHg. Flow-fraction errors in the range of $0.15 - 0.25$ are concentrated along the upstream vertebral and carotid arteries.

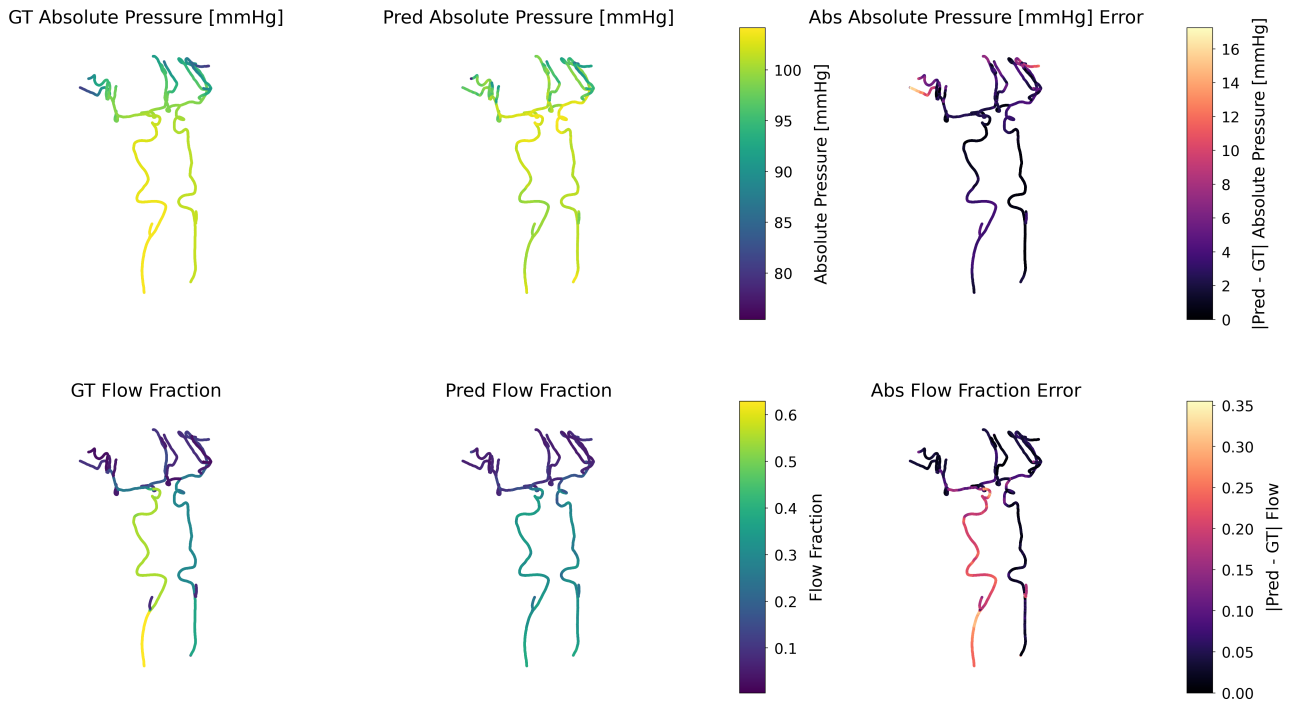

Figure 24: Additional external benchmark prediction from Policlinico di Milano, case 4. The first row reports ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error, while the second row reports ground-truth flow fraction, predicted flow fraction, and absolute flow-fraction error. The pressure error remains comparatively low over most of the central vascular region, with localized discrepancies up to approximately $14 - 16$ mmHg in distal branches, originating from the MCAs. Flow-fraction errors are mainly concentrated in the right carotid artery, reaching around $0.30 - 0.35$, consistent with the larger anatomical domain not represented in the original training geometries.

7. Academic Acknowledgements

I would like to express my sincere gratitude to Politecnico di Milano, which has been much more than a place of study throughout these years. Beyond providing me with a solid academic education, it offered opportunities that profoundly shaped both my personal and professional growth, from international experiences to research activities that ultimately led to this thesis and opened the door to future challenges.

I am deeply grateful to Professor Alberto Redaelli for the opportunity to work on this project and for his invaluable guidance throughout its development. His scientific rigor, expertise, and attention to detail continuously challenged me to look beyond immediate results, to question assumptions, and to strive for a deeper understanding of the work. His feedback and high standards have undoubtedly contributed to my growth as both a student and a researcher.

A special thank you goes to Elisabetta Criseo for her constant support throughout this journey. Her availability, patience, and willingness to answer countless questions were invaluable at every stage of this work. Beyond her technical guidance, I am particularly grateful for the trust and respect she consistently showed me, always encouraging discussion and valuing my ideas and perspectives. Working with someone who treated me not simply as a student, but as a young colleague in the research process, made this experience especially meaningful and rewarding.

This thesis represents the conclusion of an important chapter of my academic journey, and I am sincerely grateful to all those who contributed to making it such a valuable and enriching experience.