POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

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

Laurea Magistrale in Biomedical Engineering - Ingegneria Biomedica

Author: Camilla Balzarotti

Advisor: Prof. Alberto Redaelli

Co-advisor: Elisabetta Criseo

Academic year: 2025-2026

1. Introduction

Aneurysmal subarachnoid hemorrhage (SAH) is a severe cerebrovascular event caused by rupture of an intracranial aneurysm and bleeding into the subarachnoid space [1]. Despite representing only a small fraction of all strokes, it is associated with high mortality, long-term disability, and reduced quality of life among survivors [2–4]. After aneurysm securing, patients remain at risk of secondary neurological deterioration, among which delayed cerebral ischemia (DCI) is one of the most serious complications [5]. Although DCI is now recognized as a multifactorial process, cerebral vasospasm is one of the major mechanisms involved, consisting of sustained narrowing of large and medium-sized intracranial arteries, frequently affecting the proximal vessels of the Circle of Willis [6]. This process increases vascular resistance and may reduce downstream perfusion, potentially leading to irreversible ischemic injury if timely interventions are not applied.

The hemodynamic consequences of vasospasm are strongly influenced by the anatomy of the Circle of Willis [7]. This arterial network connects the anterior and posterior cerebral circulations and can provide collateral pathways when one vascular territory becomes compromised, redistributing between hemispheres and between the carotid and vertebrobasilar systems. However, the Circle of Willis is highly variable across individuals: complete and symmetric configurations are present only in a minority of subjects, while hypoplastic or absent segments are common [8]. These anatomical differences directly affect the capacity of the network to compensate for localized arterial narrowing. Therefore, the same degree of vasospasm may have different hemodynamic consequences depending on the patient-specific vascular configuration, collateral availability, and distribution of vascular resistance.

In clinical practice, a common management strategy for symptomatic vasospasm and suspected delayed cerebral ischemia is induced hypertension [9]. This approach consists of deliberately increasing systemic arterial pressure, often using vasopressor agents such as noradrenaline, with the aim of augmenting cerebral blood flow in territories at risk. However, the treatment response is intrinsically patient-

specific and heterogeneous, because it depends on the location and severity of arterial narrowing, the state of cerebral autoregulation, the available collateral routes, and the risk of systemic complications. As a result, the same pressure target may improve perfusion in one vascular territory but have limited benefit in another, while excessive blood-pressure augmentation may expose patients to cardiopulmonary adverse events. Against this background, key clinical uncertainty is whether increasing arterial pressure will effectively improve downstream supply in the threatened territories, or whether vascular anatomy and limited collateral pathways will restrict flow redistribution despite pressure support.

Current clinical tools provide essential but indirect information for this decision. Angiography describes vessel narrowing, transcranial Doppler provides velocity-based markers of vasospasm, perfusion imaging reflects tissue-level consequences, and systemic pressure targets guide treatment adjustment. However, these measurements do not directly provide a rapid and individualized quantification of how pressure and flow are redistributed throughout the Circle of Willis under alternative vasospasm or pressure-support scenarios. Computational fluid dynamics can fill this gap by simulating patient-specific pressure, velocity, and flow patterns, but its computational cost and preprocessing requirements limit its use in time-sensitive settings or for repeated evaluation of multiple possible conditions in the same patient.

For this reason, surrogate modeling has gained increasing attention as a way to approximate CFD-derived hemodynamic quantities with substantially lower computational cost [10, 11]. Purely data-driven deep learning models can learn complex mappings between vascular geometry, boundary conditions, and flow variables, but they often require large datasets and may generate physically inconsistent predictions. Physics-informed learning addresses this limitation by incorporating physical constraints, such as conservation principles, into the training process [12]. At the same time, the choice of model architecture is crucial. Graph neural networks provide a representation that is naturally aligned with vascular anatomy: vessel segments and centerline points can be represented as graph elements, and information can propagate along physiologically meaningful connections.

The objective of this thesis is therefore to develop and validate a physics-informed graph neural network surrogate for near real-time prediction of patient-specific hemodynamics in the Circle of Willis under vasospasm-inspired conditions and pressure-support regimes. Starting from vascular reconstructions derived from clinical imaging, the framework represents each anatomy as a graph and predicts relative pressure and normalized flow redistribution across the vascular network. 3D CFD simulations are used as ground truth during model development, while physics-informed constraints encourage physiologically plausible predictions, particularly with respect to flow conservation. In this context, the proposed framework aims to bridge the gap between detailed computational hemodynamics and clinically deployable decision-support tools for the management of vasospasm after aneurysmal subarachnoid hemorrhage. The final surrogate provides near-instantaneous predictions of pressure and normalized flow redistribution in the Circle of Willis, while avoiding dependence on clinically unavailable boundary-condition quantities such as outlet resistance and inlet flow fraction.

2. Materials and Methods

2.1. Dataset generation and vasospasm modeling

A patient-derived dataset of Circle of Willis geometries was generated starting from the TopCoW2024 dataset, which provides vascular segmentations and anatomical annotations from CTA and MRA acquisitions. The original volumetric geometries were converted into a centerline-based representation, providing the basis for both the definition of CFD boundary sections and the graph neural network surrogate. The dataset was designed to capture clinically relevant anatomical variability of the Circle of Willis, including complete configurations, unilateral or bilateral posterior communicating artery absence, fetal-type posterior cerebral artery configurations, A1 hypoplasia, and anterior communicating artery absence. The distribution of these variants was chosen to approx-

imately reflect population frequencies reported in the literature [8], while preserving the original patient-derived geometry whenever possible. To model post-subarachnoid hemorrhage vasospasm-inspired conditions, controlled arterial narrowing was introduced in 60 out of 125 patient-specific geometries. Since no reliable data were available on vasospasm occurrence across specific Circle of Willis variants, the modified cases were selected to approximately preserve the anatomical-variant distribution of the cohort and limit selection bias. Vasospasm was simulated by locally reducing the diameter of selected arterial segments in the anterior and posterior circulation. Narrowing severity was classified as mild, moderate, or severe using percentile-based criteria by comparing the radius of each modified geometry with the corresponding non-narrowed reference geometry along the centerline. Moderate and severe cases were intentionally enriched because they were expected to produce more pronounced hemodynamic alterations and therefore provide more informative examples for model training and evaluation.

2.2. CFD simulations and hemodynamic targets

3D CFD simulations were performed in SimVascular to generate the reference hemodynamic fields used as ground truth for the surrogate model. Blood was modeled as an incompressible Newtonian fluid and vessel walls were assumed rigid. Under these assumptions, the mathematical foundation is represented by the incompressible Navier–Stokes equation. Three steady-state hemodynamic regimes were simulated for each geometry: patient-specific baseline flow (Q_{base}), low noradrenaline-induced flow increase ($Q_{\text{NE}}^{\text{low}}$), and high noradrenaline-induced flow increase ($Q_{\text{NE}}^{\text{high}}$), mimicking increasing pressure-support conditions. Total inflow was sampled from three increasing physiological intervals and distributed among the basilar artery and the internal carotid arteries according to anatomical criteria. Outlet boundary conditions accounting for distal vasculature were prescribed through resistance-based Windkessel models. Since patient-specific flow measurements were not available for the training dataset, outlet resistances were assigned from literature values with controlled variability and kept constant across simulation rounds and between corresponding healthy and vasospasm-induced geometries, consistently with the assumption of impaired autoregulation. Simulations exhibiting non-physiological pressure or velocity values, or for which the solver failed to converge were systematically excluded from the final dataset.

2.3. AI Model Development

2.3.1 Graph representation and surrogate model

Each vascular anatomy was represented as a graph ($G = (V, E)$), where nodes corresponded to sampled centerline points and edges represented connectivity between consecutive centerline points. In addition to physical centerline edges, auxiliary edges were introduced to connect internal nodes to their nearest boundary nodes in terms of graph distance. Node features encoded local geometry, topology, and pathology, including cross-sectional area, node type, centerline tangent, area reduction, and area gradient. Edge features captured spatial and topological connectivity through normalized relative position, shortest-path distance, and edge type. A global-context vector was also incorporated to describe patient- and simulation-level information, including mean inlet pressure, anatomical variant, vasospasm severity, noradrenaline-related pressure-support condition, and mean inlet pressure increase.

As detailed in Figure 1, the surrogate model was implemented as a physics-informed graph neural network with an encode–process–decode structure. Separate encoders mapped node, edge, and global-context features into latent representations. The processor then performed 20 message-passing steps, during which edge and node embeddings were iteratively updated using neighboring information and global context. This allowed the model to learn both local geometric effects and long-range hemodynamic interactions across the vascular network. Finally, the decoder mapped the updated node embeddings to the predicted relative pressure and normalized flow fraction. Both single-decoder and double-decoder multi-task architectures were evaluated.

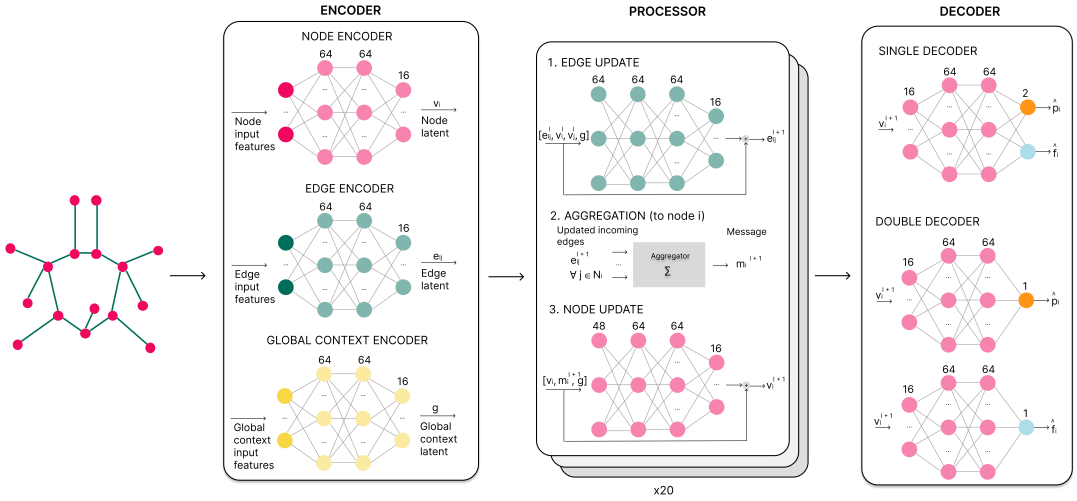

Figure 1: Schematic representation of the PiGNN architecture.

2.3.2 Training strategies, loss function and evaluation

Patients were split into training, validation, and test sets using a 60/20/20 partition. Importantly, this subdivision was performed at the patient level to avoid data leakage between anatomies or simulation rounds belonging to the same patient, thus providing a more realistic estimate of generalization to unseen vascular geometries. Similarly, input features were standardized using statistics computed only on the training set. The model was trained end-to-end using regularization strategies including early stopping, weight decay, and cosine learning-rate scheduling. As reported in Equation 1, the training objective combined a supervised data-fitting term with a physics-informed mass-conservation term:

$$\mathcal{L} = \frac{1}{B}\sum_{b=1}^{B}\left[\delta\mathcal{L}_{\text{data}}^{(b)} + (1-\delta)\mathcal{L}_{\text{phys}}^{(b)}\right]. \quad (1)$$

The data term penalized the mismatch between predicted and CFD-derived pressure and flow at the node level, using a Huber loss for pressure and a squared error for normalized flow, while the physics component enforced global mass conservation by penalizing discrepancies between the total predicted inlet and outlet flow fractions.

Model performance was evaluated using graph-wise metrics. Specifically, pressure accuracy was quantified through the pressure relative error:

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

while flow accuracy was assessed using the mean absolute error of the normalized flow fraction:

$$\text{MAE}_{\text{flow}}^{(g)} = \frac{1}{N_g}\sum_{i=1}^{N_g}\left|\hat{q}_i - q_i\right|. \quad (3)$$

Beyond pointwise pressure and flow errors, physical consistency was assessed by checking mass balance at the inlet and outlet boundaries, and a graph was considered to violate mass conservation when the absolute difference between prescribed and predicted total boundary flow exceeded 0.05.

Model definition was performed in sequential steps. First, different multi-task architectures were compared by evaluating single- and double-decoder formulations, together with the potential benefit of fine-tuning from a flow-only pretrained network, rather than training from scratch. After selecting the baseline architecture, an ablation study was performed to determine whether clinically unavailable quantities could be removed without degrading predictive performance or physical consistency, while also quantifying the contribution of individual and grouped input features. Finally, the robustness of the final model was benchmarked against four

external patient-specific geometries from Policlinico di Milano, reconstructed from CTA images and simulated with patient-specific CFD boundary conditions calibrated using available clinical and transcranial Doppler information.

3. Results

3.1. CFD-derived hemodynamic behaviour

The CFD simulations were first analyzed to verify that the generated dataset reproduced the expected hemodynamic trends across the three simulated pressure-support regimes and between physiological and vasospasm-induced geometries.

Absolute pressure distributions progressively shifted toward higher values from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, consistently with the imposed increase in inlet flow. Across all simulated geometries, mean absolute pressure increased from 128.3 ± 20.3 mmHg in Q_{base} to 147.8 ± 23.6 mmHg in $Q_{\text{NE}}^{\text{low}}$ and 166.0 ± 27.0 mmHg in $Q_{\text{NE}}^{\text{high}}$. Beyond this global pressure increase, vasospasm-induced geometries exhibited larger spatial pressure heterogeneity within the vascular network. The robust pressure range, defined as the difference between the 95th and 5th percentiles of centerline pressure within each geometry, was indeed consistently higher in vasospastic configurations than in healthy ones.

Velocity fields showed a similar controlled response with maximum velocities at inlet and outlet sections raising progressively with the increase inlet flow rate. Vasospasm-induced configurations also showed broader velocity distributions and included more high-velocity cases, especially in the MCA and ACA outlets, whereas PCA distributions remained more similar between non-narrowed and narrowed geometries.

Overall, the CFD analysis confirmed that the simulated dataset captured both the global effect of pressure-support regimes and the local hemodynamic impact of vasospasm-induced narrowing.

3.2. Model architecture comparison and ablation-driven feature selection

The architecture comparison evaluated the effect of decoder design and fine-tuning strategy on joint pressure-flow prediction.

While the flow-only model achieved the lowest flow mean absolute error, equal to $3.5 \cdot 10^{-2}$, it did not provide pressure predictions and was therefore unsuitable for the final multi-task objective. Among the multi-task models trained from scratch, single- and double-decoder architectures showed comparable pressure and flow accuracy, with pressure relative error around $4.4 \cdot 10^{-2}$ and flow MAE around 5.1–$5.2 \cdot 10^{-2}$. However, both models exhibited high outlet mass-conservation violation rates.

Fine-tuning from a flow-specialized model improved the physical consistency of the predictions. In particular, the fine-tuned double-decoder model achieved a pressure relative error of $4.7 \cdot 10^{-2}$, a flow MAE of $5.2 \cdot 10^{-2}$, and no outlet mass-conservation violations. This configuration was therefore selected as the baseline architecture for the following analyses. Nevertheless, it still relied on the full feature set, including outlet resistance and inlet flow fraction. Since these quantities are not readily available in routine clinical practice, an ablation study was performed to evaluate whether they could be removed without compromising predictive performance.

The ablation results revealed that the effect of feature removal differed between pressure and flow prediction. Flow MAE improved across the three hemodynamic regimes after several feature removals, suggesting partial redundancy among input features, while pressure was more sensitive to the removal of features describing vascular anatomy, vessel caliber, and pressure-related global context, particularly under Q_{base}. Furthermore, some ablations improved pointwise accuracy while worsening mass conservation, confirming that numerical error and physical consistency capture complementary aspects of model quality. Jointly considering target accuracy, boundary-flow consistency, and clinical applicability, the simultaneous removal of outlet resistance and inlet flow fraction represented the most appropriate compromise, since it preserved accuracy close to the full-feature baseline,

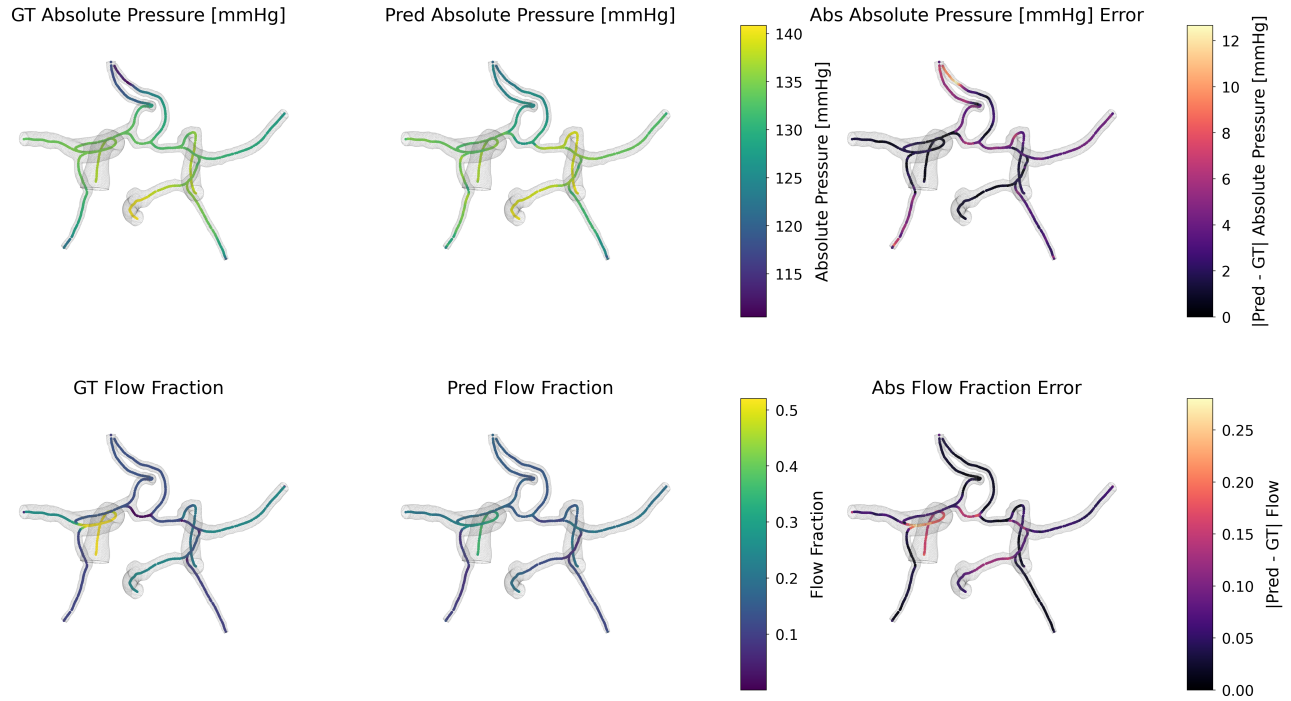

Figure 2: Representative example of final-model predictions on a test graph with error close to the median test performance. Ground-truth and predicted absolute pressure and flow fraction are shown together with the corresponding error maps. Absolute pressure errors remain low over most of the central 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.

while excluding two quantities that are not readily available in routine clinical practice.

3.3. Final model performance

Since inlet mass-conservation violations remained frequent, the model selected in the ablation was further improved by incorporating an explicit inlet conservation term in the loss function. This refinement substantially enhanced physical consistency at the inlet boundaries while preserving pressure and flow prediction accuracy. Therefore, the final model was defined as a fine-tuned double-decoder physics-informed graph neural network trained without outlet resistance and inlet flow fraction as input features. A representative prediction example is shown in Figure 2. As reported in Table 1, the final model achieved a global pressure relative error of $4.4 \cdot 10^{-2}$ and a flow MAE of $5.2 \cdot 10^{-2}$ on the internal test set. Performance remained stable across the three simulation hemodynamic regimes, indicating that the model generalized consistently across the imposed hemodynamic regimes. The average inference time was approximately $3.5 \cdot 10^{-2}$ s per sample, supporting the use of the model for near real-time scenario evaluation.

Set	PRE	Flow MAE
Q_{base}	4.23×10^{-2}	5.14×10^{-2}
$Q_{\text{NE}}^{\text{low}}$	4.43×10^{-2}	5.21×10^{-2}
$Q_{\text{NE}}^{\text{high}}$	4.62×10^{-2}	5.34×10^{-2}
Global average	4.43×10^{-2}	5.23×10^{-2}

Table 1: Final-model performance on the internal test set.

Compared with the reduced-feature model before inlet-loss refinement, the final model introduced only a limited increase in flow MAE and outlet violation rate, while reducing inlet mass-conservation violations from 100% to 1.67%. The final model was therefore selected because it provided a more balanced compromise between predictive accuracy, clinical usability, and phys-

ical consistency at both inlet and outlet boundaries.

3.3.1 Performance across anatomical variants, node types, and vasospasm severity

Performance was further analyzed across anatomical variants, node types, and vasospasm status to identify the most challenging prediction settings. Complete Circle of Willis anatomies showed the lowest pressure error, with mean PRE around $1.6 \cdot 10^{-2}$, while non-complete variants exhibited higher pressure errors and larger variability, especially when communicating or posterior circulation pathways were altered. Flow prediction was less sensitive to anatomical variation, remaining within a narrower error range across variants.
Node-wise analysis showed that outlet nodes exhibited higher and more variable pressure errors, whereas flow errors were largest at inlet nodes and junctions. This suggests that boundary regions and branching locations remain the most challenging areas, due to outlet pressure accumulation, inlet-flow partitioning, and local flow splitting.
Since vasospasm represents the main pathological scenario targeted by the framework, model robustness was assessed on 30 matched healthy and vasospasm-induced geometries. Vasospasm increased PRE by an average of $5.28 \cdot 10^{-2}$, whereas flow MAE increased by only $9.56 \cdot 10^{-3}$. The pressure error increase was severity-dependent, reaching $8.37 \cdot 10^{-2}$ in severe cases, while flow errors showed smaller variations. These results indicate that the model remains applicable to vasospasm-induced configurations, although pressure prediction is more sensitive to severe narrowing.

3.3.2 Benchmarking

On the external benchmark, PRE raised from $5.1 \cdot 10^{-2}$, while flow MAE increased from $5.2 \cdot 10^{-2}$ to $8.2 \cdot 10^{-2}$. Although outlet mass conservation deteriorated, the largest errors were mainly localized in vessels outside the original training domain, while predictive ability was better preserved in the central Circle of Willis region. A representative benchmarking prediction is shown in Figure 3.

4. Discussion

4.1. Interpretation of the CFD-derived dataset

The CFD analysis represents an essential validation step, since the simulated pressure and velocity fields constitute the reference data used to train the graph neural network surrogate. As explained in Section 3.1, the generated dataset captured two clinically meaningful hemodynamic effects: the progressive pressure increase associated with the simulated pressure-support regimes and the local redistribution effects induced by vasospasm-like narrowing.
Absolute pressure shifted progressively from Q_{base} to $Q_{\text{NE}}^{\text{high}}$, confirming that the simulation design produced controlled hemodynamic regimes with increasing pressure levels, without implying a direct prediction of the in vivo response to noradrenaline. Vasospasm-induced geometries also exhibited larger within-geometry pressure ranges, indicating increased spatial pressure heterogeneity due to narrowing-induced pressure losses, especially under higher-flow regimes.
Velocity distributions further supported the physiological consistency of the dataset. Vasospasm-induced configurations showed broader distributions and higher-velocity tails, particularly in the MCA and ACA outlets, coherently with lumen reduction. This same principle underlies transcranial Doppler monitoring of vasospasm, although CFD-derived velocities should not be interpreted as direct clinical Doppler values. Overall, the CFD dataset provided a physically meaningful ground truth for learning pressure and flow redistribution.

4.2. Target formulation and model behaviour

A central methodological choice of this work was the definition of the model outputs as relative pressure and normalized flow fraction.
Pressure was expressed relative to the mean inlet pressure, because in incompressible flow the relevant physical quantity is not the absolute pressure offset, but the pressure gradient driving flow through the network. This formulation reduces sensitivity to absolute pressure scale differences across simulations and focuses the learning task

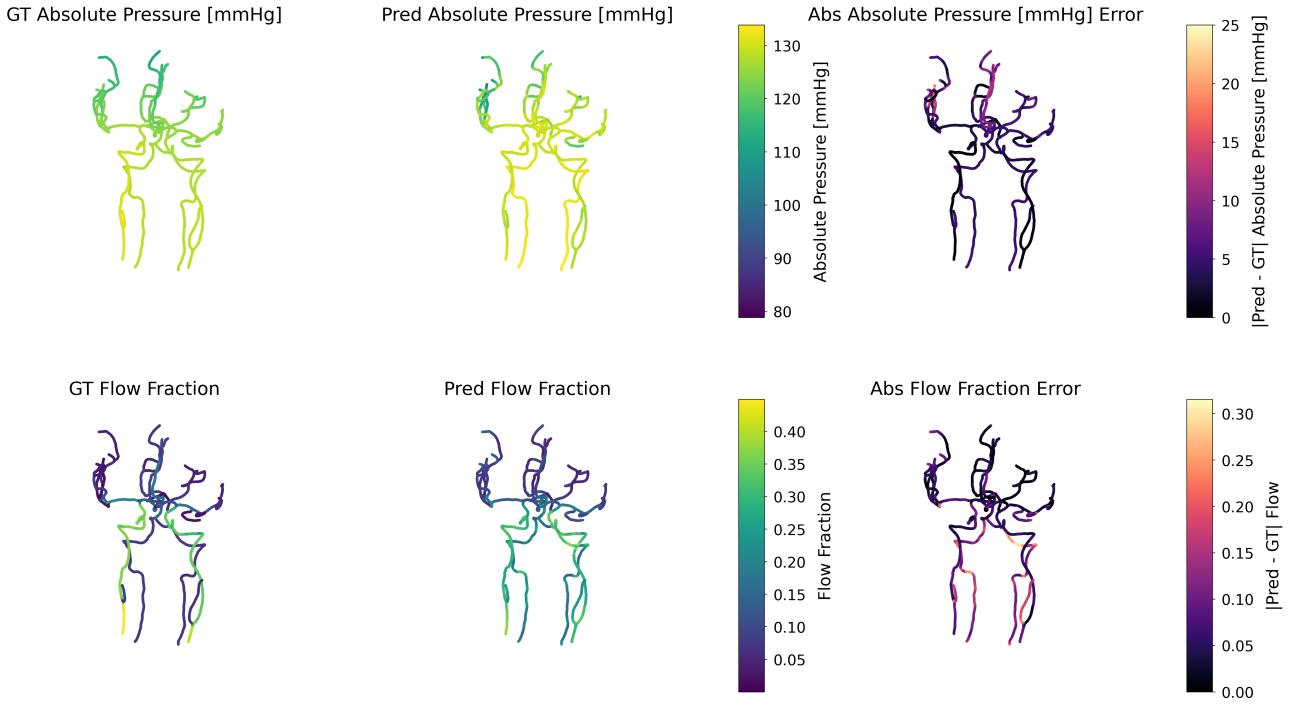

Figure 3: Representative prediction on an external benchmark geometry from Policlinico di Milano. The first row shows ground-truth absolute pressure, predicted absolute pressure, and absolute pressure error. The second row shows ground-truth flow fraction, predicted flow fraction, and absolute flow-fraction 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.

on spatial pressure variations caused by vascular geometry, vessel narrowing, and boundary conditions.

Flow was expressed as a fraction of total inlet flow rather than as an absolute volumetric flow rate. This was motivated by the clinical objective of the thesis, which is not simply to estimate the amount of blood passing through one vessel, but to assess how available inflow is redistributed across the Circle of Willis and its collateral pathways. Against this background, normalized flow fraction provides a scale-independent description of this redistribution, making predictions more comparable across patients and pressure-support regimes.

These target definitions strongly influenced both model training and interpretation of the results. Since flow fraction is strongly constrained by network topology and mass conservation, it was generally more stable across hemodynamic regimes, ablations, and vasospasm conditions. Pressure, in contrast, remained more sensitive to local resistance, vessel caliber, stenosis-induced losses, and pressure-support level. Therefore, the different behaviour observed for the two targets is consistent with their underlying physical meaning.

4.3. Selection of the baseline architecture

As detailed in Section 3.2, the comparison between model configurations was designed to assess both the decoder layout and the training strategy for multi-task pressure-flow prediction. Single- and double-decoder multi-task architectures trained from scratch achieved similar pressure and flow errors, but still produced substantial outlet mass-conservation violations. This highlights that prediction accuracy alone was not sufficient to identify the best surrogate, as locally plausible pressure and flow values did not necessarily guarantee a globally consistent flow balance at the boundaries.

Fine-tuning from a flow-only pretrained model improved this behaviour by leveraging a simpler preliminary task to initialize the subsequent multitask problem. Since flow redistribution is more directly constrained by graph topology and mass conservation, the flow-only stage allowed the network to learn a first mapping of flow redistribution. When this representation was used to initialize the multi-task model, physical consistency improved substantially, suggesting that pretraining provided a more physically meaningful starting point than random weights.
The advantage of the double-decoder design can be interpreted in terms of task specialization. Pressure and flow are related hemodynamic quantities, but they have different numerical distributions and physical sensitivities. A single shared decoder may therefore force the two targets to compete at the output level, whereas separate decoders allow the model to share geometric and topological information in the encoder and processor while using target-specific mappings for the final predictions.
For this reason, the fine-tuned double-decoder PiGNN was selected as the reference architecture. At this stage, however, the reference model still relied on the full simulation-derived feature set. This motivated the following ablation study, which assessed whether a more clinically realistic input representation could be obtained without losing the accuracy and physical consistency achieved by the selected architecture.

4.4. Ablation study and final feature set

The ablation study was performed with a twofold objective: to determine whether the model could rely only on clinically available information, and to assess feature importance. This analysis is significant because some inputs used in simulation-based frameworks, such as outlet resistance and inlet flow fraction, are meaningful from a CFD perspective but are not readily available in routine clinical practice. Their use would require assumptions, empirical estimation, or additional measurements, reducing the reliability and practical applicability of the surrogate.
The ablation results reported in Section 3.2 highlighted that flow prediction often improved after feature removal, suggesting partial redundancy among geometric, topological, and global-context inputs. This does not mean that the removed features were physically irrelevant. Rather, it indicates that similar information could often be recovered from correlated descriptors. For example, vessel area, resistance, area reduction, and topological context all encode information related to the ability of a vascular segment to conduct flow. Removing redundant or noisy inputs can therefore simplify the learning problem and improve generalization. Pressure prediction was more sensitive to feature removal, consistently with the physical nature of pressure, which depends more directly on local pressure losses and boundary conditions.
The ablation study also confirmed that pointwise accuracy and mass conservation must be evaluated separately: some feature removals improved PRE or flow MAE but worsened outlet conservation.
Therefore, the final feature set was selected by jointly considering predictive accuracy, inlet and outlet conservation behaviour, and clinical applicability. Under this combined criterion, removing outlet resistance and inlet flow fraction represented the most appropriate compromise. The resulting feature set preserved accuracy close to the full-feature baseline and improved clinical usability, avoiding dependence on quantities that would not normally be available at inference time. However, this formulation was further refined to address the remaining inlet mass-conservation limitations.

4.5. Final model interpretation

The final model consisted of a fine-tuned double-decoder PiGNN trained with the reduced, clinically oriented feature set and a physics-informed loss enforcing mass conservation not only at the outlets, but also at the inlets. This refinement was introduced because the reduced-feature model selected from the ablation study, although representing a substantial step toward clinical applicability, still exhibited high inlet mass-conservation violations. An explicit inlet mass-conservation term was therefore added to the already present outlet constraint, encouraging the model to learn inlet flow predictions that were globally consistent with the expected total inflow.
This modification substantially improved inlet

mass conservation while maintaining pressure and flow errors close to the reduced-feature baseline. This result highlights the role of the loss function as more than a numerical optimization tool: it defines which types of errors are penalized and therefore shapes the behaviour of the learned model.
The analysis by anatomical variant, node type, and vasospasm status clarified where the model performs most reliably and where limitations remain. Pressure prediction was more sensitive to anatomical variability and severe vasospasm, while flow prediction remained comparatively more stable, being more strongly constrained by topology and conservation. Node-wise errors further showed that boundary and junction regions were the most challenging locations, reflecting the difficulty of predicting inlet-flow partitioning, outlet pressure accumulation, and local flow splitting at bifurcations. The matched comparison between healthy and vasospasm-induced geometries demonstrated that the model remains applicable to the main pathological setting of the thesis, but pressure predictions in severe narrowing cases should be interpreted with greater caution.

4.6. External benchmarking

The external benchmark on four Policlinico di Milano geometries provided a particularly demanding out-of-distribution robustness test. These cases included a broader vascular domain than the training set, with upstream carotid and vertebral arteries and additional distal branches. In this context, the moderate decrease in performance should be interpreted in relation to the substantially more complex anatomical and topological setting, rather than as a loss of predictive capability. Indeed, pressure and flow estimates remained meaningful in the central Circle of Willis region, while the largest errors appeared mainly in distal or extended vessels outside the original training domain.
This occurrance indicates that the model learned transferable hemodynamic relationships within the Circle of Willis, while also highlighting the expected sensitivity of conservation and flow redistribution to outlet multiplicity and graph topology. For instance, the deterioration of outlet mass conservation in the benchmark cases is likely due to the increase in terminal branches in the external geometries, allowing small local flow-fraction errors to accumulate into a larger global residual. Overall, the external benchmark supports the robustness of the proposed surrogate within its intended anatomical domain, while indicating that future training on extended vascular reconstructions would be required to further improve performance on broader clinical geometries.

4.7. Clinical impact

The clinical value of the proposed framework lies in its potential to support near-instantaneous hemodynamic assessments once the patient-specific vascular graph is available. Rather than replacing clinical monitoring or high-fidelity CFD, the tool is intended to support rapid exploration of different patient-specific scenarios. This is particularly relevant in post-SAH management because the clinical question is often time-sensitive and dynamically evolving. For example, the surrogate could be used to rapidly compare alternative pressure-support conditions or plausible vasospasm configurations, without requiring a new CFD simulation for each hypothesis. By providing insight into pressure and flow redistribution mechanisms in the individual vascular anatomy, the model could complement clinical decision-making during monitoring and treatment assessment. In particular, it could help interpret whether collateral pathways are likely to preserve downstream supply, or whether specific vascular territories may be exposed to insufficient perfusion and therefore require closer surveillance or treatment reassessment.
A second important contribution is the generated dataset itself, which combines patient-derived Circle of Willis anatomies, systematic anatomical variants, vasospasm-inspired narrowing, and multiple pressure-support regimes. Given the limited availability of large hemodynamic datasets in cerebrovascular modelling [13], this dataset represents a valuable basis for future work on deep-learning surrogates for Circle of Willis hemodynamics.

4.8. Limitations and future work

Despite the promising performance of the proposed surrogate model, some limitations should be acknowledged.

First, dataset composition may influence model robustness. Rare anatomical variants were represented by fewer samples, since the dataset followed population-inspired frequencies, and vasospasm was synthetically imposed rather than derived from real longitudinal scans. Future datasets should therefore enrich rare anatomical variants and include patient-derived vasospasm geometries when available, especially severe or highly localized cases.
Second, the model is trained on CFD-derived ground truth and therefore inherits the assumptions of the numerical simulations, including rigid walls, Newtonian blood behaviour, steady-state flow, boundary-condition uncertainty, and the projection of three-dimensional fields onto centerline-based graph quantities. Moreover, while the graph representation is appropriate for predicting pressure and flow redistribution, it cannot recover full volumetric information such as wall shear stress, secondary flows, or cross-sectional velocity profiles.
Finally, the current framework assumes that the vascular geometry and graph have already been extracted from clinical images. Clinical translation would therefore require integration with automated image-to-graph processing tools, user-oriented interfaces, and prospective validation on independent clinical cohorts.

5. Conclusion

This thesis developed a physics-informed graph neural network surrogate for fast prediction of patient-specific pressure and flow redistribution in the Circle of Willis under vasospasm-inspired conditions and pressure-support regimes. Starting from patient-derived vascular geometries, a CFD-based dataset was generated to include anatomical variants, controlled vasospasm-induced narrowing, and multiple hemodynamic rounds. The final model, defined through architectural comparison and ablation analysis, combined a fine-tuned double-decoder structure with a clinically oriented feature set excluding outlet resistance and inlet flow fraction. This formulation preserved predictive accuracy while improving applicability to realistic clinical workflows, where such boundary-condition quantities are not directly available.
The results showed that the surrogate can approximate CFD-derived pressure and normalized flow fraction with near-instantaneous inference, while physics-informed constraints improve boundary-flow consistency. External benchmarking confirmed partial generalization beyond the training domain, while highlighting the need for broader anatomical coverage.
Overall, this work demonstrates the feasibility of graph-based physics-informed surrogates for rapid Circle of Willis hemodynamic assessment. Future developments should focus on automated image-to-graph integration, enrichment of rare variants and severe vasospasm cases, and prospective validation on independent clinical cohorts.

6. Acknowledgements

I would like to thank Politecnico di Milano for the education, opportunities, and academic environment that have accompanied my growth throughout these years and made this work possible.
My sincere gratitude is due to Professor Alberto Redaelli for his guidance, expertise, and valuable feedback throughout this project.
A special recognition goes to Elisabetta Criseo for her constant support, availability, and encouragement. Her willingness to discuss ideas, answer countless questions, and value my perspective as an active participant in the research process greatly enriched this experience and contributed significantly to both my academic and personal growth.

References

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

[2] 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.

[3] 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.

[4] 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.

[5] 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.

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

[7] Angela Straccia, Michael C Barbour, Fanette Chassagne, David Bass, Guilherme Barros, Daniel Leotta, Florence Sheehan, Deepak Sharma, Michael R Levitt, and Alberto Aliseda. Numerical modeling of flow in the cerebral vasculature: understanding changes in collateral flow directions in the circle of willis for a cohort of vasospasm patients through image-based computational fluid dynamics. *Annals of biomedical engineering*, 52(9):2417–2439, 2024.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.