

<https://doi.org/10.1038/s41746-026-02404-z>

Physics constrained graph neural network for real time prediction of intracranial aneurysm hemodynamics

Check for updates

Vincent Lanelongue¹, Paul Garnier¹, Pablo Jeken-Rico¹, Aurèle Goetz¹, Philippe Meliga¹, Yves Chau² & Elie Hachem¹ ✉

Intracranial aneurysms (IAs) are life-threatening vascular conditions requiring accurate risk assessment to guide treatment. Hemodynamic biomarkers such as wall shear stress and oscillatory shear index are promising predictors of rupture risk but remain underused clinically due to the high computational cost of traditional CFD methods. We propose a physics-constrained graph neural network (GNN) framework trained on high-fidelity CFD data to predict full 3D, time-resolved hemodynamic fields throughout the cardiac cycle. Our model incorporates enhanced node features and physics-based constraints to capture complex spatio-temporal flow behavior in near real time. It generalizes to varying inflow conditions and unseen patient-specific geometries with no fine-tuning. Additionally, we release a benchmark dataset of 105 patient-derived aneurysm geometries with CFD fields to support the machine learning (ML) community. This is the first GNN model applied to transient 3D aneurysmal flow prediction, paving the way for rapid, AI-driven hemodynamic analysis toward risk stratification and treatment planning.

Intracranial aneurysms (IAs) are pathological dilations of cerebral arteries, most frequently found in the Circle of Willis, that affect approximately 2–3% of the population. These aneurysms pose a significant clinical challenge due to their potential to rupture, resulting in subarachnoid hemorrhage, a condition with a case fatality rate over 30%^{1,2}. Current clinical decisions regarding aneurysm management rely heavily on static, macroscopic metrics such as size, shape, and location, along with patient-specific factors like age and hypertension history³. However, these metrics have limited reliability in predicting rupture risk, often leading to inconclusive or contradictory assessments⁴.

Recent advances in computational methods and imaging technologies have enabled the modeling of aneurysmal blood flow, highlighting hemodynamic biomarkers such as wall shear stress (WSS), oscillatory shear index (OSI), and intra-aneurysmal pressure gradients as key determinants of rupture risk⁴. Computational fluid dynamics (CFD) has been instrumental in providing high-fidelity, patient-specific hemodynamic simulations^{5,6}. Despite its utility, the prohibitive computational cost, time-intensive setup, and specialized expertise required for CFD limit its integration into clinical workflows, where timely decision-making is critical.

In parallel, artificial intelligence (AI) and machine learning (ML) methods have demonstrated transformative potential across a range of neurovascular interventions, as evidenced by their application in device

tracking during endovascular procedures. For instance, real-time AI-assisted platforms like Neuro-Vascular Assist (iMed technologies, Tokyo, Japan) have improved precision and reduced cognitive load during interventions such as aneurysm coiling and liquid embolization^{7,8}, demonstrating the clinical feasibility of AI-driven decision support. Inspired by these successes, researchers have begun exploring ML-based surrogates for CFD, aiming to achieve accurate, near-instantaneous predictions of hemodynamics⁹. Traditional ML approaches have largely relied on convolutional neural networks (CNNs) due to their effectiveness in image-based tasks. While CNNs have shown promise in predicting certain hemodynamic parameters from 2D or simplified 3D data^{10–12}, their reliance on structured grids makes them ill-suited for the complex, unstructured nature of aneurysm geometries. This limitation is particularly patent when attempting to model the intricate, three-dimensional flow patterns critical to understanding aneurysm pathophysiology.

Graph neural networks (GNNs) provide a compelling alternative by leveraging the graph-like structure of CFD domains, enabling the direct processing of irregular, non-Euclidean data. Going beyond point cloud-based methods like PointNet and PointNet++^{13,14}, which have been tested on simplified geometries and stationary hemodynamics predictions^{15,16}, GNNs rely on message passing mechanisms to iteratively update unstructured graphs and efficiently learn local patterns.

¹Mines Paris - PSL, University Centre for Material Forming (CEMEF) CNRS, Sophia Antipolis Cedex, France. ²Interventional Neuroradiology Department, Nice University Hospital, Nice, France. ✉e-mail: elie.hachem@minesparis.psl.eu

Recent studies have demonstrated the efficiency of GNNs in predicting 2D time-dependent flows, steady-state 3D flows, and external flows around simple geometries, particularly in physics-driven tasks involving partial differential equations and numerical simulation^{17–22}. To the best of our knowledge, while GNNs have been used in biomedical blood flow and hemodynamic studies^{23–28}, their application to time-dependent 3D blood flow predictions with intricate vortex and swirling patterns within IAs remains an unexplored frontier.

This work introduces a novel framework that integrates CFD-trained, physics-constrained graph neural networks (GNNs) to predict full 3D, time-resolved hemodynamic fields for IAs throughout the cardiac cycle. To enhance the accuracy of capturing complex spatio-temporal dynamics, the framework incorporates physics-based embedding laws and constraints directly into the learning process. Additionally, it leverages detailed node feature enhancements within the graph representation, enabling precise modeling of intricate hemodynamic variations. Therefore, the proposed learning-prediction framework achieves a balance between computational efficiency and accuracy, addressing the need for high-fidelity simulations while meeting the immediacy required for clinical applications. To further validate the generalizability of our approach, we perform two inference experiments on out-of-distribution cases: new inflow profiles unseen during training, and four fully patient-specific aneurysm geometries. These results highlight the model's robustness and its potential as a transferable foundation for future clinical extensions. Finally, to support the broader research community, we introduce a comprehensive dataset comprising 105 patient-specific aneurysm geometries coupled with high-resolution CFD simulations, referred to as BenchAnXplore. Unlike existing datasets, which often suffer from limited size or heterogeneity, this dataset enables standardized evaluation of ML models in aneurysm modeling.

This work lays the foundation for future clinical applications with significant potential impact. Indeed, real-time prediction of hemodynamics could support future advances in rupture risk assessment, enable more personalized treatment planning²⁹, assist in optimizing interventions such as flow diverters or coiling³⁰, and accelerate long-term simulations of IAs healing³¹. Moreover, the introduction of a comprehensive benchmark dataset helps address a major limitation in the field, the scarcity of large-scale, high-quality training data for ML models. By integrating modern ML architectures with domain-specific knowledge, this study contributes a foundational step toward a new generation of AI-assisted neurovascular modeling. It lays the groundwork for future innovations in clinical decision support, with the long-term goal of improving outcomes for patients with IAs.

Results

Dataset

The dataset consists of 105 aneurysm 3D meshes, created by intersecting an ideal half torus with aneurysm bulges extracted from patient cases in the *Intra* dataset³², following Goetz et al. methodology³³. We refer to them as semi-idealized geometries, as they combine an idealized toroidal parent artery with a patient-specific aneurysm bulge, as illustrated in Fig. 1.

These chosen geometries represent sidewall aneurysms located on the internal carotid artery (ICA), particularly near the ICA siphon. This choice

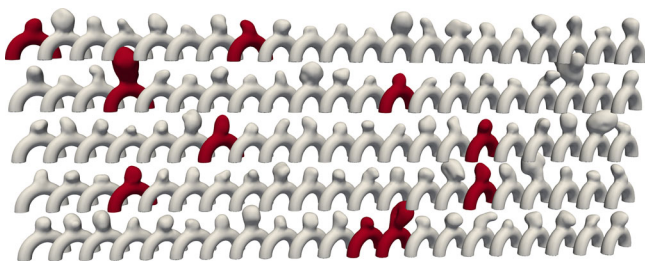


Fig. 1 | Dataset of 105 semi-idealized patient-specific geometries³³ with 10 highlighted test cases.

is motivated by its clinical representativeness: the ICA siphon region is a common site for aneurysm formation and is frequently modeled in hemodynamic studies due to its characteristic curvature³⁴. Several prior works have highlighted the ICA's susceptibility to aneurysm development at curved segments due to altered flow patterns and WSS concentrations³⁵. Notably, the siphon's curvature is often approximated as toroidal, making the torus-shaped backbone both anatomically justified and reproducible across computational settings.

The bulges range from 3 to 9 mm in diameter, covering diverse pathological shapes, and the selected vessel diameter in our dataset is 4 mm, which aligns with clinical measurements of the ICA in adult patients typically ranging from 3 to 6 mm³⁶.

We generated the **BenchAnXplore** dataset following the meshing, coarsening, and simulation processes further detailed in the Methods section, and obtained 105 aneurysm geometries, with 80 simulated timeframes for each of these cases. This dataset is now made available as a benchmark case for AI IAs applications.

Training

We trained and compared four models to evaluate the impact of node feature enhancement and physically-informed improvements separately. To ensure a fair comparison, all models were trained using the same configuration. All MLPs (multilayer perceptrons) used for encoders, decoders, and update functions within the GN blocks have two hidden layers of 128 neurons each and use a ReLU activation function. We used 15 rounds of message passing, bringing the total number of parameters to approximately 3 million.

With 105 simulations, each spanning 80 timesteps, we trained on 95 simulations and randomly selected 10 for testing and evaluation (highlighted in red in Fig. 1). This test set gathers a wide variety in shapes and flow characteristics as depicted in Figure 2. Given that our meshes contain in average, 20,000 nodes, this results in over 75 million training tokens. Recall that the number of training tokens is equal to trajectories * timesteps * data points.

We trained for 20 epochs, e.g., 144k training steps, using a fixed learning rate of 10^{-4} for 16 epochs followed by 4 epochs with exponential decay to reach a minimum learning rate of 10^{-7} . Training was conducted on L4 NVIDIA GPUs, taking approximately 20 h, with a slight (~5%) increase for the physically informed network due to gradient computations.

Boundary conditions. To handle boundary conditions over the domain, we added a node type as an input feature to separate different training behaviors according to the considered nodes. We fixed a null velocity for the wall nodes and did not compute loss nor performed learning onto these nodes. The velocity at the inlet nodes was fixed; we imposed the same pulsatile parabolic inflow profile as in the ground-truth CFD

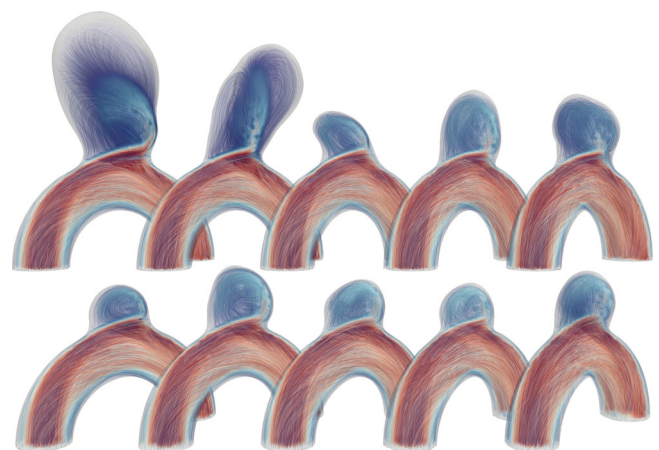


Fig. 2 | Test set of 10 representative blood flow CFD simulations.

simulations. All other nodes of the fluid domain, outflow included, were used to train the network.

Training noise. Our models are trained for 1-step prediction tasks but are intended to be used in a “rollout” scenario, where they iteratively use their own predictions to generate an entire simulation. In such scenarios, the model’s input data can deviate from the training distribution due to errors introduced in previous steps. This deviation amplifies errors over successive iterations, resulting in cumulative inaccuracies throughout the simulation. Following the same strategy as³⁷, we introduced noise to all our dynamical variables during training to enhance the robustness of our models.

Timestep. To simplify training and eliminate temporal ambiguity, we fixed the training timestep to 0.01 s, which is coarser than the CFD solver’s timestep of 0.001 s. This was done intentionally to reduce sequence length while maintaining temporal coherence. As such, the model is not required to generalize to arbitrary time horizons. While we did not test interpolation across multiple timestep sizes, the current design is adequate for clinical use cases where a consistent frame rate is expected.

Numerical performance

We evaluated all four models on the full test set for both next-step and rollout prediction with 1-step and 50-step Root Mean Square Errors (RMSE) as described in the Methods section; mean errors over all test cases are given in Table 1. While all models perform well on next-step prediction, as shown by the 1-step error in Table 1, significant differences emerge in rollout prediction in particular at time t_2 as clearly highlighted in Fig. 3.

Indeed, the In-MGN and In-PI-MGN models produce high-quality rollouts, with the 50-step error for In-PI-MGN being more than six times lower than the MGN baseline. In-MGN also performs well on longer predictions, whereas introducing the physical loss alone does not improve long-term prediction performance. Both MGN and PI-MGN models converge for single-step predictions but struggle with error accumulation in rollout predictions.

Because all geometries share the same parent artery, averaging the errors over the entire domain may obscure localized variations in model performance. To better assess the predictive accuracy of each model, we

computed the 1-step and 50-step RMSEs separately across three anatomically distinct regions: the aneurysm bulge, the neck, and the distal segment of the parent artery (see Fig. 4). The results for all four models are summarized in Table 2, confirming the same overall tendencies observed in the global error metrics.

To assess the individual impact of acceleration and inflow information on model performance, we conducted an ablation study as part of our experiments. The results showed that adding inflow features alone improved both training convergence and long-term rollout predictions, whereas adding acceleration without inflow features did not lead to model convergence. This suggests that local temporal history is insufficient for accurate prediction in the absence of global context. The combination of both inflow and acceleration features yielded the best performance in terms of error metrics and qualitative agreement with CFD fields.

Figures 5 and 6 show a temporal analysis of the full geometry volume. It clearly demonstrates that GNN predictions are closely aligned with ground truth CFD simulations, even during the systolic phase, where velocity fluctuations are pronounced. Additionally, the model effectively captures the evolution of recirculation within the aneurysm bulge. To assess the model’s performance across a diverse range of flow characteristics and bulge shapes, velocity predictions were generated for the entire test set, as illustrated in Fig. 7.

Quantitative comparisons at the aneurysm neck (Fig. 8) demonstrate good agreement in predicted flow dynamics, particularly where high-velocity jets interact with the aneurysm wall. Velocity measurements at evaluation points (Fig. 9) further illustrate model accuracy. While the baseline MGN model deviates significantly from the ground truth after a few timesteps, both inflow-enhanced In-MGN and In-PI-MGN maintain an error margin within 5% for most of the cardiac cycle at p_{neck} and p_{out} . Predictions at p_{wall} are less precise for In-MGN and exhibit some delay, whereas In-PI-MGN remains closely aligned with the ground truth, particularly after systole.

Beyond velocity evaluation, WSS results derived from velocity fields were also analyzed by comparing TAWSS in the aneurysm bulge for the In-PI-MGN model and the ground truth, as shown in Figs. 10 and 11. The prediction error remains within 10% across the entire test set,

Table 1 | Model configurations and performance

Model	Input features	Loss function	1-RMSE	50-RMSE	50-RMSE SD
MGN	u	\mathcal{L}_{data}	1.11	50.51	1.43
PI-MGN	u	\mathcal{L}_{PI}	1.06	55.62	1.27
In-MGN	$u, du/dt, inflow$	\mathcal{L}_{data}	1.09	9.22	1.47
In-PI-MGN	$u, du/dt, inflow$	\mathcal{L}_{PI}	0.85	7.58	1.02

Errors averaged on the test set over the whole geometry; next step prediction (1-RMSE) and rollout predictions (50-RMSE) errors are given in mm/s. The standard deviation (SD) of the 50-RMSE across the test set cases is also displayed. Best results are displayed in **bold**; second best in *italic*.

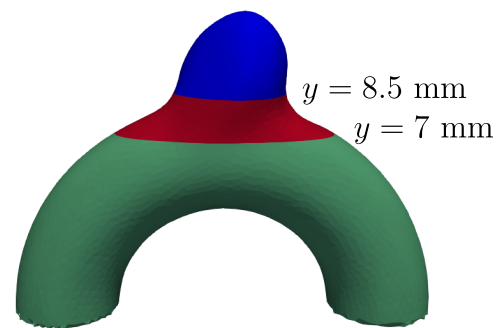
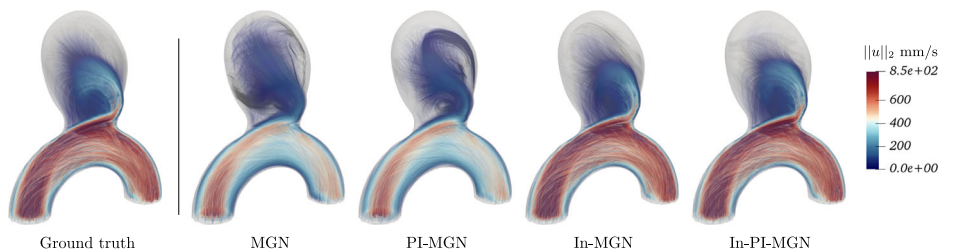


Fig. 4 | Split of our geometries at planes $y = 8.5$ mm and $y = 7$ mm to evaluate RMSE metrics in three anatomically regions: the aneurysm bulge (blue), the neck (red), and the parent artery (green).

Fig. 3 | Rollout predictions at time t_2 for all trained models.



demonstrating strong agreement between predicted and reference patterns. Finally, instant WSS and space-averaged WSS (SAWSS) were examined in Fig. 12 to track WSS variations throughout the cardiac cycle. The SAWSS curves indicate that while the In-MGN model follows the general trend, the In-PI-MGN model exhibits the closest match to the reference, particularly in capturing peak values and transient dynamics. The WSS plots at our three evaluation timesteps confirm again the accuracy of the model pattern prediction over the cardiac cycle.

Table 2 | 1-RMSE and 50-RMSE in mm/s averaged across the test set on the different regions of the geometries

Model	Bulge 1-RMSE	Neck 1-RMSE	Artery 1-RMSE
MGN	1.22	1.63	0.91
PI-MGN	1.14	1.31	0.92
In-MGN	1.11	1.47	0.83
In-PI-MGN	0.68	1.16	0.74
Model	Bulge 50-RMSE	Neck 50-RMSE	Artery 50-RMSE
MGN	38.16	62.25	51.77
PI-MGN	44.88	66.01	57.87
In-MGN	18.28	10.66	4.62
In-PI-MGN	11.87	8.38	4.66

Best results are displayed in **bold**; second best in *italic*.

Furthermore, in order to evaluate the ability of the model to conserve mass, we conducted quantitative mass flow analyses at multiple cross-sectional planes along the artery, as displayed in Fig. 13. Specifically, we measured the instantaneous volumetric flow rates at the inlet, the outlet, and at two additional intermediate planes normal to the centerline of the parent artery. These measurements were performed across the full cardiac cycle for both the CFD ground truth and the In-PI-MGN model predictions.

The results are presented in Fig. 13, illustrating the flow rates across the geometry for both the ground truth and the predicted velocity fields. The mean difference between inlet and outlet flow rates for ground truth and prediction are 2.3% and 3.7%, respectively, with maximums across the cardiac cycle at 2.9% and 5.2%. These results further support that the model captures not only the spatial and temporal dynamics of blood flow but also respects core physical laws such as continuity.

Generalization to unseen cases

In order to evaluate the generalization capability of the model trained solely on semi-idealized geometries, we conducted two experiments without any additional fine-tuning. Indeed, the proposed GNN framework operates entirely on local mesh and flow features rather than relying on global geometry descriptors; therefore, it is essential to assess its ability to adapt to new settings, provided the underlying physics and boundary conditions remain comparable.

First, we examined the model’s robustness to unseen inflow rates. We ran new CFD simulations on our three test cases using alternative inflow velocity time profiles from ICA and middle cerebral artery (MCA) as well, derived from the literature^{38,39}, shown in Fig. 14. When evaluated on these six new inflow conditions (with no retraining), the model achieved a 1-step

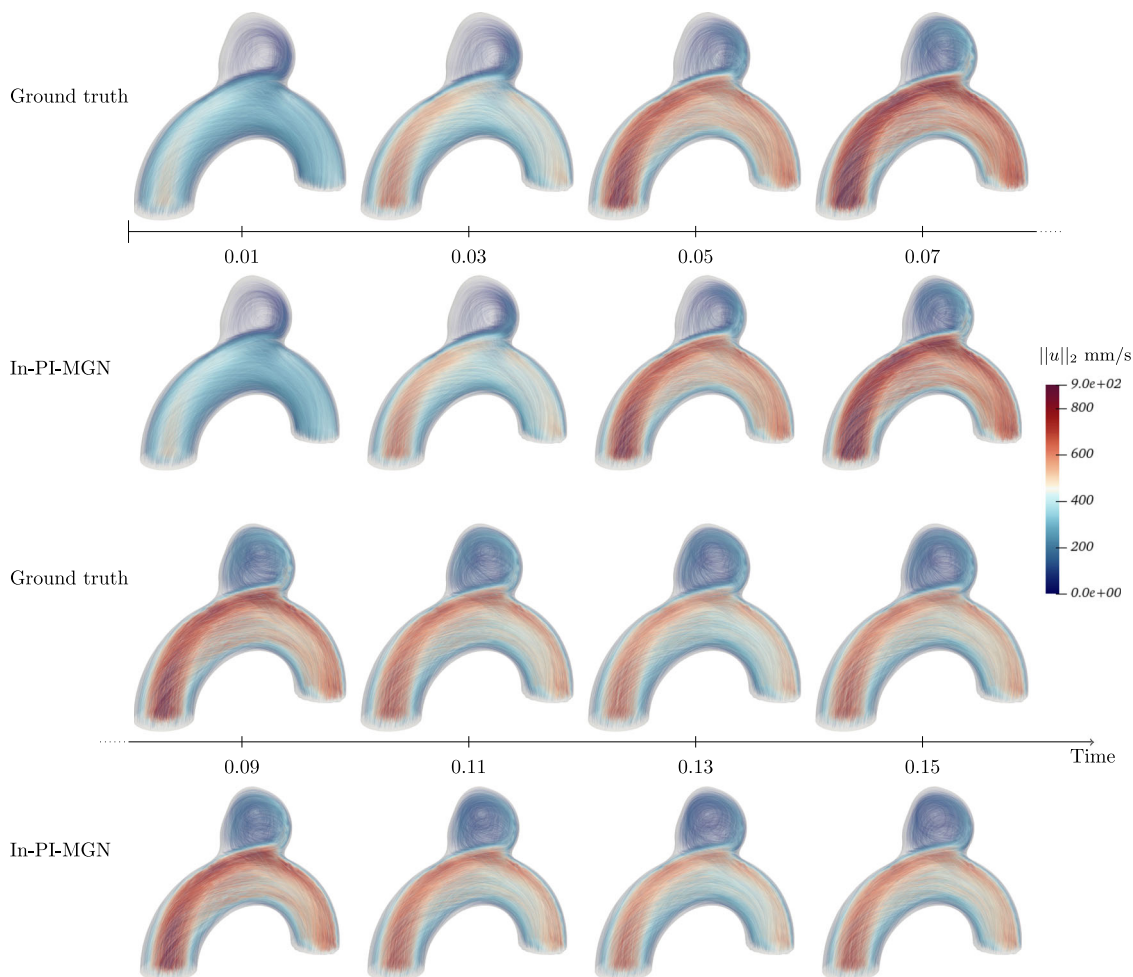


Fig. 5 | Blood flow during systole for Case B, showing predictions closely matching simulations throughout the acceleration phase.

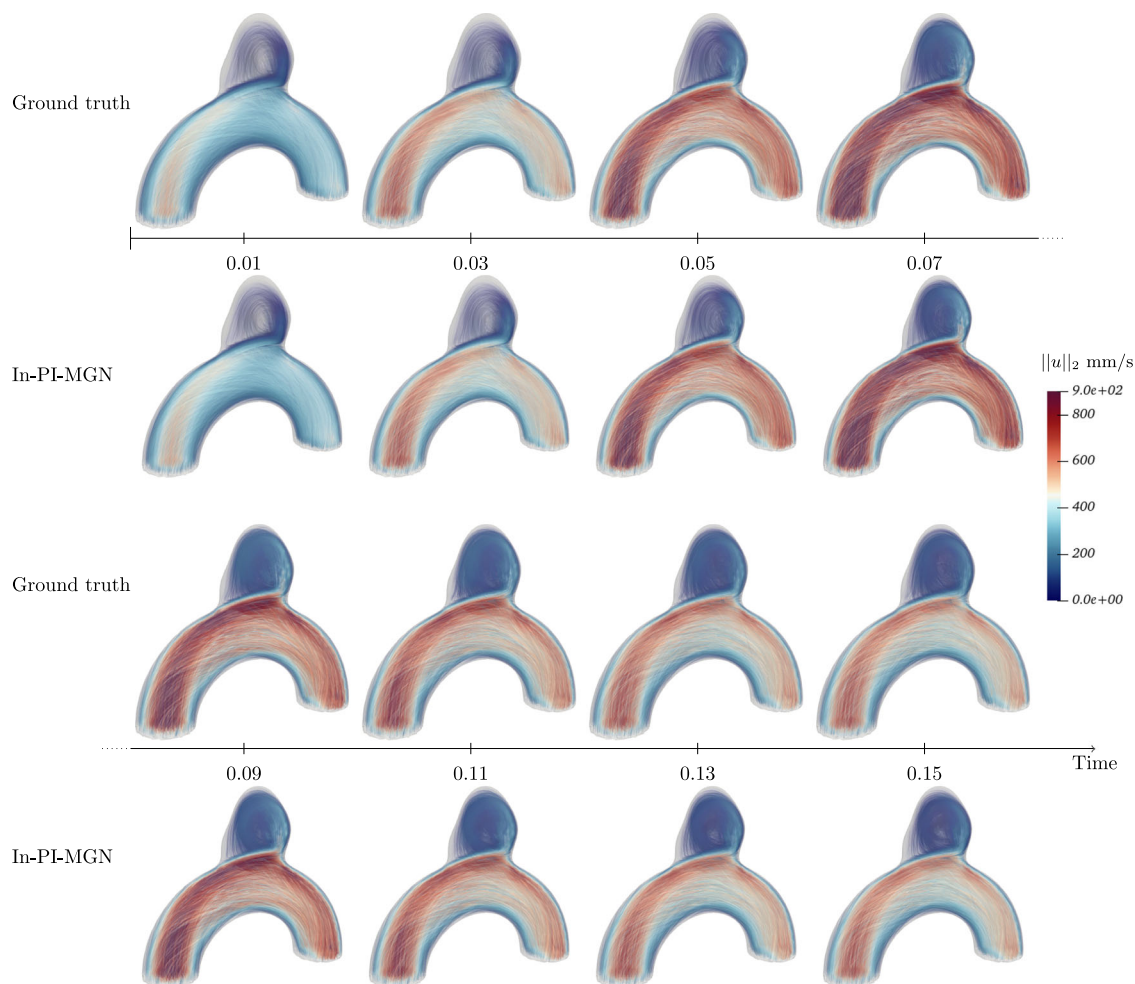


Fig. 6 | Blood flow during systole for Case C, demonstrating accurate predictions in a larger, more asymmetrical aneurysm bulge.

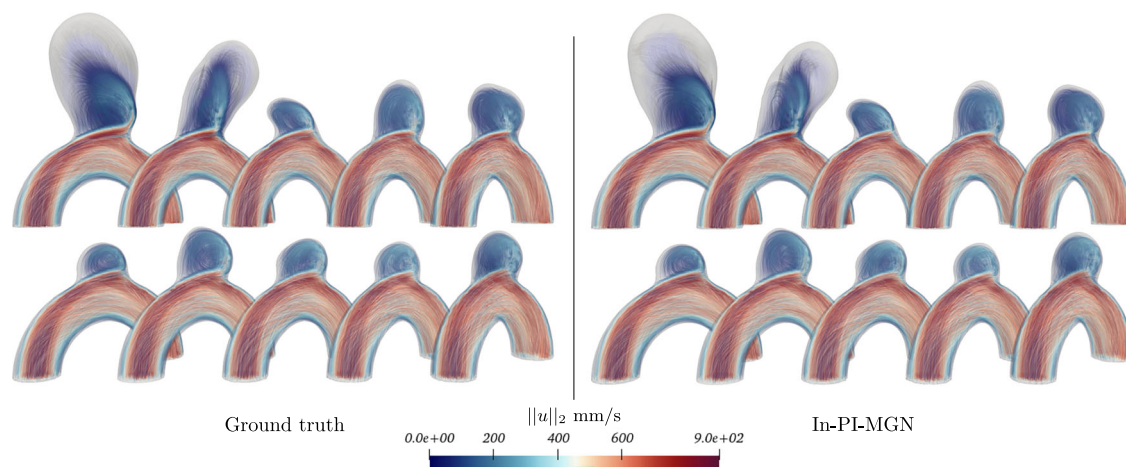


Fig. 7 | Velocity fields comparison for the whole test set at time t_2 .

RMSE of 2.4 and a 50-step RMSE of 14.3. Velocity field comparisons in Fig. 15 confirm that predictions remain close to the ground truth, despite never encountering these inflow profiles during training. This result demonstrates the model’s adaptability to physiologically realistic variations in inflow dynamics, even though it was trained on a single fixed profile.

Second, we tested the model on four zero-shot predictions on out-of-distribution, patient-specific aneurysm geometries. These new cases were randomly selected from our internal database and exhibit a range of

morphological variations, including aneurysm location, parent vessel diameter, neck size, aneurysm sac diameter, and centerline length. Their geometrical features are further described in Table 3 and highlight the difference of these cases from the BenchAnXplore training and testing data.

Remarkably, the In-PI-MGN model trained only on synthetic semi-idealized geometries was able to reproduce realistic flow behaviors in this complex clinical scenario. While the predicted velocity fields show discrepancies in order-of-magnitude for cases 3 and 4, the predicted velocity

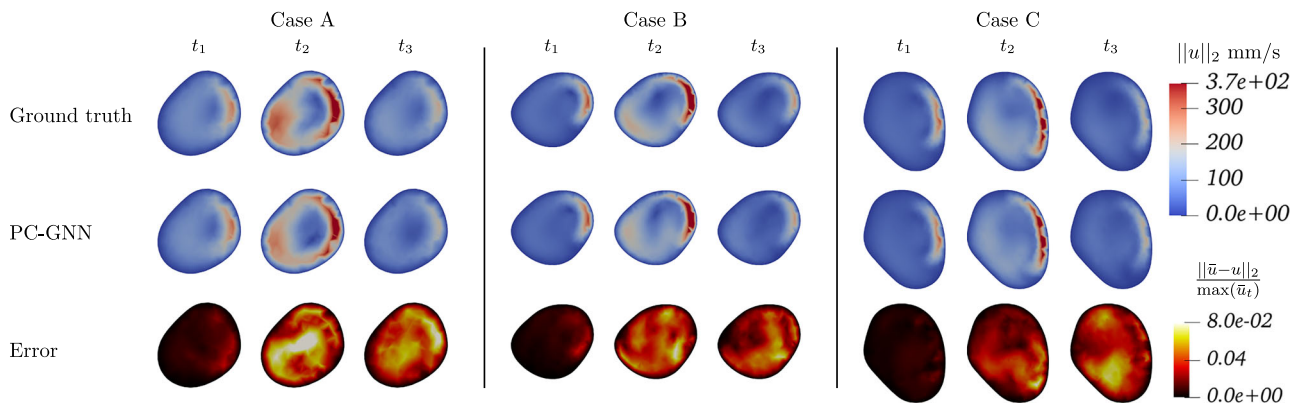


Fig. 8 | Blood velocity comparison at the aneurysm neck; displayed error is normalized by maximum ground-truth velocity at the given timestep. Maximum error is occurring at high-speed flow regions close to the wall, and remains below 8%.

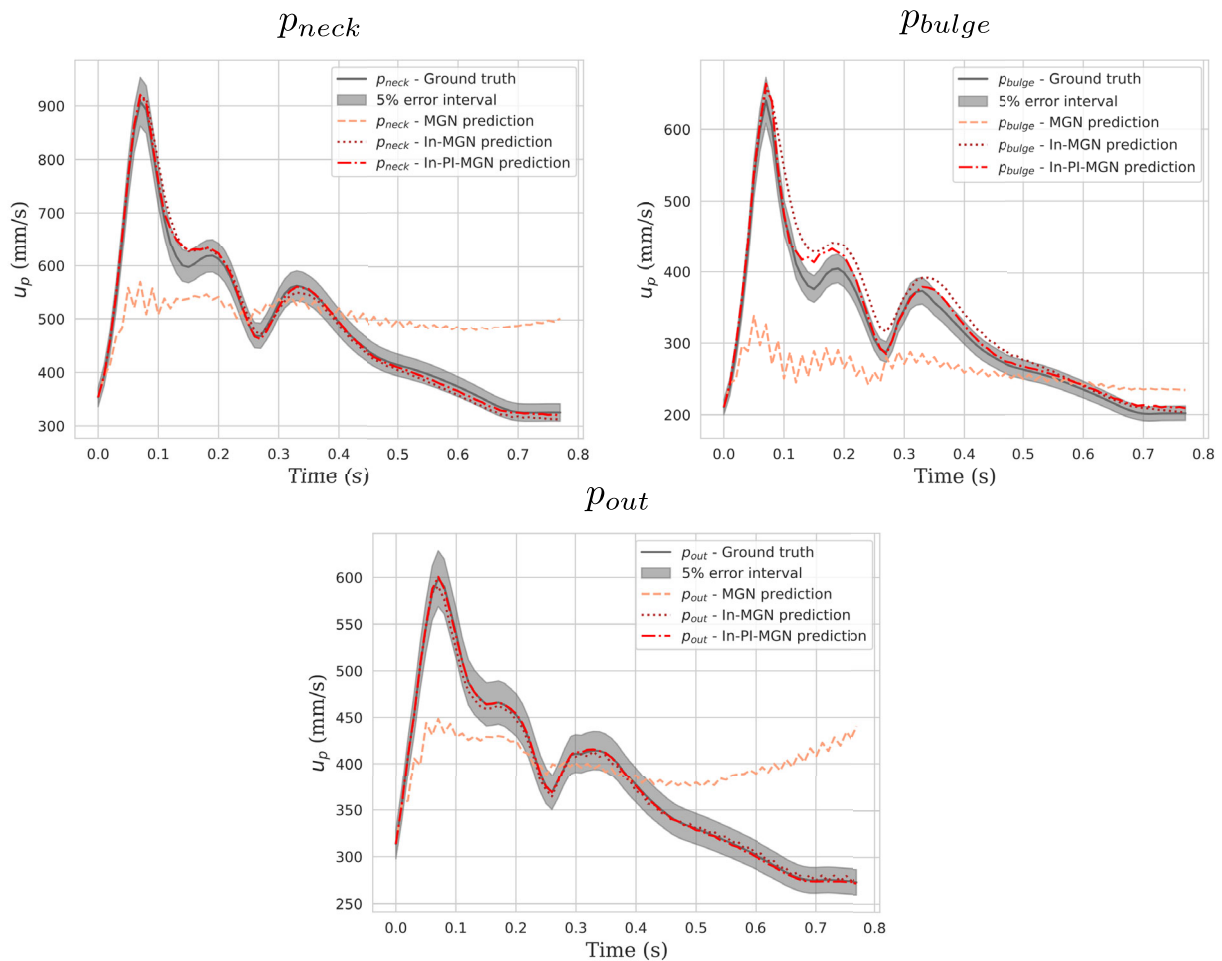


Fig. 9 | Velocity evolution plotted over time for Case A.

fields match the CFD ground truth in spatial distribution in all cases, as shown in Fig. 16.

Discussion

Our results demonstrate that In-PI-MGN, a physics-constrained GNN, can accurately predict spatio-temporal blood flow fields across a diverse range of 3D aneurysmal geometries. The model maintains close agreement with CFD ground truth while offering significant computational gains. The model can simulate 80 timesteps in under a minute, whereas traditional

CFD methods require approximately 1 h for the same task. Even accounting for the 20-h training time, the proposed approach becomes more efficient after just 20 simulations. Including dataset generation, the total model development time amounts to approximately 70 h. This leap in speed could be critical in clinical decision support, especially for time-sensitive or resource-constrained contexts.

Indeed, the ability to produce fast full-field hemodynamic predictions opens new possibilities for simulation-augmented clinical workflows. For instance, rupture risk assessments or device deployment planning could

Fig. 10 | TAWSS magnitude plotted on the aneurysm bulge, and relative error normalized by maximum ground-truth TAWSS. Error is thresholded to mask errors under 5%; maximum local error is at 9%.

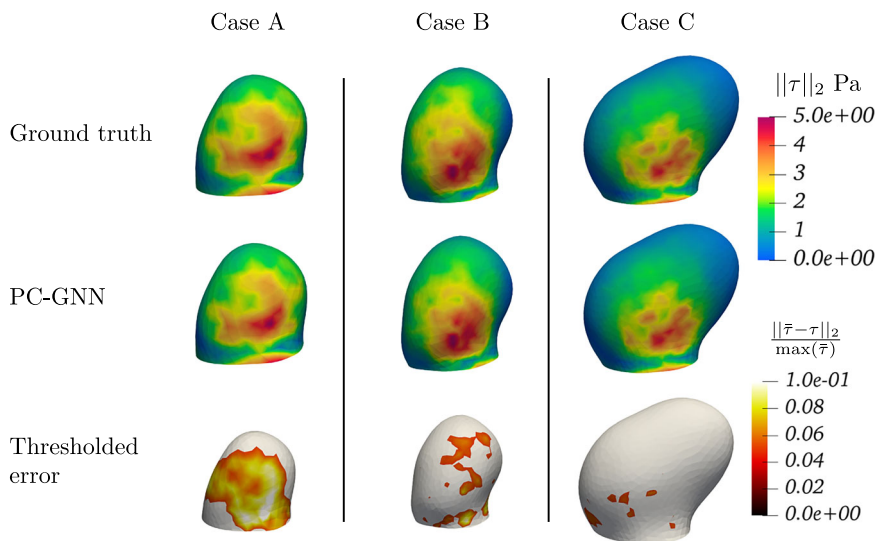
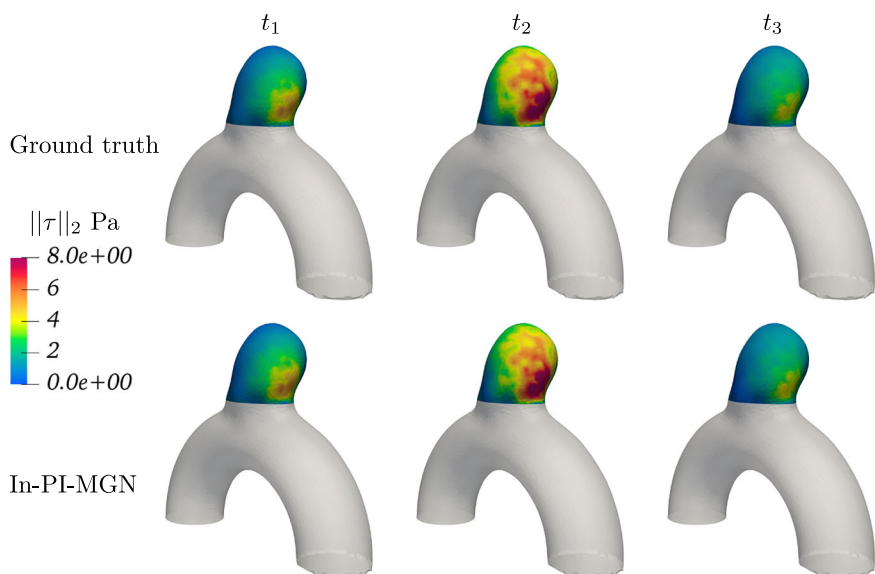


Fig. 11 | WSS plotted over the aneurysm bulge at different timeframes for Case B.



benefit from immediate feedback, without the delays of traditional CFD. We foresee this approach contributing to preoperative planning, follow-up monitoring, and even in-room guidance during interventional neuroradiology.

Our contributions are threefold. First, we tackle the challenge of predicting transient, non-Newtonian 3D blood flow, with the ability to accurately capture the evolution of recirculation and swirling patterns, within complex aneurysmal geometries, an area previously unexplored with GNNs. Second, we introduce two architectural advances: contextual inflow feature enhancement to improve long-range temporal representation, and physics-informed training to promote stability and physical plausibility across time. Third, we release BenchAnXplore, a large-scale benchmark dataset of 105 IA geometries with matched CFD simulations, establishing a new standard for AI-based hemodynamics research.

Prior ML-based approaches for hemodynamic modeling have mostly focused on reduced-order representations, simplified vascular segments, or static features^{15,23,25}. Our work departs from this by directly learning spatio-temporal dynamics from full 3D velocity fields across realistic geometries. Leveraging the expressive capabilities of GNNs, our model captures complex flow behavior—including swirling and recirculation—within

topologically diverse aneurysm shapes. The integration of physics-informed loss terms and contextual inflow features marks a methodological advance, enabling accurate long-term rollout and reducing error accumulation.

Furthermore, we have confronted our method to out-of-distribution unseen cases with different inflow rates and patient-specific cases, and observed that our model had obtained generalization capabilities and was able to uncover realistic flow patterns with no additional training. These preliminary results highlight the potential of our approach as a foundation for future fine-tuning toward patient-specific clinical applications.

In summary, this work offers a reproducible, extensible framework that brings high-fidelity hemodynamic inference closer to real-time clinical use. By enabling rapid hemodynamic inference with strong physical consistency, this work paves the way for real-time risk assessment, personalized treatment planning, and simulation-augmented clinical workflows. It offers a reproducible and extensible step toward bridging high-fidelity fluid mechanics and deep learning in cerebrovascular medicine.

However, several important limitations remain before clinical deployment can be fully realized. First and foremost, the proposed model is trained on idealized vessel geometries. Clinical routine is characterized by a wide variety of anatomical presentations in terms of size, shape, and

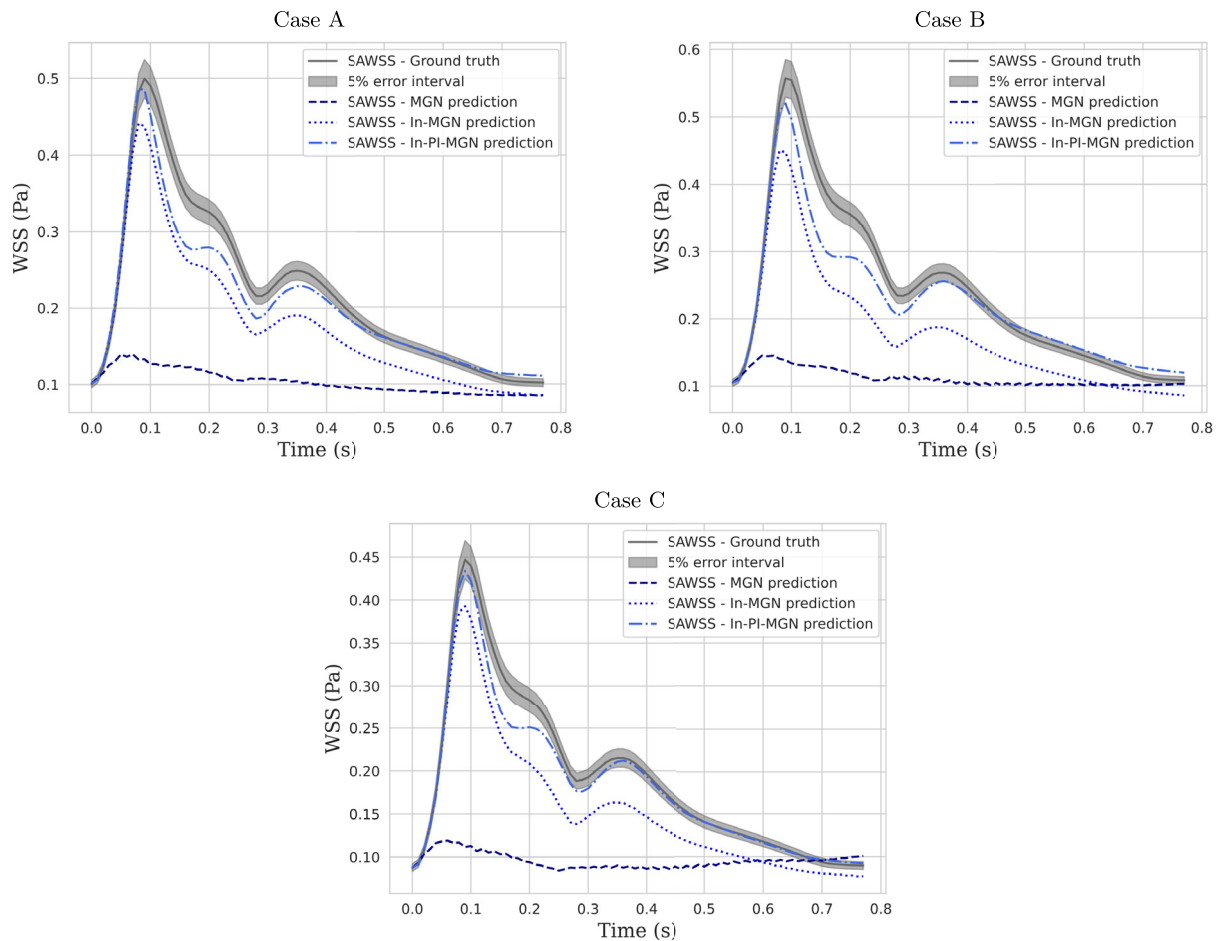


Fig. 12 | SAWSS computed over the aneurysm bulge for all 3 test cases.

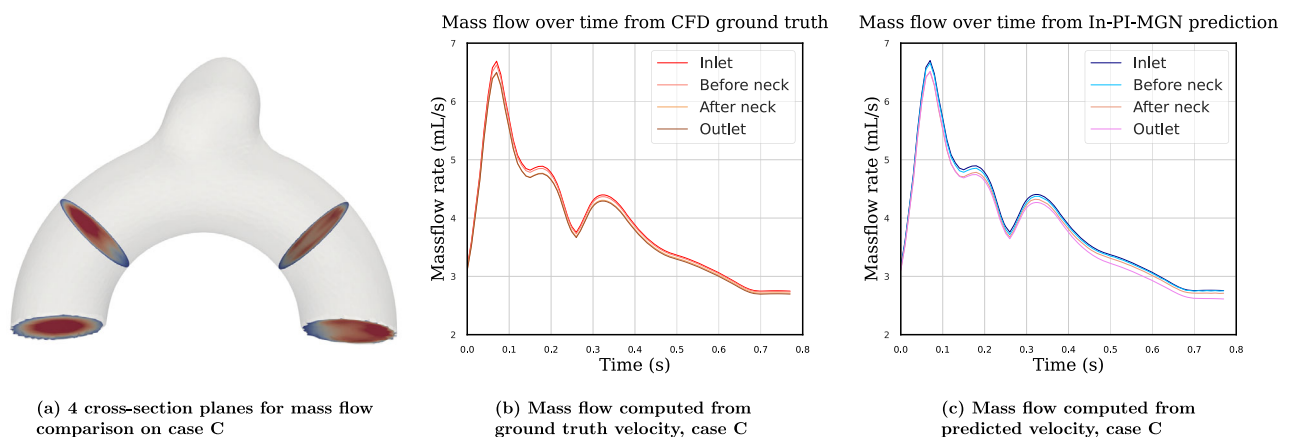


Fig. 13 | **Transient mass flow comparison between CFD ground truth and In-PI-MGN predictions.** **a** Case C geometry illustrating the four cross-sectional planes used for mass flow evaluation: inlet, before the neck, after the neck, and outlet. **b** Mass flow rates computed from the CFD ground-truth velocity field at the four

cross-sections over one cardiac cycle. **c** Mass flow rates computed from the In-PI-MGN-predicted velocity field at the same cross-sections. Mass flow is reported in mL s^{-1} as a function of time.

orientation. As highlighted by the promising yet degraded performance of our model on out-of-distribution patient-specific cases (e.g., average 50-step relative velocity error on the four cases of 11% compared to 1.6% on the original test set), generalization to clinical settings requires training on representative data. While our current dataset is varied and complex enough to serve as a valuable benchmark for model evaluation, it represents only a

first step toward patient-specific hemodynamics prediction. In future work, this will need to be complemented by a larger collection of patient-derived geometries or the generation of more realistic synthetic cases.

Second, our CFD ground-truth simulations were generated using a single, temporally and spatially idealized inflow profile. We have demonstrated the model’s ability to generalize across different inflow time profiles,

even when these were not part of the training data. Although these inflow rates are widely used in CFD studies of IAs^{40–42}, they remain statistical means extracted from literature^{38,39}, and may not reflect true patient-specific variability. It would be highly valuable to extend this approach to include clinically acquired Doppler ultrasound or 4D Flow MRI inflow profiles. Additionally, all inlet profiles in this study are prescribed as idealized parabolic flows, which—while common^{41–43}, omit higher-order effects captured in more realistic Womersley profiles⁴⁰. Fortunately, previous studies have shown that aneurysmal flow patterns remain robust across such inlet assumptions when appropriate domain lengths are used⁴⁴. Nonetheless, we aim to integrate 4D Flow MRI-derived boundary conditions in future work⁴⁵, both for fine-tuning and evaluating our model under more realistic conditions.

Lastly, we acknowledge that while hemodynamic quantities such as WSS and OSI are commonly studied, a biomarker strongly correlated with IA rupture risk remains elusive. The interpretation of these biomarkers, including proposed thresholds for risk stratification, continues to be debated within the scientific and medical communities. While the primary objective of this work is not direct rupture prediction, it is important to clarify that our model outputs are intended to serve as high-fidelity, physics-consistent inputs to further risk analysis, not as standalone diagnostic tools. We see this as a step toward augmenting clinical decision-making by providing accessible, real-time flow field estimation, which can be combined with other risk factors to support informed diagnostic.

Building on this foundation, future work may explore advanced architectures such as multigrid GNNs⁴⁶, which enable hierarchical information flow across mesh scales, and attention-based GNNs, which enhance the model’s ability to prioritize key flow structures. These approaches could

enable higher-resolution predictions while preserving computational efficiency, further advancing AI-driven hemodynamic modeling. Moreover, while we demonstrate zero-shot generalization on both unseen inflow profiles and a fully patient-specific case, we emphasize the need for further validation on larger, more heterogeneous clinical datasets to support further research.

Methods

Graph neural networks

GNNs are a class of neural networks working on sets of vertices and edges. Through different deep learning techniques such as convolution (Graph Convolutional Networks)⁴⁷, attention (Graph Attention Networks)⁴⁸, or recurrent layers (Graph Recurrent Networks)⁴⁹, they exploit message passing mechanisms to learn the relations between neighboring nodes, and have shown great performances in many areas of physics^{19,21,50–52}.

Solving complex systems of partial differential equations, such as the Navier-Stokes equations, often requires mesh-based discretisation of the computational domain. These meshes are typically adapted to enhance numerical accuracy in critical regions while minimizing computational cost in less challenging areas⁵³. However, traditional deep-learning architectures, such as convolutional, recurrent, or fully connected layers, struggle to effectively handle the anisotropic and irregular nature of meshes. They either fail to capture local dependencies properly or do so at a prohibitive computational expense. Building on the founding that the deep-learning toolbox layers did not offer the possibility of learning arbitrary relations between elements of a set with relational structure, such as a graph, Battaglia et al.⁵⁴ introduced the Graph Network (GN) layer, aiming to create a general framework for graph computation. Inspired from Message Passing Neural Networks⁵⁵ among others, this layer is designed to take a graph as input, process the graph with any subsequent model, such as a neural network, but not necessarily, and return a graph, using the nodes as entities with features, edges as relations between entities, and global attributes over the whole graph.

We note an undirected graph $G = (V, E)$ with V a set of vertices with $V = \{v_i\}_{i \in N^v}$ vertices attributes, and E a set of edges with e_{ij} linking vertices with indexes i, j .

The GN updates edge attributes with an edge-update function φ_E , gathers a vertex neighborhood information by aggregating the updated edges connected to it with an aggregation function ρ , and updates the vertex with a vertex-update function φ_V (Fig. 17). Because the update functions are shared over every node and edge, GNNs are adapted to learn local physical behavior over every relation in the graph. We purposefully left out global graph attributes here to avoid unnecessary complications of the framework, as we will not be using them.

$$\begin{aligned} \tilde{e}_{ij} &= \varphi_E(e_{ij}, v_i, v_j) \\ \tilde{e}_i &= \rho(\{\tilde{e}_{ij}\}_j) \\ \tilde{v}_i &= \varphi_V(v_i, \tilde{e}_i) \end{aligned} \tag{1}$$

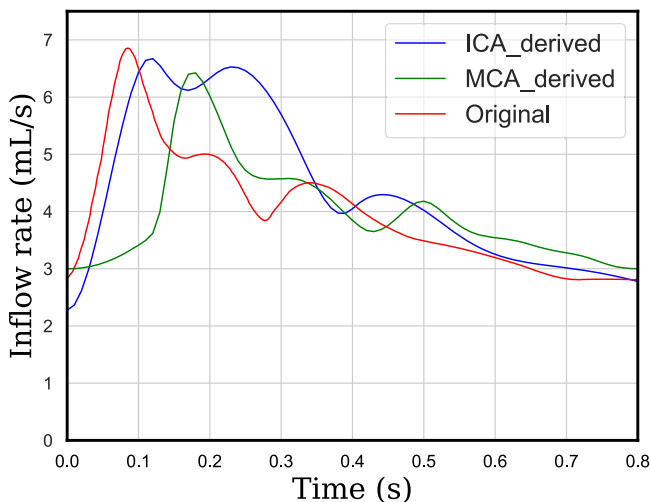


Fig. 14 | Different inflow rates derived from patient-specific ICA and MCA data^{38,39}.

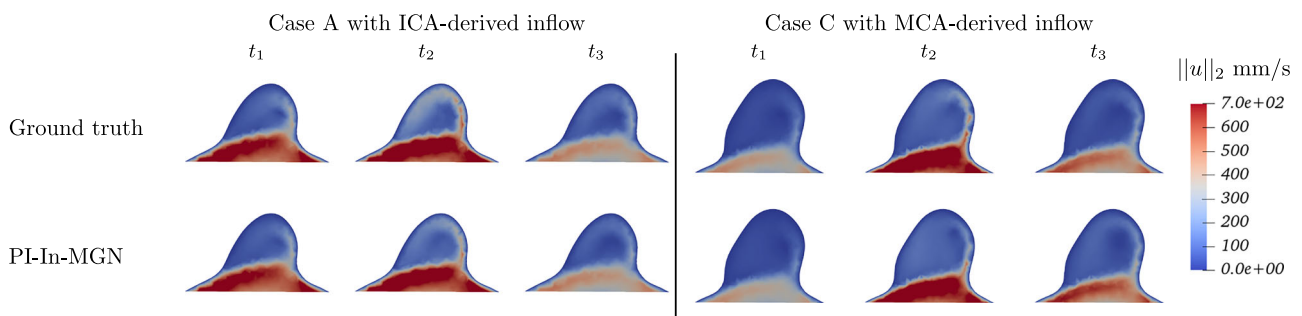


Fig. 15 | Blood velocity comparison at the aneurysm neck on unseen velocity profiles.

Table 3 | Zero-shot cases geometrical description with aneurysm location, inlet artery diameter, aneurysm diameter, aneurysm neck diameter, and artery centerline length from inlet to furthest outlet, all in mm

Case	Location	Inlet diameter	Aneurysm diameter	Aneurysm neck diameter	Centerline length
Case 1	ICA	4.8	4.5	4.2	49.1
Case 2	MCA	2.5	6.0	5.2	23.2
Case 3	MCA	2.7	7.8	7.3	21.4
Case 4	ICA	4.3	15.0	8.7	65.5
BenchAnXplore	ICA	4.0	3–9	3.2	25.6

The BenchAnXplore values are given as a comparison.

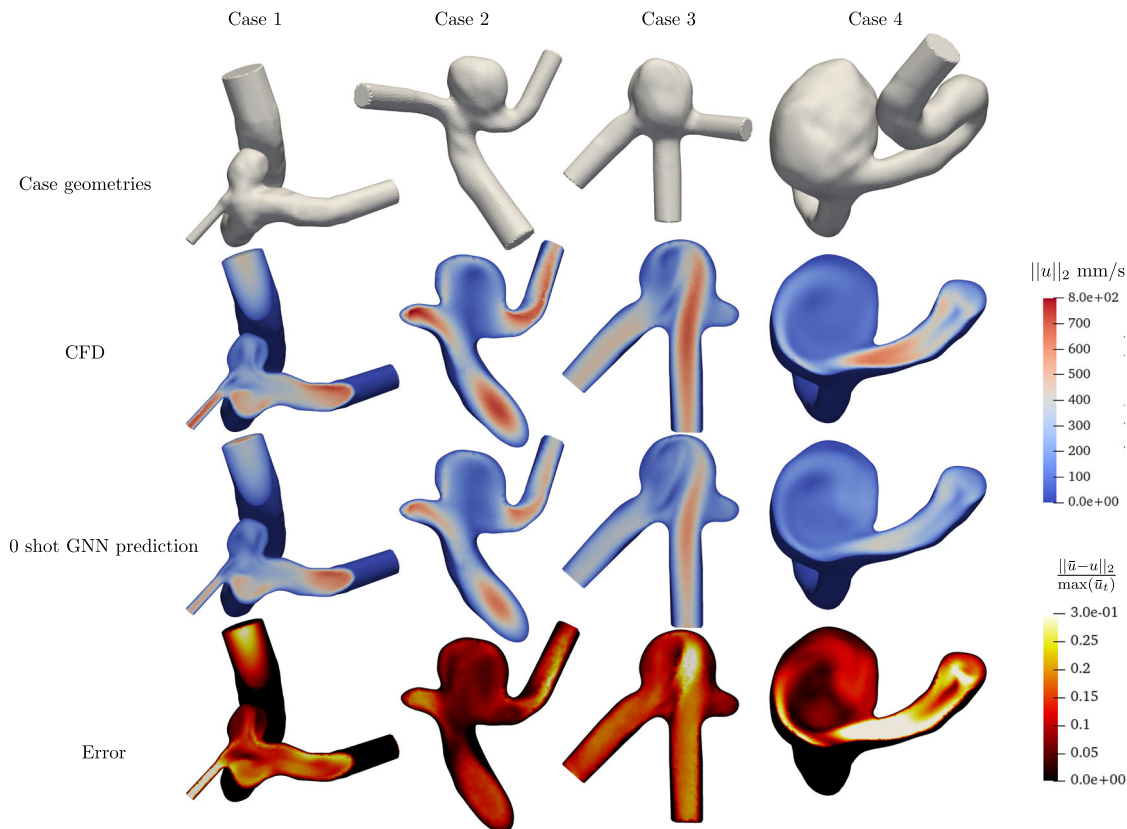


Fig. 16 | Zero-shot predictions on out-of-distribution patient-specific ICA and MCA aneurysm cases.

with \tilde{e}_{ij} being the updated edge attributes from the initial edge attributes and the connected vertices attributes, \tilde{e}_i being the aggregated information of all edges linked to vertex i , and \tilde{v}_i being the updated vertex attributes from the initial vertex attributes and the aggregated neighborhood information.

MeshGraphNet—GNNs for physics simulations

Numerical physics simulations like CFD relies on meshes. Pfaff et al.¹⁷ took advantage of this data structure and built a GNN model using physical meshes as inputs, designed for next-step prediction in physical problems such as flow field prediction around a cylinder or a wing profile. With the mesh information at t (and potentially $t - 1$), the trained network can predict the physical fields at $t + 1$. By inferring on its own predictions, the network can predict a whole simulation with only a few timesteps to infer from; the MeshGraphNet (MGN) model has achieved state-of-the-art results, being able to predict the temporal evolution of these systems during significant time periods while remaining coherent with the physical simulations. The MGN model learns the system’s forward dynamics with an encode-process-decode architecture (Fig. 18), using multilayer perceptrons (MLPs) Fig. 19.

Encoder. The encoder extracts all input features for nodes and edges from the graph G_t , which is derived from the mesh at time t . These features include physical properties such as velocity, pressure, or stress, as well as geometrical attributes like edge lengths and node positions. The node and edge features are respectively encoded into latent spaces by the encoder MLPs ϵ_V and ϵ_E into latent spaces of size 128, as graph nodes v_i and edges e_{ij} , to constitute the latent space graph Q_t .

Processor. The processor then performs rounds of message passing by applying L several successive GN layers. For each round of message passing, each node of the graph sends messages to its neighbors, which consist of the encoded and processed node’s and edge’s feature information. As in the GN layer described in Subsection GNNs, the neighbors aggregate and process these messages to gain a collective understanding of their local neighborhood and to update their own feature representation. This allows them to incorporate the new information received from their neighbors and refine their own states. Each of these GN layers has update functions ϕ_E and ϕ_V , respectively, for edges and nodes, as described in Equation (1). These functions are MLPs as well, with

different sets of parameters for each GN layer, and the edge aggregation function ρ is an element-wise summation. Performing several consecutive rounds of message passing allows the GNN to infer the evolution of the graph with information circulating over the whole domain, and to return an updated latent space graph Q_{t+1} with updated node and edge latent vectors \tilde{v}_i and \tilde{e}_{ij} .

Decoder. Finally, the decoder translates the processed node vectors contained in the latent space updated graph Q_{t+1} with a final MLP δ_V to a physical update, which is added to the input mesh in order to get the next-step physical mesh G_{t+1} with updated physical fields.

All the underlying processing and updating functions here are ReLU-activated 2 hidden-layer MLPs. The model learns the physical behavior by minimizing the differences between its predictions and the ground truth data computed with a numerical solver, with a data loss being the mean square error, for a given physical quantity a :

$$\mathcal{L}_{data} = \text{mean}_{v \in V} \| a_v - \tilde{a}_v \|_2^2 \tag{2}$$

with a_v being the ground truth and \tilde{a}_v the prediction at node v .

Hemodynamics simulations

The MeshGraphNet model has demonstrated excellent performance in predicting flow in 2D settings for benchmark fluid dynamics problems, such as incompressible steady flow around a cylinder and compressible flow around an airfoil. However, its ability to generalize to blood flow predictions in a three-dimensional environment remains an unexplored frontier and requires further investigation.

The CFD numerical framework used in hemodynamics simulations provide fully resolved velocity u and pressure p fields on the whole 3D domain by solving the mixed formulation of the transient incompressible

Navier-Stokes equations given by:

$$\begin{aligned} \rho(\partial_t u + u \cdot \nabla u) &= -\nabla p + \nabla \cdot (2\mu\epsilon(u)) \\ \nabla \cdot u &= 0 \end{aligned} \tag{3}$$

with ρ the fluid density, $\epsilon(u)$ the strain rate tensor and μ the fluid viscosity. Blood exhibits a non-Newtonian rheology, a shear-thinning behavior that can be described with constitutive models such as the Casson or the Carreau–Yasuda law⁵⁶:

$$\mu(\dot{\gamma}) = \mu_\infty + (\mu_0 - \mu_\infty) \left(1 + (\tau \dot{\gamma})^a \right)^{\frac{n-1}{a}} \tag{4}$$

with the relaxation time τ , power indices n and a , and asymptotic values of the effective viscosity $\mu(\dot{\gamma})$, μ_0 , and μ_∞ , at zero and infinite shear-rates, respectively, being known fluid constants.

At the inlet of the fluid domain, we impose a physiologically realistic pulsatile mass flow profile corresponding to one full cardiac cycle, derived from ICA measurements as reported in the literature⁵⁷. The velocity profile is applied as a fully developed parabolic inflow condition. To ensure smooth initialization of the simulation, the cardiac waveform is preceded by a 0.2 s linear ramp-up. The simulation spans two full cardiac cycles of 0.8 s each, with the first cycle serving as transient initialization. This waveform corresponds to a cardiac cycle of 0.8 s and was scaled to produce a peak systolic flow rate of approximately 6 mL/s, which, given the 4 mm vessel diameter, yields a peak Reynolds number of approximately 400. This flow regime is consistent with reported physiological conditions in the internal carotid artery^{34,36,38,57}. Note that only the second cycle is retained for analysis to ensure periodic flow has been reached. At the outlet, we apply a resistive boundary condition to emulate the downstream vasculature. The resistance is tuned to maintain a physiological pressure variation of approximately 40 mmHg over the cardiac cycle, consistent with clinical observations for ICA outflow and guided by known systolic and diastolic flow rates. The vessel walls are modeled as rigid with a no-slip boundary condition, hence zero velocity is imposed at all wall-adjacent nodes. This is a standard assumption in many CFD studies of IAs, particularly when the primary interest lies in evaluating flow patterns and shear-related biomarkers rather than wall compliance.

Several quantities of interest have been identified for risk assessment and treatment decisions for unruptured IAs. Obtaining these quantities by post-processing the simulated velocity fields inside the aneurysm is of great interest for clinical routine, and is a key feature expected from CFD frameworks for hemodynamics. Among these hemodynamical quantities is

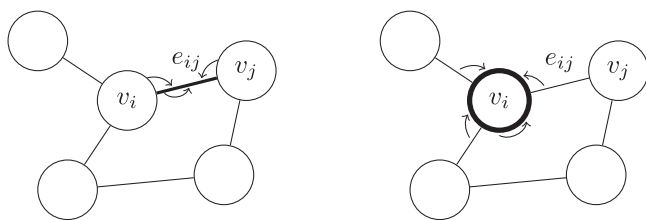
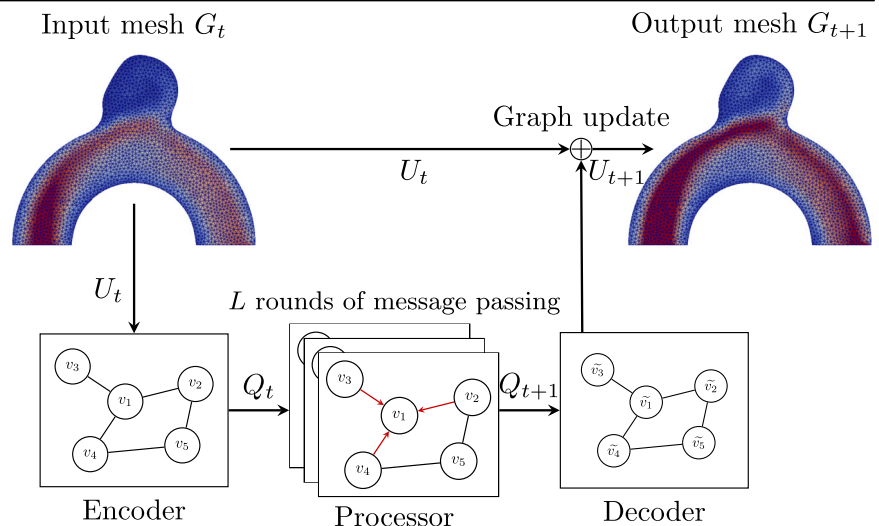


Fig. 17 | Graph network updating edges (left) and nodes (right).

Fig. 18 | Encode-process-decode (EPD) architecture¹⁷ for velocity U_{t+1} prediction from U_t .



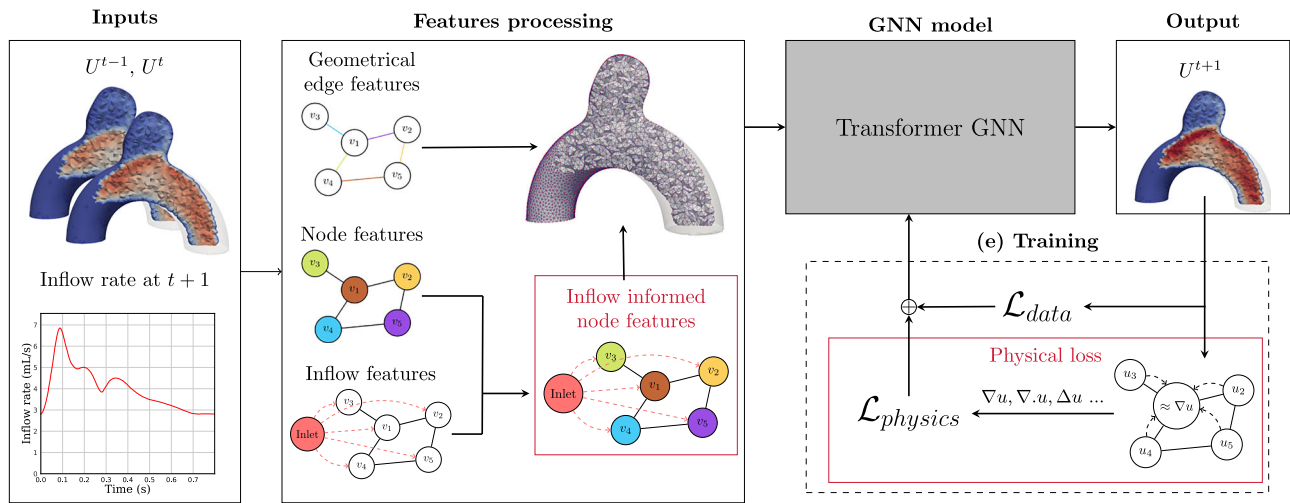


Fig. 19 | Model architecture - the input meshes and inflow data are exploited to create a feature graph, from which our GNN model can predict the next step mesh, taking into account physical and data losses combined during model training.

the WSS, defined here as:

$$\tau = n \times [(\sigma \cdot n) \times n] = \sigma \cdot n - [(\sigma \cdot n) \cdot n]n \quad (5)$$

where n is the unit normal vector at the wall and σ is the Cauchy stress tensor defined as $\sigma = -pI + \mu(\nabla u + \nabla^T u)$.

This framework unveils a higher level of complexity compared to benchmark fluid dynamics cases on which MGN was evaluated. The shift to 3D adds significant intricacy to the observed flow patterns and increases the size of the meshes. Additionally, the geometries at hand in a medical context are derived from real patients, making them inherently more intricate and variable compared to synthetic or benchmark models. Moreover, blood rheology can be considered in hemodynamics simulations, needing to account for non-Newtonian fluid properties and their effects on flow behavior. This adds another layer of complexity as blood exhibits varying viscosity under different flow conditions, as shown in Equation (4). Finally, the inflow conditions in blood flow studies follow a pulsatile cardiac cycle, which introduces time-dependent variations that must be accurately captured to understand the hemodynamics accurately. This pulsatile nature requires a deeper understanding of hemodynamics by the model to simulate the rhythmic changes in blood flow and pressure.

For a GNN-based CFD surrogate approach to be justified, the GNN counterpart to traditional numerical solvers needs to produce local physical spatio-temporal fields in a relatively thin margin of error from the CFD ground truth outputs, with a significant reduction in computing time.

Dataset generation

For each of the 105 semi-idealized aneurysm geometries, numerical fluid simulations were performed following the general fluid mechanics problem formulated in the previous subsection, with a 0.001 s timestep. All the geometries hemodynamics were simulated using the velocity profile derived from⁵⁷. The simulations ran on an in-house finite-element C++ based library by solving a mixed formulation of the incompressible Navier–Stokes equations (Equation (3)) using a variational multiscale method, as described by Hachem et al.⁵⁸.

The dataset meshes were generated with refined boundary layers around the artery walls to accurately compute velocity gradients during the CFD simulations. The resulting meshes are composed of around 200 k nodes each, which translates to about 1 M elements.

In order to build a reproducible experiment with a versatile dataset, we derived a lighter version of the dataset, allowing us to test many configurations with reasonable computational resources and time. This lighter

dataset was generated using an automated pipeline. Starting from the high-resolution CFD meshes with refined boundary layers, we extracted the outer surface of each vascular geometry and re-meshed it using isotropic triangular elements. These surfaces were then used to generate volumetric meshes consisting of unstructured tetrahedral elements with uniform resolution, using an average target element size of 0.3 mm. This mesh size was selected based on a trade-off between physical fidelity and computational efficiency, after experimenting with multiple mesh resolutions, ranging from 0.15 to 0.4 mm.

No boundary layer refinement was applied in the light dataset, as our aim was to reduce memory requirements and allow rapid GNN training iterations. To populate these lighter meshes with flow data, we applied linear interpolation of the full-resolution CFD velocity and pressure fields onto the new node sets. To validate the accuracy of this interpolation procedure, we performed round-trip interpolation experiments: interpolated flow fields from the fine mesh were mapped onto the coarse mesh and then back onto the original fine mesh. The resulting discrepancies in both velocity and pressure fields remained below a few percent across the domain, confirming the adequacy of the remeshing and interpolation steps for surrogate modeling.

With this framework, meshes with under 25 k nodes and 120 k elements were automatically generated, resulting in a dataset of 105 aneurysm geometries, with 80 simulated timeframes for each of these cases, with a 0.01 s timestep. The obtained dataset, referred to as **BenchAnXplore** dataset, is now made available as a benchmark case for AI IAs applications.

By working on this lighter dataset, we were able to increase our model size, train much faster, and run several experiments in a reasonable time; training one epoch of our models on the original dataset took around 40 times longer than on the lighter dataset (respectively 20 h and 30 min on a NVIDIA A100 GPU).

Node feature enhancement with contextual inflow information—In-GNN

We designed our model using the MeshGraphNet (MGN) model¹⁷ presented in Subsection MeshGraphNet—GNNs for physics simulations as a baseline to build upon, and to assess our obtained results. This baseline model uses the velocity field at time t as input, and returns the velocity update for time $t + 1$. However, several challenges arise compared to prior GNN applications that have largely focused on 2D scenarios, steady-state 3D flows, or external flow problems around simpler geometries. This work directly confronts the formidable challenges posed by the highly non-linear, pulsatile nature of blood flow within the 3D complex, irregular, and patient-derived geometries of aneurysms. Indeed, the dynamics of the whole

domain are deeply linked with the pulsatile nature of the inlet flow, but information from the inflow would need many rounds of message passing to reach every node of our mesh with the raw MGN architecture, fed on velocity data. Following the lines in Gladstone et al.²¹, who faced the problem of long-range interaction in a time-dependent mechanical problem, we fed our network with additional information by adding statistical information about the inflow acceleration at $t + 1$ to every node of the mesh, as the boundary conditions at the inlet are known beforehand.

This allowed the whole domain to grasp the evolution of the inflow dynamics whilst trying to predict the velocity at $t + 1$, without adding layers of message passing, which would have made the model heavier. Moreover, we also added the acceleration between timesteps $t - 1$ and t as a node feature to give more local temporal history to the model.

For the baseline MGN model, the node feature vector is composed of the following concatenated quantities:

- Velocity (3 components).
- Spatial position (3 components).
- Timestamp (1 component).
- Node type (1-hot encoded: inlet, wall, interior, outlet).

For our model, we enrich this vector by adding:

- Local acceleration computed as the velocity difference between t and $t + 1$ (3 components).
- Inflow context vector: mean, minimum, and maximum velocity magnitude of inlet nodes at $t + 1$ (3 components).

The edge feature vector, consistent across all models, contains:

- The relative coordinate vector between connected nodes (3 components).
- The Euclidean distance (1 component).

Physically-constrained GNN

Although the model is trained on datasets generated by an in-house numerical solver that computes solutions to the Navier–Stokes equations (see Subsection Dataset), the integration of physical knowledge into the training process can enhance the accuracy and reliability of its predictions. This approach aligns with the principles outlined in physics-informed machine learning, which has been shown to improve predictive performance by embedding domain-specific laws and constraints directly into the learning framework⁵⁹.

Moreover, a data-only-based learning approach can lack robustness to out-of-distribution cases. To enhance the performance of our model, we added a physics-based regularization through the addition of a more physically relevant loss (PI-GNN) by Suk et al.¹⁶.

Thus, we approximated differentiable operators in a discrete vector field using finite differences over the neighborhood of nodes, thereby preventing training slowdowns caused by gradient computations. These gradients are computed locally within each cell neighborhood on the mesh, which helps preserve topological consistency even in geometrically complex zones.

With u_i the velocity vector at node $n_i \in N$, \mathcal{N}_i the collection of its direct neighbors and x_i its position, we can approximate the gradient of u_i with:

$$\nabla u_i \approx \frac{1}{\mathcal{N}_i} \sum_{j \in \mathcal{N}_i} \frac{(u_j - u_i)(x_j - x_i)^T}{\|x_j - x_i\|_2^2} \quad (6)$$

Following the same method, we can write the conservation of mass using the discrete approximation of the divergence:

$$\nabla \cdot u_i = \frac{\partial u_{i,x}}{\partial x} + \frac{\partial u_{i,y}}{\partial y} + \frac{\partial u_{i,z}}{\partial z} = 0 \text{ thus } \nabla \cdot u_i \approx \frac{1}{\mathcal{N}_i} \sum_{j \in \mathcal{N}_i} \frac{(u_j - u_i) \cdot (x_j - x_i)}{\|x_j - x_i\|_2^2} \quad (7)$$

Furthermore, we evaluate the convective term $u \cdot \nabla u$ and the viscous term $\mu \Delta u$ using the following approximation of the Laplacian operator:

$$\Delta u \approx \frac{1}{\mathcal{N}_i} \sum_{j \in \mathcal{N}_i} \frac{(\nabla u_j - \nabla u_i)(x_j - x_i)^T}{\|x_j - x_i\|_2^2} \quad (8)$$

Using these approximations, we developed a physical loss function to augment our training by incorporating not only the discrepancies between predicted and ground truth velocities, but also a deeper understanding of the fundamental principles of fluid mechanics. This approach enables us to calculate a continuity loss based on the conservation of mass, as well as convective and viscous losses, which quantify the errors in the convective and viscous terms, respectively (Equation (3)):

$$\mathcal{L}_{continuity} = \text{mean}_{v \in V} |(\nabla \cdot u)_v| \quad (9)$$

$$\mathcal{L}_{convective} = \text{mean}_{v \in V} \|\rho u_v \nabla u_v - \rho \tilde{u}_v \tilde{u}_v\|_2^2 \quad (10)$$

$$\mathcal{L}_{viscous} = \text{mean}_{v \in V} \|\mu \Delta u_v - \mu \Delta \tilde{u}_v\|_2^2 \quad (11)$$

To mitigate the importance of each loss term during training, we combined the data and physical terms in a properly tuned weighted loss, as following:

$$\mathcal{L}_{PI} = \mathcal{L}_{data} + \alpha \mathcal{L}_{continuity} + \beta \mathcal{L}_{convective} + \gamma \mathcal{L}_{viscosity} \quad (12)$$

with α , β , and γ being scaling factors. We tuned our weighted loss by evaluating the physical loss terms during the first epochs of training the data-only equivalent models, to adjust the scaling factors in order to get comparable scales for each contributor of the global loss. The data loss term was given a weight of 0.5, while the remaining physics-informed terms were equally distributed over the remaining 0.5. While the continuity term stabilizes after only one epoch, the convection and viscosity losses decrease steadily during the whole training alongside the data loss, providing a faster model convergence.

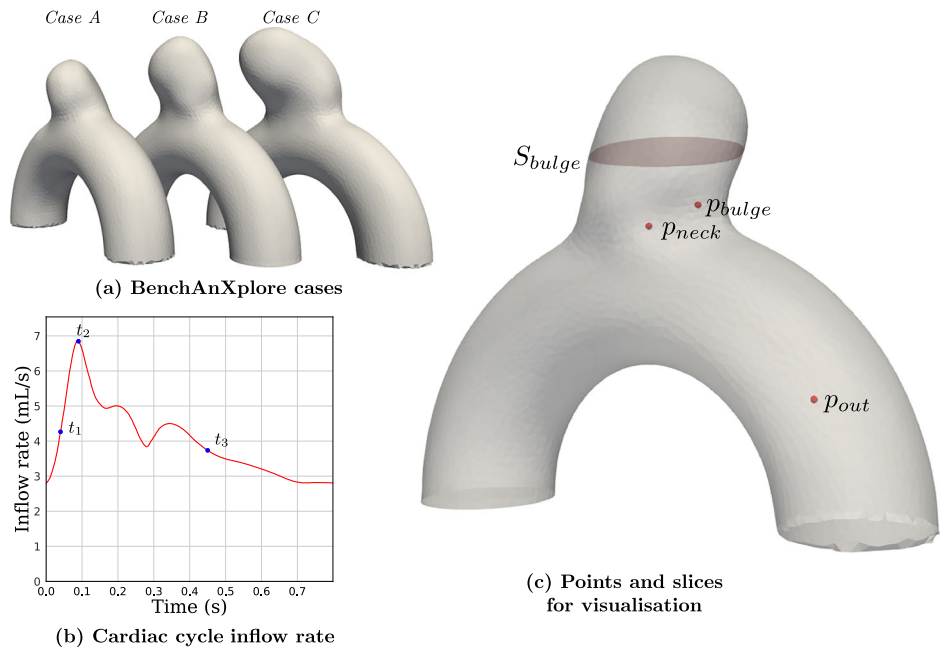
Evaluation metrics

To compare the performances of our models, we computed the L2 norm of the velocity error, and with this error, the Root Mean Square Error (RMSE) averaged across case geometry. We used the 1-RMSE to evaluate the next-step predictions of our model, as well as the 50-RMSE on a rollout prediction to determine the robustness of the model on full trajectory predictions, as formulated in Equation (13). We displayed these errors, averaged on the test set, in Table 1, as well as the standard deviation of the 50-step RMSE to evaluate the stability of the models on different test cases. These errors are in mm/s, and the values reported in Table 1 are absolute errors, not normalized by maximum velocity.

$$\begin{aligned} 1 - \text{RMSE} &= \sqrt{\frac{\sum_i^N (y_i^t - f(x^{t-1}))^2}{N}} \\ T - \text{RMSE} &= \frac{1}{T} \sum_t^T \sqrt{\frac{\sum_i^N (y_i^t - f^t(x^0))^2}{N}} \end{aligned} \quad (13)$$

Three representative cases were randomly chosen from the 10 test cases while guaranteeing coverage of three bulge sizes. Three key timesteps were also selected along the cardiac cycle, as shown in Fig. 20, to represent the systolic, peak systole, and diastolic phases. We computed WSS as well as time-averaged WSS (TAWSS) as defined in Methods to further analyze the prediction capability of the models needed for risk assessment. All WSS-based quantities were derived directly from the velocity fields and were not interpolated from the finer cases. We also plotted the aneurysm neck plane to show the velocity of the flow entering and exiting the bulge, as well as on

Fig. 20 | Evaluation cases, inflow conditions, and sampling locations. **a** Cases selected from the test dataset for detailed analysis (cases A–C). **b** Prescribed pulsatile inflow rate over one cardiac cycle, with three representative time instants (t_{1t_1} , t_{2t_2} , t_{3t_3}) selected for evaluation. **c** Slice and point locations used for velocity data extraction, including the bulge cross-section and the points p_{neck} , p_{bulge} , and p_{out} .



three selected points along the full cardiac cycle to show the velocity evolution inside the neck, the outflow artery and close to the bulge wall, as shown in Fig. 20. The mentioned velocity and WSS errors are normalized by the ground truth maximum at a given timestep; because of the highly dynamic nature of the flow, normalizing by the overall maximum value would highly reduce the error on every timestep far from the peak systole.

Data availability

The datasets used and analyzed during the current study are available at the following link: https://storage.googleapis.com/large-physics-model/datasets/aneurysm/coarse_03_dataset.zip. The custom code used for data generation and visualization are relying solely on open-source Python libraries and are available from the corresponding author upon reasonable request as well.

Code availability

The underlying code for this study are available from the corresponding author upon reasonable request. The custom code used for data generation and visualization are relying solely on open-source Python libraries and are available from the corresponding author upon reasonable request as well.

Received: 6 June 2025; Accepted: 22 January 2026;
Published online: 06 February 2026

References

1. Vlak, M. H., Algra, A., Brandenburg, R. & Rinkel, G. J. Prevalence of unruptured intracranial aneurysms, with emphasis on sex, age, comorbidity, country, and time period: a systematic review and meta-analysis. *Lancet Neurol.* **10**, 626–636 (2011).
2. Nieuwkamp, D. J. et al. Changes in case fatality of aneurysmal subarachnoid haemorrhage over time, according to age, sex, and region: a meta-analysis. *Lancet Neurol.* **8**, 635–642 (2009).
3. Baharoglu, M. I., Lauric, A., Gao, B.-L. & Malek, A. M. Identification of a dichotomy in morphological predictors of rupture status between sidewall- and bifurcation-type intracranial aneurysms: Clinical article. *J. Neurosurg.* **116**, 871–881 (2012).
4. Xiang, J. et al. Hemodynamic-morphologic discriminants for intracranial aneurysm rupture. *Stroke* **42**, 144–152 (2011).

5. Zhou, G., Zhu, Y., Yin, Y., Su, M. & Li, M. Association of wall shear stress with intracranial aneurysm rupture: systematic review and meta-analysis. *Sci. Rep.* **7**, 5331 (2017).
6. Xiang, J., Tutino, V. M., Snyder, K. V. & Meng, H. CFD: computational fluid dynamics or confounding factor dissemination? The role of hemodynamics in intracranial aneurysm rupture risk assessment. *Am. J. Neuroradiol.* **35**, 1849–1857 (2014).
7. Sakakura, Y., Kono, K. & Fujimoto, T. Real time artificial intelligence assisted carotid artery stenting: a preliminary experience. *J. NeuroInterv. Surg.* <https://jn.is.bmj.com/content/early/2025/01/25/jnis-2024-021600> (2024).
8. Sakakura, Y., Masuo, O., Fujimoto, T., Terada, T. & Kono, K. Pioneering artificial intelligence-based real time assistance for intracranial liquid embolization in humans: an initial experience. *J. NeuroInterv. Surg.* <https://jn.is.bmj.com/content/early/2024/06/27/jnis-2024-022001> (2024).
9. Arzani, A., Wang, J.-X., Sacks, M. S. & Shadden, S. C. Machine learning for cardiovascular biomechanics modeling: challenges and beyond. *Ann. Biomed. Eng.* **50**, 615–627 (2022).
10. Gharleghi, R., Sowmya, A. & Beier, S. Transient wall shear stress estimation in coronary bifurcations using convolutional neural networks. *Comput. Methods Prog. Biomed.* **225**, 107013 (2022).
11. Ferdian, E., Dubowitz, D. J., Mauger, C. A., Wang, A. & Young, A. A. WSSNet: aortic wall shear stress estimation using deep learning on 4D flow MRI. *Front. Cardiovasc. Med.* **8**, <https://www.frontiersin.org/journals/cardiovascular-medicine/articles/10.3389/fcvm.2021.769927/full> (2022).
12. Rutkowski, D. R., Roldán-Alzate, A. & Johnson, K. M. Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data. *Sci. Rep.* **11**, 10240 (2021).
13. Charles, R. Q., Su, H., Kaichun, M. & Guibas, L. J. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 77–85, <https://ieeexplore.ieee.org/abstract/document/8099499> (2017).
14. Qi, C. R., Yi, L., Su, H. & Guibas, L. J. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems 30, Annual Conference on*

- Neural Information Processing Systems* (eds Guyon, I. et al.) 5099–5108, <https://proceedings.neurips.cc/paper/2017/hash/d8bf84be3800d12f74d8b05e9b89836f-Abstract.html> (2017).
15. Li, G. et al. Prediction of cerebral aneurysm hemodynamics with porous-medium models of flow-diverting stents via deep learning. *Front. Physiol.* **12**, <https://www.frontiersin.org/journals/physiology/articles/10.3389/fphys.2021.733444/full> (2021).
 16. Suk, J. et al. Physics-informed graph neural networks for flow field estimation in carotid arteries. *CoRR* **abs/2408.07110**, <https://doi.org/10.48550/arXiv.2408.07110> (2024).
 17. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A. & Battaglia, P. W. Learning mesh-based simulation with graph networks <http://arxiv.org/abs/2010.03409> (2021).
 18. Hernández, Q., Badías, A., Chinesta, F. & Cueto, E. Thermodynamics-informed graph neural networks. *IEEE Trans. Artif. Intell.* **5**, 967–976 (2024).
 19. Lam, R. et al. Learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023).
 20. Li, S., Sun, Z., Zhu, Y. & Zhang, C. Physics-constrained and flow-field-message-informed graph neural network for solving unsteady compressible flows. *Phys. Fluids* **36**, 046123 (2024).
 21. Gladstone, R. J. et al. GNN-based physics solver for time-independent PDEs. Preprint at <https://arxiv.org/abs/2303.15681v1> (2023).
 22. Pelissier, U., Parret-Fréaud, A., Bordeu, F. & Mesri, Y. Graph Neural networks for mesh generation and adaptation in structural and fluid mechanics. *Mathematics* **12**, 2933 (2024).
 23. Sen, A. et al. Physics-informed graph neural networks to solve 1-d equations of blood flow. *Comput. Methods Prog. Biomed.* **257**, 108427 (2024).
 24. Suk, J., de Haan, P., Lippe, P., Brune, C. & Wolterink, J. M. Mesh neural networks for SE(3)-equivariant hemodynamics estimation on the artery wall. *Comput. Biol. Med.* **173**, 108328 (2024).
 25. Pegolotti, L. et al. Learning reduced-order models for cardiovascular simulations with graph neural networks. *Comput. Biol. Med.* **168**, 107676 (2024).
 26. Feiger, B. et al. Accelerating massively parallel hemodynamic models of coarctation of the aorta using neural networks. *Sci. Rep.* **10**, 9508 (2020).
 27. Li, G. et al. Prediction of 3D cardiovascular hemodynamics before and after coronary artery bypass surgery via deep learning. *Commun. Biol.* **4**, 99 (2021).
 28. Liang, L., Mao, W. & Sun, W. A feasibility study of deep learning for predicting hemodynamics of human thoracic aorta. *J. Biomech.* **99**, 109544 (2020).
 29. Hachem, E. et al. Reinforcement learning for patient-specific optimal stenting of intracranial aneurysms. *Sci. Rep.* **13**, 7147 (2023).
 30. Chong, B. et al. A multicenter pilot study on the clinical utility of computational modeling for flow-diverter treatment planning. *Am. J. Neuroradiol.* [ajnr;ajnr.A6222v1](http://www.ajnr.org/lookup/doi/10.3174/ajnr.A6222) <http://www.ajnr.org/lookup/doi/10.3174/ajnr.A6222> (2019).
 31. Pelissier, U., Meliga, P. & Hachem, E. A continuum approach with adaptive mesh refinement for platelet plug formation. *Int. J. Numer. Methods Biomed. Eng.* **41**, e70073 (2025).
 32. Yang, X., Xia, D., Kin, T. & Igarashi, T. IntraA: 3D Intracranial Aneurysm Dataset for Deep Learning. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2653–2663 <https://doi.org/10.1109/CVPR42600.2020.00273> (Seattle, WA, USA, 2020).
 33. Goetz, A. et al. AnXplore: a comprehensive fluid-structure interaction study of 101 intracranial aneurysms. *Front. Bioeng. Biotechnol.* **12**, <https://www.frontiersin.org/journals/bioengineering-and-biotechnology/articles/10.3389/fbioe.2024.1433811/full> (2024).
 34. Ramalho, S., Moura, A. B., Gambaruto, A. M. & Sequeira, A. Influence of Blood Rheology and Outflow Boundary Conditions in Numerical Simulations of Cerebral Aneurysms. In (eds Ledzewicz, U., Schättler, H., Friedman, A. & Kashdan, E.) *Mathematical Methods and Models in Biomedicine*, 149–175 (Springer, 2013).
 35. Sunderland, K., Jiang, J. & Zhao, F. Disturbed flow's impact on cellular changes indicative of vascular aneurysm initiation, expansion, and rupture: A pathological and methodological review. *J. Cell. Physiol.* **237**, 278–300 (2022).
 36. Baz, R. A., Scheau, C., Niscoveanu, C. & Bordei, P. Morphometry of the entire internal carotid artery on CT angiography. *Medicina* **57**, 832 (2021).
 37. Sanchez-Gonzalez, A. et al. Learning to simulate complex physics with graph networks. In *Proc. 37th International Conference on Machine Learning (ICML '20)*, Vol. 119, 8459–8468 (JMLR.org, 2020).
 38. Hoi, Y. et al. Characterization of volumetric flow rate waveforms at the carotid bifurcations of older adults. *Physiol. Meas.* **31**, 291–302 (2011).
 39. Reymond, P., Merenda, F., Perren, F., Rüfenacht, D. A. & Stergiopoulos, N. Validation of a one-dimensional model of the systemic arterial tree. *Am. J. Physiol. - Heart Circ. Physiol.* **297**, 208–222 (2009).
 40. Cezbral, J. R. et al. Characterization of Cerebral Aneurysms for Assessing Risk of Rupture By Using Patient-Specific Computational Hemodynamics Models. *Am. J. Neuroradiol.* **26**, 2550–2559 (2005).
 41. Larrabide, I. et al. Change in aneurysmal flow pulsatility after flow diverter treatment <https://ri.conicet.gov.ar/handle/11336/33280> (2015).
 42. Lu, X., Chen, J., Liao, Y. & Wang, X. Hemodynamics of unruptured and ruptured middle cerebral artery aneurysms based on fluid-structure interaction. *J. Clin. Neurosci.* **139**, 111433 (2025).
 43. Qian, Y., Takao, H., Umezumi, M. & Murayama, Y. Risk analysis of unruptured aneurysms using computational fluid dynamics technology: preliminary results. *Ajnr. Am. J. Neuroradiol.* **32**, 1948–1955 (2011).
 44. Saalfeld, S., Voß, S., Beuing, O., Preim, B. & Berg, P. Flow-splitting-based computation of outlet boundary conditions for improved cerebrovascular simulation in multiple intracranial aneurysms. *Int. J. Comput. Assist. Radiol. Surg.* **14**, 1805–1813 (2019).
 45. Brindise, M. C. et al. Multi-modality cerebral aneurysm haemodynamic analysis: in vivo 4D flow MRI, in vitro volumetric particle velocimetry and in silico computational fluid dynamics. *J. R. Soc. Interface* **16**, 20190465 (2019).
 46. Garnier, P., Viquerat, J. & Hachem, E. Multi-grid graph neural networks with self-attention for computational mechanics. *Phys. Fluids* **37**, 087140 (2025).
 47. Kipf, T. N. & Welling, M. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, <https://openreview.net/forum?id=SJU4ayYgl> (2017).
 48. Veličković, P. et al. Graph attention networks. *International Conference on Learning Representations*, <https://openreview.net/forum?id=rJXMpikCZ> (2018).
 49. Ruiz, L., Gama, F. & Ribeiro, A. Gated graph recurrent neural networks. *IEEE Trans. Signal Process.* **68**, 6303–6318 (2020).
 50. Janny, S. et al. EAGLE: Large-scale Learning of Turbulent Fluid Dynamics with Mesh Transformers. *The Eleventh International Conference on Learning Representations*, <https://openreview.net/forum?id=mflX4QpsARJ> (2023).
 51. Garnier, P., Lannelongue, V., Viquerat, J. & Hachem, E. MeshMask: Physics-Based Simulations with Masked Graph Neural Networks. *The Thirteenth International Conference on Learning Representations*, <https://openreview.net/forum?id=bFHR8hNk4I> (2025).
 52. Lin, F. et al. Up-sampling-only and adaptive mesh-based GNN for simulating physical systems. <https://openreview.net/forum?id=bFHR8hNk4I> (2024).
 53. Coupez, T. & Hachem, E. Solution of high-Reynolds incompressible flow with stabilized finite element and adaptive anisotropic meshing. *Comput. Methods Appl. Mech. Eng.* **267**, 65–85 (2013).
 54. Battaglia, P. W. et al. Relational inductive biases, deep learning, and graph networks. *CoRR* **abs/1806.01261**, <http://arxiv.org/abs/1806.01261> (2018).

55. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural Message Passing for Quantum Chemistry. In *Proc. 34th International Conference on Machine Learning*, 1263–1272 (PMLR, 2017).
56. Lynch, S., Nama, N. & Figueroa, C. A. Effects of non-Newtonian viscosity on arterial and venous flow and transport. *Sci. Rep.* **12**, 20568 (2022).
57. Ford, M. D., Alperin, N., Lee, S. H., Holdsworth, D. W. & Steinman, D. A. Characterization of volumetric flow rate waveforms in the normal internal carotid and vertebral arteries. *Physiol. Meas.* **26**, 477–488 (2005).
58. Hachem, E., Rivaux, B., Kloczko, T., Digonnet, H. & Coupez, T. Stabilized finite element method for incompressible flows with high Reynolds number. *J. Comput. Phys.* **229**, 8643–8665 (2010).
59. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).

Acknowledgements

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. European Union (ERC, CURE, 101045042). The views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Author contributions

V.L.: conceptualization, formal analysis, study design, investigation, data curation, software development, methodology, database generation, validation, visualization, writing original draft, review, and editing. P.G, P.J.R., and A.G.: contributed to software development, database generation, methodology, and review. P.M.: contributed to study design, data interpretation, and drafting. Y.C.: provided clinical perspectives to the

study, in addition to critically reviewing the manuscript. E.H.: developed the project concept and design, conceptualization, formal analysis, study design, investigation, software development, methodology, validation, visualization, writing original draft, review, editing, secured funding, and revised the manuscript for important intellectual content.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Elie Hachem.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026