

Original Research

Received 18 May 2026

Revised 24 June 2026

Accepted 15 July 2026

Natural Resources for Human Health 2026, Vol. 6(7s): 1147–1161

KEYWORDS:

Cerebral aneurysm, computational fluid dynamics, neural surrogate modelling, SIREN, graph neural operator, OpenFOAM, haemodynamics.

Surrogate Modelling for Patient-Specific Cerebral Aneurysm CFD Reconstruction

Neha Janu¹, Anahat Gill², Renu Kumawat³, Hemlata Goyal^{4*}

^{1,2,3,4}School of Computer Science and Engineering, Manipal University Jaipur, Jaipur

¹neha.janu@jaipur.manipal.edu

²anahatk05@gmail.com

³renu.kumawat@jaipur.manipal.edu

⁴hemlata.goyal@jaipur.manipal.edu

ABSTRACT: Patient-specific computational fluid dynamics (CFD) shows in detail how blood moves inside a cerebral aneurysm, but it is slow to run. For a single patient, preparing the geometry, generating a mesh, solving, and post-processing may take anywhere from hours to days, and this becomes a real obstacle when many cases need to be reviewed in a short time. In this work we ask a more limited question: once a CFD simulation has been run, can a neural network reconstruct its velocity and pressure fields accurately enough to be useful, and which type of network does this best? We built an automated OpenFOAM pipeline that converted 58 aneurysm geometries from the AneuriskWeb database into 10,210,671 labelled CFD points, and we trained three surrogate models on the same data. A single global model that uses Fourier coordinate encoding and a per-case embedding (CaseFourierField) reaches 74.08% velocity reconstruction accuracy from one checkpoint. Training a separate Sinusoidal Representation Network (SIREN) for each patient does better, with 85.05% mean accuracy and 97.67% on the best case, but needs one model per geometry. A graph neural operator (AneurysmGNO) is the most physically grounded of the three, although under the memory-limited evaluation we were able to run it is not yet competitive, and we explain why. We also report an early approach that did not work: a Poiseuille physics prior performed well on a small controlled dataset but failed badly across all 58 cases. We include this because it explains the design decisions that followed. All results are measured within known geometries and should be read as field reconstruction, not as prediction on unseen patients.

1. INTRODUCTION

Intracranial aneurysms are localised bulges in the wall of a cerebral artery. When one ruptures it causes a subarachnoid haemorrhage, which has a high mortality rate and often leaves survivors with lasting neurological damage [1, 2]. In clinical practice, the risk that an aneurysm will rupture is judged mainly from its shape, using measurements such as size, aspect ratio, and neck width. Haemodynamics adds a further layer of information. How blood mechanically loads the wall is captured by velocity, pressure, wall shear stress, and recirculation, and studies have connected these to both the growth and the rupture of aneurysms [3, 4].

To obtain these quantities, CFD is the standard tool. It solves the incompressible Navier-Stokes equations over a patient's vascular geometry and returns the complete three-dimensional velocity and pressure field within the aneurysm. What makes it costly is not the solve but the work surrounding it: acquiring the surface, repairing and volumetrically meshing it, assigning boundary conditions, running the solver, and post-processing the output. For one patient this can run to several days [5, 6], and reducing that cost is what we set out to do. One route to that reduction is a neural surrogate. Trained offline on a collection of existing CFD results, it predicts velocity and pressure at new locations without a fresh solve. This does not replace CFD; the surrogate can only reproduce the CFD it learned from. The geometry is where the real difficulty lies. As Figure 1 illustrates, a curved parent vessel, branching outlets, and an irregular sac drive flow that shifts abruptly from place to place. No single generic model covers every patient, and in practice each geometry has to be treated on its own.




Figure 1. A patient-specific aneurysm geometry during preparation. The irregular sac and branching vessels are exactly the features that an idealised pipe-flow model gets wrong.

Reaching our final models was not a straight path, and the attempts that failed were as instructive as the ones that worked. We started with a physics prior, reasoning that blood flow obeys laws we understand well. On a small, controlled dataset it performed well, but it broke down once we extended it to all 58 cases. Three data-driven surrogates survived that stage, and Figure 2 summarises them. CaseFourierField is a single global model trained across every case. Per-case SIREN trains one network per patient. AneurysmGNO treats the CFD points as nodes in a graph, connected to their nearest neighbours by edges. The path from a failed prior to Fourier fields, SIREN, and a graph operator is the thread that runs through the paper. Reconstructing this kind of flow well seems to need two things at once: enough spatial resolution to represent sharp gradients, and enough structure to be aware of the geometry.

1.1 Scope of the study

This is a study about reconstructing CFD fields, not about predicting rupture, and we want to be clear about the distinction from the start. Every result reported below is measured on a point-level split within geometries the model has already seen. It therefore describes interpolation inside known cases, not zero-shot prediction on new patients. A rupture-prediction study would additionally require per-patient rupture outcomes, clinical metadata, and longitudinal follow-up, none of which are part of this dataset. We treat rupture prediction as future work and are careful not to imply that our models do it.

1.2 Contributions

The main contributions of this work are threefold.

First, we provide an automated OpenFOAM pipeline that converts 58 AneuriskWeb surfaces into 10,210,671 labelled CFD points, where each point carries its coordinates together with five geometric features: wall distance, axial position, inlet distance, outlet distance, and a local radius proxy.

Second, we compare a global surrogate, a per-patient surrogate, and a graph operator on the same dataset under a common protocol, which to our knowledge has not been done for aneurysm CFD reconstruction.

Third, we document the approaches that failed, including a physics prior that reversed the flow direction at scale, a pointwise model with no sense of which patient it was looking at, and a graph-evaluation error that made a working model appear broken. These failures are reported because they directly shaped the final design.




Figure 2. The CFD-to-surrogate pipeline, from AneuriskWeb geometry to reconstructed field.

2. RELATED WORK

Physics-informed neural networks (PINNs) embed the governing equations directly into the training loss. This is useful when data are limited or noisy, and the idea has carried over successfully to cardiovascular haemodynamics, structural monitoring, and hybrid solvers [7]. Below we organise the closest work according to how it steered our own decisions..

2.1 Physics-informed networks for cardiovascular flow

In a comparison of network architectures for three-dimensional steady blood flow, Moser and colleagues [6] reported that a simple fully-connected network struck the best balance between accuracy and cost, and that pressure inside the aneurysm was the quantity hardest to recover. We later ran into the same difficulty. Zhang and colleagues [8] modelled fluid-structure interaction in elastic vessels with a mesh-free formulation. Arzani and colleagues [4] reconstructed near-wall velocity and wall shear stress from sparse, noisy measurements when the boundary conditions were unknown. Sarabian and colleagues [9] combined Doppler ultrasound, reduced-order simulation, and angiography geometry to recover velocity, pressure, and vessel area in the brain. Galazis and colleagues [10] handled uncertainty in infant arterial spin labelling MRI, Sen and colleagues [11] embedded one-dimensional blood-flow equations in a graph network, and Guan and colleagues [5] focused on speed by splitting transient problems into parallel steady-state sub-problems.

2.2 Data-efficient and physics-regularised learning

Sel and colleagues [12] embedded Taylor-series approximations of cardiovascular relationships in the loss for cuffless blood-pressure estimation and reached high accuracy with much less data. Isaev and colleagues [13] showed that adding conservation

laws improves physical consistency and convergence at arterial bifurcations. Changdar and colleagues [14] combined a PINN with symbolic regression to study nonlinear waves in arterial blood flow.

2.3 Hybrid physics-informed and reduced-order modelling

A large body of work couples PINNs with classical numerical methods. Examples include closure modelling in convection-dominated reduced-order systems [15], ResNet connections added to counter gradient vanishing [16], reduced-order equilibrium equations for structural response [17], and several combinations of the finite element method with PINNs [18, 19, 20, 21, 22, 23]. This literature is the reason we began with a reduced-order prior, and also the reason we then tested whether its pipe-flow assumption survives real sac geometry.

2.4 Neural fields, graph networks, and the remaining gap

MeshGraphNet [24] showed that message passing over irregular meshes outperforms pointwise methods for physical dynamics. GATv2 [25] makes graph attention edge-aware. SIREN [26] represents smooth signals and their derivatives well. Each of these addresses part of our problem, but none of them provides an end-to-end aneurysm pipeline that generates a large multi-case CFD dataset and compares global, per-patient, and graph surrogates under one protocol. Table 1 sets out where the closest prior work stops short of what we do here.

Table 1. Where the closest prior work stops short of the present study.

Work	Method	Contribution	What it does not do
Arzani et al. [4]	PINN reconstruction	Near-wall velocity and WSS from sparse data	No multi-case aneurysm dataset; no comparison of surrogate families
Sarabian et al. [9]	Physics-informed imaging	Cerebral haemodynamics from medical data	No neural-field or graph surrogates for 3D reconstruction
Moser et al. [6]	PINN architecture study	Architecture affects flow and pressure accuracy	No Fourier, SIREN, or GNO comparison on 58 cases
Sitzmann et al. [26]	SIREN	Smooth signals and derivatives	General method; no CFD pipeline or graph comparison
Pfaff et al. [24]	Graph simulator	Message passing on irregular domains	General framework; not aneurysm-specific
Sen et al. [11]	Physics-informed GNN	1D arterial flow	1D only, not 3D point-wise CFD
This work	CaseFourierField, SIREN, AneurysmGNO	58-case dataset with a three-model comparison	—

3. MATERIALS AND METHODS

3.1 Dataset and case selection

The geometries came from the AneuriskWeb angiography database [28]. We started from 101 cases, of which 84 had a usable surface. We removed any case that failed at meshing, solving, or field extraction, and any case with a uniformly zero velocity field, which indicated a failed extraction. Fifty-eight cases came through the full pipeline. We mention this attrition directly, because the cases that failed were not a random sample. They tended to be the ones with messier surfaces, so the accuracies we report are, if anything, a slightly optimistic view of AneuriskWeb rather than a neutral one.




Figure 3. A representative case from AneuriskWeb, with the sac forming at an arterial bifurcation. This irregular morphology is why a pipe-flow prior is a poor fit.

3.2 OpenFOAM CFD pipeline

Each surface was processed through an automated OpenFOAM [27] workflow into a CFD-ready domain with labelled wall, inlet, and outlet patches, and was then solved and exported. The stages were, in order: STL surface, mesh generation, patch assignment, OpenFOAM simulation, VTK export, and CSV extraction. A prepared mesh is shown in Figure 4 and the corresponding velocity field in Figure 5. In total the dataset holds 10,210,671 points; every point stores its coordinates, three velocity components, a scalar pressure, and a label for its boundary patch.




Figure 4. A prepared vascular surface mesh, before CFD simulation.




Figure 5. CFD velocity field inside case C0039. The fast inflow jet and the recirculation in the dome are the structures a surrogate has to reproduce.

3.3 Features and normalisation

Because coordinate ranges, velocity scales, and pressures vary between cases, we normalised the data before training. Per case, coordinates were rescaled to the interval from minus one to one. Velocity was scaled using the 99th-percentile magnitude across the dataset instead of the maximum, since a single impingement point often produces an outlying maximum. Pressure was scaled by its global minimum and range. After this, each point is represented by an eight-dimensional input vector:

$$b^{(0)} = [x', y', z', d_{wall}, a, d_{in}, d_{out}, r_{loc}] \quad (1)$$

In this vector, d_{wall} is the distance to the closest wall node, obtained with a KD-tree to keep the computation feasible across ten million points. The term a gives the projection of the point onto the case-average flow direction. The terms d_{in} and d_{out} are Euclidean distances to the estimated inlet and outlet patch centres. The term r_{loc} is a proxy for distance from the centreline, taken as the perpendicular distance from the first principal axis of the point cloud. The five geometric features were normalised per case to the range zero to one.

3.4 Physics background

Blood flow obeys the incompressible Navier-Stokes equations:

$$\rho (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \mu \nabla^2 \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0 \quad (2)$$

with $\rho = 1060 \text{ kg/m}^3$ and $\mu = 0.004 \text{ Pa}\cdot\text{s}$. The divergence-free constraint later becomes a soft penalty during graph training. Two components recur in the models below. SIREN layers use the activation $\varphi(x) = \sin(\omega x)$, whose derivatives are again sinusoidal, so smooth fields and their gradients are represented cleanly. GATv2 makes graph attention depend on the edge between two points, computing the weights as $\alpha_{ij} = \text{softmax}_j(a^T \text{LeakyReLU}(W[h_i \parallel h_j \parallel e_{ij}]))$ with the edge feature $e_{ij} = [\Delta x, \Delta y, \Delta z, \text{dist}]$. As a result, a neighbour lying downstream of the flow is weighted differently from one lying across it. Finally, a wall-distance no-slip gate, $\text{gate}(\varphi) = 1 - \exp(-\alpha \varphi)$ with $\alpha = 30$, pulls the predicted velocity towards zero at the wall and relaxes away from it.

3.5 The three surrogate models

CaseFourierField (global). This is a single network trained across all 58 cases. Raw coordinates suffer from spectral bias, meaning the network learns smooth structure quickly but fine detail slowly, so we map the coordinates through ten-level Fourier features, $\text{FF}(x') = [\sin(2^k \pi x'), \cos(2^k \pi x')]$ for $k = 0$ to 9, giving a 60-dimensional encoding. Each patient also receives a learnable embedding of dimension 32, which is what lets the model tell one geometry from another. A six-layer MLP with hidden width 512, GELU activations, LayerNorm, and 0.02 dropout, about 1.6 million parameters in total, maps the concatenated input to velocity and pressure. Training used AdamW at a learning rate of 2×10^{-4} for 30,000 epochs, with batch size 65,536 and cosine annealing.




Figure 6. CaseFourierField validation pressure RMSE over 30,000 epochs.

Per-case SIREN (per-patient). Here we train a separate SIREN for each case, so all of the model's capacity is spent on a single geometry. The architecture uses $\omega = 30$, depth 5, and hidden width 256. Training used Adam at 10^{-4} for 8,000 epochs per case, with batch size 16,384 and cosine annealing. All 58 checkpoints are kept for analysis and deployment.

AneurysmGNO (graph). Each case is represented as a directed graph in which the CFD points are nodes and each node is connected to its 12 nearest neighbours in normalised coordinate space. Twelve neighbours was a compromise: fewer reduce the local context available, while more increase memory and training time. The edge features carry the relative offset and distance between the two points. The full component list is given in Table 2, and the model has 1,278,276 parameters. All 58 case graphs are merged into one disconnected graph, and a NeighborLoader samples mini-batches while keeping each node's neighbourhood intact. Training has two phases: a data-fitting phase with loss $L1 = L_u + 0.01 L_p$, followed by a second phase that adds the incompressibility penalty and the no-slip gate. The reason for splitting training in two is given in Section 4.1.

Table 2. AneurysmGNO components.

Component	Configuration	Output dim.
Global branch	MLP 8→64→128, GELU, LayerNorm	128
Node encoder	Linear(8→256), GELU, LayerNorm	256
Graph trunk	4 × GATv2Conv, heads = 2, edge_dim = 4	256
Global injection	Add Linear(128→256) after graph layer 2	256
Skip connection	Add Linear(8→256) from input node features	256
Fusion	Linear(384→256), GELU	256
Velocity head	Linear(256→128), GELU, Linear(128→3)	3
Pressure head	Linear(256→128), GELU, Linear(128→1)	1
Total parameters	1,278,276	




Figure 7. The three surrogate architectures. Left: the global CaseFourierField. Centre: the per-patient SIREN. Right: the graph-based AneurysmGNO.

3.6 Evaluation

All models use a point-level 80/10/10 split inside each case. Because every case contributes points to training, the metric measures interpolation within known geometries rather than prediction on unseen patients. We report the relative L^2 velocity accuracy:

$$Acc_u = 1 - \|u - \hat{u}\|_2 / (\|u\|_2 + \varepsilon) \quad (3)$$

A value of 1.0 is a perfect reconstruction, 0.0 means the error is as large as the field itself, and a negative value means the prediction points in the wrong direction on average. AneurysmGNO is evaluated with a NeighborLoader that tracks node identity, which keeps the graph connected. If evaluation is done naively in chunks, edges are lost and the model looks far worse than it is, a point we return to in the next section.

4. RESULTS AND DISCUSSION

4.1 What we tried, and what it taught us

What we ended up with is the remainder after several ideas failed. We set them out here because each failure pointed to the next decision, and because a neat retrospective account would hide how the models were really arrived at.

A physics prior that looked right, then failed. Our starting point was a Poiseuille reduced-order prior: a closed-form parabolic axial profile that meets the no-slip condition with no numerical work at all. It came directly from the hybrid ROM and PINN literature [15, 17, 16] and from that literature's success in haemodynamic problems [5, 9]. On a controlled set of 18 cases it reached 89.17% accuracy on four fully held-out patients, encouraging enough that we scaled it up. Across all 58 cases, though, the same prior scored -9.52% . A negative score means more than poor accuracy; it means the prior drives velocity in the wrong direction. It presumes axial flow along the centreline, which holds roughly in a straight segment, but within a sac the flow enters from one side, slows, recirculates, and exits through the neck, leaving much of the velocity lateral or reversed. This lines up with the flow behaviour reported by Meng and colleagues [3] and Moser and colleagues [6]. The 89% and the -9.52% are not in conflict; the 18-case set simply had gentler sacs. Placing a network on top of the prior lifted accuracy only to around 21%, since the network spent its effort working against the prior rather than with it, so we discarded it.

Removing the prior revealed a problem in the architecture. Accuracy stayed near 21% even without the prior, which told us the prior was not the whole problem. A plain coordinate MLP has no way to distinguish a small sidewall sac from a large bifurcation sac when both are at the same normalised coordinate, and spectral bias stops it from resolving the sharp velocity gradients near walls and inflow jets. Closing those two gaps is exactly what the per-case embedding and the Fourier encoding in CaseFourierField were meant to do.

A model that looked broken was really a broken metric. AneurysmGNO trained normally but validated at about -20% . The cause was the evaluation, not the model. To fit in GPU memory, we had split each validation graph, which held between 150,000 and 220,000 nodes, into fixed chunks of 4,096 nodes. Any edge whose two endpoints landed in different chunks was silently dropped. With 4,096-node chunks taken from a 150,000-node graph, only about 2.7% of the edges survived, so a model built around message passing had almost no messages to pass. Replacing this with a NeighborLoader evaluation, which keeps each node's neighbourhood intact, removed the artefact. We flag it because it is an easy error to make and a hard one to spot.

The no-slip gate has to wait. Turning the gate on from the first epoch suppresses gradients in precisely the region where most of the points lie, near the wall, before the model has learned anything there. Moving the gate to the second phase, after a plain data-fitting phase, made training more stable. Table 3 collects these four findings and the change each one produced.

Table 3. Four failures and the design change each one forced.

Finding	What we observed	What we changed
Poiseuille prior fails at scale	-9.52% on 58 cases; the prior predicts axial flow while sac flow is lateral and recirculating	Drop the prior and move to a fully data-driven model
Pointwise MLP with no case identity	About 21%; the network cannot separate 58 geometries from coordinates alone	Add a case embedding and Fourier coordinate encoding
Chunked graph evaluation	-20% despite healthy training; about 97% of edges dropped at inference	Use NeighborLoader, neighbourhood-preserving evaluation
No-slip gate applied too early	Near-wall gradients suppressed before they can form	Two-phase training: fit first, apply the gate in phase 2

4.2 CaseFourierField results

On the full point-level split, with 1,021,037 points in each of the validation and test sets, CaseFourierField reaches 75.41% velocity accuracy on validation and 74.08% on test, with pressure RMSE of 0.007591 and 0.006909 respectively (Table 4). Validation and test agree closely, so the model is not badly overfit. Its main value is coverage rather than peak accuracy. One checkpoint serves all 58 cases, which is convenient for a viewer or for a quick scan across patients.

Table 4. CaseFourierField, full point-level evaluation.

Metric	Validation	Test
Velocity accuracy (%)	75.41	74.08
Pressure RMSE	0.007591	0.006909
Evaluated points	1,021,037	1,021,037




Figure 8. CaseFourierField validation velocity accuracy over 30,000 training epochs.

4.3 Per-case SIREN results

The per-patient SIRENs are the strongest reconstructors, with 85.05% mean accuracy, 87.38% median, and 97.67% on the best case, C0071 (Table 5, Figure 9). Of the 58 cases, 38 reach at least 85% and 14 exceed 90%. The six weakest cases, which sit between 56% and 71%, are C0023, C0081, C0095, C0014, C0069, and C0063. These are the same cases that were hardest under every model, which suggests the difficulty lies in the geometry and flow rather than in any one architecture. C0095, the smallest case with roughly 18,900 points, is probably under-resolved. Figure 10 places the best case against its ground truth.

Table 5. Per-case SIREN reconstruction across 58 cases.

Metric	Value
Mean velocity accuracy	85.05%
Median accuracy	87.38%
Minimum accuracy	56.35% (C0023)
Maximum accuracy	97.67% (C0071)
Cases at 90% or above	14 of 58

Metric	Value
Cases at 85% or above	38 of 58
Mean pressure RMSE	0.006040




Figure 9. Per-case SIREN accuracy across 58 cases, with reference lines at 85%, 90%, and 95%.




Figure 10a. CFD ground-truth velocity magnitude for case C0071.




Figure 10b. SIREN-predicted velocity magnitude for case C0071.

4.4 Comparison of all approaches

Table 6 lists every approach we ran. It has to be read with care, because the settings are not identical. The early ROM-PINN result comes from an 18-case controlled experiment, CaseFourierField and SIREN use point-level reconstruction on the full 58-case dataset, and AneurysmGNO was evaluated only on sampled graph patches under a memory limit. For that reason its figure of about 63% is not a like-for-like number, and we do not present it as one. Figure 11 plots predicted against ground-truth velocity, where the spread at low speed corresponds to the near-stagnant interior of the sac, which is the hardest region to reconstruct.

Table 6. All surrogate approaches. The settings differ, so this should be read as a progression rather than a controlled benchmark.

Model	Evaluation	Val. acc.	Test acc.	Note
Poiseuille prior only	All 58 cases	−9.52%	−9.52%	Predicts axial pipe flow while sac flow is lateral and recirculating
MLP with prior, 18 cases	Case-level holdout	96.09%	89.17%*	Good on a controlled set; does not scale to the full dataset
MLP with prior, 58 cases	Point-level	21.4%	n/a	Pointwise model and prior cannot handle the geometric diversity
AneurysmGNO	Sampled patches	~63%†	n/a	Memory-limited; not comparable to the full-node runs
CaseFourierField	Full point-level	75.41%	74.08%	One global checkpoint across all 58 cases
Per-case SIREN	Per-case point-level	n/a	85.05% mean	Best accuracy; one model per patient

* Cross-patient holdout on four cases that contributed no training data. † Sampled patches only; not comparable to full-node evaluation.




Figure 11. Predicted against ground-truth velocity magnitude over all test points. The scatter at low velocities reflects the near-stagnant regions of the sac.

4.5 Discussion

The gap of about ten percentage points between the global model and the per-patient SIRENs has two causes that push in the same direction. A SIREN spends all of its capacity on one geometry, whereas CaseFourierField has to divide its capacity across 58. In addition, sinusoidal activations represent smooth fields and their gradients more naturally than GELU does. The two models are therefore better seen as two operating points than as competitors. One offers a single checkpoint with reasonable accuracy, the other offers a model per patient with the best accuracy we obtained.

This raises the question of where AneurysmGNO fits. It is the only one of the three that builds in the local spatial coupling that the governing equations actually contain, which is why we believe it is the right direction for the problem we cannot yet solve, namely generalising to a patient the model has never seen. We want to be honest that this is an expectation rather than a result. Our current figure is memory-limited and measured within known geometries, and the cross-geometry claim remains untested.

On the clinical side, we reconstruct velocity and pressure and nothing beyond that. Wall shear stress, the oscillatory shear index, relative residence time, and endothelial cell activation potential can all be derived from a resolved velocity field, and a rupture model could in principle be built on top of them, but that would need outcome labels and independent validation that we do not have. Table 7 states plainly what is available now and what is not, so that the boundary is not blurred.

Table 7. Rupture-related quantities: what this study provides and what it does not.

Quantity	Requires	Status
Velocity magnitude	3D velocity field	Available
Pressure variation	Pressure field	Available
Inflow jet and recirculation	Velocity and geometry	Qualitative
Wall shear stress	Wall-normal velocity gradients	Future work
TAWSS and OSI	Time-resolved WSS	Future work
Rupture classification	Outcome labels and metadata	Future work

4.6 Limitations

Three limitations stand out. The first is that our evaluation reflects interpolation inside known geometries rather than prediction on new patients. The second is that the CFD assumes steady-state boundary conditions and treats blood as Newtonian; both are usual simplifications at this scale, yet neither captures every physiological effect. Third, the AneurysmGNO figure is memory-limited, so it should be read as a floor set by the available hardware rather than as a verdict on the architecture.

5. Conclusion and Future Work

We built a 58-case CFD dataset from AneuriskWeb geometries and used it to ask which kind of surrogate best reconstructs aneurysm velocity and pressure fields. The per-patient SIREN was the most accurate, with 85.05% mean accuracy and 97.67% on the best case. A single global Fourier model traded a few points of accuracy, reaching 74.08%, for the convenience of one checkpoint across every case. The graph operator, though memory-limited at present, is the model we would choose to pursue for the unsolved problem of generalising to new patients. The most useful outcome is not a single leaderboard number but a pattern: this flow needs both high-frequency spatial capacity, which SIREN provides, and geometry-aware coupling, which the graph operator provides. A model that combines the two, trained at full scale on a larger set of patients, is the natural next target. Throughout, we have tried to keep our claims in line with the evidence. These are within-geometry reconstructions rather than rupture predictions, and the one incomplete result is labelled as such.

Several directions follow naturally. The most immediate is to give AneurysmGNO the memory it needs, through multi-GPU training or hierarchical graph coarsening, and then run the evaluation we actually care about, which is to hold out whole cases and test true zero-shot reconstruction. Beyond that, the reconstructed fields form a basis for a second stage that computes wall shear stress, TAWSS, OSI, RRT, and ECAP, which together with rupture labels would let a downstream model learn rupture-associated patterns. Moving to pulsatile inflow, using Doppler or phase-contrast MRI waveforms with Windkessel outlet models, would make the time-resolved biomarkers available. Adding a non-Newtonian model such as Carreau-Yasuda, and re-processing the 26 excluded cases, would improve near-wall accuracy in low-shear regions and broaden the range of geometries. Finally, an interactive tool that we built alongside the models (Figure 12) already supports dataset upload, model inference, field visualisation, metric computation, and report export; extending it to encode new geometries would make the pipeline usable from end to end.




Figure 12. The interactive haemodynamic viewer, showing reconstructed three-dimensional velocity magnitude for one case.

ACKNOWLEDGEMENTS

The authors thank the maintainers of the AneuriskWeb database for making the aneurysm geometries publicly available.

AUTHOR CONTRIBUTIONS

Neha Janu: supervision, writing and review. Anahat Gill: Conceptualization, methodology, software, writing. Renu Kumawat: review and editing. Hemlata Goyal: correspondence, revised and edited.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

FUNDING

No external funding was received for this work.

ETHICS STATEMENT

This study uses publicly available, anonymised aneurysm geometry data and does not involve any direct interaction with patients.

DATA AVAILABILITY STATEMENT

The patient-specific vascular geometries used in this study are publicly available from the AneuriskWeb database[28]. AneuriskWeb[28] provides surface geometries only; the velocity and pressure labels were generated by the authors by running CFD simulations on these geometries. The generated dataset, trained model checkpoints, and source code are available from the corresponding author on reasonable request.

REFERENCES

- [1] Rinkel, G.J., Djibuti, M., Algra, A., Van Gijn, J., 1998. Prevalence and risk of rupture of intracranial aneurysms: a systematic review. *Stroke* 29(1), 251–256.
- [2] Suarez, J.I., Tarr, R.W., Selman, W.R., 2006. Aneurysmal subarachnoid hemorrhage. *New England Journal of Medicine* 354(4), 387–396.
- [3] Meng, H., Tutino, V.M., Xiang, J., Siddiqui, A., 2014. High WSS or low WSS? Complex interactions of hemodynamics with intracranial aneurysm initiation, growth, and rupture. *American Journal of Neuroradiology* 35(7), 1254–1262.
- [4] Arzani, A., Wang, J.-X., D'Souza, R.M., 2021. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. *Physics of Fluids* 33(7).
- [5] Guan, H., Dong, J., Lee, W.-N., 2023. Towards real-time training of physics-informed neural networks: applications in ultrafast ultrasound blood flow imaging. *arXiv preprint arXiv:2309.04755*.
- [6] Moser, P., Fenz, W., Thumfart, S., Ganitzner, I., Giretzlehner, M., 2023. Modeling of 3D blood flows with physics-informed neural networks: comparison of network architectures. *Fluids* 8(2), 46.
- [7] Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. 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.
- [8] Zhang, H., Chan, R.H., Tai, X.-C., 2024. A meshless solver for blood flow simulations in elastic vessels using a physics-informed neural network. *SIAM Journal on Scientific Computing* 46(4), C479–C507.
- [9] Sarabian, M., Babae, H., Laksari, K., 2022. Physics-informed neural networks for brain hemodynamic predictions using medical imaging. *IEEE Transactions on Medical Imaging* 41(9), 2285–2303.
- [10] Galazis, C., Chiu, C.-E., Arichi, T., Bharath, A.A., Varela, M., 2025. Pinning cerebral blood flow: analysis of perfusion MRI in infants using physics-informed neural networks. *Frontiers in Network Physiology* 5, 1488349.
- [11] Sen, A., Ghajar-Rahimi, E., Aguirre, M., Navarro, L., Goergen, C.J., Avril, S., 2024. Physics-informed graph neural networks to solve 1-D equations of blood flow. *Computer Methods and Programs in Biomedicine* 257, 108427.
- [12] Sel, K., Mohammadi, A., Pettigrew, R.I., Jafari, R., 2023. Physics-informed neural networks for modeling physiological time series for cuffless blood pressure estimation. *npj Digital Medicine* 6(1), 110.
- [13] Isaev, A., Dobroserdova, T., Danilov, A., Simakov, S., 2024. Physically informed deep learning technique for estimating blood flow parameters in arterial bifurcations. *Lobachevskii Journal of Mathematics* 45(1), 239–250.
- [14] Changdar, S., Bhaumik, B., Sadhukhan, N., Pandey, S., Mukhopadhyay, S., De, S., Bakalis, S., 2024. Integrating symbolic regression with physics-informed neural networks for simulating nonlinear wave dynamics in arterial blood flow. *Physics of Fluids* 36(12).
- [15] Kaya, F., Koc, B., Aygun, A., Ata, O., Karakus, A., 2026. Hybrid ROM-PINN framework for closure modeling in convection-dominated systems. *arXiv preprint arXiv:2603.01998*.
- [16] Liu, Y., Hou, J., Wei, P., Jin, J., Zhang, R., 2025. Research and application of ROM based on Res-PINNs neural network in fluid system. *Symmetry* 17(2), 163.
- [17] Sun, Z.-Y., Guo, Y.-T., Ge, K., Hou, C., Hu, Z.-Z., 2026. ROM-PINN: a physics-informed neural network with reduced-order modelling for nonlinear structural response prediction. In: *Structures* 86, Elsevier, 111350.
- [18] Ebrahimi, H., Morovati, A., 2025. Hybrid FEM–PINN framework for 2D shallow-water flow simulation and error analysis.
- [19] Beitalmal, A.O., 2025. A hybrid physics-informed neural network (PINN) and finite element method (FEM) framework for multiscale fluid-structure interaction problems. *RA Journal of Applied Research* 11(04).
- [20] Feng, X., Shangguan, H., Tang, T., Wan, X., Zhou, T., 2024. A hybrid FEM-PINN method for time-dependent partial differential equations. *arXiv preprint arXiv:2409.02810*.
- [21] Sibuet, N., Ares de Parga, S., Bravo, J.R., Rossi, R., 2025. Discrete physics-informed training for projection-based reduced-order models with neural networks. *Axioms* 14(5), 385.
- [22] He, X., Li, K., Wang, S., Lai, X., Yang, L., Kan, Z., Song, X., 2024. Toward an online monitoring of structural performance based on physics-informed hybrid modeling method. *Journal of Mechanical Design* 146(1), 011702.
- [23] Sahin, T., Wolff, D., von Danwitz, M., Popp, A., 2024. Towards a hybrid digital twin: physics-informed neural networks as surrogate model of a reinforced concrete beam. *arXiv preprint arXiv:2405.08406*.
- [24] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.W., 2020. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*.

- [25] Brody, S., Alon, U., Yahav, E., 2021. How attentive are graph attention networks? arXiv preprint arXiv:2105.14491.
- [26] Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G., 2020. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems* 33, 7462–7473.
- [27] Jasak, H., Jemcov, A., Tukovic, Z., et al., 2007. OpenFOAM: a C++ library for complex physics simulations. In: *International Workshop on Coupled Methods in Numerical Dynamics 1000*, Dubrovnik, Croatia, 1–20.
- [28] AneuriskWeb project website, <http://ecm2.mathcs.emory.edu/aneuriskweb>. Emory University, Department of Math&CS, 2012.