AneuG-Flow: A Large-Scale Synthetic Dataset of Diverse Intracranial Aneurysm Geometries and Hemodynamics

Wenhao Ding

Department of Biomedical Engineering Imperial College London United Kingdom w.ding23@imperial.ac.uk

Simão Castro

Department of Mechanical Engineering
Instituto Superior Técnico
Portugal
simao.vitoriano@tecnico.ulisboa.pt

Yiying Sheng*

Department of Biomedical Engineering National University of Singapore Singapore yiyings@nus.edu.sg

Hwa Liang Leo

Department of Biomedical Engineering National University of Singapore Singapore bielhl@nus.edu.sg

Choon Hwai Yap

Department of Biomedical Engineering Imperial College London United Kingdom c.yap@imperial.ac.uk

Abstract

Hemodynamics has a substantial influence on normal cardiovascular growth and disease formation, but requires time-consuming simulations to obtain. Deep Learning algorithms to rapidly predict hemodynamics parameters can be very useful, but their development is hindered by the lack of large dataset on anatomic geometries and associated fluid dynamics. This paper presents a new large-scale dataset of intracranial aneurysm (IA) geometries and hemodynamics to support the development of neural operators to solve geometry-dependent flow governing partial differential equations. The dataset includes 14,000 steady-flow cases and 570 pulsatile-flow cases simulated with computational fluid dynamics. All cases are computed using a laminar flow setup with more than 3 million cells. Boundary conditions are defined as a parabolic velocity profile with a realistic waveform over time at the inlet, and geometry-dependent mass flow split ratios at the two downstream outlets. The geometries are generated by a deep generative model trained on a cohort of 109 real IAs located at the middle cerebral artery bifurcation, capturing a wide range of geometric variations in both aneurysm sacs and parent vessels. Simulation results shows substantial influence of geometry on fluid forces and flow patterns. In addition to surface mesh files, the dataset provides volume data of velocity, pressure, and wall shear stresses (WSS). For transient cases, spatial and temporal gradients of velocity and pressure are also included. The dataset is tested with PointNet and graph U-Nets for WSS prediction, which showed relative L2 loss of 4.67% for normalized WSS pattern.

^{*}The first and second authors contributed equally to this work.

1 Introduction

For many vascular diseases, there has been a long history of investigating fluid mechanics behaviors aimed at improving surgical decision-making and post-operative treatment [1]. However, a significant gap remains between biomechanics research and its adoption in clinical practice. A clear example of this gap is intracranial aneurysm (IA)—a condition where a weakness in the vessel wall leads to a bulge that carries a potentially lethal risk of rupture [2]. Despite extensive research in biomechanics and several meta studies suggesting a correlation between biomechanical markers and aneurysm instability [3], these markers have yet to be accepted or utilized by physicians. In practice, physicians continue to rely primarily on morphological markers [4] to assess rupture risk, even though such systems [5] have been reported to suffer from high false-positive rates (up to 56%) [6].

One of the key reasons for this gap is that fluid dynamics simulations are time-consuming and require specialized expertise. As a result, physicians are unable to obtain fluid dynamics data in sufficiently large sample sizes to evaluate the clinical utility of biomechanics markers, preventing their adoption of it. As such, it would be useful to develop deep learning models for rapid prediction of flow dynamics in disease morphologies. Driven by this clinical demand, the concept of predicting CFD solutions directly from vascular geometries has gained significant momentum in recent years. In terms of application domains, models have been developed for coronary arteries [7], aorta [8], aortic aneurysms [9], and the left ventricle [10]. From a methodological perspective, both supervised and self-supervised models are actively being explored. Representative work in the supervised category includes a family of geometry-informed neural operators [10, 11, 12, 13], while self-supervised approaches—such as [14, 15, 16]—leverage physics-informed training. Both approaches require a substantial dataset for training and/or validation.

Unfortunately, available datasets in the bioengineering domain to support such tasks are far from satisfactory. Most existing studies—particularly those focused on theoretical model development—rely on idealized geometry datasets. For instance, simple harmonic functions are employed in [15] and [16] to represent the 3D geometries of coronary arteries. Similarly, models in the geometry-informed neural operator family are often evaluated on highly simplified and impractical geometric domains [11], limiting their translational relevance to real-world applications. This limitation largely stems from the inherent difficulty of assembling a sufficient range of patient-specific geometries. As shown in Table 1, most available datasets contain only a few hundred geometries. Moreover, generating the corresponding CFD solutions demands significant High-Performance Computing (HPC) resources and considerable human expertise to design, tune, and maintain an efficient simulation workflow.

In addition to being limited in scale, several existing datasets suffer from other notable shortcomings. For example, [17] provides real aneurysm sac geometries, but attaches them to a single idealized mother vessel, resulting in globally unphysiological vascular structures with no variation. Similarly, while [18] offers a large-scale dataset, over 95% of them are manually deformed from real shapes rather than being generated through data-driven approaches, and the parent vessels were not generated even though they are critical to flow behaviors. To address the shortage of realistic 3D hemodynamics data and to offer a more physiological and large-scale alternative, we introduce a new dataset of IA geometries with associated hemodynamics. The main contributions are summarized as follows:

- We provide the first large hemodynamics dataset of both IA 3D geometry and its detailed hemodynamics where both aneurysm pouch and parent vessels are modelled. Geometries show decent diversity representative of that of clinical data.
- A relatively large amount of geometries are provided compared to existing datasets in the bioengineering community. We include 14,000 steady cases and 200 pulsatile cases.
- In addition to the volume data, we provide wall shear stress (WSS) solution on IA sac surfaces with a consistent mesh graph structure (connectivity), allowing easier downstream deep learning processing.

2 Case Description

2.1 Geometry

Unlike existing datasets which create geometry variations by manually warping real shapes [9, 18], we use a generative model AneuG [19] that learns the geometry distribution from a real IA cohort.

Table 1: Reivew on hemodynamics datasets

Dataset	Anatomical region	Size	Geometry source		
Faisal et al.[9]	Aortic aneurysm	230	Manually generated from 23 real shapes.		
AnXplore[17]	Intracranial aneurysm	101	Real IA sac geometries merged with the same idealized mother vessel.		
Aneumo[18]	Intracranial aneurysm	10,000	9,534 synthetic shapes manually generated from 466 real shapes.		
AneuG-Flow (Ours)	Intracranial aneurysm (with parent vessels)	14,000	Synthetic shapes generated with a generative model trained on a cohort of 116 real shapes.		

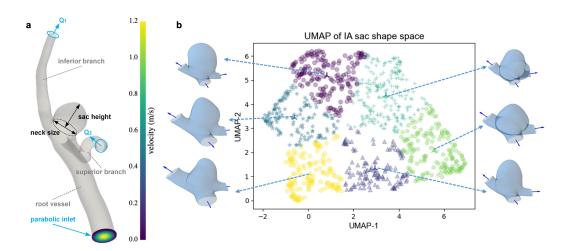


Figure 1: Geometry configuration. a: Flow configuration. b: UMAP of aneurysm sac geometries.

AneuG leverages the technique of Graph Harmonic Deform (GHD) [20] to encode the spatial warping of real IA shapes with respect to a canonical shape into a sequence of tokens. The distribution of tokens are then approximated with a two-stage Variational Autoencoder (VAE). More details can be found in [19] and [20]. AneuG generates both the aneurysm sac and its parent vessels, with the latter being conditioned on the former. This approach allows us to create diverse physiological IA shapes for CFD simulations, while existing works such as [17] fails to capture the joint distribution of aneurysm sacs and their parent vessels. AneuG uses data from the AneuX morphology database [21], an open-access, multi-centric database combining data from three European projects: AneuX project, @neurIST project and Aneurisk. As reported in [21], all patients/participants provided written informed consent to participate in the study.

As shown in Fig. 1, we create IAs located at the bifurcation of the middle cerebral artery (MCA). Such a topology contains one root vessel and two downstream vessel branches (referred to as the inferior and superior branch). Compared to the inferior branch, the superior branch generally has smaller radius and larger spatial angle with respect to the root vessel.

2.2 Mesh Generation Pipeline

We assembled a set of open-source repositories and commercial software automation scripts to develop a fully automated pipeline for volume mesh generation. Low-resolution surface meshes generated from AneuG were processed using Geomagic Wrap v2024.3.0 for remeshing and localized smoothing. We then used the Vascular Modeling Toolkit (VMTK) to generate 3D volume meshes (codes modified from [8]). For each case, 4 inflation layers were applied with a total thickness equal to 0.5 times the local edge length near the wall. The final meshes contained an average of 3.4 million volume cells, which mesh convergence studies have shown to be sufficient (see supplementary

materials for details). This pipeline is generalizable to vascular structures with a fixed number of inlets and outlets. These automation codes can be found in [22].

2.3 Boundary Conditions

As shown in Fig. 1, we apply a parabolic velocity profile at the inlet with an average velocity of $0.684\,\mathrm{m/s}$, as measured in [24]. For the outlets, several studies have demonstrated that flow split conditions more accurately represent physiological hemodynamics compared to fixed pressure boundaries [25, 26]. Following these work, we determine the mass flow split between the superior and inferior branches using a modified form of Murray's law:

$$\frac{Q_1}{Q_2} = \left(\frac{D_1}{D_2}\right)^{\gamma} \tag{1}$$

where Q_1 and Q_2 are the outlet flow rates, D_1 and D_2 are the corresponding vessel diameters, and γ is the flow split exponent. We follow [27] and choose $\gamma = 2.45$. This formulation ensures that the outlet boundary conditions reflect geometry-dependent mass flow split.

Blood were assumed as imcompressible non-Newtonian fluid. We adopt the Carreau–Yasuda model to account for the non-Newtonian behavior of blood:

$$\mu(\dot{\gamma}) = \mu_{\infty} + (\mu_0 - \mu_{\infty}) \left[1 + (\lambda \dot{\gamma})^a \right]^{\frac{n-1}{a}}$$
 (2)

The model parameters are set as follows: zero-shear viscosity $\mu_0 = 0.056\,\mathrm{Pa\cdot s}$, infinite-shear viscosity $\mu_\infty = 0.00345\,\mathrm{Pa\cdot s}$, time constant $\lambda = 3.313\,\mathrm{s}$, and power-law index n = 0.3568. The Yasuda parameter is assumed to be a = 2 and the blood density is set to $1050\,\mathrm{kg/m^3}$.

An laminar flow setup was used, as the average Reynolds numbers in the root vessel and downstream branches were approximately 410 and 330, respectively. And vessel walls were assumed as rigid with a no-slip condition. Since measurements of blood pressure and flow velocity are not routinely performed during the clinical management of intracranial aneurysms, an average waveform reported in [28] was adopted for all transient simulations. The average inlet velocity was kept same across all morphologies to ensure that aneurysms with larger parent vessels experienced higher mass flow rates. The velocity waveform signal is included in our dataset. Each transient case was simulated over two cardiac cycles using a time step of 0.001, s, which was confirmed to be sufficient through a time-step convergence study. Only the results from the second cycle were extracted.

2.4 Solver and HPC Setup

All simulations were performed using ANSYS Fluent 2023R2 (ANSYS Inc., Canonsburg, PA, USA) on the High-Performance Computing (HPC) Services at Imperial College London and National University of Singapore. Simulations ran on AMD nodes, each contains 128 cores and 1TB RAM. The Research Data Store at Imperial College London were used for data storage during the runs. Each simulation was run on a single AMD node using 64 cores. Each case took approximately 3 minutes to mesh, and 2 minutes to solve the steady simulation. Transient cases each took around 10 hours to solve.

2.5 Graph structural consistency during CFD data extraction

In addition to the raw CFD solution data, we also provide additional post-processed WSS on the surfaces of IA sacs as graphs. Each graph contains the same number of nodes and the same connectivity, allowing easier implementation of downstream deep learning tasks. AneuG gnerates surface meshes using a mesh morphing approach, where every case has exactly 3500 triangle faces for the aneursym sac [19]. As this mesh size is of low resolution, we subdivided each triangular element by adding a new node at the center of the each edge and dividing the element into four new elements, leading to 14,000 faces. WSS at nodes were then extracted through k-NN interpolation method from [23]. Further subdivision can be performed if higher resolution is desired. The standardization of node and connectivity structure allows a natural and easy node-to-node / edge-to-edge registration between different IA cases, and facilitates deep learning processing. The associated graph connectivity is also provided (see Table 2). Extraction codes are available at [22].

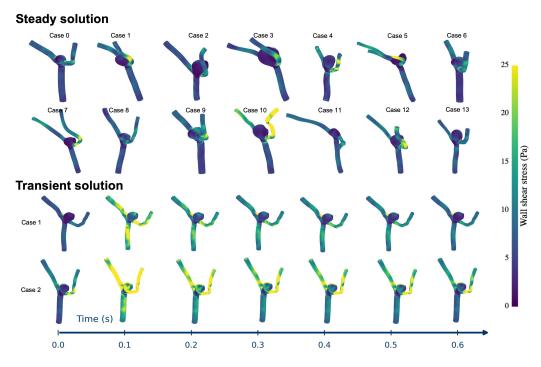


Figure 2: Example of wall shear stress data.

3 Dataset Description

3.1 Geometry Variation

We follow [19] to generate synthetic shapes using pretrained checkpoint files. Latent features were sampled from a 64-dimensional uniform distribution with a mean of 0 and a standard deviation of 2.5, and were subsequently passed through the decoder. We chose this setup because it covers over 98% of the latent space of a theoretically well-trained VAE. It is also the maximum deviation from the latent space center that still yields physiologically plausible shapes. As shown in Fig. 1 b, diverse aneurysm sac geometries are included.

3.2 Dataset Contents

A summary of the content of this dataset is provided in Table 2 and Table 3. For steady cases, we assemble the solution data of all cases into one Pytorch .pth file while keeping geometry-dependent files in separate case folders. As mentioned above, we also provide node-to-node registration for the aneurysm sac region to construct a well-structured PyTorch tensor object. This tensor has the shape of [B,N,C], where B,N, and C denote the number of cases, the number of surface nodes, and the number of physics variables, respectively.

We also provide a list of downsampled node indexes and associated edge connection of downsampled meshes obtained using the method described in [30]. This downsampling approach preserves the topological integrity of the shapes, whereas generating low-resolution k-NN graphs based solely on Euclidean coordinates may introduce connections between points that are close in space but distant in terms of surface geodesic distance. These graph structures can be used for U-Net-like structures with graph convolution layers. A visualization is provided in Fig. 3.

For transient cases, we provide time-series data over one full cardiac cycle, as summarized in Table 3. Solutions were extracted at 80 uniformly spaced time steps within the second cardiac cycle, resulting in each PyTorch tensor object having the shape [T,N,1], where T denotes the number of time steps and N the number of nodes. PyTorch .pth files were saved separately within each case folder.

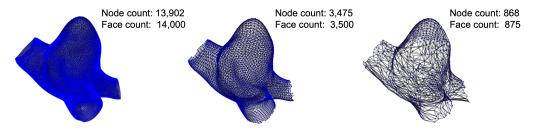


Figure 3: Mesh downsampling.

Table 2: Dataset contents of steady cases

Solution data						
File & Structure	Keys	Content				
	label	List of solution variable names.				
D 1 (4	tensor	PyTorch tensor object of solution data.				
Raw_data.pth List[Dict[str, Any]]	ghd	GHD tokens [19] and rigid alignment checkpoints.				
Elst[Diet[sti, / my]]	log	Fluent residuals log.				
	case_name	Case name.				
	tensor	PyTorch tensor object of registered solution data.				
	tensor_norm	Mean and standard deviation of registered data.				
Assembled_data.pth Dict[str, Any]	idx_list	Vertex index list of downsampled surface meshes				
Diet[su, my]	edge_index_list	Edge connections of downsampled surface meshes.				
	ds_factors	List of downsampling factors.				
Geometry data (inside each case folder)						
flowsplit_ratio.txt		Mass flow split ratio between two outlets.				
shape.obj		Surface mesh generated by AneuG.				
shape_remeshed.obj		Post-processed surface mesh (by Geomagic).				

3.3 How to access the dataset

The **AneuG-Flow** dataset is open-source under a CC BY-SA 4.0 license. It can be accessed and downloaded directly from the Hugging Face Hub. Codes for processing the raw CFD data and baseline training are available at GitHub [23].

3.4 Limitations

We aim to provide a large-scale hemodynamics dataset to support the prediction of biomechanical markers directly from patient-specific vascular geometries. However, several limitations remain. First, the dataset includes a relatively small number of transient cases (578) compared to that of steady cases (14,000), which may limit its utility for predicting temporal biomechanical markers such as time-averaged wall shear stress and oscillatory shear index. Expanding the dataset with more transient simulations is a priority for our future work. Second, the current dataset is limited to geometries with a single inlet and two outlets. To improve generalizability, additional vascular topologies — such as shapes with a single inlet and a single outlet — should be incorporated in the next phase. Finally, we do not consider variations of boundary conditions at the moment, as measuring blood pressure and mass flow waveform signals in the middle cerebral artery is not yet a routine in clinical practice.

Table 3: Dataset contents of transient cases

Solution data (inside each case folder)					
File & Structure	Keys	Content			
blood_data.pt Dict[str, Tensor]	various	Volume solution data of the blood domain, including spatial coordinates, velocity components, presure, viscosity, temporal derivative of pressure, and spatial derivatives of velocity components.			
wall_data.pt Dict[str, Tensor]	various	solution data of the surface domain, including spatial coordinates, Wall shear stress components, and the total wall shear stress magnitude.			
Geometry data (inside each case folder)					
same as Table 2					

4 ML Evaluation

We demonstrate a simple regression task as an application of the dataset. Given the geometries of IAs, we train several baseline models to predict the steady-state WSS map on the aneurysm sac. Leveraging the surface node-to-node registration described in Section 2.5, we construct input graphs with the same connectivity across different cases, each containing 13,902 nodes and 14,000 triangles. The network is designed to output WSS vector components in the x, y, and z directions. We use 80% of the steady-state cases for training and the remaining 20% for testing.

Networks & loss function design. We adopt a U-Net-like structures for the models. Specifically, we train a PointNet++ and three graph U-Nets: one using simple Graph Convolutional Networks (GCN), another using Graph Attention Networks (GAT), and a third using Chebyshev Spectral Graph Networks (ChebNet) as the core convolutional layers. For PointNet++, downsampling is performed using farthest point sampling (FPS) [29]. For graph U-Nets, a pre-computed set of downsampled node indices and associated edge connections was used, as mentioned in Section 3.2. A visualization of the mesh downsampling is provided in Fig. 3. In addition to the mean squared error (MSE) loss computed on z-score-normalized wall shear stress (WSS), we also evaluate an MSE loss defined on WSS values normalized using an power mapping:

$$\mathcal{L}_{\text{exp}} = \text{MSE}\left[f(w^{\text{pred}}), \ f(w^{\text{true}})\right], \quad f(w) = \left(\frac{\alpha \cdot w}{w_{\text{max}}}\right)^{\beta} \cdot \frac{1}{\alpha}$$
 (3)

where w denotes the WSS components and f is the normalization function. α and β are manually selected hyperparameters. For this task, we chose $\alpha=100$ and $\beta=\frac{2}{3}$. We introduce this loss term because high WSS values often appear near the junction areas between the aneurysm sac and its parent vessels. However, it is the low WSS distribution on the sac that is generally considered more clinically significant by physicians [3]. By applying such a nonlinear normalization to the WSS, the model is encouraged to learn the overall spatial pattern of WSS rather than being dominated by high-magnitude areas. As shown in Fig. 4, the normalized WSS exhibits reduced contrast between high and low values, thereby emphasizing the underlying distribution pattern.

Metrics & Results. Each model is trained on a single NVIDIA RTX 3090 GPU for 24 hours, with the learning rate decayed by a factor of 0.75 every 100 epochs. Model performance is evaluated using the root mean square error (RMSE) and mean absolute error (MAE) computed on the WSS values. In addition, both metrics and the relative L2 error are computed on WSS normalized using Eq. (3), which reflects the model's capability of predicting the spatial pattern of WSS. Prediction performances are visualized for three random cases in Fig. 4. The model captures the global WSS map on the aneurysm sac well. As expected, high WSS is observed near the junction between the sac and the parent vessels, while low WSS appears around lobulated regions on the sac. Among different network designs, the U-Net using ChebNet as the convolutional layer performs the best. During training, the ChebNet kernel size was set to 3, and all networks were configured with identical depth

Table 4: Model performance evaluated on WSS and Eq. (3)-normalize

Madal	WSS		Normalized WSS		
Model	RMSE (Pa)	MAE (Pa)	RMSE	MAE	Relative L2 (%)
PointNet++	0.204	0.122	0.0353	0.0263	5.39
U-Net (GCN)	0.199	0.114	0.0317	0.0233	4.82
U-Net (GAT)	0.213	0.120	0.0333	0.0242	5.04
U-Net (ChebNet)	0.191	0.108	0.0307	0.0223	4.67

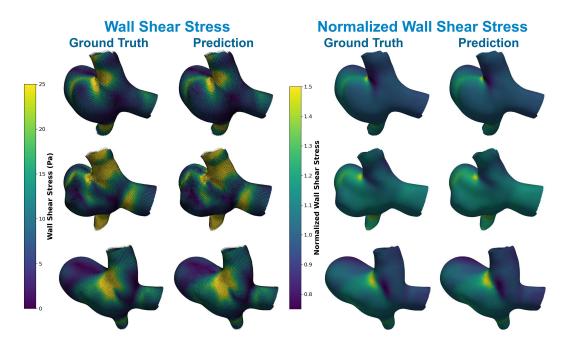


Figure 4: Prediction of WSS and WSS normalized with Eq. (3) from the trained U-Net (GCN).

and channel dimensions. The better performance is likely attributable to the enhanced receptive field of ChebNet, allowing more effective propagation of information across the mesh.

5 Relationship between morphological & biomechanical markers

We go beyond steady cases and investigate the relationship between morphological and biomechanical markers with our transient data. For morphological markers, we include Aspect Ratio (AR), Size Ratio (SR), and maximum sac height (H_{max}) as they are generally accepted in existing physiology research [4, 5]. Further, We include Lobulation Index (LI) defined as the sac's surface area divided by its volume. We also compute the radius of a sphere with the same volume of the sac as it reflects the 3D size of the aneurysm, which we refer to as Equivalent Radius (ER). A detailed graphic definition of these markers can be found in the supplementary material. For biomechanical markers, we consider two of them: Oscillatory Shear Index (OSI) and Relative Residence Time (RRT) [31]. The Oscillatory Shear Index (OSI) measures the directional changes of wall shear stress throughout a cardiac cycle, defined as:

$$OSI = \frac{1}{2} \times \left(1 - \frac{\left| \int_0^T \boldsymbol{\tau_w}(t) dt \right|}{\int_0^T |\boldsymbol{\tau_w}(t)| dt} \right)$$
(4)

with values ranging from 0 (unidirectional flow) to 0.5 (oscillatory flow). Here $\tau_w(t)$ is the wall shear stress vector. Relative Residence Time (RRT) combines both parameters to evaluate flow stagnation:

$$RRT = \frac{1}{(1 - 2 \times OSI) \times TAWSS}$$
 (5)

where Time-Averaged Wall Shear Stress (TAWSS) is the average magnitude of wall shear stress over a complete cardiac cycle, calculated as:

$$TAWSS = \frac{1}{T} \int_0^T |\boldsymbol{\tau_w}(t)| dt$$
 (6)

It is generally considered dangerous when large OSI and RRT are observed on the aneurysm sac's surface. As shown in Fig. 5 and Fig. 6, most morphological markers demonstrated weak or statistically insignificant correlations with both OSI and RRT averaged on the aneurysm sac. This suggests that traditional morphological markers alone may be insufficient to capture the complex hemodynamic behavior within aneurysmal sacs.

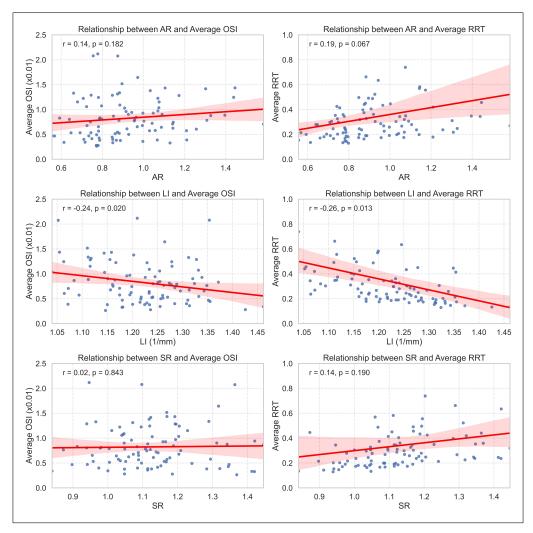


Figure 5: Morphological-biomechanical marker correlations (AR, SR, LI).

An interesting observation is the moderate but statistically significant negative correlation between the lobulation index (LI) and both OSI ($r=-0.24,\,p=0.020$) and RRT ($r=-0.26,\,p=0.013$). While this trend is statistically supported, it is counterintuitive. One would expect aneurysms with higher LI to be associated with more chaotic and oscillatory flow patterns. A possible explanation is that true daughter sacs—often linked to rupture risk and flow complexity—were relatively rare. And the LI marker as defined here, may be elevated in aneurysms with elongated but smooth morphologies. These cases could exhibit high surface-to-volume ratios without necessarily possessing complex internal flow structures. Therefore, the specificity of LI as a marker may be limited. As this is an preliminary investigation, further studies with larger and more diverse datasets are needed. In contrast, the equivalent radius (ER), which reflects the size of the aneurysm sac based on its volume, showed a

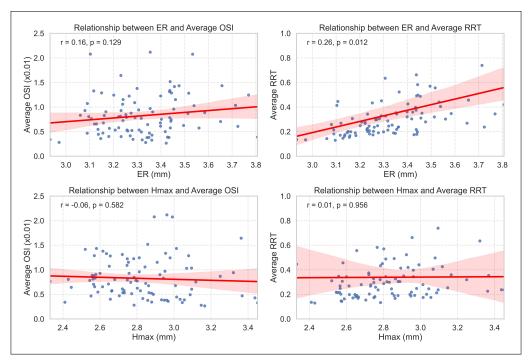


Figure 6: Morphological-biomechanical marker correlations (ER and H_{max})

moderate positive correlation with both OSI (r=0.16, p=0.129) and RRT (r=0.26, p=0.012). This trend is intuitive, as a larger ER corresponds to a larger sac volume, increasing the likelihood of developing slow and recirculating flow. Such regions are associated with disturbed hemodynamics behaviors, including elevated oscillatory shear and prolonged residence time.

6 Conclusion

In this paper, we present a new large-scale and open-source dataset designed to support the development of data-driven models for predicting hemodynamics from geometries. The dataset includes 14,000 steady-flow cases and 200 pulsatile-flow cases, each computed using high-resolution CFD simulations with anatomically physiological IA geometries. These geometries model both the aneurysm sacs and their parent vessels as a joint distribution, addressing key limitations in previous datasets that relied on idealized or manually deformed shapes [9, 18].

By leveraging a deep generative model trained on a real IA cohort, we capture a broad range of physiologically plausible geometries representative of real-world anatomical variations. We provide solution data including pressure, velocity, and WSS. Spatial gradients for velocity components and temporal gradient for pressure are also provided. Initial experiments using PointNet and graph U-Nets demonstrate the dataset's utility in enabling WSS pattern prediction, achieving a best relative L2 error of 4.67% on normalized WSS.

We hope this dataset will contribute to the biomechanics and machine learning communities, accelerating the development of neural operators and other data-driven solvers for geometry-conditioned partial differential equations.

Acknowledgments and Disclosure of Funding

This study was supported by the Imperial College startup funding and MOE-AcRF-Tier1-FRC-FY2024 Grant. We would also like to acknowledge that computational work involved in this study is partly supported by the National University of Singapore's IT Research Computing group under grant number NUSREC-HPC-00001.

References

- [1] Carpenter, H.J., Gholipour, A., Ghayesh, M.H., Zander, A.C., Psaltis, P.J.: A review on the biomechanics of coronary arteries. International Journal of Engineering Science **147**, 103201 (2020).
- [2] Lee, K.S., Zhang, J.J.Y., Nguyen, V., Han, J., Johnson, J.N., Kirollos, R., Teo, M.: The Evolution of Intracranial Aneurysm Treatment Techniques and Future Directions. Neurosurgical Review **45**(1), 1–25 (2022).
- [3] Zhou, G., Zhu, Y., Yin, Y., Su, M., Li, M.: Association of Wall Shear Stress with Intracranial Aneurysm Rupture: Systematic Review and Meta-Analysis. Scientific Reports 7(1), 5331 (2017).
- [4] Merritt, W.C., Berns, H.F., Ducruet, A.F., Becker, T.A.: Definitions of intracranial aneurysm size and morphology: A call for standardization. Surgical Neurology International **12**, 506 (2021).
- [5] Etminan, N., Brown, R.D., Jr., Beseoglu, K., et al.: The Unruptured Intracranial Aneurysm Treatment Score: A Multidisciplinary Consensus. Neurology **85**(10), 881–889 (2015).
- [6] Mocco, J., Brown, R.D., Torner, J.C., Capuano, A.W., Fargen, K.M., Raghavan, M.L., Piepgras, D.G., Meissner, I., Huston, J.: Aneurysm Morphology and Prediction of Rupture: An International Study of Unruptured Intracranial Aneurysms Analysis. Neurosurgery 82(4), 491–496 (2018).
- [7] Li, G., Wang, H., Zhang, M. et al.: Prediction of 3D Cardiovascular Hemodynamics Before and After Coronary Artery Bypass Surgery via Deep Learning. Communications Biology **4**, 99 (2021).
- [8] Pajaziti, E., Montalt-Tordera, J., Capelli, C., Sivera, R., Sauvage, E., Quail, M., Schievano, S., Muthurangu, V.: Shape-Driven Deep Neural Networks for Fast Acquisition of Aortic 3D Pressure and Velocity Flow Fields. PLOS Computational Biology **19**(4), e1011055 (2023). doi: 10.1371/journal.pcbi.1011055.
- [9] Faisal, M.A.A., Mutlu, O., Mahmud, S. et al.: Rapid Wall Shear Stress Prediction for Aortic Aneurysms Using Deep Learning: A Fast Alternative to CFD. Medical & Biological Engineering & Computing (2025).
- [10] Yin, M., Charon, N., Brody, R. et al.: A Scalable Framework for Learning the Geometry-Dependent Solution Operators of Partial Differential Equations. Nature Computational Science **4**, 928–940 (2024).
- [11] Li, Z., Huang, D.Z., Liu, B., Anandkumar, A.: Fourier Neural Operator with Learned Deformations for PDEs on General Geometries. Journal of Machine Learning Research **24**(1), Article No. 388, 18593–18618 (2023).
- [12] Quackenbush, B., Atzberger, P.J.: Transferable Foundation Models for Geometric Tasks on Point Cloud Representations: Geometric Neural Operators. arXiv preprint arXiv:2503.04649 (2025).
- [13] Quackenbush, B., Atzberger, P.J.: Geometric Neural Operators (GNPs) for Data-Driven Deep Learning of Non-Euclidean Operators. arXiv preprint arXiv:2404.10843 (2024).
- [14] Zhong, W., Meidani, H.: Physics-Informed Geometry-Aware Neural Operator. Computer Methods in Applied Mechanics and Engineering 434, 117540 (2025).
- [15] Wong, H.S., Chan, W.X., Li, B.H. et al.: Strategies for Multi-Case Physics-Informed Neural Networks for Tube Flows: A Study Using 2D Flow Scenarios. Scientific Reports 14, 11577 (2024).
- [16] Chan, W.X., Ding, W., Li, B., Wong, H.S., Yap, C.H.: Role of Physics-Informed Constraints in Real-Time Estimation of 3D Vascular Fluid Dynamics Using Multi-Case Neural Network. Computers in Biology and Medicine 190, 110074 (2025).
- [17] Goetz, A., Jeken-Rico, P., Pelissier, U., Chau, Y., Sédat, J., Hachem, E.: AnXplore: a comprehensive fluid-structure interaction study of 101 intracranial aneurysms. Frontiers in Bioengineering and Biotechnology **12** (2024).

- [18] Li, X., Zhou, Y., Xiao, F., Guo, X., Zhang, Y., Jiang, C., Ge, J., Wang, X., Wang, Q., Zhang, T., Lin, C., Cheng, Y., Qi, Y.: Aneumo: A Large-Scale Comprehensive Synthetic Dataset of Aneurysm Hemodynamics. arXiv preprint arXiv:2501.09980 (2025).
- [19] Ding, W., Ji, K., Castro, S., Luo, Y., Roi, D., Yap, C.H.: Two-Stage Generative Model for Intracranial Aneurysm Meshes with Morphological Marker Conditioning. Medical Image Computing and Computer Assisted Intervention – MICCAI 2025, Lecture Notes in Computer Science 15969, 595–604 (2025).
- [20] Luo, Y., Sesija, D., Wang, F., Wu, Y., Ding, W., Huang, J., et al.: Explicit Differentiable Slicing and Global Deformation for Cardiac Mesh Reconstruction. arXiv preprint arXiv:2409.02070v2 (2024).
- [21] Juchler, N., Schilling, S., Bijlenga, P., Kurtcuoglu, V., Hirsch, S.: AneuX: Shape Trumps Size Image-Based Morphological Analysis Reveals that the 3D Shape Discriminates Intracranial Aneurysm Disease Status Better than Aneurysm Size. Frontiers in Neurology 13, 809391 (2022).
- [22] Sheng, Y.: AneuG-CFD: Automatic mesh generation code for intracranial aneurysms. GitHub repository, https://github.com/yiyingsheng07/AneuG_CFD.git (2025).
- [23] Ding, W.: AneuG-Flow: Dataset and mesh generation code for intracranial aneurysm flow prediction. GitHub repository, https://github.com/WenHaoDing/AneuG-Flow.git (2025).
- [24] Weston, M.E., Barker, A.R., Tomlinson, O.W., Coombes, J.S., Bailey, T.G., Bond, B.: Middle Cerebral Artery Blood Velocity and End-Tidal Carbon Dioxide Responses to Moderate Intensity Cycling in Children, Adolescents, and Adults. Journal of Applied Physiology 137(5), 1117–1129 (2024).
- [25] Chnafa, C., Brina, O., Pereira, V.M., Steinman, D.A.: Better Than Nothing: A Rational Approach for Minimizing the Impact of Outflow Strategy on Cerebrovascular Simulations. American Journal of Neuroradiology **39**(2), 337–343 (2018).
- [26] 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. International Journal of Computer Assisted Radiology and Surgery **14**(10), 1805–1813 (2019).
- [27] Chnafa, C., Bouillot, P., Brina, O., Delattre, B.M.A., Vargas, M.I., Lovblad, K.O., Pereira, V.M., Steinman, D.A.: Vessel calibre and flow splitting relationships at the internal carotid artery terminal bifurcation. *Physiological Measurement* **38**(11), 2044–2057 (2017).
- [28] 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. Physiological Measurement 26(4), 477–488 (2005).
- [29] Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. arXiv preprint arXiv:1706.02413 (2017).
- [30] Tan, Q., Gao, L., Lai, Y.-K., Xia, S.: Variational Autoencoders for Deforming 3D Mesh Models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5841–5850 (2018).
- [31] Trenti, C., Ziegler, M., Bjarnegård, N., et al.: Wall shear stress and relative residence time as potential risk factors for abdominal aortic aneurysms in males: a 4D flow cardiovascular magnetic resonance case—control study. Journal of Cardiovascular Magnetic Resonance 24, 18 (2022). 10.1186/s12968-022-00848-2

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",
- Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Main contributions are covered in the introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are discussed in Section 3.4.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Full assumptions are covered in Section 2.3. Mesh and time-step convergence analyses are covered in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [No]

Justification: We do not describe all details of the network design. However, we do provide them in the codes.

Guidelines:

• The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Links can be found in Section 3.3.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

• Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [No]

Justification: We do not mention all details of the training. However, we provide them in the codes.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Our dataset is simulated with synthetic geometries, we therefore do not have statistical significance of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Related information are reported in Section 2.4.

Guidelines:

• The answer NA means that the paper does not include experiments.

- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]
Justification: [NA]
Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]
Justification: [NA]
Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- · For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLM is used for minor grammar checking, but not for any methodology parts. Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.