

École polytechnique de Louvain

Hemodynamic Characterization of Intracranial Aneurysms

Author: **Edouard BARETTE**

Supervisor: **Benoît MACQ**

Readers: **Colin VANDEN BULCKE, John LEE**

Academic year 2025–2026

Master [120] in Computer Science and Engineering

Abstract

Intracranial aneurysms affect approximately 3.2% of the general population. Rupture causes subarachnoid hemorrhage, a life-threatening emergency with a case fatality rate approaching 40%. Existing risk scores aggregate morphological and clinical variables but do not capture patient-specific blood flow patterns inside the aneurysm sac. This work pursues two complementary objectives: (i) identifying which hemodynamic indices separate growing from stable aneurysms in a matched cohort, and (ii) assessing the ability of a physics-informed neural network (PINN) to reproduce the hemodynamic fields obtained from a reference finite element (FEM) simulation on the same geometries, together with the associated computational cost.

Both objectives are addressed on six patient-specific intracranial aneurysm geometries drawn from the Vascular Model Repository, arranged into three anatomically matched growing–stable pairs. FEM simulations with svMultiPhysics serve as the ground truth. PINNs trained with the NVIDIA PhysicsNeMo framework are evaluated against this reference on identical geometries and boundary conditions.

Absolute wall-averaged time-averaged wall shear stress (TAWSS) does not discriminate the growing from the stable group at this sample size, because the two anatomical sites carry markedly different baseline shear levels. Normalising dome shear stress by the paired parent-vessel value removes this confound. The dome-to-parent TAWSS ratio is consistently lower in the growing group across all three pairs (0.26, 0.40, 0.39 versus 0.79, 0.58, 0.58), with a large effect size (Cohen’s $d = -2.99$). This provides cross-sectional matched-pair evidence that relative dome shielding is associated with growth status.

The pure pulsatile PINN achieves a cycle-averaged relative ℓ_2 velocity error between 0.81 and 0.90, with TAWSS underestimated by 65 to 89%. Training requires approximately 24 hours on a single A100 GPU per case, which is 8 to 37 times more expensive than the FEM solve for five of the six cases. Adding sparse FEM supervision reduces the velocity error to 0.10–0.30 and closes the TAWSS gap to within $\pm 32\%$ of the FEM reference. In the best case, on the smallest mesh, about 2% of the FEM mesh nodes are enough to drive most of this gain. The hybrid approach is a practical strategy when even a small number of reference measurements is available.

Acknowledgements

I would like to thank my supervisor, Benoît Macq, for the guidance, the time invested in reviewing this work, and the conversations that shaped the direction of this thesis.

I am grateful to Colin Vanden Bulcke, the reader of this thesis, for his careful reading and constructive feedback.

I also thank Simon François for his feedback on this work and for introducing me to SLURM job scripting.

This work relied heavily on the CÉCI computing infrastructure, in particular the Lyra cluster. I thank the CÉCI support team for their remarkably prompt responses to technical issues and for the quality of their documentation and video tutorials, which proved invaluable throughout this project.

I thank Dr Patrice Finet, neurosurgeon at Cliniques universitaires Saint-Luc, for taking the time to answer my clinical questions.

Finally, I thank the SimVascular Development Team for developing and maintaining both the SimVascular software suite and the Vascular Model Repository dataset, and for making both freely available.

List of Abbreviations

A

ACA Anterior Cerebral Artery
ACoA Anterior Communicating Artery

C

CÉCI Consortium des Équipements de Calcul Intensif
CFD Computational Fluid Dynamics
CISM Calcul Intensif et Stockage de Masse
CT Computed Tomography
CTA CT Angiography

D

DSA Digital Subtraction Angiography

E

ECAP Endothelial Cell Activation Potential

F

FEM Finite Element Method
FSI Fluid-Structure Interaction

G

GPU Graphics Processing Unit

H

HPC High-Performance Computing

I

IA Intracranial Aneurysm
ICA Internal Carotid Artery

L

LSA Low Shear Area

M

MCA Middle Cerebral Artery
MLP Multi-Layer Perceptron
MPI Message Passing Interface
MRA Magnetic Resonance Angiography
MRI Magnetic Resonance Imaging

O

OSI Oscillatory Shear Index

P

PCoA Posterior Communicating Artery
PDE Partial Differential Equation
PHASES Population, Hypertension, Age, Size, Earlier SAH, Site
PINN Physics-Informed Neural Network
PSWSS Peak-Systolic Wall Shear Stress

R

ReLoBRaLo Relative Loss Balancing with Random Lookback
RRT Relative Residence Time

S

SAH Subarachnoid Hemorrhage
SCI Shear Concentration Index
SLURM Simple Linux Utility for Resource Management
STL Stereolithography (file format)
SV SimVascular

T

TAWSS Time-Averaged Wall Shear Stress
TOF Time of Flight

V ---

- VMR** Vascular Model Repository
- VTK** Visualization Toolkit (software for 3D)
- VTP** VTK PolyData (file format)
- VTU** VTK Unstructured Grid (file format)

W ---

- WSS** Wall Shear Stress
- WSSG** Wall Shear Stress Gradient

Contents

Abstract	i
Acknowledgements	ii
List of Abbreviations	iii
List of Figures	x
List of Tables	xii
Introduction	1
1 State of the Art	2
1.1 Medical Context of Intracranial Aneurysms	2
1.1.1 Context and Epidemiology	2
1.1.2 Pathogenesis and Morphological Classification	3
1.1.3 Detection and Imaging Modalities	4
1.1.4 Clinical Decision-Making and Treatment Options	5
1.2 Hemodynamics of Intracranial Aneurysms	6
1.2.1 Blood Flow Fundamentals	6
1.2.2 Hemodynamic Parameters and Indices	7
1.2.3 Hemodynamic-Based Rupture Risk Assessment	9
1.2.4 Growing vs. Stable Aneurysms: Hemodynamic Signatures	10
1.3 Computational Fluid Dynamics for Cerebrovascular Research	11
1.3.1 Governing Equations	11
1.3.2 Patient-Specific CFD Simulation Pipeline	12
1.3.3 SimVascular	13
1.3.4 Other CFD Tools and Limitations	14

1.4	Physics-Informed Neural Networks	14
1.4.1	Fundamentals of PINNs	14
1.4.2	PINNs for Fluid Dynamics	15
1.4.3	NVIDIA PhysicsNeMo	16
1.4.4	PINNs vs. Traditional CFD	17
2	Materials and Methods	18
2.1	Dataset: Growing versus Stable Aneurysms	18
2.2	SimVascular CFD Simulation Pipeline	19
2.2.1	Model Preparation	19
2.2.2	Meshing	19
2.2.3	Boundary Conditions	21
2.2.4	Solver Configuration	21
2.2.5	Post-Processing and Parameter Extraction	22
2.3	PINN-Based Simulation with PhysicsNeMo	23
2.3.1	Mesh Processing and STL Generation	23
2.3.2	Problem Formulation	24
2.3.3	Network Architecture and Training	24
2.3.4	Adaptive Loss Weighting	25
2.3.5	Hybrid PINN with Sparse FEM Supervision	26
2.3.6	PINN Post-Processing and WSS Extraction	27
2.3.7	HPC Deployment	28
2.4	Experimental Design	28
2.4.1	E1: Steady-State PINN Architecture Ablation	28
2.4.2	E2: Pulsatile PINN on the Full Cohort	29
2.4.3	E3: Pure PINN vs FEM Accuracy and Cost	29
2.4.4	E4: Hybrid PINN with Sparse FEM Supervision	30
2.4.5	E5: Data-Sparsity Sweep	30
2.4.6	E6: Growing versus Stable Hemodynamic Comparison	30
2.5	Hemodynamic Analysis Pipeline	30
2.5.1	Parameter Definitions and Computation	30
2.5.2	Dome and Parent Vessel Analysis	31
2.5.3	Statistical Analysis	31

3	Results	33
3.1	SimVascular CFD Results	33
3.1.1	Simulation Overview and Convergence	33
3.1.2	Velocity and Pressure Fields	34
3.1.3	Wall Shear Stress and OSI Maps	34
3.2	PINN-Based Simulation Results	39
3.2.1	Steady-State PINN: Preliminary Exploration	39
3.2.2	Pulsatile PINN Training Convergence	39
3.2.3	Predicted Pulsatile Flow Fields	41
3.3	Pulsatile PINN vs FEM Comparison	43
3.3.1	Cycle-Averaged Velocity Error	43
3.3.2	Wall TAWSS and OSI Agreement	43
3.3.3	Per-Case Error Distribution	43
3.3.4	Per-Cycle Visual Comparison	45
3.4	Hybrid PINN: Sparse FEM Supervision	46
3.4.1	Per-Patient Training	46
3.4.2	Data-Sparsity Sweep	47
3.4.3	Per-Patient Hybrid Cycle Animations	48
3.5	Hemodynamic Analysis	48
3.5.1	Per-Case Hemodynamic Indices	48
3.5.2	Growing vs Stable Group Comparison	49
3.5.3	Paired Analysis	50
4	Discussion	52
4.1	Evaluation of the PINN Approach	52
4.1.1	Root Causes of the Velocity Discrepancy	52
4.1.2	Effect of Sparse FEM Supervision	53
4.1.3	Accuracy and Computational Cost	53
4.1.4	Potential for Clinical Application and Future Directions	54
4.2	Interpretation of Hemodynamic Differences	54
4.3	Limitations of This Study	57
4.4	Perspectives for Future Work	58
	Conclusion	59

Bibliography	65
A Use of Artificial Intelligence	66
B Code and Configuration Details	67
C Supplementary Hemodynamic Data	68

List of Figures

1.1	Example of an intracranial aneurysm (also called cerebral aneurysm) [1].	2
1.2	Main intracranial aneurysm morphologies	3
1.3	Cerebral vasculature and common aneurysm locations.	4
1.4	Surgical and endovascular treatments for intracranial aneurysms. (a) Endovascular coiling of the aneurysm sac. (b) Surgical clipping of the aneurysm neck. (c) Endovascular treatment combining coils and a stent. (d) Endovascular treatment with a flow diverter [2].	5
1.5	Illustrative TAWSS distribution on an aneurysm surface: low WSS (blue) concentrates in the dome while elevated WSS (beige-orange) appears at the inflow neck. This pattern is consistently associated with increased rupture risk.	9
1.6	SimVascular GUI [3].	13
1.7	Example of a PINN architecture.	15
2.1	Meshes for all six cases. Each row shows one matched pair (growing left, stable right). Aneurysm sacs are visible as bulging regions off the parent vessel. Mesh density statistics are given in Table 2.2.	20
2.2	Population-averaged internal carotid artery inflow waveform used as the inlet boundary condition for all six cases (VMR case 0207, period $T = 0.95$ s, mean $Q_{\text{mean}} = 4.084 \text{ cm}^3/\text{s}$). Red markers: the tabulated VMR samples. Solid line: the 10-mode Fourier reconstruction. The dashed line marks the cycle mean.	22
3.1	FEM velocity magnitude over one cardiac cycle, all six cases (20-frame orbit loop). Each row: one matched pair. Left: growing, right: stable. Per-case colour scale. Static fallback: peak-systole frame.	35
3.2	Gauge pressure over one cardiac cycle, same paired layout as Figure 3.1. Per-case symmetric scale centred on the mean. Static fallback: peak-systole frame.	36
3.3	TAWSS surface maps, all six cases (20-frame 360° orbit). Per-case colour scale. Same paired layout as Figure 3.1. Static fallback: single still.	37

3.4	OSI surface maps, all six cases (20-frame 360° orbit, 0 = unidirectional, 0.5 = fully oscillatory). Same paired layout as Figure 3.1. Static fallback: single still.	38
3.5	Steady-state velocity magnitude on case 0203: FEM reference (left) vs Modified Fourier $4 \times 256 + \text{ReLoBRaLo PINN}$ (right). Shared $[0, 0.51]$ cm/s colour scale.	40
3.6	Pulsatile PINN total training loss for all six cases. Faint lines: raw loss sampled every 100 steps. Solid lines: running-median smooth over 100 samples.	41
3.7	Pulsatile PINN velocity magnitude on case 0203 over one cardiac cycle (20-frame orbit loop). Static fallback: peak-systole frame.	42
3.8	PINN vs FEM agreement on the aneurysm wall. Left: TAWSS (Pa). Right: OSI. The dashed line is the $y = x$ identity.	44
3.9	Case 0203 over the analysis cycle (20-frame orbit loop). Left: SV velocity. Centre: PINN velocity (shared colour bar). Right: pointwise error $ \mathbf{u}_{\text{PINN}} - \mathbf{u}_{\text{SV}} $. Static fallback: peak-systole frame.	45
3.10	Three-way velocity comparison on case 0203 (20-frame orbit loop): SV reference (left), pure PINN (centre), hybrid PINN (right), shared colour bar. Static fallback: peak-systole frame.	46
3.11	Cycle-averaged velocity (blue) and pressure (red) ℓ_2 errors on case 0203 vs N_s . Dashed line: pure-PINN baseline.	47
3.12	Hybrid PINN vs SimVascular, case 0203 (20-frame orbit loop): SV velocity (left), hybrid PINN velocity (centre), pointwise error (right), shared colour scale. Static fallback: single still.	48
3.13	Box plots of wall hemodynamic indices for the growing and stable groups ($n = 3$ each). Individual case values are overlaid as black dots.	50

List of Tables

1.1	PHASES score: factors, categories, and assigned points [4]. The total score maps to a 5-year rupture risk ranging from 0.4% (score ≤ 2) to 17.8% (score ≥ 12).	6
2.1	Dataset overview with clinical metadata. Each pair consists of one growing and one stable aneurysm at the same anatomical site. Outlets: number of outflow boundaries. HTN: hypertension, Hx: family history of aneurysm, M: male, F: female, Y: yes, N: no.	19
2.2	Mesh statistics of the reused VMR meshes. BL: prismatic boundary-layer sublayers.	21
2.3	Summary of the six experiments. Each experiment is reported in a dedicated subsection of Chapter 3.	28
3.1	Pulsatile FEM simulation overview. Four MPI ranks, $\Delta t = 5 \times 10^{-4}$ s, four cardiac cycles (7600 steps), analysis window: steps 5719–7600.	34
3.2	Steady-state PINN ablation on case 0203. Final total loss, wall time, and Adam steps completed within the 24-hour budget.	39
3.3	Pulsatile PINN training duration per patient on a single A100 GPU.	40
3.4	Cycle-averaged relative ℓ_2 velocity error (Eq. 2.10) between PINN and FEM, restricted to lumen voxels.	43
3.5	Wall TAWSS and OSI from PINN and FEM, with relative difference (PINN – FEM)/FEM on TAWSS.	44
3.6	Per-cycle peak and volume-mean velocity magnitude (PINN and FEM), averaged over 20 frames. Last column: mean gauge-pressure relative ℓ_2 error.	45
3.7	Hybrid PINN vs pure-PINN agreement on the analysis cycle, averaged over the same 20 frames per case as Table 3.6.	46
3.8	Data-sparsity sweep on case 0203. Cycle-averaged ℓ_2 errors on velocity and pressure, and TAWSS error, vs supervised samples per frame N_s	47
3.9	Group comparison between growing and stable aneurysms ($n = 3$ per group). Mean $\pm$ SD and Cohen’s d effect size (growing – stable). r_X denotes the dome-to-parent ratio of index X . Per-case values are in Appendix C.	49

3.10	Within-pair differences (growing – stable) for all indices. “Cons.” is “Y” when the sign is the same across all three pairs.	51
C.1	Per-case wall area-averaged hemodynamic indices. TAWSS, PSWSS, Δ WSS in Pa. OSI dimensionless. RRT and ECAP in Pa^{-1} . WSSG in Pa/cm . LSA ^{0.4} is the wall-area fraction with TAWSS < 0.4 Pa.	68
C.2	Per-case dome (subscript d) and parent-vessel (subscript p) area-averages, with the dome-to-parent TAWSS ratio r_{TAWSS} . Units as in Table C.1. . . .	68
C.3	Per-case low-shear area fraction at three TAWSS thresholds. Subscript d = dome region, no subscript = full wall.	69
C.4	Supplementary group-level statistics for secondary hemodynamic indices not reported in Table 3.9. Mean $\pm$ SD and Cohen’s d (growing – stable). Units: WSSG in Pa/cm . PSWSS and Δ WSS in Pa. PSWSS/TAWSS and dome-to-parent LSA ratios dimensionless.	69

Introduction

Intracranial aneurysms are localised dilations of the cerebral arterial wall whose rupture causes subarachnoid hemorrhage, an acute cerebrovascular event with a case fatality rate close to 40% [5]. They are far from rare in the general population, with a pooled prevalence around 3.2% [6], making incidental detection increasingly common. The clinical question is therefore rarely whether an aneurysm exists, but whether and when it should be treated. Most lesions are managed conservatively and require an estimate of the rupture risk from the available patient data. Risk tools such as the PHASES score aggregate clinical and morphological variables but do not capture the patient-specific blood flow patterns inside the sac [4].

Patient-specific computational fluid dynamics (CFD) provides access to these hemodynamic patterns [7, 8]. Indices such as time-averaged wall shear stress (TAWSS) and oscillatory shear index (OSI) computed from CFD have shown consistent statistical associations with rupture status in large cohorts [9]. The main limitation of CFD is computational cost. A single pulsatile simulation typically requires several hours on a multi-core machine.

Physics-informed neural networks (PINNs) offer a mesh-free alternative [10]. Rather than discretizing the Navier-Stokes equations on a volumetric mesh, a PINN encodes the governing equations into its training loss and learns the flow solution at scattered collocation points. To date, however, few studies have systematically benchmarked PINNs against finite element (FEM) references for intracranial aneurysm applications specifically [11]. A hybrid approach that augments the physics constraints with sparse reference data could further improve accuracy, making PINNs a practical complement to CFD when limited reference measurements are available [12].

This thesis pursues two complementary objectives. The first one is clinical: identifying which hemodynamic indices, computed from patient-specific pulsatile simulations, separate growing intracranial aneurysms from stable aneurysms in a matched cohort. The second one is methodological: assessing the ability of a PINN to reproduce the hemodynamic fields obtained from a reference finite element (FEM) simulation on the same patient geometries, and quantifying the associated computational cost.

Both objectives are addressed on the same six patient-specific aneurysm geometries drawn from the Vascular Model Repository [13]: three growing and three stable cases arranged into three anatomically matched pairs. FEM simulations are run with svMultiPhysics and serve as the ground truth for both the clinical comparison and the PINN evaluation. PINNs are trained with the NVIDIA PhysicsNeMo framework [14]. Running both solvers on identical geometries and boundary conditions ensures output differences are attributable to the solver, not the geometry.

Chapter 1

State of the Art

1.1 Medical Context of Intracranial Aneurysms

1.1.1 Context and Epidemiology




Figure 1.1: Example of an intracranial aneurysm (also called cerebral aneurysm) [1].

Intracranial aneurysms (IAs) are pathological dilations of cerebral artery walls, most commonly forming at arterial bifurcations subject to elevated hemodynamic stress. Their prevalence in the general population is approximately 3.2%, based on a systematic review by Vlak et al. of 68 studies [6]. The majority of aneurysms remain asymptomatic and are discovered incidentally during imaging performed for other reasons, such as headache evaluation or post-traumatic screening.

The primary clinical concern is rupture. When an aneurysm ruptures, blood escapes into the space surrounding the brain, causing subarachnoid hemorrhage (SAH). SAH is a life-threatening emergency. The case fatality rate for aneurysmal SAH approaches 40%, and a significant fraction of survivors sustain permanent neurological deficits [5]. The overall incidence of SAH is around 9 per 100 000 person-years, with marked geographic variation ranging from about 4 per 100 000 in South and Central America to over 20 per

100 000 in Japan [15].

Known risk factors for IA formation, growth, and rupture include hypertension, cigarette smoking, family history of intracranial aneurysm, and being a female. Women represent roughly 61% of cases [6]. Certain genetic conditions, notably autosomal dominant polycystic kidney disease and connective tissue disorders, are associated with significantly elevated prevalence, pointing to a hereditary component in aneurysm susceptibility.

1.1.2 Pathogenesis and Morphological Classification

Aneurysm formation involves an interplay of hemodynamic loading, arterial wall degeneration, and inflammatory remodeling. Repetitive hemodynamic stress is thought to damage the endothelium and weaken the internal elastic lamina, leading to localized outpouching of the vessel wall [16]. Once an aneurysm forms, its continued growth and eventual rupture depend on the local wall mechanics and the hemodynamic environment within the sac.

The dominant type is the **saccular** (or berry) aneurysm, accounting for the vast majority of intracranial cases. It is characterized by a roughly spherical dome connected to the parent vessel by a neck, and typically arises at bifurcations. Less common forms, fusiform and dissecting aneurysms, also exist but are considerably rarer.




Figure 1.2: Main intracranial aneurysm morphologies

The most frequent locations are along the Circle of Willis and its branches: the internal carotid artery (ICA) and its bifurcation, the anterior communicating artery (ACoA), the posterior communicating artery (PCoA), and the middle cerebral artery (MCA) bifurcation [7, 17]. Together these sites account for the large majority of saccular aneurysms.




Figure 1.3: Cerebral vasculature and common aneurysm locations.

Geometrically, aneurysms are characterized by parameters such as maximum dome diameter, neck diameter, aspect ratio (dome height to neck width), and dome-to-neck ratio. Larger aneurysms are generally considered at higher rupture risk, but size alone is a poor predictor: rupture risk also depends strongly on aneurysm location and patient clinical characteristics [4, 20].

1.1.3 Detection and Imaging Modalities

Three imaging modalities are routinely used in the clinical workup of intracranial aneurysms:

Computed Tomography Angiography (CTA) is fast, widely available, and provides high-resolution three-dimensional reconstructions of the cerebrovascular anatomy. It is typically the first-line tool in acute settings, particularly for suspected SAH. Its limitations include radiation exposure and the use of iodinated contrast agents [21].

Magnetic Resonance Angiography (MRA), usually acquired with time-of-flight (TOF) sequences, is a non-invasive alternative without ionizing radiation. It is particularly useful for follow-up imaging of known aneurysms. It achieves high sensitivity but more variable specificity, with aneurysms most often missed near the skull base and the middle cerebral artery [22].

Digital subtraction angiography (DSA) remains the gold standard for high-resolution, dynamic arterial imaging. It provides the most detailed view of aneurysm morphology and the parent vessel but is invasive and carries a small procedural risk, so it is typically reserved for pre-treatment planning or inconclusive cases. Less invasive modalities such as CTA are compared against it as the reference standard [23].

Patient-specific CFD workflows rely on CTA or MRA images to reconstruct the vascular geometry used for simulation.

1.1.4 Clinical Decision-Making and Treatment Options

Managing an unruptured IA involves balancing the risk of rupture against the risk of intervention. Many small incidentally found aneurysms carry a low annual rupture risk (below 1%), and may be followed conservatively with periodic imaging rather than treated [20].

When intervention is indicated, two main approaches exist. **Surgical clipping** places a metal clip across the aneurysm neck to exclude the sac from the circulation [24]. **Endovascular coiling** fills the aneurysm with platinum coils via catheter, promoting thrombosis and progressive occlusion [24]. More recently, **flow diverters**, mesh stents placed in the parent artery, have been used to redirect flow away from the sac and induce gradual aneurysm occlusion without direct sac access [25].




Figure 1.4: Surgical and endovascular treatments for intracranial aneurysms. (a) Endovascular coiling of the aneurysm sac. (b) Surgical clipping of the aneurysm neck. (c) Endovascular treatment combining coils and a stent. (d) Endovascular treatment with a flow diverter [2].

Risk stratification tools such as the **PHASES score** provide an estimate of 5-year rupture risk to guide clinical decisions [4]. The acronym stands for Population origin, Hypertension, Age, aneurysm Size, Earlier SAH from a different aneurysm, and aneurysm Site. Each factor is assigned an integer point value, and the total score maps to a 5-year rupture risk ranging from below 1% for low scores to above 10% for the highest. These tools are practical and widely used, but they rely entirely on clinical and morphological variables, without incorporating patient-specific hemodynamics. Two aneurysms with identical PHASES scores may differ substantially in their intra-aneurysmal flow patterns and wall shear stress, information that is invisible to such scoring systems. Furthermore, treatment selection between clipping, coiling, and flow diversion is itself influenced by the local flow patterns: flow diverters, for instance, operate by modifying the inflow jet entering the sac. CFD-derived hemodynamic indices can complement these clinical scores by capturing patient-specific flow information that morphological predictors alone cannot provide, potentially improving both rupture risk assessment and treatment planning.

Table 1.1: **PHASES score**: factors, categories, and assigned points [4]. The total score maps to a 5-year rupture risk ranging from 0.4% (score ≤ 2) to 17.8% (score ≥ 12).

Factor	Category	Points
(P) Population	North American, European (other than Finnish)	0
	Japanese	3
	Finnish	5
(H) Hypertension	No	0
	Yes	1
(A) Age	< 70 years	0
	≥ 70 years	1
(S) Size	< 7.0 mm	0
	7.0–9.9 mm	3
	10.0–19.9 mm	6
	≥ 20.0 mm	10
(E) Earlier SAH	No	0
	Yes	1
(S) Site	ICA	0
	MCA	2
	ACoA / PCoA / posterior	4

1.2 Hemodynamics of Intracranial Aneurysms

1.2.1 Blood Flow Fundamentals

Accurate CFD simulation of intracranial aneurysms requires a precise characterization of blood as a fluid. Blood is a suspension of cellular elements, red blood cells, white blood cells, and platelets, carried in a liquid plasma phase [26]. These components interact with each other and with the vessel wall, giving blood mechanical properties that differ markedly from simple Newtonian fluids [27]. Understanding these properties is necessary to choose appropriate modeling assumptions and to interpret the limitations of the resulting simulations.

Blood is a non-Newtonian fluid whose apparent viscosity decreases with increasing shear rate (shear-thinning behaviour), described by models such as Carreau or Casson. However, in large vessels with diameters above approximately 1 mm, such as the cerebral arteries carrying intracranial aneurysms, the shear rates are high enough that viscosity approaches a nearly constant asymptotic value. For this reason, the Newtonian approximation is widely used in patient-specific CFD studies of intracranial aneurysms, allowing for simpler fluid modeling [8]. A sensitivity analysis by Cebal et al. found that the choice of viscosity model does not alter the characterization of intra-aneurysmal flow patterns, supporting this simplification for the large cerebral vessels considered here [28]. This assumption has nonetheless been questioned: a meta-analysis by Saqr et al. argued that the Newtonian model, used in nearly 90% of published studies, may be insufficient to fully capture

intracranial aneurysm hemodynamics [29].

Blood flow in cerebral arteries is pulsatile, driven by the cardiac cycle. The unsteady, oscillatory nature of this flow determines whether it remains laminar or transitions to turbulence, and how strongly the velocity profile deviates from a quasi-steady shape during the cycle. Two dimensionless numbers are used to characterize this regime and inform the modeling assumptions adopted in CFD simulations. The **Reynolds number**,

$$Re = \frac{\rho U D}{\mu}, \quad (1.1)$$

where ρ is the blood density, U a characteristic velocity, D the vessel diameter, and μ the dynamic viscosity, quantifies the ratio of inertial to viscous forces. In cerebral vessels, Re typically ranges from a few hundred in the parent artery to around 1000 at peak systole, placing the flow firmly in the laminar regime. The **Womersley number**,

$$Wo = \frac{D}{2} \sqrt{\frac{\omega \rho}{\mu}}, \quad (1.2)$$

where ω is the angular frequency of the cardiac cycle, describes the relative importance of oscillatory inertial effects versus viscous diffusion. For typical cerebral arteries ($D \approx 3$ -5 mm at 1 Hz heart rate), Wo is on the order of 3-5, meaning that the velocity profile departs significantly from a quasi-steady Poiseuille profile throughout the cardiac cycle.

Inside the aneurysm sac, the inflow jet from the parent artery drives a primary recirculating vortex whose structure depends strongly on the sac geometry, orientation, and inlet conditions. The flow may concentrate into a narrow jet impinging on the dome or spread more diffusely, and the resulting intrasaccular flow can range from simple, stable patterns to complex, unstable ones with one or more vortices [30].

1.2.2 Hemodynamic Parameters and Indices

The hemodynamic environment of an aneurysm is most directly characterized by the **wall shear stress** (WSS), the tangential force per unit area exerted by the flowing blood on the vessel wall:

$$\boldsymbol{\tau}_w = \mu \left. \frac{\partial \mathbf{u}}{\partial n} \right|_{\text{wall}}, \quad (1.3)$$

where $\mathbf{u}$ is the velocity field and n is the coordinate normal to the wall pointing inward. Because WSS is pulsatile, several time-integrated scalar indices are used in practice to summarize the near-wall hemodynamic conditions over a full cardiac cycle of period T .

Time-Averaged Wall Shear Stress (TAWSS) is the time average of the WSS magnitude:

$$\text{TAWSS} = \frac{1}{T} \int_0^T |\boldsymbol{\tau}_w| \, dt. \quad (1.4)$$

In healthy cerebral arteries, TAWSS is typically on the order of a few Pa. Values below 1 Pa in the aneurysm dome are consistently associated with increased rupture risk in the literature [9, 31].

Oscillatory Shear Index (OSI) quantifies temporal directional variation in the WSS vector:

$$\text{OSI} = \frac{1}{2} \left(1 - \frac{\left| \int_0^T \boldsymbol{\tau}_w dt \right|}{\int_0^T |\boldsymbol{\tau}_w| dt} \right). \quad (1.5)$$

OSI ranges from 0 (purely unidirectional shear) to 0.5 (purely oscillatory, zero net shear). High OSI (> 0.2) reflects unstable, disturbed near-wall flow and is associated with endothelial activation and wall remodeling [32, 33]. Elevated OSI at the aneurysm dome has been linked to sites of subsequent growth and to increased rupture probability in longitudinal studies [34].

Relative Residence Time (RRT) combines TAWSS and OSI to estimate how long blood particles spend near the wall in low-shear, oscillatory conditions:

$$\text{RRT} = \frac{1}{(1 - 2 \cdot \text{OSI}) \cdot \text{TAWSS}}. \quad (1.6)$$

High RRT corresponds to slow, oscillatory near-wall flow, which promotes platelet accumulation and inflammatory cell adhesion. Regions of high RRT are considered at elevated risk for wall degeneration and aneurysm rupture [33, 35].

Endothelial Cell Activation Potential (ECAP) is defined as:

$$\text{ECAP} = \frac{\text{OSI}}{\text{TAWSS}}, \quad (1.7)$$

and serves as a combined indicator of hemodynamic conditions likely to activate the endothelium and promote wall dysfunction. High ECAP values mark regions where both oscillatory and low-shear conditions coincide, a pattern associated with aneurysm wall instability [33, 36].

Low Shear Area (LSA) is the fraction of the aneurysm dome surface on which TAWSS falls below a threshold τ_{th} :

$$\text{LSA} = \frac{\text{area}_{\text{dome}}(\text{TAWSS} < \tau_{\text{th}})}{\text{area}_{\text{dome}}}, \quad (1.8)$$

with typical thresholds in the range 0.1–1 Pa. A large LSA indicates that a substantial portion of the dome surface is exposed to chronically low shear, a configuration consistently associated with wall thinning and elevated rupture risk [9, 31]. Together, TAWSS, OSI, RRT, ECAP, and LSA form the standard hemodynamic fingerprint computed from a CFD simulation and are the primary outputs analyzed in this work.

Five further indices are reported in the recent CFD literature on intracranial aneurysms and complement the standard fingerprint above. The **peak-systolic wall shear stress** is the instantaneous WSS magnitude at the systolic frame, $\text{PSWSS} = \max_t |\boldsymbol{\tau}_w(t)|$. It captures the maximal mechanical load on the wall during the cycle and is therefore a candidate marker for sites where the wall is intermittently overloaded rather than chronically understimulated [37]. The **systolic-to-diastolic WSS range** is the peak-to-trough amplitude of the WSS magnitude over the cycle, $\Delta\text{WSS} = \max_t |\boldsymbol{\tau}_w(t)| - \min_t |\boldsymbol{\tau}_w(t)|$, and quantifies the pulsatile component of the loading. The $\text{PSWSS}/\text{TAWSS}$ **peaking factor**

measures how concentrated the shear loading is around systole. Values close to unity correspond to nearly steady shear, whereas larger values mark cycles with a pronounced systolic spike. The **wall-shear-stress gradient magnitude**, $WSSG = \|\nabla_s TAWSS\|$, is the surface gradient of the TAWSS field. It quantifies the spatial sharpness of low-shear or high-shear patches on the sac and has been linked to mechanotransduction disturbances at the endothelium [33]. Finally, because the LSA depends on the choice of the low-shear cut-off τ_{th} , recent studies report it at several thresholds in the literature range. In this work LSA is reported at 0.4 Pa, 1.0 Pa and 1.5 Pa to make the threshold sensitivity explicit [9, 31].

For each of the indices above, the area-weighted average over the aneurysm sac is sensitive to the absolute inlet flow rate, which is itself uncertain because patient-specific inlet waveforms are rarely available [38]. To partly absorb this confound, **dome-to-parent ratios** of the form $r_X = \bar{X}_{dome}/\bar{X}_{parent}$ are also reported. The parent vessel acts as an internal reference of the upstream flow conditions, so r_X characterises the relative shielding (or amplification) of index X inside the dome compared with the immediate upstream segment.




Figure 1.5: Illustrative TAWSS distribution on an aneurysm surface: low WSS (blue) concentrates in the dome while elevated WSS (beige-orange) appears at the inflow neck. This pattern is consistently associated with increased rupture risk.

1.2.3 Hemodynamic-Based Rupture Risk Assessment

The connection between intracranial aneurysm hemodynamics and rupture has been an active research topic since Cebal et al. demonstrated in 2005 that patient-specific CFD flow characteristics could distinguish ruptured from unruptured aneurysms in a clinical cohort [7]. Subsequent work established quantitative hemodynamic differences: ruptured aneurysms tend to exhibit more complex and unstable flow patterns, with stronger and more concentrated inflow jets impinging on the sac [37]. The direction of the wall shear stress association, by contrast, remains debated, as discussed below.

Two competing hypotheses have been proposed regarding the direction of hemodynamic risk. The **low WSS hypothesis** holds that chronically low wall shear stress impairs endothelial function, triggers inflammatory remodeling, and progressively weakens the

aneurysm wall until rupture. A meta-analysis by Zhou et al. of multiple CFD studies found low TAWSS to be consistently and significantly associated with rupture status across the pooled population [9]. Suzuki et al. corroborated this at the sub-aneurysm level, identifying the exact rupture point in a series of ruptured cases and showing that those locations had markedly lower WSS than the surrounding dome [31]. The underlying cellular mechanism has been investigated by Wei et al., who found that low shear stress activates macrophage infiltration via the CCL7/CCR1/TAK1/NF- κ B signaling axis, leading to extracellular matrix degradation and wall thinning [39]. The alternative **high WSS hypothesis** attributes bleb formation and local wall thinning to concentrated high-shear regions. It finds support in studies such as Cebal et al., who reported elevated maximum WSS and concentrated inflow streams in aneurysms with a history of rupture [37]. Across pooled cohorts, however, the low-WSS association has been the more consistently reproduced of the two [9].

Beyond TAWSS, oscillatory indices such as OSI and RRT are considered complementary markers. High-OSI regions co-localize with flow stagnation and recirculation zones, which favor thrombus formation and sustained inflammatory cell exposure. Detmer et al. analyzed nearly 2000 aneurysms and found that, while the adverse profile of ruptured aneurysms (faster, more oscillatory, and more complex flow) is broadly consistent across locations, rupture rates themselves vary markedly with aneurysm site, patient age and sex, and bifurcation versus sidewall configuration [40]. This underscores that hemodynamic markers must be interpreted alongside anatomical and clinical context rather than in isolation.

The translation of these hemodynamic markers to clinical practice is limited by several factors. CFD-derived indices show population-level statistical associations with rupture status but carry substantial inter-patient variability. Steinman and Pereira have provided a careful analysis of the principal sources of uncertainty in patient-specific hemodynamic modeling, including image segmentation, inlet waveform prescription, and mesh resolution, and noted that these can individually produce significant variability in computed WSS values [38]. These considerations motivate both improved modeling standards and access to well-controlled, longitudinal datasets.

1.2.4 Growing vs. Stable Aneurysms: Hemodynamic Signatures

While most published CFD studies compare ruptured with unruptured aneurysms, a smaller body of work focuses on aneurysm growth as an intermediate endpoint. Longitudinal enlargement is considered a marker of hemodynamic instability and a risk factor for eventual rupture, making it an important target for surveillance and intervention decisions.

Dabagh et al. analyzed the hemodynamic and morphological evolution of a growing cerebral aneurysm over time and found that areas of low WSS and high OSI tended to spatially coincide with regions of subsequent dome enlargement [41]. This spatial correlation suggests that hemodynamic loading drives wall remodeling locally rather than uniformly across the sac.

Karnam et al. examined a larger series and identified multiple competing hemodynamic growth pathways: in some cases, localized growth was driven by high-velocity inflow jets producing elevated WSS at the neck, while in others, dome expansion occurred in low-WSS stagnation zones [42]. This heterogeneity indicates that aneurysm growth cannot

be attributed to a single dominant hemodynamic mechanism.

Nishiwaki et al. specifically compared hemodynamics in growing aneurysms that subsequently ruptured against those that remained stable after growth. For aneurysms with homothetic (diffuse) growth patterns, maximum OSI after enlargement was significantly higher in ruptured cases than in stable ones (0.455 vs. 0.319, $p = 0.003$), while TAWSS differences were not statistically significant. A newly emerged focal high-OSI area after growth consistently corresponded to the rupture point [34], suggesting that the appearance of concentrated oscillatory flow, rather than global WSS reduction, may be the key hemodynamic signature discriminating rupture-prone from stable growing aneurysms.

1.3 Computational Fluid Dynamics for Cerebrovascular Research

1.3.1 Governing Equations

Patient-specific CFD simulations of intracranial aneurysms are built on the incompressible Navier-Stokes equations. For a Newtonian fluid with constant density ρ and dynamic viscosity μ :

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u}, \quad (1.9)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (1.10)$$

where $\mathbf{u}$ is the velocity vector field, p is the pressure, and body forces are neglected. Equation (1.10) enforces mass conservation (incompressibility).

Three standard modeling assumptions are adopted in most patient-specific studies of intracranial aneurysms:

- **Rigid walls:** The vessel wall is treated as non-deforming. This eliminates the need for fluid-structure interaction (FSI) coupling and reduces computational cost substantially. Wall compliance alters near-wall hemodynamics, with the largest effects on near-dome recirculation zones. A fluid-structure interaction study of 101 sidewall aneurysms reported sac-averaged OSI changes ranging from -36% to $+674\%$ relative to the rigid-wall value [43].
- **Newtonian fluid:** Dynamic viscosity μ is taken as constant. This approximation is supported by sensitivity analyses showing that the viscosity model does not alter the characterization of intra-aneurysmal flow patterns [28], with any residual non-Newtonian effects remaining small relative to the dominant sources of uncertainty in patient-specific models, such as boundary conditions and wall compliance [38]. This simplification has nonetheless been questioned for intracranial aneurysm flows [29].
- **No-slip boundary condition:** The fluid velocity at the vessel wall is set to zero.

Standard blood properties used in the literature are $\rho = 1060 \text{ kg/m}^3$ and $\mu = 0.004 \text{ Pa}\cdot\text{s}$ [26, 28].

1.3.2 Patient-Specific CFD Simulation Pipeline

Patient-specific cerebrovascular CFD follows a standardized multi-step pipeline [3, 7]:

1. **Image acquisition:** A CTA or MRA scan provides the three-dimensional anatomical data.
2. **Segmentation and geometric modeling:** The arterial lumen is segmented from the images to obtain a surface model of the vasculature, including the aneurysm sac and a sufficient length of the parent and daughter vessels to allow proper flow development.
3. **Mesh generation:** The enclosed volume is discretized into a tetrahedral mesh. Accurate WSS computation requires a fine boundary layer mesh at the vessel wall. Typical mesh sizes range from 10^5 to 10^6 elements per model. Mesh convergence, verifying that hemodynamic indices are independent of resolution, is an essential quality check [38].
4. **Boundary condition specification:** A pulsatile velocity waveform is imposed at the inlet, either derived from patient measurements or from a population-averaged waveform scaled to the local vessel diameter. A zero-pressure or resistance condition is applied at each outlet.
5. **Transient simulation:** The Navier-Stokes equations are solved over multiple cardiac cycles until the flow reaches a periodic steady state. At least two to three cycles are needed to eliminate initial transients.
6. **Post-processing:** Velocity and pressure fields from the converged solution are used to compute hemodynamic indices (TAWSS, OSI, RRT, ECAP, LSA) over the aneurysm surface.

The choice of boundary conditions is one of the most influential parameters: the inlet waveform directly controls computed WSS levels, and patient-specific flow measurements are rarely available in clinical practice. This introduces a systematic source of uncertainty that must be acknowledged when comparing results across studies [38].

1.3.3 SimVascular




Figure 1.6: SimVascular GUI [3].

SimVascular is an open-source software platform designed specifically for the patient-specific cardiovascular simulation workflow [3]. It integrates all steps of the pipeline into a single graphical environment:

- **Path planning:** Interactive definition of vessel centerlines through the anatomy of interest.
- **2D segmentation:** Delineation of the lumen contour at cross-sections along each centerline, using spline-based or level-set contours fitted to image intensity gradients.
- **3D model construction:** Lofting the 2D contours into a solid surface model, with boolean union operations to merge vessel branches.
- **Meshing:** Tetrahedral mesh generation with boundary layer refinement, supported by TetGen and MeshSim backends.
- **CFD solver - svMultiPhysics:** A finite element solver for the incompressible Navier-Stokes equations using a stabilized variational multiscale (VMS) formulation, which provides good accuracy on unstructured tetrahedral meshes.
- **Post-processing:** Simulation results are stored in VTK format (VTP for surface data, VTU for volume data) and can be visualized directly in SimVascular or in external tools such as ParaView.

The Vascular Model Repository (VMR), developed and maintained by the SimVascular team, provides a publicly accessible collection of patient-specific cardiovascular models [13]. It covers a range of anatomical regions and pathologies, including intracranial aneurysms, and is a community resource for benchmarking and validation.

1.3.4 Other CFD Tools and Limitations

General-purpose CFD solvers including ANSYS Fluent, OpenFOAM, and COMSOL Multiphysics have also been applied to vascular hemodynamics. These platforms offer greater flexibility in numerical schemes and mesh types, but require significantly more manual configuration to set up a patient-specific vascular workflow. A multi-group challenge found that pressure predictions were remarkably consistent across a wide range of solvers and solution strategies for identical geometries and boundary conditions, although intra-sac flow patterns and WSS-derived quantities were considerably more variable [44].

Traditional mesh-based CFD has several limitations that constrain its use at scale and motivate alternative approaches. First, the full pipeline from image to result involves multiple manual steps (segmentation, model cleanup, mesh generation, boundary condition tuning), each of which introduces operator-dependent variability and is time-consuming. Second, a single simulation can take hours to days on a workstation, limiting throughput for large cohort studies. Third, the rigid wall assumption neglects wall compliance and the mechanical coupling between blood pressure and wall deformation, FSI approaches address this but at substantially higher computational cost [43]. Fourth, as discussed above, boundary conditions remain a major source of uncertainty [38]. These limitations motivate the exploration of physics-informed neural network approaches as a complementary simulation strategy.

1.4 Physics-Informed Neural Networks

1.4.1 Fundamentals of PINNs

Physics-informed neural networks (PINNs) are a class of machine learning methods that embed physical laws, expressed as partial differential equations (PDEs), directly into the training objective of a neural network [10, 11, 45]. This framework was formalized by Raissi, Perdikaris, and Karniadakis in 2019, building on earlier work from the 1990s on neural network approximation of PDEs.

The central idea is straightforward. A neural network $\mathcal{N}(\mathbf{x}; \boldsymbol{\theta})$ with trainable parameters $\boldsymbol{\theta}$ is used to approximate the solution of a PDE over a spatial (and possibly temporal) domain. Rather than discretizing the domain on a mesh, the network is evaluated at a set of *collocation points* sampled from the domain interior and boundaries. Automatic differentiation computes the spatial derivatives of the network output, and the PDE residuals at each collocation point are included as penalty terms in the total training loss:

$$\mathcal{L}(\boldsymbol{\theta}) = \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}} + \lambda_{\text{data}} \mathcal{L}_{\text{data}}, \quad (1.11)$$

where $\mathcal{L}_{\text{PDE}}$ penalizes PDE residuals at interior collocation points, $\mathcal{L}_{\text{BC}}$ enforces boundary and initial conditions, and $\mathcal{L}_{\text{data}}$ optionally fits the network to available measurement data. The weights λ balance the contribution of each term. The network is trained by minimizing $\mathcal{L}$ over $\boldsymbol{\theta}$ using gradient descent.

For the incompressible Navier-Stokes equations in the steady-state case, the PDE residuals

at a collocation point $\mathbf{x} = (x, y, z)$ are:

$$r_{\text{momentum}} = \rho(\mathbf{u} \cdot \nabla \mathbf{u}) + \nabla p - \mu \nabla^2 \mathbf{u}, \quad (1.12)$$

$$r_{\text{continuity}} = \nabla \cdot \mathbf{u}, \quad (1.13)$$

and training minimizes $\|r_{\text{momentum}}\|^2 + \|r_{\text{continuity}}\|^2$ over the sampled points.

The standard network architecture for PINN flow problems is a fully connected multilayer perceptron (MLP). For a three-dimensional steady flow, the input is $\mathbf{x} = (x, y, z)$ and the output is the flow state (u, v, w, p) , where (u, v, w) are the velocity components and p is the pressure. Network depth and width are hyperparameters that control expressiveness and training cost.




Figure 1.7: Example of a PINN architecture.

Several properties distinguish PINNs from classical CFD methods: (i) they are **mesh-free**, requiring only sampled points rather than a structured or unstructured grid. (ii) they produce a **continuous** solution representation that can be evaluated at any point in the domain. (iii) they are **differentiable** with respect to inputs, allowing direct computation of velocity gradients and thus wall shear stress without finite differencing and (iv) they can incorporate sparse measurement data as additional loss terms, making them suitable for data assimilation problems.

1.4.2 PINNs for Fluid Dynamics

Early PINN demonstrations for fluid dynamics used simple canonical flows (Poiseuille flow, lid-driven cavity, Kovasznay flow) to validate the approach against known solutions. Jin et al. extended this to the full incompressible Navier-Stokes system in their NSFnets framework, showing that PINNs could reproduce benchmark flows at moderate Reynolds numbers with reasonable accuracy [46]. A key practical finding was that careful loss weighting and network initialization are essential for convergence, since the PDE loss landscape tends to be ill-conditioned when momentum and continuity residuals are combined naively.

Two adaptive strategies have become standard. **Learning Rate Annealing** rescales each loss weight at regular intervals so that every constraint contributes gradients of comparable magnitude [47]. **ReLoBRaLo** (Relative Loss Balancing with Random Lookback) tracks both long-term and short-term ratios of each loss component and adjusts weights accordingly, with a stochastic reset mechanism to avoid ratio saturation [48]. Both strategies are

designed to prevent any single constraint from dominating the parameter updates during training.

In the cardiovascular context, Kissas et al. demonstrated a particularly relevant application: they trained a PINN on noisy in-vivo flow and wall-displacement measurements near the aorta-carotid bifurcation and used 1D pulsatile-flow conservation laws to infer the unmeasured pressure field [12]. This approach effectively enriches sparse clinical imaging data with physically consistent, complete hemodynamic fields, a task that is impossible with standalone CFD, which requires full boundary conditions to be prescribed a priori.

Applying PINNs to patient-specific aneurysm geometries introduces specific challenges. The geometry is complex and non-convex, which complicates collocation point sampling and makes it harder for the network to resolve fine near-wall features such as the boundary layer responsible for WSS. The no-slip condition must be enforced on an arbitrary curved surface defined by an STL mesh. Coordinate normalization, scaling the domain to a unit cube, is important for numerical stability and balanced gradient magnitudes. Integral mass flux constraints at the inlet and outlet cross-sections are an effective strategy to enforce physically consistent flow rates and to stabilize training, since they provide global conservation information that point-wise residuals alone may not capture reliably.

1.4.3 NVIDIA PhysicsNeMo

PhysicsNeMo (formerly NVIDIA Modulus) is an open-source Python framework developed by NVIDIA for physics-informed and physics-constrained machine learning [14]. It provides a high-level interface for defining PDE systems symbolically, specifying geometries, and training PINN models on GPU hardware.

In PhysicsNeMo, governing equations are specified symbolically using SymPy expressions. The framework automatically differentiates the network outputs to compute PDE residuals at runtime. Geometry is defined via STL files that describe the domain boundaries: for an aneurysm, separate STL surfaces are provided for the inlet, outlet(s), the no-slip vessel wall, and a closed surface enclosing the interior volume. PhysicsNeMo samples collocation points on each boundary type and in the interior according to user-specified densities, and assembles the corresponding loss terms automatically.

Training uses gradient descent (typically Adam followed by L-BFGS) on GPU hardware. PhysicsNeMo supports multi-GPU training and is designed to integrate with SLURM-based high-performance computing (HPC) clusters, which is important given the scale of training runs required for realistic three-dimensional geometries. The framework also provides checkpointing, loss monitoring, and model export utilities.

A reference aneurysm tutorial is included in the PhysicsNeMo documentation and provides a complete working example of training a Navier-Stokes PINN on a patient-specific cerebral aneurysm geometry [14]. This tutorial uses a fully connected MLP with input (x, y, z) and output (u, v, w, p) , trained on five STL-defined surfaces (inlet, outlet, no-slip wall, closed interior, and an integral constraint surface), and forms the starting point for the implementation carried out in this work.

1.4.4 PINNs vs. Traditional CFD

PINNs and classical finite element or finite volume solvers approach the same problem through fundamentally different paradigms. Understanding their respective strengths and limitations is important for interpreting the results of this work, where both approaches are applied to the same aneurysm geometries.

Accuracy and convergence. Classical FEM solvers have well-established convergence theory and achieve high accuracy with mesh refinement. PINNs lack comparable guarantees: training can stall in local minima, and accuracy depends on architecture choices, collocation point density, and training duration. For complex three-dimensional flows at realistic Reynolds numbers, PINNs currently struggle to match the pointwise accuracy of well-converged FEM solutions [11]. The near-wall gradient resolution needed for reliable WSS computation is particularly challenging.

Geometric workflow. FEM requires mesh generation, which is a time-consuming and often manual step. PINNs replace this with point sampling from an STL surface, which can be automated more easily. On the other hand, PINN training is itself a computationally intensive process that requires GPU hardware and careful hyperparameter tuning.

Data assimilation. A key advantage of PINNs is their native ability to incorporate sparse measurement data (e.g., 4D flow MRI velocity fields) as additional loss terms, yielding a physically consistent flow reconstruction that seamlessly integrates simulation and measurement [12]. This capability has no direct equivalent in classical CFD.

Inference speed. Once a PINN is trained, evaluating the flow field at arbitrary points is essentially instantaneous, without re-running a simulation. This makes trained models potential surrogates for rapid parametric studies or real-time hemodynamic assessment, provided the geometry and boundary conditions are consistent with the training configuration.

Current limitations. The main barriers to practical use of PINNs for patient-specific aneurysm hemodynamics are: (i) difficulty capturing the fine boundary layer structure needed for accurate WSS, particularly for higher Reynolds number flows. (ii) sensitivity to hyperparameters including loss weights, learning rate schedule, and network depth. (iii) absence of validated open-source pipelines for complex patient-specific vascular geometries comparable in maturity to SimVascular and (iv) limited comparative validation against high-fidelity FEM references for intracranial aneurysm applications specifically. Addressing these limitations is part of the broader research effort in which this work participates.

Chapter 2

Materials and Methods

This chapter describes the materials and methods used to address the two objectives of this thesis: (i) identifying which hemodynamic indices, computed from patient-specific simulations, separate growing aneurysms from stable aneurysms in a matched cohort, and (ii) assessing the ability of a physics-informed neural network (PINN) to reproduce the hemodynamic fields obtained from a reference finite element (FEM) simulation on the same patient geometries, together with the associated computational cost.

The two objectives are addressed jointly. The FEM pipeline (Section 2.2) produces the reference fields used both for the clinical comparison and for the validation of the PINN. The PINN pipeline (Section 2.3) is built on the same six VMR surface models, so that any difference at the output is attributable to the solver and not to the geometry. The hemodynamic post-processing common to both pipelines is described in Section 2.5. Section 2.4 closes the chapter with a concise list of the experiments that produce the results reported in Chapter 3.

2.1 Dataset: Growing versus Stable Aneurysms

Six patient-specific intracranial aneurysm geometries were drawn from the Vascular Model Repository (VMR) [13], a publicly accessible collection of cardiovascular models maintained by the SimVascular team. The VMR includes a subset of cases with longitudinal follow-up imaging that classifies each aneurysm as *growing* or *stable* on the basis of serial measurements. Three growing and three stable cases were selected, arranged into three matched pairs sharing the same anatomical location. Pairing by site controls for the distinct baseline flow environments of the cavernous internal carotid artery (ICA) and the ICA terminus, so within-pair differences are more attributable to growth status than to anatomy.

The six cases are summarised in Table 2.1.

For each case, the VMR provides a pre-built surface model of the vascular lumen derived from computed tomography angiography (CTA), supplied as a triangulated VTK surface mesh covering the aneurysm sac and a sufficient length of the parent and daughter vessels. No additional segmentation was performed. The VMR surfaces served directly as the geometric input to both simulation pipelines. The volumetric meshes generated from these

Table 2.1: Dataset overview with clinical metadata. Each pair consists of one growing and one stable aneurysm at the same anatomical site. Outlets: number of outflow boundaries. HTN: hypertension, Hx: family history of aneurysm, M: male, F: female, Y: yes, N: no.

Pair	Case	Status	Location	Out.	Sex	Age	HTN	Smk.	Hx
1	0203	Growing	Cavernous ICA	1	F	65	Y	N	N
	0204	Stable	Cavernous ICA	1	F	65	N	N	N
2	0207	Growing	Left ICA Terminus	2	F	61	N	N	N
	0208	Stable	ICA Terminus	2	F	71	N	N	N
3	0209	Growing	Left Cavernous ICA	1	F	48	N	N	Y
	0210	Stable	Cavernous ICA	1	M	50	N	Y	Y

surfaces are shown in Figure 2.1.

2.2 SimVascular CFD Simulation Pipeline

This thesis compares two simulation approaches applied to the same six patient geometries: a finite element method (FEM) solver and a physics-informed neural network (PINN). Running both on identical geometries and boundary conditions ensures that any difference in output is attributable to the solver and not to the geometry. The FEM pipeline described in this section serves two purposes. It produces the reference hemodynamic fields against which the PINN is evaluated (Section 2.3), and it provides the ground-truth indices for the growing-versus-stable hemodynamic comparison of Experiment E6 (Section 2.4.6).

Patient-specific CFD simulations were run for all six cases using the SimVascular platform [3]. The pipeline below follows the standard image-to-WSS workflow introduced in Section 1.3.2.

2.2.1 Model Preparation


The VMR surface models were imported directly into SimVascular [3], the pre-existing segmentation required no modification. Each surface was inspected to verify geometric integrity (absence of self-intersections, correct identification of inlet and outlet faces), and the enclosed fluid domain was extracted before meshing.

2.2.2 Meshing

The volumetric tetrahedral meshes were not generated from scratch for this work. For each case, the mesh distributed with the corresponding SimVascular model in the Vascular Model Repository was reused directly [13]. These meshes were produced by the SimVascular team with the TetGen-based meshing module and include prismatic boundary-layer elements at the vessel wall to resolve the near-wall velocity gradients needed for WSS computation. They are the same meshes used for the steady-state reference simulations shipped with each VMR case (`mesh-complete.mesh.vtu`). Reusing the published meshes keeps the




(a) 0203, Growing, Cavernous ICA



(b) 0204, Stable, Cavernous ICA



(c) 0207, Growing, ICA Terminus



(d) 0208, Stable, ICA Terminus



(e) 0209, Growing, Cavernous ICA



(f) 0210, Stable, Cavernous ICA

Figure 2.1: Meshes for all six cases. Each row shows one matched pair (growing left, stable right). Aneurysm sacs are visible as bulging regions off the parent vessel. Mesh density statistics are given in Table 2.2.

geometry identical to the reference dataset and removes one source of operator-dependent variability, at the cost of not having performed an independent mesh-convergence study. The node and element counts for each case are given in Table 2.2.

Table 2.2: Mesh statistics of the reused VMR meshes. BL: prismatic boundary-layer sublayers.

Case	Nodes	Tetrahedral elements	BL layers
0203	4 972	25 201	3
0204	23 391	128 688	3
0207	13 873	73 833	3
0208	9 570	50 363	3
0209	8 782	36 996	3
0210	5 468	27 678	3

2.2.3 Boundary Conditions

Blood was modeled as an incompressible Newtonian fluid with density $\rho = 1060 \text{ kg/m}^3$ and dynamic viscosity $\mu = 0.004 \text{ Pa s}$, in line with the standard physiological values adopted in patient-specific intracranial aneurysm CFD [8, 28]. Three boundary conditions were applied:

- **Inlet.** A pulsatile parabolic velocity profile was prescribed at the proximal ICA inlet. The waveform shape is a population-averaged internal carotid artery flow rate waveform reconstructed from 10 Fourier modes with period $T = 0.95 \text{ s}$ and mean volumetric flow rate $Q_{\text{mean}} = 4.084 \text{ cm}^3/\text{s}$, derived from case 0207 of the VMR dataset [13]. All six cases share the same waveform shape. The peak inlet velocity U_{max} for each case is set so that the mean parabolic flow rate matches Q_{mean} : for a circular inlet of area A_{in} , $U_{\text{max}} = 2 Q_{\text{mean}}/A_{\text{in}}$. The waveform is shown in Figure 2.2. It is strictly unidirectional: the volumetric flow rate stays positive throughout the cardiac cycle and never reverses sign, so the bulk inflow does not reverse direction at any phase.
- **Outlets.** A zero-pressure (traction-free) condition was imposed at each outlet face. This is standard practice for isolated aneurysm models where patient-specific downstream resistance data are unavailable [38].
- **Vessel wall.** A rigid no-slip condition ($\mathbf{u} = \mathbf{0}$) was applied at the luminal surface. Wall compliance and fluid-structure interaction were not modeled, consistent with common practice in patient-specific aneurysm CFD [8].

2.2.4 Solver Configuration

Simulations were run with svMultiPhysics, the finite element solver embedded in SimVascular. It uses a stabilized variational multiscale (VMS) formulation for the incompressible




Figure 2.2: Population-averaged internal carotid artery inflow waveform used as the inlet boundary condition for all six cases (VMR case 0207, period $T = 0.95$ s, mean $Q_{\text{mean}} = 4.084$ cm³/s). Red markers: the tabulated VMR samples. Solid line: the 10-mode Fourier reconstruction. The dashed line marks the cycle mean.

Navier-Stokes equations on unstructured tetrahedral meshes [3]. Each simulation ran for four cardiac cycles at a time step of $\Delta t = 5 \times 10^{-4}$ s (0.5 ms), with the first cycle discarded as transient initialization and the last cycle retained for analysis. All runs were executed on the UCLouvain HPC cluster (CISM) as MPI-parallel jobs using four cores. Wall-time results for each case are reported in Table 3.1 in Chapter 3.

2.2.5 Post-Processing and Parameter Extraction

The svMultiPhysics solver writes the velocity and pressure fields every 19 solver steps to VTK surface files (VTP format), producing 100 stored frames per cardiac cycle (400 files in total over the four simulated cycles. The last 100, steps 5719–7600, form the analysis window). These files were loaded with a custom Python post-processing script built on PyVista. For each case, the simulation surface was split into an *aneurysm dome* and a *parent vessel* region by reusing the region labels distributed with each SimVascular case in the VMR dataset. These labels originate from a separate steady-state rigid-wall simulation included with the dataset. The dome and parent-vessel point coordinates were extracted from those labels and transferred to the pulsatile wall mesh by nearest-neighbour lookup. The reference surface is geometrically identical to the pulsatile wall in all six cases, with a maximum nearest-neighbour distance below 0.02 cm, so the mapping is unambiguous. Hemodynamic indices (TAWSS, OSI, RRT, ECAP, LSA, WSSG, PSWSS, Δ WSS and the PSWSS/TAWSS peaking factor) were computed as area-weighted surface averages over each region. Their definitions and operational expressions are given in Section 2.5.1. The WSS vector $\boldsymbol{\tau}_w$ at each surface node was obtained from the FEM velocity gradient projected onto the local surface tangent plane. Dome-to-parent ratios were computed from the per-region averages as described in Section 2.5.2.

2.3 PINN-Based Simulation with PhysicsNeMo

A PINN was trained for each aneurysm geometry using NVIDIA PhysicsNeMo [14]. Unlike a mesh-based FEM solver, a PINN encodes the solution as a neural network whose parameters are optimised by minimising the residuals of the governing equations at scattered collocation points. No volumetric mesh is required, which makes the approach naturally adaptable to the complex surface geometries produced by medical image segmentation. Once trained, the network evaluates the flow field at any point in space and time at the cost of a single forward pass.

Two difficulties drive the rest of this section. The first one is enforcing global mass conservation across the domain. The second one is resolving the steep near-wall velocity gradients that govern the accuracy of the wall shear stress. The architecture, the loss decomposition, the adaptive loss-weighting strategies, and the WSS extraction procedures described below each target one of these two difficulties. The corresponding design choices are evaluated experimentally in Section 2.4.

2.3.1 Mesh Processing and STL Generation

PhysicsNeMo represents the fluid domain through triangulated surfaces in STL format. Each VMR surface model was converted into five STL files using a custom PyVista script (`generate_aneurysm_stls.py`):

1. `*_inlet.stl` - planar cap disk closing the proximal inlet, identified as the boundary loop with the largest enclosed area.
2. `*_outlet_XX.stl` - one numbered planar cap disk per outlet (e.g. `_01`, `_02`), automatically identified from the remaining boundary loops after inlet detection.
3. `*_noslip.stl` - the open vessel wall surface.
4. `*_integral.stl` - an internal cross-sectional disk placed at distance $d = 1.0 D_{\text{in}}$ from the inlet centroid along the vessel axis (D_{in} : inlet diameter), used as a flux monitoring surface during training.
5. `*_closed.stl` - the watertight surface (wall + both caps), used for sampling interior collocation points.

For open mesh inputs, caps were generated by detecting boundary edge loops and fitting triangulated planar disks. A companion script (`inspect_stl_scales.py`) reports the bounding-box extent and inlet geometry of each STL set, allowing the physical parameters (`nu`, `inlet_vel`) to be calibrated consistently across geometries before training.

2.3.2 Problem Formulation

The PINN targets the incompressible Navier-Stokes equations. In the pulsatile regime used for all production runs the residual form is

$$\rho(\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}) + \nabla p - \mu \nabla^2 \mathbf{u} = \mathbf{0}, \quad (2.1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.2)$$

and the network maps space-time coordinates (x, y, z, t) to the four flow variables (u, v, w, p) . The preliminary architecture ablation of Experiment E1 (Section 2.4.1) drops the time input and the time derivative, so the network reduces to a map $(x, y, z) \mapsto (u, v, w, p)$ and the residual reduces to the steady form. The spatial coordinates are taken directly from the VMR surface model without additional scaling. Blood is modelled with the same physiological properties used for the FEM reference (Section 2.2.3): density $\rho = 1060 \text{ kg/m}^3$ and dynamic viscosity $\mu = 4 \times 10^{-3} \text{ Pa} \cdot \text{s}$, equivalently a kinematic viscosity $\nu = \mu/\rho \approx 3.77 \times 10^{-6} \text{ m}^2/\text{s}$. The per-case peak parabolic inlet velocity at the cycle-mean flow rate, $U_{\max} = 2Q_{\text{mean}}/A_{\text{in}}$ with $Q_{\text{mean}} = 4.084 \text{ cm}^3/\text{s}$, ranges from 0.55 m/s (case 0203, $A_{\text{in}} = 0.148 \text{ cm}^2$, $D_{\text{in}} \approx 4.3 \text{ mm}$) to 1.38 m/s (case 0208, $D_{\text{in}} \approx 2.7 \text{ mm}$). The corresponding bulk Reynolds numbers $\text{Re} = \bar{U}_{\text{in}} D_{\text{in}}/\nu$, based on the cross-sectional mean velocity $\bar{U}_{\text{in}} = Q_{\text{mean}}/A_{\text{in}} = U_{\max}/2$, span roughly 300 to 500 across the six geometries, consistent with the physiological range for cerebral arteries, where Re is typically of the order of a few hundred [8]. The flow remains nominally laminar but, unlike the steady ablation of Section 3.2.1, is time-dependent and reaches instantaneous peak-systolic Reynolds numbers roughly $1.5\times$ higher during systole, exhibiting sharper near-wall gradients and stronger advective effects that the network must resolve.

2.3.3 Network Architecture and Training

Three architectures were carried through the steady ablation of Experiment E1 (Section 2.4.1). The first is the PhysicsNeMo default fully connected MLP FC 6×512 : six hidden layers of 512 neurons each with SiLU activations, around 1.6 million trainable parameters. SiLU is smooth and non-saturating, providing more stable gradient flow through deep networks than classical choices such as tanh or sigmoid [49]. The second is a smaller fully connected MLP FC 4×256 with the same activation, four layers of 256 neurons and roughly 400k parameters. The third is a Modified Fourier 4×256 network [50] that replaces the standard linear input projection with trainable sinusoidal basis functions applied independently to each input coordinate across multiple frequency levels. This directly targets the spectral bias of standard MLPs, which tend to learn low-frequency components first and therefore under-resolve the steep near-wall velocity gradients that drive WSS [50].

The Modified Fourier 4×256 network with ReLoBRaLo loss balancing was the winning configuration of the E1 ablation (Section 3.2.1) and is the only architecture used for the pulsatile production runs of Experiments E2–E5 (Sections 2.4.2–2.4.5). The other two MLPs and the LR-annealing loss-balancing variant appear in the present chapter only because they define the search space of the ablation.

All networks share the composite training loss:

$$\mathcal{L} = \lambda_{\text{int}} \mathcal{L}_{\text{int}} + \lambda_{\text{in}} \mathcal{L}_{\text{in}} + \lambda_{\text{out}} \mathcal{L}_{\text{out}} + \lambda_{\text{wall}} \mathcal{L}_{\text{wall}} + \lambda_{\text{flux}} \mathcal{L}_{\text{flux}}, \quad (2.3)$$

where:

- $\mathcal{L}_{\text{int}}$ penalizes the momentum and continuity residuals (Eq. 2.1–2.2) at interior collocation points sampled from the closed STL volume.
- $\mathcal{L}_{\text{in}}$ enforces a circular parabolic inlet velocity profile. In the steady E1 ablation the inlet amplitude is fixed at the cardiac-cycle mean. In all pulsatile runs the same profile is modulated in time by the same ICA inflow waveform $W(t)$ used for the FEM runs (Section 2.2.3), so that the inlet constraint reads

$$\mathbf{u}_{\text{in}}(r, t) = W(t) U_{\text{max}} \max\left(1 - \frac{r^2}{R_{\text{in}}^2}, 0\right) \hat{\mathbf{n}}_{\text{in}}, \quad (2.4)$$

with r the distance from the inlet centroid, $R_{\text{in}} = \sqrt{A_{\text{in}}/\pi}$ the inlet radius inferred from the cap mesh area, and $\hat{\mathbf{n}}_{\text{in}}$ the inward inlet normal.

- $\mathcal{L}_{\text{out}}$ enforces zero gauge pressure at the outlet ($p = 0$).
- $\mathcal{L}_{\text{wall}}$ enforces no-slip at the vessel wall ($\mathbf{u} = \mathbf{0}$).
- $\mathcal{L}_{\text{flux}}$ penalizes deviations of the volumetric flux, integrated over the outlet cap and the internal cross-section disk, from the theoretical parabolic inlet flux $Q = \frac{1}{2}\pi R_{\text{in}}^2 U_{\text{max}}$. These integral constraints provide global mass-conservation information that pointwise residuals alone do not enforce reliably [10]. Pointwise collocation of $\nabla \cdot \mathbf{u} = 0$ at a finite set of interior samples controls the local divergence but does not directly constrain the net flux across the boundary faces. By the divergence theorem, zero net boundary flux requires the continuity equation to hold globally, not merely at a finite sample. A network can therefore drive $\mathcal{L}_{\text{int}}$ to a small value while still carrying a non-zero outflow imbalance at the outlet caps. The weight was fixed at $\lambda_{\text{flux}} = 0.1$.

Weights were optimised using the Adam optimizer with an exponential learning rate schedule (initial rate 10^{-3} , decay factor 0.95 every 15,000 steps) for up to 1,500,000 iterations for the steady ablation and for the pure pulsatile runs, and for up to 800,000 iterations for the hybrid runs of Section 2.3.5, where the supervised term accelerates the early convergence. Each training step draws from fixed per-constraint batch sizes: $N_{\text{interior}} = 3500$, $N_{\text{no-slip}} = 3000$, $N_{\text{inlet}} = 700$, $N_{\text{outlet}} = 400$ per outlet, and $N_{\text{flux}} = 200$ quadrature points per integral surface. Training was stopped either at convergence, assessed by monitoring the mean inlet pressure logged every 10,000 steps, or at the 24-hour wall-time limit. After training, the field was evaluated on a uniform 128^3 voxel grid (batch size 8192), producing volumetric arrays of (u, v, w, p) .

2.3.4 Adaptive Loss Weighting

With fixed unit weights ($\lambda_k = 1$ for all constraints), the gradients contributed by different loss components to the shared network parameters can differ by several orders of magnitude, a consequence of the varying scales of physics residuals, boundary terms, and the integral flux constraint. When one component dominates parameter updates the others stagnate, often manifesting as a persistently large $\mathcal{L}_{\text{flux}}$ despite overall loss reduction. This imbalance

has a formal interpretation in terms of the Neural Tangent Kernel (NTK) eigenvalue spectrum and motivates dynamic weight adjustment [51].

Two adaptive aggregation strategies were evaluated in combination with the Modified Fourier 4×256 architecture:

Learning Rate Annealing (LR Annealing) [47]. At every $N_{\text{upd}} = 200$ training steps, the ℓ_2 norm of the gradient of each loss component with respect to the network parameters is computed. Each weight λ_k is then adjusted so that its gradient magnitude matches that of the continuity residual, used as the reference constraint:

$$\hat{\lambda}_k = \frac{\|\nabla_{\theta} \mathcal{L}_{\text{continuity}}\|}{\|\nabla_{\theta} \mathcal{L}_k\|}, \quad \lambda_k \leftarrow (1 - \alpha) \lambda_k + \alpha \hat{\lambda}_k, \quad (2.5)$$

where $\alpha = 0.01$ is a slow exponential moving average factor that prevents abrupt weight changes between steps.

ReLoBRaLo [48]. Each weight is updated proportionally to two complementary loss ratios: a long-term ratio comparing the current loss to its value at initialisation, and a short-term ratio comparing it to the previous step. An exponential moving average with $\alpha = 0.95$ blends the two signals. At each step, the long-term reference point is retained with probability $\beta = 0.99$ and reset to the current losses with probability $1 - \beta = 0.01$. This occasional random reset re-anchors the long-term ratio, preventing it from becoming trivially close to 1 as training progresses while still preserving a stable memory of the training history.

2.3.5 Hybrid PINN with Sparse FEM Supervision

The pure PINN of the previous subsections is trained from physics losses only. The network never sees a velocity or a pressure value from a reference solver. In a clinical setting, however, a small number of FEM snapshots are often available. They may come from a CFD reference already run on a similar geometry, or from a coarse FEM solution that is cheap to produce. To exploit this, a hybrid variant was implemented. It augments the physics losses of Equation 2.3 with a sparse data term constructed from the FEM solution of Section 2.2.

For each case, a supervised dataset is built from $N_t = 20$ uniformly spaced frames over one cardiac cycle. Within each frame, N_s samples are drawn at random from the FEM volumetric mesh nodes, with $N_s \in \{100, 500, 2000, \text{all}\}$. Each sample stores the spatio-temporal coordinates (x, y, z, t) and the FEM velocity and pressure $\mathbf{q}_{\text{FEM}} = (u, v, w, p)$. The training loss becomes

$$\mathcal{L}_{\text{hybrid}} = \mathcal{L}_{\text{PDE+BC}} + \lambda_{\text{data}} \frac{1}{|\mathcal{S}|} \sum_{(\mathbf{x}, t) \in \mathcal{S}} \|\mathbf{q}_{\theta}(\mathbf{x}, t) - \mathbf{q}_{\text{FEM}}(\mathbf{x}, t)\|_2^2, \quad (2.6)$$

where $\mathcal{L}_{\text{PDE+BC}}$ groups the interior, inlet, outlet, wall, and flux terms of Equation 2.3, $\mathbf{q}_{\theta}$ is the network output, and $\mathcal{S}$ is the sparse supervised set. The data weight λ_{data} is set to 1.0 at initialisation in every hybrid run and is then rebalanced by the ReLoBRaLo scheme of Section 2.3.4 alongside the other constraint weights, so its initial value acts only as a starting point.

The intent is not to fit the FEM solution pointwise. The supervised term anchors the network in regions where the physics losses alone converge slowly, typically near the dome and the bifurcation. The experiment also measures how rapidly the agreement with FEM improves as the supervision budget N_s grows. The corresponding experiments are E4 and E5 in Section 2.4. Results are reported in Section 3.4.

FEM velocity fields are stored in cm/s. To bring them into the PINN’s kinematic-pressure unit system, pressure values are divided by $\rho_{\text{blood}} = 1.06 \text{ g cm}^{-3}$ and then gauge-shifted per time-step to remove the absolute offset introduced by the resistance-free outlet boundary condition, which the PINN cannot recover from its soft Dirichlet $p = 0$ constraint. The architecture, optimiser, learning rate schedule, and boundary condition implementation are identical to the pure-PINN reference, so the supervised term is the only experimental variable between the two runs.

2.3.6 PINN Post-Processing and WSS Extraction

WSS computation from a trained PINN follows two distinct stages: an approximate in-training monitor used as a convergence diagnostic, and an exact post-hoc extraction based on automatic differentiation.

In-training monitor (finite difference). During training, a lightweight WSS estimate is tracked via a `PointwiseMonitor` evaluated every 10,000 steps. Because the no-slip condition enforces $\mathbf{u}(\mathbf{x}_w) = \mathbf{0}$, the network is queried at off-wall points displaced a distance $\varepsilon = 10^{-3}$ inward from $N_{\text{wall}} = 2000$ surface points along the outward normal $\hat{\mathbf{n}}$:

$$|\boldsymbol{\tau}_w|_{\text{FD}} \approx \nu \frac{|\mathbf{u}(\mathbf{x}_w - \varepsilon \hat{\mathbf{n}})|}{\varepsilon}. \quad (2.7)$$

This estimate is inexpensive. It requires only a single forward pass, and provides a practical diagnostic: stagnation of the mean wall WSS despite overall loss reduction indicates that boundary-layer resolution has not improved, motivating early stopping or architecture changes.

Post-training extraction (automatic differentiation). The finite-difference estimate is sensitive to the choice of ε and to numerical noise in the velocity field near the wall. After training, exact WSS values are recovered using the script `compute_wss.py`, which exploits a key property of PINNs: the solution network $\mathbf{u}(x, y, z)$ is a differentiable function of its inputs, so its spatial Jacobian can be computed to machine precision by reverse-mode automatic differentiation via `torch.autograd.grad`, without any mesh or finite-difference stencil.

The procedure is as follows. A set of N_w points is sampled from the no-slip STL surface together with their outward normals $\hat{\mathbf{n}} = (n_x, n_y, n_z)$. For each point the network is evaluated with `requires_grad=True`, and all nine components of the velocity gradient tensor $\partial u_i / \partial x_j$ are obtained by automatic differentiation. The viscous stress tensor is then

$$\sigma_{ij} = \nu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad (2.8)$$

and the traction vector at the wall is $t_i = \sigma_{ij} n_j$. The WSS magnitude is the norm of the tangential component of the traction:

$$|\boldsymbol{\tau}_w| = |\mathbf{t} - (\mathbf{t} \cdot \hat{\mathbf{n}}) \hat{\mathbf{n}}|. \quad (2.9)$$

Points are processed in batches of 4096 to limit GPU memory usage. The script saves the wall coordinates, outward normals, WSS magnitude, and full traction vector to a compressed `.npz` archive for downstream hemodynamic analysis.

2.3.7 HPC Deployment

All PINN runs were submitted to the UCLouvain CISM cluster under SLURM, the job scheduler that manages resource allocation on the cluster nodes. Each job was allocated a single GPU and a wall-time ceiling of 24 hours. The environment consisted of Python 3.12, CUDA 12.6, PyTorch, and `nvidia-physicsnemo` in a dedicated virtual environment on the cluster scratch storage.

2.4 Experimental Design

The pipelines described above support a sequence of six experiments that together address the two objectives of this thesis. Each experiment isolates one design choice and uses the same six patient geometries. Table 2.3 summarises them and points to the corresponding results section. The rest of this section describes each experiment in turn.

Table 2.3: Summary of the six experiments. Each experiment is reported in a dedicated subsection of Chapter 3.

ID	Pipeline	Question	Results
E1	PINN	Which architecture and loss weight-	§3.2.1
E2	PINN	Does the chosen configuration train	§3.2.2
E3	PINN vs FEM	How accurate and how costly is the	§3.3
E4	Hybrid PINN	What does sparse FEM supervision	§3.4
E5	Hybrid PINN	How many FEM samples are needed?	§3.4.2
E6	FEM	Which indices separate growing from	§3.5
		stable?	

2.4.1 E1: Steady-State PINN Architecture Ablation

Four PINN configurations were trained on the reference case 0203 under a steady-state Navier-Stokes setup. The steady setting removes the time dimension from the network input, so the network reduces to a map $(x, y, z) \mapsto (u, v, w, p)$ and the momentum residual

reduces to the steady form. It was chosen for the ablation because it halves the training degrees of freedom and allows faster iteration: the only remaining sources of difficulty are the patient-specific geometry, the loss component weighting, and the representation capacity of the network. The inlet is a fixed parabolic profile with peak velocity $U_{\text{in}}^{\text{max}} = 0.375 \text{ cm/s}$, the outlet is held at zero pressure, and the wall is no-slip. The fluid properties are taken from the NVIDIA PhysicsNeMo aneurysm tutorial used as the starting point for the ablation: kinematic viscosity $\nu = 2.5 \times 10^{-4} \text{ cm}^2/\text{s}$ and density $\rho = 1.0 \text{ g/cm}^3$. These differ from the physiological blood properties used in the pulsatile runs ($\nu \approx 3.77 \times 10^{-2} \text{ cm}^2/\text{s}$, $\rho = 1.06 \text{ g/cm}^3$). With $D_{\text{in}} \approx 0.43 \text{ cm}$ for case 0203, the steady Reynolds number is $\text{Re} \approx 325$, in the laminar range. The absolute loss values of this ablation are not comparable to the pulsatile runs. The ablation is used solely to rank the four configurations under matched conditions.

The first configuration is the baseline FC 6×512 with uniform loss weights. The second is the smaller FC 4×256 with uniform loss weights and isolates the effect of model capacity. The third is the Modified Fourier 4×256 network combined with LR Annealing. The fourth is the same Modified Fourier network combined with ReLoBRaLo. All other settings are held fixed across the four runs. These include the collocation budgets, the optimiser, the learning rate schedule, and the 24-hour wall-time limit. The configuration retained for the rest of the study is the one that reaches the lowest total loss within the wall-time budget.

2.4.2 E2: Pulsatile PINN on the Full Cohort

The configuration selected in E1 is extended to a pulsatile setting. Time is added as a network input and the inlet boundary condition is modulated with the population-averaged ICA waveform used for the FEM simulations. The same network is trained independently for each of the six cases under identical optimisation settings. Convergence is assessed from the inlet-pressure trace and the wall TAWSS monitor described in Section 2.3.6.

2.4.3 E3: Pure PINN vs FEM Accuracy and Cost

The converged pulsatile PINN of E2 is compared to the FEM reference for each of the six cases. Three quantities are reported. The first one is the relative ℓ_2 velocity error on a common volumetric grid restricted to the lumen. Specifically, the PINN cycle-averaged velocity is sampled on a regular 48^3 voxel grid covering the FEM bounding box; the relative error

$$\epsilon_{\ell_2} = \frac{\|\bar{\mathbf{u}}_{\text{PINN}} - \bar{\mathbf{u}}_{\text{FEM}}\|_2}{\|\bar{\mathbf{u}}_{\text{FEM}}\|_2} \quad (2.10)$$

is evaluated only on grid points falling inside the FEM lumen to avoid contaminating the metric with extrapolation outside the vessel. The second one is the agreement on the wall TAWSS and OSI averaged over the aneurysm dome. The third one is the wall-clock cost of each solver, that is, the PINN training time on a single GPU and the FEM solve time on the cluster nodes used in Section 2.2.4. A fourth diagnostic quantifies inlet boundary-condition satisfaction: the PINN axial velocity is integrated over 8000 surface points sampled from the inlet cap to recover the predicted volumetric flow rate $Q_{\text{PINN}}(t)$,

and the time-mean ratio $\overline{|Q_{\text{PINN}}|}/Q_{\text{target}}$ is reported for each case. A per-cycle visual comparison is also produced for case 0203. Twenty VTU frames are selected uniformly over the analysis cycle and the PINN velocity and pressure fields are evaluated at the exact SimVascular mesh nodes at the matching cardiac phase $t_{\text{PINN}} = t_{\text{phys}} \bmod T$, giving a three-panel side-by-side view of the SV field, the PINN field, and the pointwise error at each phase. This experiment answers whether a pure PINN is competitive with FEM on patient-specific aneurysm geometries.

2.4.4 E4: Hybrid PINN with Sparse FEM Supervision

The hybrid loss of Equation 2.6 is trained on all six cases with a fixed supervision budget of $N_t = 20$ frames and $N_s = 2000$ samples per frame, drawn at random from the FEM mesh nodes. The same accuracy and cost metrics as in E3 are reported. The gain attributable to the sparse FEM supervision can therefore be read directly against the pure-PINN baseline.

2.4.5 E5: Data-Sparsity Sweep

On the reference case 0203, the hybrid PINN is retrained for $N_s \in \{100, 500, 2000, \text{all}\}$ samples per frame, with $N_t = 20$ frames in every run. At $N_s = 100$ the supervised set contains 2000 tuples in total, covering about 2 % of the FEM mesh nodes per frame. At $N_s = 500$ it rises to 10 000 tuples; at $N_s = 2000$ to 40 000 tuples, roughly 40 % of the mesh. The *all* configuration uses every FEM node across all 20 frames, exactly 99 440 tuples in total ($4972 \text{ nodes} \times 20 \text{ frames}$, see Table 2.2), and is equivalent to a physics-regularised neural interpolant of the FEM solution. All other settings are held fixed. The aim is to characterise the trade-off between supervision budget and accuracy, and to identify the regime in which a small number of FEM samples already brings most of the benefit.

2.4.6 E6: Growing versus Stable Hemodynamic Comparison

The FEM solutions are used as ground truth. The hemodynamic indices of Section 2.5.1 are computed on the dome and on the parent vessel of each case. Dome-level values, dome-to-parent ratios, and within-pair differences are reported. Effect sizes are quantified using Cohen’s d (Section 2.5.3). This experiment answers the clinical question of which indices separate growing aneurysms from stable aneurysms in this matched cohort.

2.5 Hemodynamic Analysis Pipeline

2.5.1 Parameter Definitions and Computation

The hemodynamic indices analysed in this work are defined in Section 1.2: TAWSS (Eq. 1.4), OSI (Eq. 1.5), RRT (Eq. 1.6), ECAP (Eq. 1.7) and LSA (Eq. 1.8), together with the additional indices PSWSS, ΔWSS , the PSWSS/TAWSS peaking factor and the WSSG introduced in the same section. This subsection only documents the numerical choices made when evaluating these definitions on the FEM output.

All indices are computed from the time-resolved WSS vector $\boldsymbol{\tau}_w$ at each surface node of the analysis cycle (steps 5719–7600, 100 stored frames, cycle 4 of 4, see Section 2.2.4). The WSS vector at each surface node is obtained from the FEM velocity gradient projected onto the surface tangent plane. Time integrals in Eq. 1.4 and Eq. 1.5 are evaluated by the composite trapezoidal rule on the 100-frame analysis cycle.

The RRT and ECAP fields are evaluated per node from the per-node TAWSS and OSI, with TAWSS floored at 0.05 Pa to avoid the $1/\text{TAWSS}$ singularity at boundary points where the numerical WSS magnitude is essentially zero. The LSA threshold τ_{th} is reported at three values: $\tau_{\text{th}} = 0.4$ Pa, 1.0 Pa and 1.5 Pa, all in the range commonly used in the literature [9, 37]. Unless stated otherwise, LSA values quoted in the text use $\tau_{\text{th}} = 0.4$ Pa. PSWSS uses the cardiac frame at which the area-averaged $|\boldsymbol{\tau}_w|$ on the wall reaches its maximum. The cycle phase is then shared across all surface nodes. The WSSG is computed with PyVista’s surface derivative on the triangulated wall mesh and expressed in Pa/cm. All scalar indices are reported as area-weighted surface averages over the region of interest.

2.5.2 Dome and Parent Vessel Analysis

The aneurysm dome was separated from the parent vessel by reusing the region labels distributed with each SimVascular case in the VMR dataset, as described in Section 2.2.5. The parent vessel region is the remainder of the simulation surface. All indices were computed separately on each region. In addition to the per-region averages, the dome-to-parent ratio $r_X = \bar{X}_{\text{dome}}/\bar{X}_{\text{parent}}$ defined in Section 1.2 was computed for every index (TAWSS, OSI, RRT, ECAP, WSSG, PSWSS, ΔWSS and LSA at each threshold).

2.5.3 Statistical Analysis

With only three growing and three stable cases, standard parametric tests are not appropriate. Given the small sample size, the test is severely underpowered and results should be interpreted as exploratory rather than confirmatory.

Effect sizes were quantified using Cohen’s d :

$$d = \frac{\bar{x}_{\text{growing}} - \bar{x}_{\text{stable}}}{s_p}, \quad (2.11)$$

where s_p is the pooled standard deviation. By convention, an effect size is labelled $|d| < 0.2$ negligible, 0.2–0.5 small, 0.5–0.8 medium, and $|d| > 0.8$ large [52]. These labels are reported throughout for reference, but the convention was established for adequately powered samples and does not transfer to a comparison of three versus three cases. With $n = 3$ per group the pooled standard deviation is estimated from two degrees of freedom per group, so d is itself a high-variance quantity: a single case shifting can move it by a full unit, and a nominally “large” $|d| > 0.8$ here carries far less evidential weight than the same value computed on a large cohort. The d values in this work are therefore used only to rank indices by the size of their between-group separation, not as evidence that any separation is established. The consistency of the sign across the three matched pairs (Section 3.5.3) is the more trustworthy signal at this sample size.

For the PINN accuracy metrics, where the six cases serve as the sample, the Pearson correlation coefficient r was used to assess whether the case-to-case ordering of a metric is

preserved between the PINN and the FEM reference:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}}, \quad (2.12)$$

where x_i and y_i denote the PINN and FEM values for case i and bars denote means over the $n = 6$ cases. The two-tailed p -value is computed from the t -distribution with $n - 2 = 4$ degrees of freedom.

Comparisons were run on both dome-level indices and dome-to-parent ratios. The within-pair differences (growing minus stable) were also examined to identify signatures consistent across all three anatomical locations.

Chapter 3

Results

Several figures in this chapter contain orbit animations (20-frame loops covering one cardiac cycle or a 360° surface rotation). These animations play automatically in PDF viewers that support the `animate` package, such as Adobe Acrobat and Foxit Reader. In other environments, including most web browsers and preview applications, the figure displays a static frame, typically the peak-systole snapshot. All animations are also available as standalone GIF files in the supplementary dataset at <https://doi.org/10.5281/zenodo.20402867>.

3.1 SimVascular CFD Results

This section presents the FEM reference results for all six cases. These results serve as the ground truth for the PINN accuracy comparison (Experiments E3 and E4) and for the clinical hemodynamic analysis (Experiment E6).

3.1.1 Simulation Overview and Convergence

Pulsatile FEM simulations were carried out with `svMultiPhysics` on all six aneurysm cases following the solver settings of Section 2.2.4. Case 0204 required substantially more wall time (16.23 h) than the other five cases (0.65 h to 2.93 h), which is consistent with its much denser mesh: the 0204 volume mesh carries about 2.3×10^4 nodes and 1.3×10^5 tetrahedra, roughly four to five times more than the meshes used for the other cases. The per-step cost scales accordingly.

For each case, the last 100 VTU files (one cardiac period) were used for the hemodynamic analysis, corresponding to cycle 4 (steps 5719–7600) in every patient.

For the six cases the cycle-to-cycle change in the area-averaged TAWSS and OSI on the wall dropped below 1% from cycle 2 onward, confirming that the analysis window captures a periodic flow field.

Table 3.1: Pulsatile FEM simulation overview. Four MPI ranks, $\Delta t = 5 \times 10^{-4}$ s, four cardiac cycles (7600 steps), analysis window: steps 5719–7600.

Case	Status	Pair	Wall time (h)
0203	Growing	1	0.65
0204	Stable	1	16.23
0207	Growing	2	2.93
0208	Stable	2	1.96
0209	Growing	3	1.74
0210	Stable	3	1.00

3.1.2 Velocity and Pressure Fields

Velocity magnitude over the cardiac cycle is shown in Figure 3.1 for the six cases. In every case, the parent ICA delivers an organized inflow jet that impinges on the aneurysm wall and drives a recirculating vortex inside the sac. The strength and spatial coherence of this vortex vary between cases. The animation also reveals the pulsatile modulation of the jet: it intensifies sharply through early systole and becomes weaker and more diffuse during diastole.

The gauge-pressure field over the cycle (Figure 3.2) shows a high-pressure region at the impingement site on the dome wall and a lower-pressure core in the recirculation zone. The animation makes the pulse wave visible: pressure rises throughout the parent ICA during systole, then decreases during diastole.

3.1.3 Wall Shear Stress and OSI Maps

TAWSS surface maps over the analysis cycle are shown in Figure 3.3 for all six cases. The spatial pattern is consistent across cases: TAWSS is highest along the impingement zone at the dome neck and on the proximal wall, while the distal dome and the recirculation core carry markedly lower TAWSS.

The area-averaged TAWSS on the wall ranges from 1.34 Pa (case 0203) to 7.17 Pa (case 0208). The variability between cases is larger than the difference between paired growing and stable aneurysms. Quantitative group comparisons are reported in Section 3.5.

OSI surface maps over the analysis cycle are shown in Figure 3.4. Elevated OSI values concentrate in the recirculation zone of the distal dome and overlap spatially with the low-TAWSS regions identified above.

The absolute OSI values are very low across the whole cohort. The wall area-averages lie between 0.0008 and 0.0045, and even the dome area-averages stay below 0.02 (Appendix C). This follows from the prescribed inlet condition. The population-averaged ICA waveform used here is strictly unidirectional and never reverses sign over the cardiac cycle (Figure 2.2), so the bulk flow keeps a near-constant direction and the near-wall shear changes little in orientation. The OSI maps are therefore reported for their spatial pattern rather than their absolute magnitude.




Figure 3.1: FEM velocity magnitude over one cardiac cycle, all six cases (20-frame orbit loop). Each row: one matched pair. Left: growing, right: stable. Per-case colour scale. Static fallback: peak-systole frame.




Figure 3.2: Gauge pressure over one cardiac cycle, same paired layout as Figure 3.1. Per-case symmetric scale centred on the mean. Static fallback: peak-systole frame.




Figure 3.3: TAWSS surface maps, all six cases (20-frame 360° orbit). Per-case colour scale. Same paired layout as Figure 3.1. Static fallback: single still.




Figure 3.4: OSI surface maps, all six cases (20-frame 360° orbit, $0 =$ unidirectional, $0.5 =$ fully oscillatory). Same paired layout as Figure 3.1. Static fallback: single still.

3.2 PINN-Based Simulation Results

This section reports the results of Experiments E1 and E2: the architecture ablation on the steady-state problem (E1) and the pulsatile PINN training on all six cases (E2). The accuracy comparison against FEM (E3) and the hybrid runs (E4 and E5) are reported in Sections 3.3 and 3.4.

3.2.1 Steady-State PINN: Preliminary Exploration

Before moving to the full pulsatile setup, a steady-state PINN was trained on the case-0203 geometry to identify a baseline architecture and loss-weighting strategy. Results are summarised in Table 3.2: the Modified Fourier 4×256 network with ReLoBRaLo reached the lowest total loss by an order of magnitude over the smaller fully-connected baseline, by two orders of magnitude over the default PhysicsNeMo 6×512 network and by nearly two orders of magnitude over the same Modified Fourier architecture trained with learning-rate annealing. The default 6×512 MLP is the worst of the four, which confirms that on this geometry blindly scaling up a vanilla MLP is not a substitute for either a better architecture or a smarter loss balancing.

Table 3.2: Steady-state PINN ablation on case 0203. Final total loss, wall time, and Adam steps completed within the 24-hour budget.

Configuration	Final total loss	Steps reached	Wall time (h)
FC 6×512 + uniform weights (default)	6.8×10^{-3}	1.48×10^6	24.0 (timeout)
FC 4×256 + uniform weights	5.4×10^{-4}	1.5×10^6	13.6
Modified Fourier 4×256 + LR annealing	1.6×10^{-3}	1.04×10^6	24.0 (timeout)
Modified Fourier 4×256 + ReLoBRaLo	4.3×10^{-5}	1.5×10^6	20.9

Predicted steady flow field. The Modified Fourier + ReLoBRaLo configuration was rolled out on the same 48^3 voxel grid as the pulsatile runs. The resulting velocity magnitude field, displayed as a clipped volume in Figure 3.5, reproduces the qualitative features of the steady FEM reference. The cycle-averaged speed inside the lumen is in the same range as the FEM reference, which is the simplest sanity check that the boundary conditions and the geometry have been encoded consistently. The near-wall velocity gradient is again smoother than in the FEM solution. The steady setting is nevertheless a useful checkpoint because it confirms that the global flow topology (jet direction, recirculation centre, stagnation location) is captured before any time-dependent term is introduced.

3.2.2 Pulsatile PINN Training Convergence

Training durations for the six pulsatile PINN runs, using the Modified Fourier 4×256 architecture with ReLoBRaLo selected in Section 3.2.1, are summarised in Table 3.3. The four cavernous ICA cases (0203, 0204, 0209, 0210) completed the 1.5×10^6 steps inside the 24-hour budget. The two ICA terminus cases 0207 and 0208 hit the 24-hour ceiling at around 1.49×10^6 steps and were resubmitted from the latest checkpoint for the remaining iterations (0.3 h and 0.2 h respectively).




Figure 3.5: Steady-state velocity magnitude on case 0203: FEM reference (left) vs Modified Fourier $4 \times 256 + \text{ReLoBRaLo}$ PINN (right). Shared $[0, 0.51]$ cm/s colour scale.

Table 3.3: Pulsatile PINN training duration per patient on a single A100 GPU.

Case	Status	Training time (h)
0203	Growing	23.5
0204	Stable	23.6
0207	Growing	24.3 (24.0 + 0.3 after resubmission)
0208	Stable	24.2 (24.0 + 0.2 after resubmission)
0209	Growing	23.8
0210	Stable	23.6

The loss trajectory follows a consistent pattern across cases (Figure 3.6). The total loss drops sharply during the first few thousand iterations as the network learns the boundary conditions, then enters a slow refinement regime dominated by the interior physics residual. The boundary loss components $\mathcal{L}_{\text{in}}$, $\mathcal{L}_{\text{out}}$ and $\mathcal{L}_{\text{wall}}$ all settle to low values within the first 10^5 steps. The interior residual $\mathcal{L}_{\text{int}}$ decreases steadily under ReLoBRaLo reweighting throughout training, without flattening at the wall-time limit, which indicates that further iterations would still improve the field.




Figure 3.6: Pulsatile PINN total training loss for all six cases. Faint lines: raw loss sampled every 100 steps. Solid lines: running-median smooth over 100 samples.

The two-outlet ICA terminus cases 0207 and 0208 present an additional difficulty: the network must split the inflow between the ACA and the MCA branches, and the single integral flux constraint in the training loss cannot enforce per-outlet mass conservation. This translates into a persistently elevated $\mathcal{L}_{\text{flux}}$ component for those two cases, while the cavernous ICA cases (0203, 0204, 0209, 0210) reach clearly lower flux residuals.

3.2.3 Predicted Pulsatile Flow Fields

Velocity and pressure fields were extracted from each trained checkpoint on a regular voxel grid covering the FEM domain, at multiple time samples spanning one cardiac period (Figure 3.7). For the cavernous ICA cases (0203, 0204, 0209, 0210), the PINN reproduces the dominant flow structure observed in the FEM solution: a high-velocity jet in the parent artery and a slower recirculation region inside the sac. The amplitude of the unsteady response is captured between systole and diastole, with peak in-sac velocities reached within a few time samples of the FEM peak.

For the ICA terminus cases (0207, 0208), the predicted fields exhibit non-physical velocity concentrations at the outlet faces, a direct consequence of the unresolved per-outlet flux imbalance described in Section 3.2.2. The quantitative comparison reported in Section 3.3




Figure 3.7: Pulsatile PINN velocity magnitude on case 0203 over one cardiac cycle (20-frame orbit loop). Static fallback: peak-systole frame.

reflects this.

3.3 Pulsatile PINN vs FEM Comparison

This section reports Experiment E3: the quantitative comparison between the pure pulsatile PINN and the FEM reference on all six cases. Accuracy is assessed through the cycle-averaged relative ℓ_2 velocity error, the wall TAWSS and OSI agreement, and the per-cycle velocity and pressure fields.

3.3.1 Cycle-Averaged Velocity Error

Table 3.4 reports the cycle-averaged relative ℓ_2 velocity error (Eq. 2.10, defined in Section 2.4.3) for all six cases.

Table 3.4: Cycle-averaged relative ℓ_2 velocity error (Eq. 2.10) between PINN and FEM, restricted to lumen voxels.

Case	Status	ϵ_{ℓ_2}
0203	Growing	0.81
0204	Stable	0.90
0207	Growing	0.83
0208	Stable	0.89
0209	Growing	0.89
0210	Stable	0.87

The relative error is very large. It sits between 0.81 and 0.90 in every case.

3.3.2 Wall TAWSS and OSI Agreement

Table 3.5 reports the area-averaged TAWSS and OSI on the entire wall surface, computed independently from the PINN and from the FEM fields. Figure 3.8 shows the corresponding scatter plots, with each marker labelled by patient identifier and the $y = x$ identity line drawn for reference.

The PINN and FEM TAWSS values are strongly correlated across the six cases (Pearson $r = 0.92$, $p = 0.009$), confirming that the network preserves the relative ranking of wall stress across geometries despite the systematic underestimation. OSI shows no statistically meaningful correlation ($r = 0.18$, $p = 0.73$), consistent with the wide scatter visible in the right panel of Figure 3.8.

3.3.3 Per-Case Error Distribution

The cycle-averaged relative ℓ_2 velocity error stays between 0.81 and 0.90 across the six cases. The wall TAWSS is underestimated by 65 to 89% in every patient, and the OSI

Table 3.5: Wall TAWSS and OSI from PINN and FEM, with relative difference $(\text{PINN} - \text{FEM})/\text{FEM}$ on TAWSS.

Case	Status	TAWSS _{FEM} (Pa)	TAWSS _{PINN} (Pa)	OSI _{FEM}	OSI _{PINN}	ΔTAWSS (%)
0203	Growing	1.34	0.47	0.0016	0.0037	-65.1
0204	Stable	2.76	0.31	0.0025	0.0141	-88.9
0207	Growing	6.29	1.50	0.0008	0.0071	-76.2
0208	Stable	7.17	1.62	0.0045	0.0094	-77.4
0209	Growing	4.75	0.81	0.0021	0.0045	-83.0
0210	Stable	4.26	0.69	0.0041	0.0044	-83.7




Figure 3.8: PINN vs FEM agreement on the aneurysm wall. Left: TAWSS (Pa). Right: OSI. The dashed line is the $y = x$ identity.

values differ from the FEM reference by up to a factor of six in the ICA terminus cases (0207 and 0208). The underestimation is consistent in sign across the full cohort. The smallest TAWSS error is at case 0203 (-65%) and the largest at case 0204 (-89%). The ICA terminus cases (0207, 0208) also show the largest OSI divergence, while the cavernous ICA cases show smaller OSI discrepancies.

3.3.4 Per-Cycle Visual Comparison

Figure 3.9 shows the per-cycle side-by-side comparison for case 0203 across 20 cardiac-phase samples.




Figure 3.9: Case 0203 over the analysis cycle (20-frame orbit loop). Left: SV velocity. Centre: PINN velocity (shared colour bar). Right: pointwise error $|\mathbf{u}_{\text{PINN}} - \mathbf{u}_{\text{SV}}|$. Static fallback: peak-systole frame.

Table 3.6 summarises the per-cycle peak and volume-mean velocity magnitudes for the two fields. The picture is consistent across cases: the PINN does capture the jet structure (peak ratio in the range 0.28–0.66), but distributes the kinetic energy more uniformly across the lumen than the FEM solution does (volume-mean ratio in the range 0.13–0.22). The gauge pressure relative error is above unity for every case, meaning the predicted spatial pressure distribution carries no useful information beyond the mean.

Table 3.6: Per-cycle peak and volume-mean velocity magnitude (PINN and FEM), averaged over 20 frames. Last column: mean gauge-pressure relative ℓ_2 error.

Case	Peak $ U _{\text{SV}}$ (cm/s)	Peak $ U _{\text{PINN}}$ (cm/s)	Peak ratio	Mean $ U _{\text{SV}}$ (cm/s)	Mean $ U _{\text{PINN}}$ (cm/s)	$\overline{\epsilon_{\ell_2,p}}$
0203	101.4	67.0	0.66	14.6	3.2	1.16
0204	255.3	72.3	0.28	22.7	3.6	1.03
0207	195.1	96.0	0.49	37.5	8.2	1.14
0208	258.4	137.4	0.53	41.6	5.2	1.19
0209	213.0	91.8	0.43	22.0	4.7	1.04
0210	222.7	112.2	0.50	33.8	5.6	1.07

The inlet boundary condition is well satisfied: the time-mean flux ratio $|\overline{Q_{\text{PINN}}}|/Q_{\text{target}}$ lies between 0.85 and 0.94 across the six cases. The velocity discrepancy is concentrated inside the lumen: the volume-mean speed ratio $\langle |U| \rangle_{\text{PINN}} / \langle |U| \rangle_{\text{SV}}$ lies in the range 0.13–0.22 across the six cases (Table 3.6). Across the six cases, the PINN and FEM volume-mean

velocities are moderately correlated (Pearson $r = 0.75$, $p = 0.09$), while the case-to-case ordering of peak velocities is less consistent ($r = 0.57$, $p = 0.24$).

3.4 Hybrid PINN: Sparse FEM Supervision

This section reports Experiments E4 and E5, which test whether augmenting the physics losses with sparse FEM reference data closes the accuracy gap. The physics losses are kept intact throughout, so the hybrid network satisfies the governing equations everywhere, not only at the supervised points. The hybrid training setup is described in Section 2.3.5.

3.4.1 Per-Patient Training

Table 3.7: Hybrid PINN vs pure-PINN agreement on the analysis cycle, averaged over the same 20 frames per case as Table 3.6.

Case	$\overline{\epsilon_{\ell_2,U}}$ (pure)	$\overline{\epsilon_{\ell_2,U}}$ (hybrid)	TAWSS error % (pure)	TAWSS error % (hybrid)
0203	0.810	0.116	−65	+32
0204	0.896	0.299	−89	−12
0207	0.834	0.095	−76	−2
0208	0.894	0.109	−77	−10
0209	0.872	0.147	−83	−8
0210	0.873	0.131	−84	+5




Figure 3.10: Three-way velocity comparison on case 0203 (20-frame orbit loop): SV reference (left), pure PINN (centre), hybrid PINN (right), shared colour bar. Static fallback: peak-systole frame.

Across the six cases the supervised term acts as intended. The cycle-averaged relative velocity error $\overline{\epsilon_{\ell_2,U}}$ drops from 0.81–0.90 for the pure PINN to 0.10–0.30 for the hybrid, and the wall TAWSS gap closes from a uniform −65 to −89 % underestimation down to within ± 32 % of the FEM reference. The closest matches are −2 % for case 0207 and −8 % for case 0209, while case 0203 overshoots by +32 %. The overshoot on case 0203 is a sign that the supervised samples are dense enough to anchor the interior and that the network now reproduces sharper near-wall gradients than the smooth pure-PINN baseline. Across the six cases, the hybrid PINN and FEM TAWSS values are very strongly correlated (Pearson $r = 0.99$, $p < 0.001$), a marked improvement over the pure-PINN correlation of $r = 0.92$.

3.4.2 Data-Sparsity Sweep

Table 3.8: Data-sparsity sweep on case 0203. Cycle-averaged ℓ_2 errors on velocity and pressure, and TAWSS error, vs supervised samples per frame N_s .

N_s per frame	$\overline{\epsilon}_{\ell_2, U}$	$\overline{\epsilon}_{\ell_2, p}$	TAWSS error (%)
0 (pure PINN)	0.810	1.16	−65
100	0.374	0.38	−5
500	0.209	0.26	+20
2000 (main hybrid)	0.116	0.02	+32
all (≈ 5000)	0.117	0.002	+37




Figure 3.11: Cycle-averaged velocity (blue) and pressure (red) ℓ_2 errors on case 0203 vs N_s . Dashed line: pure-PINN baseline.

The sparsity sweep on case 0203 already exposes a clear monotonic trend. The cycle-averaged velocity error $\overline{\epsilon}_{\ell_2, U}$ falls from 0.81 (pure PINN) to 0.37 at $N_s = 100$, 0.21 at $N_s = 500$ and 0.12 at $N_s = 2000$, a near-linear decay in $\log_{10}(N_s)$. The gauge pressure error follows the same monotonic trend, from 1.16 down to 0.02. The wall TAWSS reacts differently: at $N_s = 100$ the hybrid is closest to the FEM reference (−5 %), at $N_s = 500$ it overshoots (+20 %) and at $N_s = 2000$ it overshoots further (+32 %). The picture is consistent: the supervised term first removes the pure-PINN bias, then gradually replaces the smooth interior reconstruction by one with sharper near-wall gradients, eventually pushing TAWSS slightly above the FEM value. The $N_s = \text{all}$ run closes the sweep at $\overline{\epsilon}_{\ell_2, U} = 0.117$ and a TAWSS error of +37 %, essentially indistinguishable from the $N_s = 2000$ run: with the FEM mesh of case 0203 carrying about 5 000 nodes per frame, the $N_s = 2000$ configuration already supervises roughly 40 % of the mesh and is close to the saturation of what additional samples can bring on this geometry.

3.4.3 Per-Patient Hybrid Cycle Animations

For each patient, one cardiac cycle of the hybrid PINN against the SimVascular reference is rendered as a 20-frame orbit animation with a per-case fixed colour scale on velocity. Figure 3.12 gives the animation for case 0203. The five other patients are available in the supplementary dataset at <https://doi.org/10.5281/zenodo.20402867>. Compared with the pure-PINN animation in Figure 3.9, the hybrid panel recovers the parent-artery jet amplitude and the near-wall velocity gradient that drives the wall shear stress in every case. The same panel layout makes it possible to compare the three fields side by side at every cardiac phase.




Figure 3.12: Hybrid PINN vs SimVascular, case 0203 (20-frame orbit loop): SV velocity (left), hybrid PINN velocity (centre), pointwise error (right), shared colour scale. Static fallback: single still.

3.5 Hemodynamic Analysis

This section reports Experiment E6: the growing-versus-stable hemodynamic comparison based on the FEM results, which serve as the ground-truth reference. PINN hemodynamic values are reported in Section 3.3.

3.5.1 Per-Case Hemodynamic Indices

The full set of per-case hemodynamic indices is reported in Appendix C: wall area-averages for all indices (Table C.1), dome and parent-vessel area-averages with the resulting dome-to-parent ratios (Table C.2), and the LSA fraction at the three reported thresholds (Table C.3). This section reports group-level statistics in Section 3.5.2 and within-pair differences in Section 3.5.3. Per-case values cited in the text refer to the appendix tables.

Across the six cases the dome-to-parent TAWSS ratio is 0.26, 0.40 and 0.39 for the three growing aneurysms and 0.79, 0.58 and 0.58 for the three stable controls. The dome-to-parent OSI ratio exceeds 10 in every case and exceeds 100 in three of the six. The dome low-shear area at 0.4 Pa reaches 25.6 % on case 0203 and 8.2 % on case 0204, an order of magnitude above the corresponding wall-averaged values.

3.5.2 Growing vs Stable Group Comparison

Group-level descriptive statistics are reported in Table 3.9. The mean, standard deviation, and Cohen’s d effect size are computed as described in Section 2.5.3 of Chapter 2.

Table 3.9: Group comparison between growing and stable aneurysms ($n = 3$ per group). Mean $\pm$ SD and Cohen’s d effect size (growing – stable). r_X denotes the dome-to-parent ratio of index X . Per-case values are in Appendix C.

Index	Growing	Stable	d
<i>Wall area-averages</i>			
TAWSS (Pa)	4.13 ± 2.53	4.73 ± 2.24	-0.25
OSI	0.0015 ± 0.0006	0.0037 ± 0.0010	-2.58
RRT (Pa^{-1})	3.53 ± 3.01	1.06 ± 1.11	$+1.09$
ECAP (Pa^{-1})	0.0055 ± 0.0014	0.0116 ± 0.0083	-1.02
WSSG (Pa/cm)	17.12 ± 10.35	18.59 ± 13.11	-0.12
PSWSS (Pa)	6.92 ± 4.23	7.87 ± 3.70	-0.24
ΔWSS (Pa)	4.59 ± 2.70	5.22 ± 2.37	-0.25
PSWSS/TAWSS	1.675 ± 0.017	1.668 ± 0.018	$+0.42$
LSA ^{0.4}	0.040 ± 0.057	0.022 ± 0.016	$+0.43$
LSA ^{1.0}	0.156 ± 0.201	0.077 ± 0.043	$+0.54$
LSA ^{1.5}	0.245 ± 0.278	0.145 ± 0.097	$+0.48$
<i>Dome area-averages</i>			
LSA _d ^{0.4}	0.091 ± 0.143	0.029 ± 0.045	$+0.59$
LSA _d ^{1.0}	0.393 ± 0.486	0.111 ± 0.126	$+0.79$
LSA _d ^{1.5}	0.498 ± 0.468	0.200 ± 0.186	$+0.84$
<i>Dome-to-parent ratios</i>			
r_{TAWSS}	0.35 ± 0.08	0.65 ± 0.12	-2.99
r_{OSI}	117.2 ± 71.1	75.4 ± 117.9	$+0.43$
r_{RRT}	3.52 ± 1.02	1.58 ± 0.38	$+2.51$
r_{ECAP}	743 ± 585	421 ± 707	$+0.50$
r_{WSSG}	0.75 ± 0.50	1.14 ± 0.30	-0.95
r_{PSWSS}	0.39 ± 0.09	0.69 ± 0.11	-3.12
$r_{\Delta\text{WSS}}$	0.49 ± 0.05	0.78 ± 0.13	-3.03

Box plots of the wall indices are shown in Figure 3.13. At the wall level, OSI carries the largest effect size ($d = -2.58$), followed by RRT ($d = +1.09$) and ECAP ($d = -1.02$). Wall TAWSS, PSWSS, ΔWSS , WSSG and the PSWSS/TAWSS peaking factor all have small effect sizes ($|d| < 0.5$). The three LSA thresholds carry small-to-medium effect sizes on the wall and medium-to-large effect sizes on the dome. Normalising by the parent vessel sharpens the picture: the dome-to-parent ratios of TAWSS, RRT, PSWSS and ΔWSS all reach $|d| \geq 2.5$ with no overlap between the two groups, whereas the dome-to-parent ratios of OSI, ECAP and WSSG carry only small-to-large effect sizes.




Figure 3.13: Box plots of wall hemodynamic indices for the growing and stable groups ($n = 3$ each). Individual case values are overlaid as black dots.

3.5.3 Paired Analysis

Pairing by anatomical location removes the location-driven baseline differences in WSS magnitude that confound the unpaired group comparison. Table 3.10 reports the within-pair differences (growing – stable) for all indices. An index is flagged as consistent when the sign of the difference is the same across all three pairs.

Seven indices flip sign consistently across the three pairs in this cohort: wall OSI, wall ECAP, the dome low-shear area at 1.5 Pa, and the dome-to-parent ratios of TAWSS, RRT, PSWSS and Δ WSS. The dome-to-parent ratios of OSI, ECAP and WSSG reverse sign on pair 3 (0209/0210) and are therefore not consistent in this cohort, even though their unpaired effect sizes in Table 3.9 are non-negligible.

Table 3.10: Within-pair differences (growing – stable) for all indices. “Cons.” is “Y” when the sign is the same across all three pairs.

Index	Pair 1 (0203–0204)	Pair 2 (0207–0208)	Pair 3 (0209–0210)	Cons.
<i>Wall area-averages</i>				
TAWSS (Pa)	–1.42	–0.87	+0.49	N
OSI	–0.0009	–0.0037	–0.0021	Y
RRT (Pa ^{–1})	+1.51	+6.11	–0.20	N
ECAP (Pa ^{–1})	–0.0021	–0.0010	–0.0151	Y
WSSG (Pa/cm)	–3.29	–5.87	+4.75	N
PSWSS (Pa)	–2.41	–1.45	+0.99	N
Δ WSS (Pa)	–1.66	–1.04	+0.81	N
PSWSS/TAWSS	–0.019	–0.0005	+0.041	N
LSA ^{0.4}	+0.077	+0.0001	–0.023	N
LSA ^{1.0}	+0.273	+0.013	–0.050	N
LSA ^{1.5}	+0.315	+0.028	–0.043	N
<i>Dome area-averages</i>				
LSA _d ^{0.4}	+0.174	+0.015	–0.001	N
LSA _d ^{1.0}	+0.686	+0.164	–0.006	N
LSA _d ^{1.5}	+0.585	+0.294	+0.014	Y
<i>Dome-to-parent ratios</i>				
r_{TAWSS}	–0.52	–0.18	–0.19	Y
r_{OSI}	+141.4	+160.1	–176.0	N
r_{RRT}	+2.47	+2.87	+0.48	Y
r_{ECAP}	+1042	+1093	–1170	N
r_{WSSG}	–0.77	+0.50	–0.90	N
r_{PSWSS}	–0.52	–0.20	–0.19	Y
$r_{\Delta\text{WSS}}$	–0.47	–0.22	–0.17	Y

Chapter 4

Discussion

4.1 Evaluation of the PINN Approach

4.1.1 Root Causes of the Velocity Discrepancy

Three factors explain the large velocity errors of the pure PINN (relative ℓ_2 between 0.81 and 0.90, TAWSS underestimated by 65 to 89 %).

The first factor is the unresolved boundary layer. The Modified Fourier MLP produces a globally smooth velocity field. It cannot represent the steep transverse velocity gradient in the thin near-wall region that carries the wall shear stress. This is a well-documented limitation of vanilla PINNs on convection-dominated incompressible flow [10, 47]: the spectral bias of standard MLPs lets them learn low-frequency components first and under-resolve high-frequency near-wall structures, which is precisely the motivation for the Fourier-feature embedding adopted here following the analysis of Tancik et al. [50]. These error levels are comparable to the values commonly reported for vanilla MLP-based PINNs on patient-specific vascular geometries at comparable training budgets [10, 11]. The consequence is a near-wall velocity that is too high, which in turn means that the gradient evaluated at the wall is too shallow, and TAWSS is underestimated. The severity of this effect is compounded by the flow regime: the pulsatile runs reach bulk Reynolds numbers of 300–500 (Section 2.3.2), within the physiological range for cerebral arteries, with instantaneous peak-systolic values roughly $1.5\times$ higher. During systole the viscous Stokes layer is thinner and the near-wall velocity gradient is steeper, sharpening the resolution demand on the network beyond what a cycle-mean estimate would suggest.

The second factor is jet diffusion in the lumen. The same smoothness prior spreads the parent-artery inflow jet across the lumen cross-section instead of keeping it confined. Mass is conserved in an integral sense (the inlet flux ratio is 0.85–0.94), but the kinetic energy is redistributed into a larger, slower volume. This explains why the volume-mean speed ratio $\langle |U| \rangle_{\text{PINN}} / \langle |U| \rangle_{\text{SV}}$ lies in the range 0.13–0.22 even though the inlet boundary condition is approximately satisfied.

The third factor affects only the ICA terminus cases (0207 and 0208). The integral flux constraint in the training loss sums over both outlets, so any per-outlet flow split imbalance is invisible to the loss. This compounds with the smoothness issue and is the reason the

OSI error is largest for those two cases. The absence of per-outlet flux control breaks the case-to-case ordering of OSI: the Pearson correlation between pure PINN and FEM OSI is $r = 0.18$ ($p = 0.73$). This per-outlet limitation is an intrinsic consequence of constraining only the total outflow rather than each outlet individually. The importance of careful, per-component constraint weighting in multi-objective incompressible Navier-Stokes PINNs was emphasized by Jin et al. [46].

The loss convergence curves (Figure 3.6) reveal a distinction between the two terminus cases. Case 0207 reaches a final total loss of 22, indistinguishable from the cavernous ICA cases (12–25). Case 0208 plateaus at 56, more than twice as high, at the end of training. Both cases share the same outlet architecture, so the flux imbalance alone does not explain the gap. Case 0208 also has the highest peak-systolic velocity of the entire cohort (258 cm/s), which means the steepest near-wall gradients and the most demanding optimization landscape. The combination of an unresolvable outlet constraint and an inherently high-Re geometry prevented the optimizer from finding a solution of comparable quality within the fixed training budget.

4.1.2 Effect of Sparse FEM Supervision

The hybrid PINN substantially reduces all three error components. The cycle-averaged ℓ_2 velocity error drops from the pure-PINN level to 0.10–0.30. The TAWSS gap closes from a uniform underestimation to within $\pm 32\%$ of the FEM reference. The Pearson correlation between hybrid PINN and FEM TAWSS rises to $r = 0.99$ ($p < 0.001$) from $r = 0.92$ for the pure PINN: sparse supervision not only corrects the mean bias but also tightens the case-to-case ranking. The supervised term anchors the interior field and allows the network to learn sharper near-wall gradients than the physics losses alone can produce.

The data-sparsity sweep on case 0203 shows that a large share of the improvement is already achieved at $N_s = 100$ per frame (2000 total supervised tuples, about 2% of the FEM mesh). The velocity error falls from 0.81 at $N_s = 0$ to 0.37 at $N_s = 100$, and continues to drop nearly linearly in $\log_{10}(N_s)$. The TAWSS responds differently: it reaches its closest value to the FEM reference at $N_s = 100$ (-5%), then overshoots progressively as more samples are added. This non-monotone characteristic indicates a transition from physics-dominated to data-dominated reconstruction of the near-wall gradient, and suggests that an intermediate supervision budget (100–500 samples per frame) gives the best wall-shear accuracy for this geometry. The benefit of injecting a small amount of reference data into a physics-only loss has also been documented in the cardiovascular setting by Kissas et al., who used sparse 4D-flow MRI measurements to constrain a PINN of arterial pulse propagation [12], and is a recurrent theme of the physics-informed machine-learning review of Karniadakis et al. [45]. The present sweep extends this observation by quantifying the saturation point at which adding more FEM samples no longer improves the near-wall reconstruction on a patient-specific aneurysm geometry.

4.1.3 Accuracy and Computational Cost

The FEM simulations took between 0.65 and 2.93 hours for five of the six cases on four CPU cores. The sixth case (0204) took 16.23 hours because of its much denser mesh

(1.3×10^5 tetrahedra). The pulsatile PINN training required approximately 24 hours on a single A100 GPU for every case, regardless of geometry size. For the five fast FEM cases the PINN training is therefore 8–37 $\times$ more expensive than the FEM solve. This ratio compares different hardware, a single A100 GPU against four CPU cores, so it is indicative rather than a like-for-like measure. Once trained, the PINN evaluates the velocity and pressure field at 128^3 grid points in a single forward pass, which takes seconds. For queries at arbitrary points the trained network is orders of magnitude faster than a new FEM run.

In practical terms, the pure PINN does not provide a cost-competitive alternative to FEM for the hemodynamic indices studied here. The large TAWSS error renders the pure-PINN results unusable as a hemodynamic surrogate. The hybrid PINN reduces the cost gap for the dense-mesh cases like 0204 (FEM is slow while PINN training time is fixed), but the accuracy improvement still falls short of what would be needed for clinical use of TAWSS or OSI as risk predictors.

4.1.4 Potential for Clinical Application and Future Directions

Several technical remedies could close the remaining accuracy gap. Hard-constrained boundary conditions, implemented by multiplying the network output by a signed-distance function that enforces exact no-slip at the wall, would remove the smoothness issue at the wall and improve near-wall gradient resolution. Per-outlet integral flux constraints would eliminate the flow-split error in the ICA terminus cases. Importance sampling near the wall, allocating a larger fraction of collocation points to the boundary layer, would focus optimization on the region that matters most for WSS. Causal training in time, which trains the network sequentially from early to late cardiac phases rather than sampling all phases simultaneously, has been shown to accelerate convergence for time-dependent PINNs [53].

Beyond these architectural remedies, the hybrid approach is promising for settings where a small number of reference samples is available. In a clinical context those samples could come from a 4D phase-contrast MRI acquisition, which would allow a PINN to enrich sparse in-vivo measurements with a physically consistent, complete hemodynamic field. That task is not achievable with standalone CFD, which requires complete boundary conditions specified a priori. Transfer learning across patient geometries, pre-training on a large corpus of VMR cases and fine-tuning on a new patient, is a further direction that could reduce the per-patient training cost to a clinically acceptable scale.

4.2 Interpretation of Hemodynamic Differences

The results in Chapter 3 show that absolute TAWSS on the full wall does not separate growing from stable aneurysms in this cohort. The group means differ by only 0.6 Pa with a Cohen’s d of -0.25 , which is negligible. This is not surprising. The six cases span two anatomical sites and three different absolute flow regimes. Cavernous ICA aneurysms (pairs 1 and 3) carry substantially less wall shear than ICA terminus aneurysms (pair 2) simply because the parent artery is less straight and the inflow velocity is lower. Using an unpaired group mean to look for a growth-status signal conflates this anatomical effect with the hemodynamic phenotype of interest.

Normalising the dome shear by the paired parent-vessel value removes this confound. The dome-to-parent TAWSS ratio is 0.35 ± 0.08 in the growing group and 0.65 ± 0.12 in the stable group, with no overlap between the two distributions and $d = -2.99$. This ratio is consistent across all three pairs: in every case the growing aneurysm has a more shielded dome, that is, a dome that receives proportionally less shear than the parent artery of the same patient. This pattern holds even for pair 3 (0209 vs 0210), where the absolute TAWSS comparison reverses sign. Pairing by anatomy and normalising by the parent vessel are both required to expose a robust separator on this small cohort.

The low-WSS association with adverse aneurysm behaviour is well established in the rupture literature. Zhou et al. pooled multiple CFD studies and found low TAWSS to be consistently and significantly associated with rupture status across a large pooled population [9]. Suzuki et al. provided complementary sub-aneurysm evidence by showing that the exact rupture point in a series of ruptured cases coincided with the region of lowest local WSS [31]. The present results extend this evidence to the growth direction: in all three pairs the growing aneurysm has a more shielded dome relative to its own parent artery. Dabagh et al. tracked hemodynamic evolution longitudinally and found areas of low WSS and high OSI coinciding with regions of subsequent dome enlargement [41], consistent with the dome-shielding pattern observed here. Karnam et al. identified multiple competing hemodynamic growth pathways and cautioned against attributing growth to a single dominant mechanism [42]. The present cohort is too small to distinguish growth pathways, but the consistency of the dome-to-parent TAWSS ratio across all three pairs suggests that relative dome shielding is the most robust signal available at this sample size.

The co-localisation of elevated OSI with the low-TAWSS dome and recirculation regions is expected: where the mean WSS vector is small in magnitude, even modest temporal reversals push the OSI close to its theoretical maximum [33].

The OSI result goes in the opposite direction to what the classical high-OSI rupture-risk narrative would predict. Growing aneurysms show lower wall-averaged OSI than their stable counterparts in all three pairs. Two factors make this comparison weak, and the present work does not draw conclusions from it. First, the absolute OSI magnitudes are anomalously low across the whole cohort (wall averages below 0.005, dome averages below 0.02), because the prescribed ICA inflow waveform is strictly unidirectional and never reverses over the cardiac cycle (Section 2.2.3, Figure 2.2). With a non-reversing bulk inflow and a rigid wall, OSI is driven only by weak secondary recirculation and sits near the noise floor, so its sign is easily flipped by small numerical differences rather than by physiology. A waveform with a genuine diastolic reversal, or a compliant wall, would be needed to generate the dome OSI levels reported in the rupture literature. Second, even setting magnitude aside, the wall-averaged comparison is a scale effect. Wall-averaged OSI pools the aneurysm dome and the parent vessel. In every case the dome contributes an OSI one to two orders of magnitude above the parent vessel, but the dome is a small fraction of the total wall area. When dome and parent are averaged together, the signal from the dome is diluted. The paired analysis on dome-only OSI actually reverses sign in pair 3, which means OSI at the dome-only level is not a consistent separator either. With three cases per group and OSI values this close to zero, no OSI-based claim is defensible here, especially since the hemodynamic profile that distinguishes aneurysms is modulated by aneurysm location and patient characteristics [40]. Nishiwaki et al. found that OSI differences between growing ruptured and growing stable aneurysms were significant after

enlargement, with higher OSI in the ruptured group [34]. The present comparison is cross-sectional at a single time point, uses a non-reversing inflow, and does not track the post-growth OSI elevation that Nishiwaki observed. The two observations are compatible.

The dome low-shear area (LSA_d) points in the expected direction for two of the three pairs but not for pair 3. The same underpowering argument applies. The paired dome-to-parent RRT ratio is consistent across all three pairs ($d = +2.51$), because it inverts and amplifies the TAWSS ratio signal.

In summary, the hemodynamic signature most robustly associated with growing status in this cohort is a relatively shielded dome that receives less shear than its parent artery. This is the dome-level expression of the low-WSS hypothesis for aneurysm wall degeneration. It is not detectable from absolute TAWSS alone but emerges clearly when anatomy is controlled for by within-pair normalisation.

The additional indices reported in Table 3.9 and Table 3.10 sharpen this picture without changing its conclusion. The dome-to-parent ratios of the peak-systolic WSS and of the systolic-to-diastolic WSS range (ΔWSS) flip sign consistently across the three pairs in the same direction as the TAWSS ratio, with effect sizes of $d = -3.12$ and $d = -3.03$ respectively. These three ratios are not independent: the inflow waveform is shared across the six cases, so the temporal scaling between TAWSS, PSWSS and ΔWSS is approximately constant across patients. The PSWSS/TAWSS peaking factor makes this explicit. It sits in a narrow band of 1.67 ± 0.02 on the wall and 1.80 ± 0.05 on the dome (Table C.4), with $|d| < 0.5$ in both cases, and does not discriminate the two groups. The same shielding signature is therefore visible in three temporal windows of the cardiac cycle, but the time average remains the most informative scalar.

The WSSG dome-to-parent ratio reaches $|d| = 0.95$ at the unpaired group level but does not flip sign consistently within pairs. WSSG measures the spatial sharpness of the TAWSS field, which is sensitive to local curvature and to mesh resolution near the wall, and is known to show large inter-patient variability [38]. At $n = 3$ per group this metric should be considered exploratory.

Explicit dome-to-parent ratios were also computed for OSI and ECAP (Table 3.9). Both reach medium unpaired effect sizes ($d = +0.43$ and $d = +0.50$) but reverse sign in pair 3 (the 0209 / 0210 cavernous ICA pair) and are therefore not consistent discriminators on this cohort. The reversal is driven by a parent vessel almost free of oscillation in both members of pair 3, which pushes the ratio of two very small numbers into a regime where the sign is dominated by numerical noise rather than by anatomy. This underlines a practical limit of dome-to-parent normalisation for indices that take values close to zero on the parent vessel: the TAWSS, RRT, PSWSS and ΔWSS ratios remain robust because their parent-vessel denominators are well separated from zero, while the OSI and ECAP ratios are not.

Raising the LSA threshold from 0.4 Pa to 1.5 Pa brings the dome low-shear area into a consistent paired pattern: the growing aneurysm has a larger dome low-shear area than its stable counterpart in every pair (Table 3.10). At the conventional 0.4 Pa threshold, two of the three growing aneurysms (0207, 0209) already exhibit dome TAWSS values close to the threshold, so the 0.4 Pa indicator saturates and the contrast with the stable control is lost. Raising the threshold restores sensitivity. This threshold dependence is in line with the broader rupture-risk literature, in which no single universal LSA cut-off has

been established [37]: the appropriate threshold depends on the absolute shear regime of the parent artery, and a fixed 0.4 Pa value chosen for ICA-scale flows may understate the low-shear footprint on smaller or lower-flow geometries.

4.3 Limitations of This Study

The primary limitation is sample size. With three growing and three stable cases, the study is severely underpowered for group-level hypothesis testing. The consistency analysis across the three matched pairs is the strongest inference available at this scale, and even that is sensitive to the choice of which index is examined. Any quantitative claim about the direction or magnitude of hemodynamic differences should be treated as hypothesis-generating rather than confirmatory.

The rigid wall assumption eliminates fluid-structure interaction. Wall compliance alters near-wall hemodynamics, with the strongest effects on near-dome recirculation zones. A fluid-structure interaction study of 101 sidewall aneurysms reported sac-averaged OSI changes ranging from -36% to $+674\%$ relative to the rigid-wall value [43]. All six cases have relatively small sacs, which limits the expected compliance effect, but it cannot be excluded entirely without FSI simulations.

Blood was modeled as a Newtonian fluid with constant viscosity. Sensitivity analyses indicate that the viscosity model does not alter the characterization of intra-aneurysmal flow patterns [28], and any residual non-Newtonian effects are small relative to the dominant uncertainties in patient-specific models, though this simplification has been questioned for intracranial aneurysm flows [29]. In regions of very low shear, such as the dome stagnation zone, the shear-thinning characteristics of blood would increase the apparent viscosity and reduce local WSS slightly below the Newtonian value.

A population-averaged ICA waveform was used for all six cases. Patient-specific flow measurements were not available. Inlet waveform prescription is a major source of uncertainty in patient-specific hemodynamics [38]. Using a shared waveform introduces a shared bias in the absolute WSS values across the cohort and means that flow-rate differences between patients are not captured, but it reduces the confounding effect of inter-patient flow-rate differences on within-pair comparisons. All six cases share the same inflow waveform, so the dome-to-parent ratios reported here reflect geometric differences between the aneurysm sacs rather than flow-rate differences.

The geometric models were drawn from the VMR and were not re-segmented. Any segmentation error or smoothing applied during the original VMR processing propagates directly into the computed hemodynamic indices. No validation against in-vivo measurements was performed for the simulation results.

The volumetric meshes were reused from the VMR rather than generated and converged for this study (Section 2.2.2). They range from about 5 000 to 23 000 nodes, which is at the lower end of, or below, the 10^5 element range commonly cited for patient-specific aneurysm CFD [38]. No independent mesh-convergence study was carried out: the meshes are accepted as published because they are the ones the SimVascular team distributes with each case and used for their own reference simulations. Boundary-layer WSS is the quantity most sensitive to near-wall resolution, so the absolute TAWSS and OSI values

reported here should be read with this caveat in mind. The within-pair and dome-to-parent analyses are more robust to this limitation than the absolute values, because every case is meshed by the same published pipeline and the normalisation cancels a shared resolution bias to first order, but a dedicated convergence study on at least one case would be needed to bound the residual mesh-induced error.

The hybrid PINN was evaluated against the same FEM solution that provided its supervision. The accuracy metrics in Section 3.4 are computed on the full voxel grid, the bulk of which is unsupervised for the sparse budgets, so the sparse runs do test spatial generalisation to unseen locations within the field. They do not, however, constitute a fully held-out evaluation: no unseen cardiac cycle, no unseen patient geometry, and no in-vivo reference were used to test the hybrid network. The $N_s = \text{all}$ configuration in particular is best understood as a physics-regularised interpolation of the FEM field rather than an independent prediction. Quantifying how the hybrid network generalises to a geometry it was not trained on, for instance through the transfer-learning protocol suggested in Section 4.1.4, is left to future work.

4.4 Perspectives for Future Work

Several extensions would strengthen the clinical and methodological conclusions of this work. A larger cohort, ideally drawn from the same VMR longitudinal subset, would provide sufficient power for formal statistical testing and allow subgroup analysis by aneurysm location. The three pairs studied here already suggest that the dome-to-parent TAWSS ratio is a robust discriminator, but confirming this across a broader population requires more cases.

On the modeling side, patient-specific inlet waveforms derived from phase-contrast MRI measurements would replace the shared population-averaged flow rate and reduce one of the main sources of uncertainty [38]. Introducing FSI coupling, at least for cases with large or thin-walled sacs, would allow the effect of wall compliance to be quantified.

For the PINN methodology, the most impactful improvements are the hard-constrained no-slip boundary, per-outlet integral flux constraints, and near-wall importance sampling described in Section 4.1.4. Implementing these changes and repeating the accuracy comparison would clarify how much of the remaining error is architectural and how much is addressable without a fundamentally different network design. A systematic study of the supervised sample budget, extended to all six cases rather than only case 0203, would provide better guidance on the minimum data requirement for the hybrid approach.

Finally, comparing the two simulation approaches on a morphological index such as maximum aneurysm diameter or neck ratio, computed from the same VMR surface models, would place the hemodynamic discriminators in context alongside the clinical predictors already integrated into risk scores such as PHASES [4].

Conclusion

Intracranial aneurysms are a common incidental finding with rare but devastating rupture consequences, and clinical risk scores still rely on morphological and demographic variables that overlook the patient-specific blood-flow patterns inside the sac. Against this background, this thesis pursued two complementary objectives on a matched cohort of six patient-specific aneurysm geometries drawn from the Vascular Model Repository: identifying hemodynamic indices that separate growing from stable aneurysms, and assessing whether a physics-informed neural network can stand in for a reference finite element simulation of the same flow.

The first objective produced a clear take-home message. Absolute wall-averaged shear-stress metrics, taken in isolation, do not separate the growing from the stable group in this cohort, because the location-driven variability of the parent-artery baseline obscures any growth-related signal. Normalising the dome shear stress by the paired parent-vessel value removes that confound. Across all three matched pairs the growing aneurysm consistently exhibits a more shielded dome, that is, a dome that receives proportionally less shear than its own parent artery. Several closely related dome-to-parent indicators, built on the peak-systolic and systolic-to-diastolic WSS or on a higher low-shear-area threshold, reinforce the same shielding pattern. This is consistent with the low-WSS hypothesis for aneurysm wall degeneration and provides cross-sectional matched-pair evidence that relative dome shielding tracks growth status as well as the better-studied rupture risk.

The second objective produced a more nuanced message about the practicality of PINNs in this application. A pure pulsatile PINN trained on each geometry reproduces the dominant flow topology but systematically under-resolves the near-wall boundary layer and over-diffuses the parent-artery jet, leaving a large residual error on the wall-averaged TAWSS and OSI. Despite this error, the predicted case-to-case TAWSS ranking remains highly correlated with the FEM reference, indicating that the residual is dominated by a near-uniform scaling offset rather than by a random reshuffling of the cases. The training cost is also markedly higher than the corresponding FEM solve for all but the densest mesh. Augmenting the physics losses with a sparse set of FEM reference samples closes most of the accuracy gap, and a data-sparsity sweep on a single case shows that most of the gain is already achieved with a small fraction of the FEM mesh. The hybrid approach is therefore the more promising direction whenever even a limited amount of reference data is available, while the pure PINN remains a useful conceptual tool rather than a quantitative surrogate at this level of accuracy.

Several extensions naturally follow. A larger cohort, ideally drawn from the same longitudinal subset of the VMR, would move the consistency analysis from hypothesis-generating

to formally testable, and could clarify whether the dome-to-parent shielding pattern generalises across locations. On the modelling side, patient-specific inlet waveforms from phase-contrast MRI and FSI coupling for the larger sacs would address the two main sources of uncertainty in the present setup. For the PINN methodology, hard-constrained no-slip walls, per-outlet integral flux conditions, and near-wall importance sampling are the most direct architectural avenues to reduce the residual error. The hybrid framework is particularly promising for data-assimilation settings, where the supervised inputs would come from 4D-flow MRI rather than from FEM, allowing a physically consistent hemodynamic field to be reconstructed from sparse in-vivo measurements. Such directions position PINNs not as a replacement for FEM, but as a complement that can integrate simulation and measurement in ways that neither approach can achieve alone.

Bibliography

- [1] National Heart, Lung, and Blood Institute, “Cerebral aneurysm diagram,” Wikimedia Commons, 2011, public domain (US government work). Accessed May 2026. [Online]. Available: https://commons.wikimedia.org/wiki/File:Cerebral_aneurysm_NIH.jpg
- [2] R. D. Perrone, A. M. Malek, and T. Watnick, “Vascular complications in autosomal dominant polycystic kidney disease,” *Nature Reviews Nephrology*, vol. 11, no. 10, pp. 589–598, 2015, figure 6 used pending permission from Nature Publishing Group.
- [3] A. Updegrove, N. M. Wilson, J. Merkow, H. Lan, A. L. Marsden, and S. C. Shadden, “SimVascular: an open source pipeline for cardiovascular simulation,” *Annals of Biomedical Engineering*, vol. 45, no. 3, pp. 525–541, 2017.
- [4] J. P. Greving, M. J. H. Wermer, R. D. Brown, A. Morita, S. Juvela, M. Yonekura, T. Ishibashi, J. C. Torner, T. Nakayama, G. J. E. Rinkel, and A. Algra, “Development of the PHASES score for prediction of risk of rupture of intracranial aneurysms: a pooled analysis of six prospective cohort studies,” *The Lancet Neurology*, vol. 13, no. 1, pp. 59–66, 2014.
- [5] A. Ziebart, J. Dremel, S. Hetjens, D. J. Nieuwkamp, F. H. H. Linn, N. Etminan, and G. J. E. Rinkel, “Case fatality and functional outcome after spontaneous subarachnoid haemorrhage: a systematic review and meta-analysis of time trends and regional variations in population-based studies,” *European Stroke Journal*, vol. 9, no. 3, pp. 555–565, 2024.
- [6] M. H. M. Vlak, A. Algra, R. Brandenburg, and G. J. E. Rinkel, “Prevalence of unruptured intracranial aneurysms, with emphasis on sex, age, comorbidity, country, and time period: a systematic review and meta-analysis,” *The Lancet Neurology*, vol. 10, no. 7, pp. 626–636, 2011.
- [7] J. R. Cebral, M. A. Castro, J. E. Burgess, R. S. Pergolizzi, M. J. Sheridan, and C. M. Putman, “Characterization of cerebral aneurysms for assessing risk of rupture by using patient-specific computational hemodynamics models,” *American Journal of Neuroradiology*, vol. 26, no. 10, pp. 2550–2559, 2005.
- [8] D. M. Sforza, C. M. Putman, and J. R. Cebral, “Computational fluid dynamics in brain aneurysms,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 28, no. 6-7, pp. 801–808, 2012.
- [9] G. Zhou, Y. Zhu, Y. Yin, M. Su, and M. Li, “Association of wall shear stress with intracranial aneurysm rupture: systematic review and meta-analysis,” *Scientific Reports*, vol. 7, p. 5394, 2017.

- [10] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [11] S. Cuomo, V. S. D. Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, “Scientific machine learning through physics-informed neural networks: where we are and what’s next,” *Journal of Scientific Computing*, vol. 92, no. 3, p. 88, 2022.
- [12] G. Kissas, Y. Yang, E. Hwuang, W. R. Witschey, J. A. Detre, and P. Perdikaris, “Machine learning in cardiovascular flows modeling: predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks,” *Computer Methods in Applied Mechanics and Engineering*, vol. 358, p. 112623, 2020.
- [13] A. L. Marsden, N. M. Wilson, S. C. Shadden, M. Pfaller *et al.*, “Vascular model repository,” <https://www.vascularmodel.com>, 2023, open-source database of patient-specific cardiovascular models, Stanford University / UC Berkeley.
- [14] NVIDIA Corporation, “NVIDIA PhysicsNeMo: physics-informed machine learning framework,” <https://github.com/NVIDIA/physicsnemo>, 2024, formerly NVIDIA Modulus. Open-source framework for physics-informed and physics-constrained deep learning.
- [15] N. K. de Rooij, F. H. H. Linn, J. A. van der Plas, A. Algra, and G. J. E. Rinkel, “Incidence of subarachnoid haemorrhage: a systematic review with emphasis on region, age, gender and time trends,” *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 78, no. 12, pp. 1365–1372, 2007.
- [16] P. Texakalidis, A. Sweid, N. Mouchtouris, E. C. Peterson, C. Sioka, L. Rangel-Castilla, J. Reavey-Cantwell, and P. Jabbour, “Aneurysm formation, growth, and rupture: the biology and physics of cerebral aneurysms,” *World Neurosurgery*, vol. 130, pp. 277–284, 2019.
- [17] G. J. E. Rinkel, M. Djibuti, A. Algra, and J. van Gijn, “Prevalence and risk of rupture of intracranial aneurysms: a systematic review,” *Stroke*, vol. 29, no. 1, pp. 251–256, 1998.
- [18] OpenStax College, “Arteries of the brain,” Wikimedia Commons, 2013, licensed under CC BY 3.0. Accessed May 2026. [Online]. Available: https://commons.wikimedia.org/wiki/File:2123_Arteries_of_the_Brain.jpg
- [19] Rhcastilhos, “Circle of Willis (annotated english diagram),” Wikimedia Commons, 2007, licensed under CC BY-SA 3.0. Accessed May 2026. [Online]. Available: https://commons.wikimedia.org/wiki/File:Circle_of_Willis_en.svg
- [20] D. O. Wiebers, J. P. Whisnant, J. Huston, I. Meissner, R. D. Brown, D. G. Piepgras, G. S. Forbes, K. Thielen, D. Nichols, W. M. O’Fallon, J. Peacock, L. Jaeger, N. F. Kassell, G. L. Kongable-Beckman, and J. C. Torner, “Unruptured intracranial aneurysms: natural history, clinical outcome, and risks of surgical and endovascular treatment,” *The Lancet*, vol. 362, no. 9378, pp. 103–110, 2003.

- [21] H. E. Westerlaan, J. M. C. van Dijk, M. C. Jansen-van der Weide, J. C. de Groot, R. J. M. Groen, J. J. A. Mooij, and M. Oudkerk, "Intracranial aneurysms in patients with subarachnoid hemorrhage: CT angiography as a primary examination tool for diagnosis – systematic review and meta-analysis," *Radiology*, vol. 258, no. 1, pp. 134–145, 2011.
- [22] A. M. H. Sailer, B. A. J. M. Wagemans, P. J. Nelemans, R. de Graaf, and W. H. van Zwam, "Diagnosing intracranial aneurysms with MR angiography: systematic review and meta-analysis," *Stroke*, vol. 45, no. 1, pp. 119–126, 2014.
- [23] M. Uricchio, S. Gupta, N. Jakowenko, M. Levito, N. Vu, J. Doucette, A. Liew, S. Papatheodorou, A. M. Khawaja, L. S. Aglio, M. A. Aziz-Sultan, H. Zaidi, T. R. Smith, and R. A. Mekary, "Computed tomography angiography versus digital subtraction angiography for postclipping aneurysm obliteration detection: A meta-analysis," *Stroke*, vol. 50, no. 2, pp. 381–388, 2019.
- [24] A. Molyneux, R. Kerr, I. Stratton, P. Sandercock, M. Clarke, J. Shrimpton, and R. Holman, "International Subarachnoid Aneurysm Trial (ISAT) of neurosurgical clipping versus endovascular coiling in 2143 patients with ruptured intracranial aneurysms: a randomised trial," *The Lancet*, vol. 360, no. 9342, pp. 1267–1274, 2002.
- [25] W. Brinjikji, M. H. Murad, G. Lanzino, H. J. Cloft, and D. F. Kallmes, "Endovascular treatment of intracranial aneurysms with flow diverters: a meta-analysis," *Stroke*, vol. 44, no. 2, pp. 442–447, 2013.
- [26] O. K. Baskurt and H. J. Meiselman, "Blood rheology and hemodynamics," *Seminars in Thrombosis and Hemostasis*, vol. 29, no. 5, pp. 435–450, 2003.
- [27] A. S. Popel and P. C. Johnson, "Microcirculation and hemorheology," *Annual Review of Fluid Mechanics*, vol. 37, pp. 43–69, 2005.
- [28] J. R. Cebal, M. A. Castro, S. Appanaboyina, C. M. Putman, D. Millan, and A. F. Frangi, "Efficient pipeline for image-based patient-specific analysis of cerebral aneurysm hemodynamics: technique and sensitivity," *IEEE Transactions on Medical Imaging*, vol. 24, no. 4, pp. 457–467, 2005.
- [29] K. M. Saqr, S. Rashad, S. Tupin, K. Niizuma, T. Hassan, T. Tominaga, and M. Ohta, "What does computational fluid dynamics tell us about intracranial aneurysms? A meta-analysis and critical review," *Journal of Cerebral Blood Flow & Metabolism*, vol. 40, no. 5, pp. 1021–1039, 2020.
- [30] J. R. Cebal, F. Mut, J. Weir, and C. M. Putman, "Association of hemodynamic characteristics and cerebral aneurysm rupture," *American Journal of Neuroradiology*, vol. 32, no. 2, pp. 264–270, 2011.
- [31] T. Suzuki, C. J. Stapleton, M. J. Koch, K. Tanaka, S. Fujimura, M. Yamamoto, Y. Fujii, Y. Murayama, and A. B. Patel, "Decreased wall shear stress at high-pressure areas predicts the rupture point in ruptured intracranial aneurysms," *Journal of Neurosurgery*, vol. 132, no. 4, pp. 1116–1122, 2020.
- [32] D. N. Ku, D. P. Giddens, C. K. Zarins, and S. Glagov, "Pulsatile flow and atherosclerosis in the human carotid bifurcation: positive correlation between plaque location and low and oscillating shear stress," *Arteriosclerosis*, vol. 5, no. 3, pp. 293–302, 1985.

- [33] O. Mutlu, H. E. Salman, H. Al-Thani, A. El-Menyar, U. A. Qidwai, and H. C. Yalcin, “How does hemodynamics affect rupture tissue mechanics in abdominal aortic aneurysm: focus on wall shear stress derived parameters, time-averaged wall shear stress, oscillatory shear index, endothelial cell activation potential, and relative residence time,” *Computers in Biology and Medicine*, vol. 154, p. 106609, 2023.
- [34] T. Nishiwaki, T. Ikedo, Y. Kushi, K. Shimonaga, H. Kobayashi, T. Itazu, R. Otsuka *et al.*, “Hemodynamic differences determining rupture and non-rupture in middle cerebral aneurysms after growth,” *PLOS ONE*, vol. 19, no. 8, p. e0307495, 2024.
- [35] H. A. Himburg, D. M. Grzybowski, A. L. Hazel, J. A. LaMack, X.-M. Li, and M. H. Friedman, “Spatial comparison between wall shear stress measures and porcine arterial endothelial permeability,” *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 286, no. 5, pp. H1916–H1922, 2004.
- [36] P. Di Achille, G. Tellides, C. A. Figueroa, and J. D. Humphrey, “A haemodynamic predictor of intraluminal thrombus formation in abdominal aortic aneurysms,” *Proceedings of the Royal Society A*, vol. 470, no. 2172, p. 20140163, 2014.
- [37] J. R. Cebral, F. Mut, J. Weir, and C. Putman, “Quantitative characterization of the hemodynamic environment in ruptured and unruptured brain aneurysms,” *American Journal of Neuroradiology*, vol. 32, no. 1, pp. 145–151, 2011.
- [38] D. A. Steinman and V. M. Pereira, “How patient specific are patient-specific computational models of cerebral aneurysms? an overview of sources of error and variability,” *Neurosurgical Focus*, vol. 47, no. 1, p. E14, 2019.
- [39] H. Wei, G. Wang, Q. Tian, C. Liu, W. Han, J. Wang, P. He, and M. Li, “Low shear stress induces macrophage infiltration and aggravates aneurysm wall inflammation via CCL7/CCR1/TAK1/NF- κ B axis,” *Cellular Signalling*, vol. 117, p. 111122, 2024.
- [40] F. J. Detmer, B. J. Chung, C. Jimenez, F. Hamzei-Sichani, D. Kallmes, C. Putman, and J. R. Cebral, “Associations of hemodynamics, morphology, and patient characteristics with aneurysm rupture stratified by aneurysm location,” *Neuroradiology*, vol. 61, no. 3, pp. 275–284, 2019.
- [41] M. Dabagh, P. Nair, J. Gounley, D. Frakes, L. F. Gonzalez, and A. Randles, “Hemodynamic and morphological characteristics of a growing cerebral aneurysm,” *Neurosurgical Focus*, vol. 47, no. 1, p. E13, 2019.
- [42] Y. Karnam, F. Mut, A. M. Robertson, N. Kaneko, and J. R. Cebral, “Competing pathways of intracranial aneurysm growth: linking regional growth distribution and hemodynamics,” *Journal of Neurosurgery*, vol. 142, no. 6, pp. 1741–1750, 2025.
- [43] A. Goetz, P. Jeken-Rico, U. Pelissier, Y. Chau, J. Sédat, and E. Hachem, “AnXplore: a comprehensive fluid-structure interaction study of 101 intracranial aneurysms,” *Frontiers in Bioengineering and Biotechnology*, vol. 12, p. 1433811, 2024.
- [44] D. A. Steinman, Y. Hoi, P. Fahy, L. Morris, M. T. Walsh, N. Aristokleous, A. S. Anayiotos, Y. Papaharilaou, A. Arzani, S. C. Shadden, P. Berg, G. Janiga, J. Bols, P. Segers, N. W. Bressloff, M. Cibis, F. H. Gijsen, S. Cito, J. Pallàres, L. D. Browne, J. A. Costelloe, A. G. Lynch, J. Degroote, J. Vierendeels, W. Fu, A. Qiao, S. Hodis,

- D. F. Kallmes, H. Kalsi, Q. Long, V. O. Kheyfets, E. A. Finol, K. Kono, A. M. Malek, A. Lauric, P. G. Menon, K. Pekkan, M. E. Moghadam, A. L. Marsden, M. Oshima, K. Katagiri, V. Peiffer, Y. Mohamied, S. J. Sherwin, J. Schaller, L. Goubergrits, G. Usera, M. Mendina, K. Valen-Sendstad, D. F. Habets, J. Xiang, H. Meng, Y. Yu, G. E. Karniadakis, N. Shaffer, and F. Loth, “Variability of computational fluid dynamics solutions for pressure and flow in a giant aneurysm: the ASME 2012 Summer Bioengineering Conference CFD Challenge,” *Journal of Biomechanical Engineering*, vol. 135, no. 2, p. 021016, 2013.
- [45] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [46] X. Jin, S. Cai, H. Li, and G. E. Karniadakis, “NSFnets (Navier-Stokes flow nets): physics-informed neural networks for the incompressible Navier-Stokes equations,” *Journal of Computational Physics*, vol. 426, p. 109951, 2021.
- [47] S. Wang, Y. Teng, and P. Perdikaris, “Understanding and mitigating gradient flow pathologies in physics-informed neural networks,” *SIAM Journal on Scientific Computing*, vol. 43, no. 5, pp. A3055–A3081, 2021.
- [48] R. Bischof and M. Kraus, “Multi-objective loss balancing for physics-informed deep learning,” arXiv preprint arXiv:2110.09813, 2021.
- [49] S. Elfving, E. Uchibe, and K. Doya, “Sigmoid-weighted linear units for neural network function approximation in reinforcement learning,” *Neural Networks*, vol. 107, pp. 3–11, 2018.
- [50] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng, “Fourier features let networks learn high frequency functions in low dimensional domains,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 7537–7547.
- [51] S. Wang, X. Yu, and P. Perdikaris, “When and why PINNs fail to train: a neural tangent kernel perspective,” *Journal of Computational Physics*, vol. 449, p. 110768, 2022.
- [52] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1988.
- [53] S. Wang, S. Sankaran, and P. Perdikaris, “Respecting causality for training physics-informed neural networks,” *Computer Methods in Applied Mechanics and Engineering*, vol. 421, p. 116813, 2024.
- [54] DeepL SE, “DeepL translator,” <https://www.deepl.com>, 2026, neural machine translation service used for language assistance.
- [55] Anthropic, “Claude,” <https://www.anthropic.com/claude>, 2026, large language model used for text rephrasing and writing assistance.

Appendix A

Use of Artificial Intelligence

Artificial intelligence tools were used in the context of this thesis solely to assist with non-scientific tasks. DeepL [54] and Claude by Anthropic [55] were used to rephrase and improve the flow of certain text passages. Under no circumstances was artificial intelligence employed to generate scientific or technical content, to interpret results, or to draw conclusions.

Appendix B

Code and Configuration Details

The source code for this thesis, including PhysicsNeMo PINN scripts, Hydra configurations, SimVascular project files, SLURM job scripts, HPC environment setup, ..., is publicly available on GitHub:

`https://github.com/EBarette/hemodynamic-characterization-intracranial-aneurysms`

Appendix C

Supplementary Hemodynamic Data

This appendix gathers the per-case hemodynamic indices used to compute the group-level statistics in Chapter 3, Section 3.5. Tables are grouped by sampling region: full-wall area-averages (Table C.1), dome and parent-vessel area-averages with the dome-to-parent TAWSS ratio (Table C.2), and multi-threshold low-shear area (Table C.3).

Table C.1: Per-case wall area-averaged hemodynamic indices. TAWSS, PSWSS, Δ WSS in Pa. OSI dimensionless. RRT and ECAP in Pa^{-1} . WSSG in Pa/cm . $\text{LSA}^{0.4}$ is the wall-area fraction with $\text{TAWSS} < 0.4 \text{ Pa}$.

Case	Status	TAWSS	OSI	RRT	ECAP	WSSG	PSWSS	Δ WSS	$\text{LSA}^{0.4}$
0203	Growing	1.34	0.0016	3.84	0.0068	7.18	2.24	1.58	0.106
0204	Stable	2.76	0.0025	2.33	0.0090	10.47	4.65	3.24	0.029
0207	Growing	6.29	0.0008	6.37	0.0040	27.84	10.47	6.81	0.004
0208	Stable	7.17	0.0045	0.26	0.0050	33.71	11.92	7.84	0.004
0209	Growing	4.75	0.0021	0.38	0.0057	16.34	8.04	5.38	0.010
0210	Stable	4.26	0.0041	0.58	0.0208	11.59	7.05	4.57	0.033

Table C.2: Per-case dome (subscript d) and parent-vessel (subscript p) area-averages, with the dome-to-parent TAWSS ratio r_{TAWSS} . Units as in Table C.1.

Case	Status	TAWSS_d (Pa)	TAWSS_p (Pa)	r_{TAWSS}	OSI_d	OSI_p	PSWSS_d (Pa)	WSSG_d (Pa/cm)
0203	Growing	0.54	2.05	0.26	0.0040	2.6×10^{-5}	0.95	1.19
0204	Stable	1.87	2.38	0.79	0.0069	6.4×10^{-4}	3.35	9.06
0207	Growing	2.45	6.08	0.40	0.0029	1.8×10^{-5}	4.49	22.88
0208	Stable	4.28	7.38	0.58	0.0174	4.5×10^{-3}	7.84	28.47
0209	Growing	3.23	8.35	0.39	0.0049	1.4×10^{-4}	5.81	11.94
0210	Stable	2.73	4.70	0.58	5.5×10^{-4}	2.6×10^{-6}	4.68	9.64

Table C.3: Per-case low-shear area fraction at three TAWSS thresholds. Subscript d = dome region, no subscript = full wall.

Case	Status	LSA ^{0.4}	LSA ^{1.0}	LSA ^{1.5}	LSA _d ^{0.4}	LSA _d ^{1.0}	LSA _d ^{1.5}
0203	Growing	0.106	0.388	0.566	0.256	0.942	0.990
0204	Stable	0.029	0.114	0.251	0.082	0.256	0.405
0207	Growing	0.004	0.044	0.089	0.015	0.218	0.444
0208	Stable	0.004	0.030	0.061	0.000	0.054	0.150
0209	Growing	0.010	0.037	0.080	0.003	0.019	0.059
0210	Stable	0.033	0.087	0.123	0.004	0.025	0.044

Table C.4: Supplementary group-level statistics for secondary hemodynamic indices not reported in Table 3.9. Mean $\pm$ SD and Cohen's d (growing – stable). Units: WSSG in Pa/cm. PSWSS and Δ WSS in Pa. PSWSS/TAWSS and dome-to-parent LSA ratios dimensionless.

Index	Growing	Stable	d
<i>Dome area-averages</i>			
WSSG _d (Pa/cm)	12.0 $\pm$ 10.8	15.7 $\pm$ 11.0	−0.34
PSWSS _d (Pa)	3.75 $\pm$ 2.51	5.29 $\pm$ 2.30	−0.64
Δ WSS _d (Pa)	2.94 $\pm$ 1.76	3.93 $\pm$ 1.72	−0.57
PSWSS/TAWSS _d	1.796 $\pm$ 0.041	1.778 $\pm$ 0.059	+0.35
<i>Parent-vessel area-averages</i>			
WSSG _p (Pa/cm)	14.5 $\pm$ 10.6	16.4 $\pm$ 15.9	−0.14
PSWSS _p (Pa)	9.02 $\pm$ 5.24	8.00 $\pm$ 4.02	+0.22
Δ WSS _p (Pa)	5.79 $\pm$ 3.32	5.18 $\pm$ 2.45	+0.21
<i>Dome-to-parent LSA ratios</i>			
$r_{\text{LSA}^{1.0}}$	14.7 $\pm$ 17.5	1.20 $\pm$ 0.62	+1.09
$r_{\text{LSA}^{1.5}}$	4.38 $\pm$ 3.31	1.64 $\pm$ 0.14	+1.17

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain

Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | www.uclouvain.be/epl