

©Copyright 2026

Bahetihazi Maidu

Physics-Informed and Data-Driven Inference of Cardiovascular Hemodynamics, Thrombosis Risk, and Model Parameters

Bahetihazi Maidu

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Juan Carlos del Álamo, Chair

Alberto Aliseda

Patrick Boyle

Nazem Akoum

Program Authorized to Offer Degree:
Mechanical Engineering

University of Washington

Abstract

Physics-Informed and Data-Driven Inference of Cardiovascular Hemodynamics, Thrombosis Risk, and Model Parameters

Bahetihazi Maidu

Chair of the Supervisory Committee:
Juan Carlos del Álamo
Department of Mechanical Engineering

Cardiovascular hemodynamics are governed by blood transport through complex cardiac chambers, veins, and arteries. Quantities such as velocity, pressure, residence time, and vascular resistance are closely linked to cardiac function, thrombosis risk, and clinical decision-making, yet many are difficult or invasive to measure directly in vivo. Medical imaging provides complementary but incomplete information, while computational fluid dynamics (CFD) can simulate realistic blood flow but often requires accurate boundary conditions, additional measurements, and substantial computational resources. These limitations motivate approaches that combine physical laws, medical imaging, and data-driven inference to recover cardiovascular hemodynamics from sparse and indirect measurements.

This thesis develops physics-informed neural network frameworks and data-driven methods for patient-specific inference of cardiovascular flow, thrombosis risk, and model parameters. The central idea is to embed governing physical principles, including the incompressible Navier–Stokes equations, indicator dilution theory, and lumped-parameter models, into machine learning and model-based inference frameworks. These constraints enable reconstruction of hemodynamic quantities that are not directly measurable from standard medical imaging, without requiring comprehensive and computationally expensive direct numerical simulations.

A major focus of this work is AI-VFM, a physics-informed framework for left ventricular (LV) vector flow mapping (VFM) from color-Doppler echocardiography. AI-VFM uses mass conservation, momentum balance, and boundary conditions to infer 2D intraventricular velocity and pressure fields from

standard Doppler acquisitions. Validation against CFD data and application to clinical echocardiographic acquisitions demonstrate its ability to recover velocity and pressure fields, perform phase unwrapping, fill missing data, and generate super-resolved hemodynamic maps. This framework is further extended to time-resolved 3D reconstruction using triplane color-Doppler imaging, where sparse velocity measurements from three intersecting planes are used to infer volumetric LV velocity and pressure fields. Validation against CFD simulations shows that 3D AI-VFM recovers dominant flow features, including vortical structures, velocity patterns, and pressure fluctuations.

This thesis also investigates left atrial flow and thrombosis risk using contrast-enhanced computed tomography (CT). In atrial fibrillation, blood stasis in the left atrial appendage (LAA) contributes to thrombus formation and stroke risk, but blood stasis remains difficult to quantify non-invasively. To address this challenge, a physics-informed neural network is used to reconstruct LA flow and residence time fields from contrast dynamics, while an indicator dilution theory framework provides rapid residence time estimation from contrast concentration curves. Benchmarking against patient-specific CFD simulations demonstrates that contrast-enhanced 4D CT contains sufficient transport information to infer atrial flow, LAA stasis, and elevated thrombosis-related risk.

Finally, this physics-informed inference paradigm is extended to reduced-order modeling of post-Norwood single-ventricle circulation. A lumped-parameter-model-informed neural network is developed to infer patient-specific hemodynamic states and model parameters from clinically measurable quantities. Using synthetic data from a physiologically derived Norwood circulation model, the framework recovers flow, pressure, and volume waveforms while estimating sensitive vascular parameters such as pulmonary vascular resistance and shunt resistance.

Together, these methods advance a unified framework for cardiovascular hemodynamic inference from sparse, noisy, and indirect clinical data. By combining medical imaging, physics-informed neural networks, indicator dilution theory, and lumped-parameter modeling, this work demonstrates how physical constraints can transform limited measurements into quantitative biomarkers of flow, pressure, residence time, thrombosis risk, and patient-specific cardiovascular parameters.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	ix
Chapter 1: Introduction	1
1.1 Cardiovascular system overview	3
1.2 Heart anatomy and cardiac cycle	4
1.3 Medical imaging for cardiac flow	7
1.3.1 Color-Doppler ultrasound	7
1.3.2 Computed tomography	8
1.3.3 Magnetic resonance imaging	9
1.4 Cardiovascular fluid dynamics	9
1.4.1 Left ventricular flow	9
1.4.2 Left atrial flow	11
1.4.3 Capillary and circulatory flows	11
1.5 Clinical implication of intracardiac flows	12
1.6 Thesis outline	14
Chapter 2: Background	15
2.1 Image-based data-driven computational predictive modeling in cardiovascular flow . .	15
2.1.1 Computational fluid dynamics	15
2.1.2 Intraventricular vector flow mapping	17
2.1.3 Indicator dilution theory	18
2.1.4 Lumped-parameter modeling	20
2.1.5 Machine learning approaches	21
2.2 Physics-informed neural network	24
2.2.1 Methodology overview	24
2.2.2 State-of-the-art	26
2.2.3 Applications in cardiovascular flow	28

Chapter 3:	Super-resolution Left ventricular flow and pressure mapping by Navier-Stokes-informed neural network (AI-VFM)	30
3.1	Introduction	30
3.2	Materials and Methods	32
3.2.1	Governing Equations of LV flow	32
3.2.2	LV geometrical domain for AI-VFM	33
3.2.3	Color-Doppler Phase Unwrapping in AI-VFM	33
3.2.4	Neural Network Architecture and Super-Resolution Loss Function	35
3.2.5	Adaptive Loss Function Weight Adjustment	38
3.2.6	PINN runs	40
3.2.7	LV geometrical domain for CFD	40
3.2.8	CFD runs	42
3.2.9	Validation and Sensitivity Analysis on CFD-generated LV flow	42
3.2.10	Echocardiographic Imaging	43
3.2.11	Reference (Vanilla) VFM and Phase-Unwrapping Methods	44
3.3	Results	45
3.3.1	AI-VFM Validation on CFD-generated LV flow	45
3.3.2	Application to Clinical Acquisitions	52
3.3.3	Phase unwrapping in AI-VFM	55
3.4	Discussion	57
3.5	Conclusion	62
3.6	Acknowledgements	62
3.7	Supporting Information	63
Chapter 4:	Time-resolved Three-dimensional AI-VFM by Triplane Color-Doppler Imaging	70
4.1	Introduction	70
4.2	Materials and Methods	71
4.2.1	Synthetic triplane Doppler generation	71
4.2.2	3D AI-VFM architecture & training strategy	73
4.3	Results	75
4.3.1	Flow and pressure reconstruction visualization and comparison	76
4.3.2	Point-wise and global comparisons against CFD ground-truth data	77
4.4	Discussion	80
Chapter 5:	Left atrial flow and thrombosis risk from 4D CT contrast dynamics by physics-informed neural networks and indicator dilution theory	82

5.1	Introduction	82
5.2	Materials and Methods	84
5.2.1	Image acquisition and segmentation	84
5.2.2	Computational fluid dynamics analysis	86
5.2.3	Physics-informed neural network	88
5.2.4	Indicator dilution theory	93
5.3	Results	98
5.3.1	Spatiotemporal Comparison of T_R Fields: CFD vs. PINN and IDT	100
5.3.2	Statistical Validation of Inferred LAA T_R	101
5.3.3	Temporal Dynamics of Inferred LAA T_R	104
5.3.4	Left Atrial Flow Reconstruction Using PINN	105
5.4	Discussion	107
5.5	Conclusion	112
5.6	Acknowledgments	113
5.7	Supporting Information	113
Chapter 6:	Patient-specific parameter inference in Norwood circulation by Lumped-Parameter-Model-informed neural network	122
6.1	Introduction	122
6.2	Materials and Method	124
6.2.1	Lumped-Parameter Model (LPM)	124
6.2.2	Synthetic Data Generation	126
6.2.3	Physics-Informed Neural Network	129
6.2.4	Parameter Identifiability Analysis	132
6.3	Results and Discussions	133
6.3.1	LPM Inference with Fixed Parameters	133
6.3.2	LPM Parameter Inference – Pulmonary Artery Resistance	136
6.3.3	Sensitivity Analysis	138
6.3.4	Conclusion	139
Chapter 7:	Conclusion	141
	Bibliography	144

LIST OF FIGURES

Figure Number	Page
1.1 Leonardo da Vinci’s drawings of human heart and lung, and an ox’s heart	2
1.2 Cardiovascular system and circulation overview	3
1.3 Human heart anatomy	4
1.4 Left ventricle pressure-volume loop	6
1.5 Medical imaging examples of color-Doppler, CT scan, and 4D flow MRI	8
3.1 AI-VFM flowchart	34
3.2 AI-VFM loss curves, loss weight evolution, and eigenvalues of NTK	39
3.3 LV geometrical domain for CFD	41
3.4 Velocity and pressure maps inferred by AI-VFM compared to ground-truth CFD data at early filling, late filling, and ejection	46
3.5 Scatter plots. Point-wise comparison of velocity and fluctuating pressure inferred by AI-VFM with ground-truth CFD data.	48
3.6 AI-VFM sensitivity analysis via CFD	51
3.7 Example of clinical application of AI-VFM in patient #1	53
3.8 Example of clinical application of AI-VFM in patient #1 to color-Doppler sequences with scan sector truncation near the LV boundary	54
3.9 Leave-one-out validation of AI-VFM on the clinical color-Doppler sequence of patient #1	56
3.10 Color-Doppler dealiasing of radial velocity v_r analysis in CFD model	57
3.11 Example of clinical application of AI-VFM in patient #2	64
3.12 Example of clinical application of AI-VFM in patient #3	65
3.13 Example of clinical application of AI-VFM in patient #1 to color-Doppler sequences with holes inside the LV domain	66
3.14 Velocity and pressure maps inferred by AI-VFM compared to ground-truth CFD data at early filling, late filling, and ejection. Data from two AI-VFM implementations, using PINNs trained with $Re = 4, 550$ and $Re = \infty$, are shown	67
3.15 Point-wise comparison of de-aliased radial velocity from One-step Laplacian method (left) and AI-VFM (right) with ground-truth CFD data	68
3.16 Point-wise comparison of velocity and fluctuating pressure values inferred by AI-VFM with CFD ground-truth data for different number of training iterations and batch sizes .	69

4.1	Synthetic triplane Doppler data acquisition on an idealized CFD data. Left panel depicts the triplane separated by 60° . Right panel shows a representative radial velocity component v_r snapshot during early filling for the triplane.	71
4.2	Flowchart of the proposed framework in this chapter. Spatiotemporal coordinates of the triplane Doppler data (top left) are used as inputs to the 3D AI-VFM (top right). After training, it reconstructs full volumetric flow (bottom left) and pressure fields (bottom right).	72
4.3	Velocity and vorticity maps inferred by 3D AI-VFM (bottom row) compared to CFD ground-truth data (top row) at early filling (left column), late filling (mid column), and ejection (right column). The blue volumes are the vortex regions computed by the Q-criterion ($3000/s^{-2}$ threshold).	75
4.4	Pressure fluctuation (iso-surface) maps overlay with velocity vectors (green arrows) inferred by 3D AI-VFM (bottom row) compared to CFD ground-truth data (top row) at early filling (left column), late filling (mid column), and ejection (right column).	76
4.5	Point-wise comparison of velocity and fluctuating pressure inferred by 3D AI-VFM with CFD ground-truth data.	77
4.6	Volume-averaged kinetic energy at each time instant comparison inferred by 3D AI-VFM with CFD ground-truth data.	78
4.7	Mean vorticity ($[s^{-1}]$) at each time instant comparison inferred by 3D AI-VFM with CFD ground-truth data.	79
4.8	Loss terms and weights vs. training iterations.	79
5.1	Patient-specific left atrial meshes used in Chapter 4 study	86
5.2	Overview of the study. (A) Patient-specific pulmonary vein (PV) contrast inflow profiles from 4D contrast CT are modeled with gamma-variate fits. (B) Contrast concentration maps obtained from CFD simulations. (C) Physics-informed residual adaptive neural network (PirateNet) framework: spatiotemporal inputs are encoded with Fourier features and passed through four residual blocks to infer contrast concentration, velocity, pressure, and residence time. Each block consists of three multilayer perceptrons (MLPs) and an adaptive skip connection step with a trainable parameter α initialized at 0. Training minimizes a composite loss including data, residence time inflow conditions, Navier-Stokes, and transport equations. (D) Indicator dilution theory (IDT) framework: chamber transport is modeled through an impulse response inferred from the PV inflow profiles and simulated contrast maps. Both inflow and response curves are parameterized with gamma-variate models, and residence time is obtained from the closed-form integral of the corresponding survival function (a hypergeometric function).	89

5.3	Contrast concentration (c) and residence time (T_R) maps comparison of case 3 (top) and case 5 (bottom) between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.	99
5.4	Comparison of LAA T_R by PirateNet and IDT with CFD ground-truth data. The top panels (A–F) display joint probability density functions of CFD residence times (x -axis) and IDT-inferred residence time (y -axis, filled black contours), $p(T_{R,CFD}, T_{R,IDT})$, or PirateNet-inferred residence time (y -axis, blue line contours), $p(T_{R,CFD}, T_{R,PirateNet})$. 15 cycles of residence time are pooled to display the joint pdfs. Panels (A) through (F) respectively represent cases 1 through 6. Each contour, from dark to light, contains 75%, 90%, 95%, and 99% of the data, respectively. The bottom right corner of each panel shows the correlation coefficients $R_{PirateNet}$ and R_{IDT} . The bottom panels (G–I) are scatter plots of mean LAA T_R from PirateNet and IDT vs. with mean T_R from CFD, averaged within the LAA over three different cardiac cycles: G), $4T \leq t \leq 5T$; H), $9T \leq t \leq 10T$; I) $14T \leq t \leq 15T$. Hollow and solid symbols respectively represent LAAT/TIA-negative and cases LAAT/TIA-positive cases.	102
5.5	Comparison of volume-averaged LAA residence time T_R inferred by the PirateNet (blue) and IDT (black) against CFD ground-truth (red). Panels (A)–(F) correspond to cases 1–6. Insets show representative training contrast concentration time curves at inlet and chamber locations (lighter/darker shades indicate PVs/LAA). For the PINN, the LAA volume-averaged contrast and the cross-section-averaged contrast at each PV inlet are shown. For IDT, the LAA volume-averaged contrast and the cross-section-averaged contrast pooled across the four PV inlets are shown. In each inset, only the time interval used to train the corresponding method is shown.	104
5.6	Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 3 (LAAT/TIA-negative, normal LA function), whereas right panels (B,D,F) show case 5 (LAAT/TIA-positive, impaired LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.0 in case 3 and 1.8 in case 5, respectively.	106
5.7	Loss curves for each loss component of a representative case 1 from PirateNet training.	114

5.8	Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 1 (LAAT/TIA-negative, normal LA function), whereas right panels (B,D,F) show case 2 (LAAT/TIA-negative, normal LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.5 in case 1 and 2.0 in case 2, respectively.	115
5.9	Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 4 (LAAT/TIA-positive, impaired LA function), whereas right panels (B,D,F) show case 6 (LAAT/TIA-positive, impaired LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.0 for both cases.	116
5.10	Contrast concentration (c) and residence time (T_R) maps comparison of case 1 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.	117
5.11	Contrast concentration (c) and residence time (T_R) maps comparison of case 2 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.	118

5.12	Contrast concentration (c) and residence time (T_R) maps comparison of case 4 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.	119
5.13	Contrast concentration (c) and residence time (T_R) maps comparison of case 6 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.	120
5.14	Violin plots of the probability density function with mean and median of the residence time T_R from CFD, PirateNet, and IDT inside the left atrial appendage are shown for the $N = 6$ studied cases.	121
6.1	Post-Norwood circulation model used to generate training and validation data in LPM-PINN.	124
6.2	Lumped-parameter-model-informed neural network training schematic	129
6.3	Flow rate, pressure, and volume inferred by LPM-PINN	134
6.4	Loss term weights vs. training iterations for the fixed parameter analysis in LPM-PINN	135
6.5	Eigenvalues of neural kernel associated with each loss term calculated at the last training iteration versus eigenvalue index in LPM-PINN	136
6.6	Flow rate, pressure, and volume inferred by LPM-PINN compared to ground-truth with all fixed params but trainable pulmonary artery resistance	137
6.7	Convergence of the inferred parameter vs. training iterations compared to the ground-truth value in LPM-PINN	138
6.8	Sensitivity analysis results in LPM-PINN	139

LIST OF TABLES

Table Number	Page
3.1 Characteristic scales used to non-dimensionalize inputs in Chapter 3 AI-VFM applications . .	37
5.1 Anatomical and functional parameters of the cases reported in Chapter 4 study	85
5.2 Global velocity comparison metrics between PINN and CFD for the six studied cases. PINN results are reported for both PirateNet and vanilla MLP architectures. ϵ_M represents the normalized magnitude discrepancy and ϵ_O is the global orientation error (in degrees). LAA residence time T_R correlations are shown for PINN vs. CFD and IDT vs. CFD.	103
6.1 List of subscripts in the lumped-parameter model.	125
6.2 List of lumped-parameter model parameters	127
6.3 List of lumped-parameter model initial conditions	127

ACKNOWLEDGMENTS

I want to thank Professor Juan Carlos del Álamo for his mentorship and guidance throughout my Ph.D. studies. Professor del Álamo has been my mentor and advisor since I began my Masters program at the University of California San Diego. He has always been an inspiring advisor, with the ability to come up with interesting research ideas and a sharp vision for important scientific directions. He has also shown great patience with the process of research discovery and progress, which I have deeply appreciated. In the lab, Professor del Álamo gives me and other trainees the freedom to manage our research progress and maintain a very healthy research-life balance. He is always calm, positive, humorous, and supportive, and he has set an example from which I have learned a great deal. I would also like to thank my committee members, Professor Alberto Aliseda, Professor Patrick Boyle, and Professor Nazem Akoum, for their valuable suggestions and recommendations on my research projects, as well as for years of collaboration and discussion during research meetings. I am also grateful to Professor Krithika Manohar and Professor Michelle DiBenedetto for their participation and support during my qualifying exam.

I sincerely thank Professor David Santillan and Professor Padmini Rangamani for serving on my advisory committee and supporting me during my time at UCSD. I would also like to thank Professor Mohammad Bahadori for his support and for providing valuable lessons on teaching and mentoring students in academia. I also want to extend my deepest appreciation to Professor Juan Lasheras for his wisdom, vision, and guidance when I first joined the del 'Alamo & Lasheras lab at UCSD. I am also thankful to the current and past members of the del 'Alamo lab, as well as my peers in the Mechanical Engineering departments at both UW and UCSD, for their contributions to my research projects and for their support throughout my Ph.D. journey, especially during challenging times. These include, but are certainly not limited to, Tu, Clarissa, Alex, Akshay, Saveer, Cathleen, Alejandro, Adithan, Manu, Amy, Antoni, Stephanie, Ernesto, Lorenzo, Obed, Hossein, and Xiaokang. I will miss our time spent in rushing coursework deadlines, conference travels, and gatherings.

I also want to extend my gratitude to all the co-authors of my conference presentations and journal publications, who come from many countries and institutions around the world, specifically Professor Pablo Martinez-Legazpi, Professor Oscar Flores, Professor Andrew M Kahn, Professor Javier Bermejo, Professor Manuel García-Villalba, and Professor Elliot McVeigh. My Ph.D. thesis would not have been possible without your contributions, advice, and collaboration. I would also like to thank my academic advisor, Wanwisa, in the UW ME Department. She has always been welcoming, kind, and responsive to my questions regarding my graduate study progress.

Moreover, I would like to thank my family and relatives back in my hometown for their unconditional support of me and my decisions. Without their support, I would not even have reached the point of continuing higher education, let alone pursuing a Ph.D. degree.

Finally, I want to thank my wife, Munira Kairat, who herself is almost at the finish line of her own Ph.D. journey. Your emotional support, constant kindness, and strength have made me more determined and have made this journey meaningful. I feel incredibly fortunate to have had you by my side throughout this journey, and I look forward to our future together, both in life and in our careers.

DEDICATION

To my Altai people and the land.

Chapter 1

INTRODUCTION

Human efforts to understand the cardiovascular system span thousands of years, beginning with early anatomical and physiological interpretations recorded in ancient medical texts and historical records. The earliest known documentation of vascular structures driven by a pulsatile force dates back to 1550 BC [1], while parallel traditions in other ancient civilizations offered formative perspectives on the heart and circulation [2–4]. Some early accounts were remarkably advanced, including the *Nei Jing*'s description around 2600 BC of continuous heart-driven blood flow in ancient China [5], and a 13th century description of pulmonary circulation [6] that contradicted Galenic physiology [7]. During the Renaissance, Leonardo da Vinci created detailed anatomical drawings of the human lung and heart vascular network and ox's heart, depicting its chambers, valves and major vessels with exceptional precision (Fig. 1.1). A decisive shift in physiological understanding occurred in the 17th century when circulation was demonstrated to operate as a closed loop propelled by the heart [8]. This seminal foundation enabled further developments, including early measurements of blood pressure in the 18th century and the identification of the cardiac conduction system in the 19th and early 20th centuries [9–11]. These milestones collectively mark the transition from speculative interpretations to a scientifically grounded understanding of cardiovascular system and physiology.

The 20th century witnessed rapid expansion in cardiovascular science driven by technology advances in measurement, imaging methods and biomedical engineering. The electrocardiograph developed in 1901 was used to systematically assess cardiac electrical activity and was applied in early forms of telemedicine [5]. Breakthroughs in pacemakers, defibrillation, valve prostheses, artificial hearts and ventricular assist devices reshaped therapeutic practice, while modalities such as echocardiography, computed tomography and magnetic resonance imaging opened new windows into cardiac anatomy and function. These developments highlighted the cardiovascular system as a complex, multiscale dynamical system, motivating continued efforts to understand how an organ weighing less than a kilogram can drive a vascular network spanning nearly 10^5 km and sustain billions of heartbeats over a human lifetime.




Figure 1.1: Leonardo da Vinci's drawings of (A) human heart and lung, and (B) an ox's heart. Adapted from [12].

In the remainder of this chapter, we introduce the core concepts relevant to this thesis. Section §1.1 provides an overview of the cardiovascular system. Section §1.2 then reviews cardiac anatomy with emphasis on chamber structure. Section §1.3 summarizes the medical imaging modalities pertinent to this work, with a focus on ultrasound and computed tomography. Section §1.4 introduces key principles of cardiovascular fluid dynamics, including characteristic flow regimes, hemodynamic metrics and their physiological implications. Section §1.5 outlines the clinical implications driving this thesis, particularly the need for improved biomarkers and quantitative flow-based assessment tools. Finally, Section §1.6 concludes with an outline of the thesis structure.

1.1 Cardiovascular system overview

The cardiovascular system consists of a complex vascular network and interconnected organs (Fig. 1.2), each contributing essential functions required to sustain life and unified by the common element of blood flow. At the center of this system is the heart, which drives and directs blood throughout the body. In a broad perspective, the circulatory system is divided into pulmonary and systemic circuits. The pulmonary circulation involves the right side of the heart (from the patient's point-of-view), which pumps deoxygenated blood to the lungs for gas exchange before it returns to the left side of the heart. The systemic circulation begins when oxygenated blood is ejected from the left heart and distributed to the body, delivering oxygen and nutrients to tissues and organs. After exchanging gases and metabolic products, the deoxygenated blood returns to the heart, completing the cycle and then repeats again.




Figure 1.2: Cardiovascular system and circulation: (A) pulmonary circulation; (B) systemic circulation.

Adapted from Fig. 20.2 in [13].

1.2 Heart anatomy and cardiac cycle

The heart is a hollow muscular organ composed of the endocardium, myocardium, and epicardium, enclosed within the serous and fibrous layers of the pericardial sac (Fig. 1.3). It contains four chambers: right atrium, right ventricle, left atrium, and left ventricle. They are organized into a unidirectional pumping system regulated by four valves: tricuspid, mitral, pulmonary, and aortic. The synchronized chamber contractions are driven by the electrical activation originating in the sinoatrial node and transmitted through the cardiac conduction system. The atrial and ventricular septa structurally separate the right and left circulations. To better understand cardiac hemodynamics, it is essential to first consider how blood moves through the chambers. Below, we introduce the anatomy and cardiac cycle of each chamber by tracing the path of a blood particle through the heart, beginning with its return of oxygen-rich blood from the lungs.




Figure 1.3: Human heart anatomy. Adapted from Fig. 1 in [14].

Left atrium

After returning from the lungs, the oxygenated blood enters the left atrium (LA) through four pulmonary veins. The LA interior wall is smooth, with the exception of the left atrial appendage (LAA), a cul-de-sac-like pouch where most clots form [15, 16]. During left ventricular (LV) systole, when the mitral valve is closed, the LA serves as a filling reservoir. As the ventricle relaxes in early diastole, LV pressure falls below LA pressure, the mitral valve opens, and blood begins to move toward the valve. Near the

end of diastole, a brief atrial contraction (often called the “atrial kick”) propels the remaining blood into the LV.

Left ventricle

Upon passing through the mitral valve (MV), the blood flows into the LV, a thick-walled, conical chamber responsible for systemic circulation. Early LV relaxation creates a pressure gradient that draws blood inward, followed by continuous passive filling until late diastole, when atrial contraction provides the final increment of volume. Once filling is complete, the MV closes and the LV begins isovolumetric contraction, rapidly increasing intraventricular pressure. When LV pressure exceeds aortic pressure, the aortic valve opens and the blood is ejected into the systemic circulation. At the end of systole, the ventricle relaxes, the aortic valve closes, and the LV begins another filling period. This sequence of filling, contraction, ejection, and relaxation produces the characteristic LV pressure–volume loop (Fig. 1.4), providing a concise summary of left ventricular function.

Right atrium

After completing its systemic journey, the blood eventually returns to the heart deoxygenated and enters the right atrium (RA) through the superior or inferior vena cava. During right ventricular systole, when the tricuspid valve is closed, the RA accumulates venous return and acts as a reservoir. As the RV relaxes and its pressure falls below that of the RA, the tricuspid valve opens and the blood particle moves toward the RV. A short atrial contraction at late diastole helps complete RV filling. Structurally, the RA resembles the LA but includes a larger right atrial appendage, whose broader junction is associated with a lower likelihood of thrombus formation compared with the LAA [18].

Right ventricle

Passing through the tricuspid valve, the blood then enters the right ventricle (RV), a thin-walled, crescent-shaped chamber adapted for pumping blood into the pulmonary circulation. During early diastole, ventricular relaxation enables rapid filling, followed by the contribution of atrial contraction. The tricuspid valve then closes, and the RV undergoes isovolumetric contraction until its pressure exceeds that in the pulmonary artery. Once the pulmonary valve opens, the blood is ejected toward the lungs for




Figure 1.4: Left ventricle pressure-volume loop. DBP, diastolic blood pressure; E_a , effective arterial elastance; EDPVR, end-diastolic pressure–volume relationship; E_{es} , end-systolic elastance; ESPVR, end-systolic pressure–volume relationship; PE, potential energy; P_{ed} , end-diastolic pressure; P_{es} , end-systolic pressure; SBP, systolic blood pressure; SW, stroke work; V_0 , volume at a P_{es} of 0 mmHg. Stroke volume (SV) is EDV - ESV. Adapted from Fig.1 in [17].

reoxygenation, completing the cardiopulmonary loop.

Remark

While all four chambers contribute to coordinated cardiac function, the LA and LV are of particular physiological and clinical importance. Their structural complexity, their central roles in governing systemic perfusion, and their involvement in a wide range of cardiovascular diseases (CVD) make them especially relevant to image-based modeling and computational analysis [19–23]. Accordingly, this thesis emphasizes on developing predictive computational models of LA and LV hemodynamics reconstruction and CVD risk inference, while also extending the analysis to broader aspects of global circulation.

1.3 Medical imaging for cardiac flow

Medical imaging technologies have existed since the late 19th century, yet many did not become clinically relevant for cardiac function evaluations until decades later. Echocardiography, first conceived in the 1880s but not widely adopted until the 1950s, introduced real-time, radiation-free visualization of cardiac structure and hemodynamics through ultrasound [24, 25]. Over time, multiple echo modalities such as transthoracic, transesophageal, color-Doppler, and 3D echocardiography, expanded diagnostic capabilities for assessing chamber function, wall motion, valve disease, and common cardiac pathologies. The emergence of computed tomography (CT) in the 1960s and its clinical use in the early 1970s allowed rapid acquisition of cross-sectional and 3D cardiac and vascular images, enabling detailed assessment of coronary stenosis and structural abnormalities [26]. Magnetic resonance imaging (MRI), first demonstrated for the heart in the early 1980s [27], has since evolved into a powerful, radiation-free modality capable of high-resolution structural and functional imaging, owing to improvements in hardware, pulse sequences, and tissue characterization techniques.

The remainder of this section introduces on imaging modalities most relevant to computational cardiac flow analysis and to this thesis, with an emphasis on color-Doppler ultrasound, and contrast-enhanced computed tomography. MRI is also discussed where pertinent, particularly in relation to its strengths in cardiac structure and tissue characterization.

1.3.1 Color-Doppler ultrasound

Color-Doppler imaging leverages ultrasound principles to transmit high-frequency sound waves that reflect off moving blood cells and return to the transducer, allowing real-time estimation of the radial component of blood velocity (e.g., v_r) within a 2D plane sector (Fig. 1.5A). Modern scanners provide relatively high temporal resolution for this 1D radial velocity field; however, the modality remains limited by comparatively low spatial resolution, phase-wrapping artifacts, and the fact that it captures only one of the three velocity components. A prominent computational method that builds on color-Doppler data is the left ventricular vector flow mapping technique [31–37], which integrates the continuity equation using the measured radial velocities to recover the azimuthal component and thereby reconstruct a physiologically meaningful 2D flow field (see Section §2.1.2). Beyond flow assessment, echocardiography offers radiation-free imaging with superior structural detail compared




Figure 1.5: Medical imaging examples: (A) 2D color-Doppler ultrasound in early filling. Adapted from [28]; (B) 4D-CT scan of left atrium and left ventricle. Adapted from [29]; (C) Three-chamber view of 4D flow MRI velocity streamlines at peak systole. Adapted from [30].

to standard radiography, providing dynamic visualization of chamber motion and wall mechanics. It is routinely used to evaluate conditions such as abdominal or thoracic aortic aneurysm, intracardiac thrombi, pericarditis, pericardial effusion, and valvular heart disease [38].

1.3.2 Computed tomography

Computed tomography (CT) creates detailed cross-sectional images of the heart (Fig. 1.5B) by rotating x-ray beams and measuring how cardiac tissues attenuate the radiation [39]. Dense structures such as coronary calcium appear bright, allowing clear visualization of plaque and other pathology. Multiple reconstructed slices form high-resolution 3D volumes, offering far greater anatomical detail than standard radiographs. Cardiac CT is valued for its speed, reliability, and excellent spatial resolution [40], and is frequently used to assess various cardiovascular diseases and complications [41]. Specialized modalities such as coronary CT angiography, calcium-score scans, and whole-body CT, serve distinct diagnostic purposes [42]. These techniques typically rely on iodine-based contrast to enhance vascular and myocardial visualization; although contrast nephropathy is a concern, overall risk is low with proper precautions [43, 44]. While CT does not directly measure cardiac flow, it has become a preferred non-invasive method for initial evaluation of suspected coronary artery disease, providing information on

anatomy, myocardial perfusion, contrast enhancement, and valve morphology [45,46].

1.3.3 Magnetic resonance imaging

Magnetic resonance imaging (MRI) is considered the gold standard for noninvasive cardiac assessment because of its precision and excellent tissue characterization [47]. The reconstruction works by using a strong magnetic field and brief radiofrequency pulses to align and disturb hydrogen protons in the body; the signals they release as they return to alignment are reconstructed into detailed images [48–58]. MRI enables comprehensive evaluation of ventricular volumes and function, myocardial viability, tissue composition, perfusion, and vascular or structural abnormalities [59–61]. MRI techniques, including MRA, MRV, and dedicated cardiac protocols, often use gadolinium contrast to enhance visualization of vessels, myocardium, and areas of scar or inflammation, although their spatial resolution is slightly lower than that of CT [62]. Importantly, cardiac MRI can also measure blood-flow velocity through phase-contrast and 4D flow imaging (Fig. 1.5C), which encode velocity information in the MR signal to generate time-resolved maps of intracardiac and vascular flow, making it a powerful tool for studying cardiac hemodynamics.

1.4 Cardiovascular fluid dynamics

In this section, we introduce the fundamentals of cardiovascular fluid dynamics with emphasis placed on the left ventricle and left atrium, as well as on models that describe global blood circulation. While cardiovascular performance arises from the complex interactions among hemodynamics, tissue mechanics, molecular pathways, and systemic physiology, focus of this thesis is on flow phenomena most relevant to fluid dynamics and computational modeling. We highlight key features of intracardiac flow and outline multiscale modeling strategies that have advanced our understanding of cardiac and circulatory hemodynamics.

1.4.1 Left ventricular flow

Flow inside the left ventricle (LV) is highly unsteady, driven by rapid volume and velocity changes throughout each cardiac cycle. During early diastole, the strong mitral jet emerging from the mitral orifice corresponding to a Reynolds number greater than 3000, a dimensionless measure of the ratio

between inertial and viscous forces. Thus, it's nominally within a turbulent regime. As the jet enters the ventricle, the shear layer forms an annular vortex downstream of the mitral valve, which evolves into a large, asymmetric vortex ring by the end of diastole [63–65]. This vortex reduces convective energy losses [66] and helps redirect flow toward the outflow tract during systole [67]. The question of whether ventricular flow becomes turbulent remains an active area of study: both imaging [68] and simulations [69, 70] indicate that late diastolic vortex breakdown can induce transitional or turbulent states. Turbulence modeling approaches such as large eddy simulation and Reynolds-averaged Navier-Stokes have shown improved agreement with experimental measurements by capturing transient and small-scale flow dynamics absent in laminar formulations [71]. Although turbulence appears intrinsic to cardiovascular flow environments, fully developed turbulence is generally not reached under physiological conditions [69, 70].

Considering this intricate hemodynamics, both in vivo imaging and computational modeling provide complementary methods to quantify LV flow. Doppler echocardiography (Fig. 1.5A) and, more recently, 4D flow MRI (Fig. 1.5C) allow time-resolved, three-dimensional visualization of blood flow [72, 73]. Although resolution limitations restrict the capture of fine-scale features and smaller vessels, ongoing advances continue to expand their clinical utility. Computational simulation offers a viable alternative for detailed analysis, especially when imaging resolution is insufficient.

Computational simulations of LV hemodynamics require patient-specific geometries, physiological wall motion, and fully coupled fluid-solid interaction with valve dynamics. Model simplifications vary depending on the specific question, ranging from lumped-parameter models that provide efficient pressure-volume predictions [74] to idealized [70, 75–77] or patient-specific geometries [64, 78]. Fine anatomical structures such as papillary muscles and trabeculae introduce additional geometric complexity, and contribute functionally to mitral valve closure, efficient filling/ejection and reduced myocardial stress [71, 78–80]. Valve modeling strategies range from omission or simplified resistance-based valves [70] to fully resolved 3D dynamic valves using Arbitrary Lagrangian-Eulerian or immersed boundary methods [77, 78, 81]. Ventricular wall motion may also be prescribed from imaging or modeled using fluid-solid interaction frameworks that incorporate electrophysiology and active contraction [82].

1.4.2 Left atrial flow

The left atrium (LA) is a highly compliant chamber that contributes to ventricular filling in three distinct ways: it stores pulmonary venous return during ventricular systole, passes blood directly into the left ventricle (LV) during early diastole, and actively boosts filling through atrial contraction in late diastole. These varied mechanical functions, coupled with pronounced geometric diversity across individuals, result in complex, 3D flow fields that remain less extensively characterized than LV haemodynamics. Early computational fluid dynamics (CFD) efforts involving the LA were often embedded within whole-left-heart simulations, and primarily explored how LA shape influenced downstream LV flow organization [69, 78, 79].

Experimental and computational studies consistently show that inflow from the pulmonary veins is markedly asymmetric: left pulmonary veins tend to direct flow into a rotational loop, whereas right pulmonary veins align their jets more closely along the atrial wall [83]. This produces uneven flow pathways that help maintain blood momentum throughout the cardiac cycle. Distinct vortical structures often emerge during systole and mid-diastole—features observed in both imaging studies [83, 84] and numerical simulations [85–88]. These vortices are thought to improve atrial washout by reducing regions where flow might otherwise stagnate [83].

Among all the structural and flow features of the LA, the left atrial appendage (LAA) stands out. The LAA protrudes from the superior aspect of the atrium and contains pectinate muscles that introduce pronounced internal ridges. Its geometry makes it prone to slow flow and stasis. Patient-specific CFD simulations have shown that atrial fibrillation (AF) significantly reduces LAA velocities, leading to incomplete washout. The typical pattern of several small vortices in the healthy atrium is often replaced by a single, less efficient vortex in AF patients [84]. CFD investigations of characteristic LAA shapes have demonstrated clear links between morphology, local haemodynamics and thrombotic risk [89, 90]. Although some LA simulations explore transitional or turbulent flow regimes, turbulence modeling is not always required for AF cases [91].

1.4.3 Capillary and circulatory flows

Blood flow in the circulation spans nearly six orders of magnitude in scale, from large arteries in centimeters to small capillaries in microns, creating tightly coupled multi-scale interactions. While

blood behaves approximately as a Newtonian fluid in large arteries, this continuum description breaks down in the microcirculation, where viscosity varies nonlinearly with vessel diameter [92]. These disparate behaviours motivate multi-physics modeling approaches that avoid the prohibitive cost of fully resolving all vessels.

Reduced-order descriptions such as lumped-parameter and one-dimensional (1D) flow models enable efficient representation of global haemodynamics. Lumped-parameter models treat the vasculature as networks of resistors, capacitors and inductors [93], while 1D models are derived by cross-sectional averaging of the Navier–Stokes equations [94, 95]. Despite their simplicity, both frameworks capture essential flow and pressure dynamics and are routinely coupled with localized 3D simulations in multi-scale modeling [96–99]. These hybrid approaches have been used to study how physiological conditions, disease or surgical interventions influence circulation across disparate vascular regions.

At the smallest scales, capillary Reynolds numbers are $O(10^{-3})$, and blood behaves as a suspension of deformable cells rather than a continuum fluid [92]. Red blood cells migrate away from vessel walls, form single-file arrangements and alter viscosity through deformation, giving rise to shear-thinning effect [100]. These microstructural dynamics directly influence resistance, oxygen transport and local haemodynamics.

1.5 Clinical implication of intracardiac flows

Left ventricular flow

During left ventricle (LV) early diastole, the interaction of the E-wave jet with the mitral leaflets and residual end-systolic flow generates a vortex ring that typically pinches off during E-wave deceleration and becomes increasingly asymmetric as it travels toward the apex [63, 101, 102]. In healthy hearts, this E-wave vortex usually decays during diastasis so that atrial contraction generates a new A-wave vortex, whereas in dilated cardiomyopathy (DCM) the E-wave vortex persists, grows in circulation and dominates diastolic flow due to the more spherical chamber geometry [102]. Changes in heart rate, filling time and pacing can further modulate vortex strength and organization, and may explain differences in response to resynchronization therapy [103, 104].

The extent to which diastolic vortices influence systolic ejection remains debated. Vortical structures may help redirect flow toward the outflow tract, potentially reducing energy dissipation and facilitating the

transition from filling to ejection [67]. Clinical studies report higher diastolic kinetic energy losses and reduced vortex circulation in heart failure, both associated with adverse outcomes [105–107]. However, because intraventricular kinetic energy represents $\leq 5\%$ of LV mechanical work, the absolute energetic contribution of vortices is small [108, 109]. Even in such cases, ventricles with limited myocardial reserve, such as DCM or hypertrophy, modest inefficiencies in vortex-mediated redirection may still have clinically meaningful effects, motivating further mechanistic study.

LV flow profoundly influences blood transport. In healthy LVs, $\approx 70\%$ of inflow leaves during the same cardiac cycle, and most of the remainder exits in the next cycle; approximately 30% end diastolic volume persists near the apex with long residence times but limited mixing [106]. In DCM, direct flow falls to $\approx 50\%$, retained and delayed ejection volumes increase, and the residual volume expands and shifts toward the cavity center, producing complex residence time distributions and greater stasis [106]. These flow alterations contribute to the elevated risk of intraventricular thrombosis in DCM and other structural diseases [110, 111]. Device therapies can further disrupt natural flow paths; for example, continuous-flow left ventricle assist device increase stasis risk, whereas pulsatile operation may improve washout [112, 113]. Understanding these transport pathways is therefore important for thrombus prediction, risk stratification and device optimization.

Left atrial flow

The left atrium (LA) plays a key role in reservoir, conduit, and booster-pump function, and disturbances in its haemodynamics, particularly within the left atrial appendage (LAA) have major clinical complications. The LAA's compliant, trabeculated geometry predisposes it to slow flow and stasis, making it the primary site of thrombus formation in atrial fibrillation (AF), where abnormal electrical activation impairs atrial contraction and reduces washout. More than 90% of atrial thrombi originate in the LAA [114], and early clinical and computational studies have shown that local flow patterns strongly correlate with thromboembolic risk [115, 116].

In AF, decreased mean and peak velocities in the LAA, the loss of discrete left atrial vortices, and the emergence of a single large recirculating structure lead to increased blood residence time and promote clot formation [84, 87, 88]. Patient-specific CFD simulations have shown that LAA morphology strongly influences these patterns, with certain shapes associated with markedly higher thrombosis risk-

findings consistent with clinical observations [89, 90]. These sensitivities underscore the importance of incorporating detailed LAA geometry into personalized risk assessment and treatment planning, including consideration of LAA occlusion devices.

Beyond AF, impaired LA flow is implicated in other conditions where elevated filling pressures, atrial dilation, or reduced atrial compliance produce regions of stasis that predispose to thrombus formation. The integration of advanced imaging with computational modeling continues to improve our ability to map atrial flow, identify high-risk regions and refine therapeutic decision making in disorders affecting LA structure and function.

1.6 Thesis outline

This general exam report and thesis are organized as follows. In Chapter 2, we discuss a range of image-based, data-driven computational approaches for cardiovascular flow applications, along with other predictive modeling frameworks, and provides an in-depth review of these approaches and physics-informed neural networks (PINN), the core methodological foundation of this thesis. In Chapter 3, we present the application of PINN to 2D left ventricular color-Doppler vector flow mapping (VFM) to reconstruct LV velocity and pressure fields (e.g., the AI-VFM framework). In Chapter 4, we extend AI-VFM framework to 3D by using triplane color-Doppler imaging to reconstruct full-volume LV flow fields. In Chapter 5, we demonstrate and compare the use of PINN and indicator dilution theory to reconstruct flow fields and thrombosis risk from 4D contrast-CT in the left atrium. In Chapter 6, we further expand the PINN framework to a post-Norwood circulation, using a lumped-parameter model to infer patient-specific parameters such as vascular resistances and hemodynamic quantities. Finally, in Chapter 7, we outline the remaining work prior to the final defense, discuss potential challenges and alternative solutions, and provide a proposed timeline for completing the thesis.

Chapter 2

BACKGROUND

2.1 Image-based data-driven computational predictive modeling in cardiovascular flow

2.1.1 Computational fluid dynamics

Computational fluid dynamics (CFD) simulations emerged in the early 1970s as increasing computational power enabled the numerical simulation of fluid flows, first through simplified transonic potential-flow models and later through full 2D and 3D Euler and Navier-Stokes equations [117–120]. By the mid 1980s, advances in numerical schemes, turbulence modeling, and supercomputing made it possible to simulate increasingly complex flows, including hypersonic vehicles, combustion, and turbomachinery. As problem complexity grew, so did the need for more sophisticated computational grid generation techniques, evolving from structured meshes to multiblock strategies, and eventually unstructured meshes with solution-based adaptation to better capture flow features. These developments greatly expanded CFD's capability and reduced the manual effort required to create high-quality computational grids.

As CFD evolved into a robust framework for simulating complex, unsteady, and geometry dependent flows, it became particularly well suited for modeling intracardiac flows, where a direct measurement is challenging or invasive, and controlled experimental studies are often impractical. CFD simulations enable systematic studies of how adverse cardiac remodeling, such as changes in chamber and valve geometry, affect intraventricular flow, while allowing individual parameters to be varied independently, unlike in clinical settings where pathological factors are tightly coupled [102, 121]. This capability has been leveraged to explore whether cardiac anatomy has evolved to optimize intraventricular flow [67, 101, 109], to predict the hemodynamic impact of surgical or device-based interventions [122, 123], and to couple flow solvers with growth and remodeling models to study the interplay between hemodynamics and myocardial mechanobiology [124–126]. Although such simulations are computationally feasible with modern hardware, they remain challenging due to complex, time-varying geometries and strong fluid-structure interactions, which motivated the development of immersed boundary and moving-mesh

methods [127, 128].

A central challenge in intracardiac CFD simulations is the accurate representation of fine anatomical structures, including inflow valves, papillary muscles, and trabeculae, all of which can substantially influence intracardiac flow dynamics. For example, diastolic vortex rings are shed from the inflow valve leaflets in the left ventricle, and valve kinematics have been shown to affect intraventricular flow [66, 129, 130]. Although less is known about the role of papillary muscles, their characteristic dimensions in the adult heart are comparable to the viscous core of the diastolic vortex ring, suggesting a potential role in enhancing viscous dissipation and mixing [102, 121, 131]. Similarly, ventricular trabeculae, often neglected or simplified in simulations, have long been hypothesized to condition inflow and reduce the energetic cost of ejection by increasing dissipation [132], a mechanism that is difficult to test in vivo but accessible through CFD simulations [133].

Beyond left ventricular flow, CFD simulations of left atrial (LA) hemodynamics have emerged more recently and remain comparatively limited. Early studies prescribed synthetic atrial wall motion to examine the role of atrial contraction in promoting left atrial appendage (LAA) washout [134, 135], while subsequent patient-specific simulations reported that reduced atrial or appendage contractility can promote blood stasis [87, 88, 136]. However, conflicting findings indicate that LAA stasis does not depend on atrial function alone; for example, subjects with better atrial function have also been shown to experience increased stasis under certain conditions [137]. Although clinical studies suggest a moderate correlation between LAA morphology and stroke risk in atrial fibrillation [138–141], CFD analyses indicate that LAA stasis depends not only on geometry but also on additional factors such as pulmonary vein inflow patterns [89, 90].

Building upon prior CFD investigations of left ventricular and left atrial flow and their implications for intracardiac transport and thrombosis risk, this thesis focuses on leveraging patient-specific CFD simulations of LV and LA hemodynamics as inputs to physics-informed neural networks (PINN) for data-driven inference. In particular, CFD-derived flow fields are used to inform the learning of hemodynamic quantities and thrombosis-related metrics such as residence time that are difficult to measure directly in vivo. All simulations are performed using the in-house CFD solver TUCAN [28, 88, 142, 143], which solves the incompressible Navier-Stokes equations for a Newtonian fluid with constant kinematic viscosity using a fractional-step method. Spatial discretization employs a staggered Cartesian grid with second-order centered finite differences, while time integration is carried out using a low-storage, three-

stage semi-implicit Runge-Kutta scheme. Detailed descriptions of the patient-specific CFD geometry, boundary conditions, and parameters are provided in the subsequent chapters.

2.1.2 Intraventricular vector flow mapping

This section reviews one of the classical computational approaches based on color-Doppler ultrasound imaging, namely vector flow mapping (VFM), which has gained substantial traction and undergone significant development over the past two decades for assessing left ventricular flow. Briefly, VFM uses the one-dimensional radial velocity measured by color-Doppler to infer the azimuthal velocity component in a polar coordinate system, thereby reconstructing a two-dimensional velocity vector field. Importantly, this thesis further develops and extends the VFM framework using state-of-the-art physics-informed machine learning to infer additional flow quantities (e.g., pressure) that were not immediately accessible in previous implementations. Details of these extensions are presented in Chapters 3 and 4.

The most widely used VFM approach is based on the two-dimensional continuity equation expressed in a polar coordinate system,

$$\nabla \cdot \tilde{\mathbf{v}} = \frac{1}{r} \frac{\partial}{\partial r}(r v_r) + \frac{1}{r} \frac{\partial v_\theta}{\partial \theta} = 0 \quad (2.1)$$

where $[r, \theta]$ denote the polar coordinates of the color-Doppler imaging sector, and $[v_r, v_\theta]$ represent the radial velocity measured by color-Doppler and the unknown azimuthal velocity component, respectively. First proposed by Garcia et al. [31], this VFM formulation directly integrates Eq. 2.1 using a finite-difference spatial discretization to recover v_θ , assuming either no-slip or free-slip boundary conditions derived from speckle-tracking analysis of the left ventricular boundaries. Itatani et al. [32] subsequently improved this formulation by refining a weighted combination of two azimuthal velocity fields computed from the anteroseptal and posterior wall boundaries, and by proposing an optimal smoothing strategy to obtain stable spatial derivatives for post-processing analyses.

Several alternative VFM formulations were later developed based on potential functions [144], stream-function approaches [37], and least-squares solution methods [34]. Although these variants address certain limitations of the planar flow assumption and improve numerical robustness, they do not incorporate full momentum conservation (e.g., the Navier-Stokes equations) and therefore require secondary analyses to derive pressure fields [145, 146]. Extensions to three-dimensional, full-volume

flow reconstruction have also been proposed using continuity-based formulations combined with triplane color Doppler imaging to recover all three velocity components [147]. More recently, machine learning-based and physics-informed neural network frameworks have been introduced to reconstruct velocity and pressure fields in both two-dimensional and three-dimensional VFM settings [28, 148–150].

2.1.3 Indicator dilution theory

Indicator dilution theory originates from late 19th century efforts to quantify blood flow based on conservation of mass. While electrocardiography emerged as a dominant clinical tool, a parallel and equally foundational development occurred with the formulation of blood flow estimation based on diffusive and convective transport [151]. This principle was extended through continuous infusion techniques [152] and later transformed into a dynamic framework, which introduced the single-injection indicator dilution method [153]. Henriques' approach departed from steady-state assumptions by relating blood flow to the ratio of injected indicator mass and the time-integrated downstream concentration curve, implicitly assuming constant or rapidly oscillating flow during sampling [154], although Henriques' work initially received limited attention [155].

Subsequent refinements of the technique by Hamilton led to broader recognition of the method, which became widely known as the Stewart-Hamilton indicator dilution technique [156]. The introduction of spectrophotometrically measurable dyes enabled practical implementation in both experimental and clinical settings [157]. With advances in cardiac catheterization and circulatory physiology during the mid 20th century, indicator dilution methods were increasingly applied to quantify cardiac output, organ perfusion, and circulation [158, 159]. Extensions of the method, such as dual-site measurements across organ beds, enabled estimation of regional blood volumes and organ-specific hemodynamics [154].

From the late 1950s onward, the theoretical interpretation of indicator dilution curves matured through tracer kinetic analysis and differential equation based modeling. Rapid bolus injection followed by high-resolution sampling revealed that dilution curves could be interpreted as impulse response functions of the cardiovascular system, encoding the distribution of transit times through a vascular bed [160]. This insight led to the formulation of the central kinetic relationship linking flow, mean transit time, and volume, enabling quantitative estimation of vascular and extravascular volumes [154]. Multi-indicator techniques further extended the framework by comparing intravascular reference tracers with diffusible

indicators to assess capillary permeability, extracellular space, and metabolic exchange, with applications spanning pulmonary physiology, hepatic circulation, and neurovascular transport [160, 161].

The indicator dilution theory also provided the foundation for compartmental analysis and modern tracer kinetics. Formulations incorporating convolution integrals and explicit input functions enabled analysis of complex transport processes and laid the groundwork for tomographic imaging techniques [162]. These models were later adapted to radioactive tracers, microspheres, and inert gases to quantify regional perfusion and metabolic activity across organs, including the brain and myocardium [163, 164]. The advent of positron emission tomography (PET) further operationalized indicator dilution principles by combining dynamic tracer injection with spatially resolved detection, allowing in vivo estimation of perfusion, metabolism, and receptor kinetics using mathematically rigorous deconvolution and compartmental analysis [165, 166].

More recently, indicator dilution theory persists in both simplified and advanced forms. Thermodilution using cold saline remains a standard method for cardiac output measurement in intensive care settings [167], while nuclear medicine and PET-based tracer techniques are routinely employed in cardiology, oncology, and neurology. Magnetic resonance imaging has adopted contrast agent and endogenous tracer strategies that, despite limited diffusibility, leverage high temporal resolution and model-based analysis to estimate perfusion, blood volume, and permeability [168–170]. Techniques such as arterial spin labeling represent a noninvasive extension of classic indicator dilution concepts by magnetically labeling blood water as an intrinsic tracer [171]. Collectively, these developments underscore the enduring relevance of indicator dilution theory as a unifying framework linking physiology, transport theory, and modern medical imaging.

Building upon these foundational works and motivated by recent work in deploying gamma-variate model to compute thrombosis risk using contrast-enhanced CT imaging [172], this thesis employs indicator dilution theory in combination with compartmental analysis and a gamma-variate model to derive a closed-form expression for left atrial residence time from 4D CT contrast dynamics. The proposed framework is used to assess thrombosis risk and is systematically compared with state-of-the-art computational fluid dynamics and physics-informed neural network approaches. Details are presented in Chapter 5.

2.1.4 *Lumped-parameter modeling*

Lumped-parameter modeling provides a simplified framework for describing cardiovascular hemodynamics in which pressure, flow rate, and volume vary as functions of time. By exploiting the hydraulic-electrical analogy where fluid flow corresponds to electric current and pressure to voltage, cardiovascular circulation can be represented using equivalent electrical circuits composed of resistors, capacitors, and inductors. Solving the coupled ordinary differential equations (ODEs) of these circuits yields a mathematical representation capable of predicting global blood flow behavior throughout the circulation system with very little computational cost [173, 174]. This concept traces back to the seminal two-element Windkessel model proposed in 1899, which modeled the arterial system using a resistor and capacitor to capture vascular resistance and compliance [175, 176]. Subsequent extensions addressed frequency dependent limitations of the original formulation, including the three-element Windkessel model incorporating characteristic impedance [93, 177] and the four-element model adding inertance to better represent arterial wave propagation across a broader frequency range, albeit at the expense of increased parameter complexity [176, 178].

Over time, numerous Windkessel variants and lumped-parameter models have been developed to capture additional physiological features such as microcirculatory effects [179–182]. These models have been widely applied to study artificial heart devices [183–185], surgical interventions [186–188], and patient-specific diagnostics [189–199]. A typical configuration consists of cardiac chamber compartments acting as pumps coupled to vascular networks that transport blood throughout the body [174]. In the electrical domain, resistors model viscous losses, capacitors represent vascular compliance, and inductors account for blood inertia [173]. The ODEs arise from mass and momentum conservation in the hemodynamic domain via the continuity equation, Poiseuille’s law, and the Navier-Stokes equations, and are translated into the electrical domain [200, 201]. The resulting system comprises coupled ODEs and algebraic pressure-volume relationships, which may be linear or nonlinear depending on model fidelity [201]. Due to their low computational cost and structural simplicity, lumped-parameter models are well suited for automation and integration into high-throughput clinical workflows [202, 203].

A key strength of lumped-parameter modeling lies in its ability to simultaneously capture cardiac function and global circulatory hemodynamics within a unified framework, which is essential given the tightly coupled nature of cardiovascular physiology. These models can serve both as standalone

diagnostic tools and as components within multiscale frameworks coupled to higher order 1D or 3D models to incorporate local flow phenomena [191, 192, 204–210]. However, a major challenge remains parameter identification, as many model elements do not correspond directly to measurable physiological quantities and often require invasive measurements or indirect estimation [200]. Consequently, parameter values are frequently derived from literature, population averages, animal studies, or prior experiments [201, 211–215]. Increasingly, non-invasive patient-specific tuning approaches have been proposed to address this limitation, enabling personalized simulations that improve diagnostic and predictive capability [189–194, 216–219]. Continued refinement and personalization of lumped-parameter models are therefore critical for advancing cardiovascular modeling as a practical tool to study disease mechanisms and support clinical decision-making.

Chapter 6 of this thesis builds upon these concepts by modeling the post-Norwood circulation using a lumped-parameter framework, in which patient-specific parameters and flow related quantities are inferred from clinically accessible measurements through a physics-informed neural network approach.

2.1.5 Machine learning approaches

Machine learning (ML) has seen rapid advancement and widespread adoption across a range of domains and is emerging as a powerful tool in healthcare. ML approaches are well suited for modeling complex and latent relationships among numerous clinical and physiological variables and their associated outcomes. In this section, we review applications of ML frameworks for predicting cardiovascular disease (CVD) trained on large-scale clinical characteristics and comorbidities data, as well as the use of medical imaging data for CVD predictions.

Cardiovascular disease prediction

The term CVD refers to a wide variety of commonly linked pathologies, including coronary artery disease, stroke, peripheral artery disease, deep vein thrombosis and pulmonary embolism. CVD prediction is currently based on risk scores that use patients' clinical characteristics and comorbidities to approximate the likelihood of future CVD, thus facilitating clinical decision-making. Models widely used in the US include the Framingham Risk Score (FRS) and the American College of Cardiology/American Heart Association (ACC/AHA) guidelines [220, 221]. Based on the risk profile calculated by these algorithms,

the physician may choose to start a patient on statin therapy, for example, in addition to recommending adequate lifestyle modifications.

Machine learning models compared with traditional risk scores

The 2013 ACC/AHA Pooled Cohort Equations Risk Calculator is based on traditional risk factors to predict cardiovascular events over 10 years. It can overestimate or underestimate the overall risk of cardiovascular events in certain groups [222]. Kakadiaris et al. used the same risk variables collected in the Multi-Ethnic Study of Atherosclerosis cohort to train an SVM ML algorithm [223]. The resulting predictive tool outperformed the ACC/AHA calculator, recommending less drug therapy and missing fewer cardiovascular events. The 2017 Framingham Stroke Risk Profile was compared with ML techniques for stroke risk prediction (random survival forest [RSF], SVM, GBT and multilayer perceptron) in a Chinese cohort of 503,842 adults [224]. These ML models improved risk prediction over the conventional Cox model-based approach, with GBT providing the best discrimination and calibration performance. However, the best means of identifying individuals at high risk of stroke was the use of an ensemble model combining the GBT and Cox methods (accuracy of 76% in men and 80% in women). Weng et al. compared four ML algorithms with an established evaluation rubric (ACC guidelines) to predict the first cardiovascular event over 10 years in a prospective cohort of 378,256 patients from UK family practices free from CVD at the study's outset [225]. All ML-based approaches tested were better at identifying individuals at risk of CVD than the established algorithm, with neural networks providing the best risk prediction (area under the receiver operating characteristic curve [AUC] 0.764 versus 0.728 for neural networks versus the conventional method, respectively).

Machine learning models can analyse more complex features

In addition to patient characteristics and demographics, ML can use blood biomarker levels or imaging data as a basis for predicting CVD. AmbaleVenkatesh et al. found that imaging, ECG and serum biomarkers were among the top predictors of CVD as opposed to traditional cardiovascular risk factors and identified RSF as the most predictive cardiovascular risk (including stroke) model out of nine tested models in an initially asymptomatic multiracial population [226]. In another study, an ML-based, 27-protein prognostic cardiovascular model was trained and validated to predict the 4-year likelihood of MI,

stroke, heart failure or death. The proteomic model performed better than the clinical model in predicting CVD (AUC 0.73 versus 0.64) [227]. Preliminary results from an ML-based model for predicting white matter lesions as a surrogate for CVD from retinal microvasculature changes (arteriovenous nicking, opacity, focal arteriolar narrowing) are also compelling [228]. Four ML classifiers were tested (RF, artificial neural network [ANN], SVM and linear regression), and both the RF and ANN achieved optimal classification accuracy. Thus, retinal imaging is proposed as a screening tool at the primary care level. Sridhar et al. developed a rapid and inexpensive triage tool based on a DNN to predict major adverse cardiovascular outcomes from ECG data in patients with COVID-19 [229]. This approach yielded comparable results to traditional statistical models incorporating extensive clinical data with the advantage of not requiring clinical expertise to gather medical history.

Clinical stroke risk assessment in atrial fibrillation

Prior work aiming to characterise ischaemic stroke risk in AF patients has focused on clinical scores, such as CHADS₂, CHA₂DS₂-VASc and ATRIA. CHADS₂ was limited by its difficulty in accurately evaluating low-risk groups [230]. To improve the predictive performance in this subset of patients, the CHA₂DS₂-VASc score was implemented. The ATRIA score outperforms CHA₂DS₂-VASc in some cases, but both scores remain imperfect predictors of stroke with modest predictive accuracy [231,232].

Early efforts to develop ML algorithms for predicting stroke risk in AF patients have shown some promise, and have achieved an AUC as high as 0.892 in one cohort analysis [233]. Whereas CHADS₂ and CHA₂DS₂-VASc use 6-7 features to stratify stroke risk, an attention-based deep neural network model identified up to 48 features that influenced stroke risk using the Korean National Health Insurance data, and provided better prediction of ischaemic stroke in AF patients compared with CHA₂DS₂-VASc scores (AUC 0.727 versus 0.651, respectively) [234]. ML was used in a German cohort to predict ischaemic strokes in AF patients started on non-vitamin K oral antagonists. The Subpopulation Optimisation and Modelling Solutions (SOMS) tool identified age, male sex, ischaemic heart disease, urinary tract infection, and dementia as predictors of ischaemic stroke in this population [235].

Remark

The reviews presented in Section §2.1.5 are, in part, a published manuscript [236]: Chahine, M. J. Magoon, B. Maidu, J. C. del Álamo, P. M. Boyle, and N. Akoum. Machine learning and the conundrum of stroke risk prediction. *Arrhythmia & Electrophysiology Review*, 12:e07, 2023. The dissertation author is a co-author in this paper.

2.2 *Physics-informed neural network*

Physics-informed neural networks (PINN) is a deep learning-based computational predictive framework designed to infer quantities governed or well approximated by differential equations with respect to spatial-temporal coordinates [237]. In the simplest formulation, PINN employs a fully-connected neural network to learn nonlinear mappings between input coordinates and output variables, which are subsequently used to construct residuals of the underlying governing equations, typically ordinary differential equations (ODEs) or partial differential equations (PDEs). These residuals are combined with additional constraints, such as measurement data, initial conditions, and boundary conditions, to form a composite loss function. Training then proceeds by minimizing this composite loss, ensuring that the predicted variables both fit the available data and satisfy the imposed physical laws.

This section introduces the fundamentals of the PINN framework using the inviscid Burgers' equation as an illustrative example and discusses the state-of-the-art in PINN formulations, including network settings and training strategies. Finally, we review current PINN implementations in cardiovascular flow applications.

2.2.1 *Methodology overview*

Consider the inviscid Burgers' equation, which is a fundamental partial differential equation occurring in various areas of applied mathematics, such as fluid mechanics, nonlinear acoustics, gas dynamics, traffic flow, and mathematical physics.

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = 0, \quad t \in [0, T], \quad x \in \Omega, \quad (2.2)$$

subject to the initial condition,

$$u(x, 0) = f(x). \quad (2.3)$$

We then represent the unknown solution $u(x, t)$ by a fully-connected neural network $[x, t] \rightarrow [u_\theta(x, t)]$, where $[x, t]$ are the spatial-temporal coordinates used as inputs, u_θ is the output, and θ denotes the neural network trainable weights and biases. This then allows us to define a residual equation of the Burgers' equation as,

$$R_\theta(x, t) = \frac{\partial u_\theta}{\partial t} + u_\theta \frac{\partial u_\theta}{\partial x}. \quad (2.4)$$

Suppose we also have some sparse measurement data of this unknown solution $u(x, t)$ as,

$$u(x_{\text{data}}, t_{\text{data}}) = u_{\text{data}}, \quad (2.5)$$

and our objective is to use PINN to predict the solution in the whole spatial-temporal space including those where we don't have prior knowledge of the solution. Then, a PINN model can be trained by minimizing the following weighted mean squared error loss function i.e., $\text{argmin}(L_{\text{total}}, \theta)$,

$$L_{\text{total}} = \lambda_{\text{data}} L_{\text{data}} + \lambda_{\text{ic}} L_{\text{ic}} + \lambda_{\text{r}} L_{\text{r}} \quad (2.6)$$

where

$$L_{\text{data}} = \frac{1}{N_{\text{data}}} \sum_{i=1}^{N_{\text{data}}} \|u_\theta(x_{\text{data}}^i, t_{\text{data}}^i) - u_{\text{data}}\|^2, \quad (2.7)$$

$$L_{\text{ic}} = \frac{1}{N_{\text{ic}}} \sum_{i=1}^{N_{\text{ic}}} \|u(x_{\text{ic}}^i, 0) - f(x_{\text{ic}}^i)\|^2, \quad (2.8)$$

$$L_{\text{r}} = \frac{1}{N_{\text{r}}} \sum_{i=1}^{N_{\text{r}}} \|R_\theta(x_{\text{r}}, t_{\text{r}})\|^2. \quad (2.9)$$

In Eq. 2.6, λ_{data} , λ_{ic} , and λ_{r} represent the weight coefficients of each loss term dictating their contribution to the training and optimization process. And in above loss terms $\{x_{\text{data}}^i, t_{\text{data}}^i\}_{i=1}^{N_{\text{data}}}$, $\{x_{\text{ic}}^i\}_{i=1}^{N_{\text{ic}}}$, and $\{x_{\text{r}}^i, t_{\text{r}}^i\}_{i=1}^{N_{\text{r}}}$ denote the sampled spatial-temporal coordinates, and they can be the whole spatial-temporal domain (e.g., full-batch) or be randomly drawn a subsample during each training

iteration (e.g., mini-batch). After the training is complete, the PINN will be able to infer the $u(x, t)$ solution inside the whole domain.

2.2.2 *State-of-the-art*

Physics-informed neural network (PINN) has emerged as a powerful paradigm within scientific machine learning by embedding known physical laws and constitutive relations, typically expressed as partial differential equations (PDEs), directly into the training of neural networks [237]. One of the principal advantages of PINN is their ability to learn physically consistent solutions with little or no labeled data, relying on governing equations and sparse measurements. This makes PINN particularly attractive for inverse and data assimilation problems, where direct observations of all state variables or parameters are unavailable or expensive to obtain. Unlike purely data-driven approaches, PINN incorporates prior physical knowledge as an inductive bias, which improves interpretability and robustness in data-scarce regimes [238].

Another key strength of PINN is its flexibility and generality. PINN provides a unified framework capable of addressing forward problems, inverse problems, and parameter identification by modifying only the loss function while retaining the same neural network architecture [237, 238]. Moreover, PINN frameworks are mesh-free, which means collocation points can be arbitrarily sampled in the spatial-temporal domain, eliminating the need for structured meshes and simplifying applications to complex geometries, moving boundaries, and high-dimensional domains. The continuous and differentiable representation offered by neural networks also allows direct evaluation of solution fields and their derivatives via automatic differentiation [REF], which is advantageous for sensitivity analysis, uncertainty quantification, and coupling with optimization or control tasks.

Despite these advantages, PINN faces several fundamental challenges that currently limit their reliability and scalability. One major difficulty lies in optimization, PINN training involves minimizing a highly nonconvex loss function that couples PDE residuals, boundary and initial conditions, and possibly data fidelity terms. As a result, convergence is sensitive to initialization, optimizer choice, and hyperparameter tuning, and low training loss does not necessarily imply an accurate physical solution [239, 240]. Standard gradient-based optimizers such as Adam and L-BFGS may stagnate in poor local minima, particularly for stiff, convection-dominated, or nonlinear PDEs [241]. To mitigate these issues,

hybrid optimization strategies combining Adam and L-BFGS-B [242], careful nondimensionalization and normalization of governing equations [243], and improved initialization strategies based on low-fidelity solutions [244] have been proposed.

Another challenge is the balancing and enforcement of multiple physical constraints inside the loss terms. PINN typically impose PDEs and boundary conditions as soft penalties, and neural networks tend to prioritize minimizing the dominant loss terms. There is currently no universal, principled method for selecting loss weights, and improper balancing can lead to violation of boundary conditions or incomplete satisfaction of the governing equations [238, 245]. To address this limitation, self-adaptive and gradient-normalized PINN have been introduced, in which loss weights are dynamically adjusted during training to prevent any single constraint from dominating [246, 247]. Alternative formulations, such as variational PINN and conservative PINN, seek to enforce physics more strongly by leveraging weak or integral forms of the governing equations [248].

PINN also struggles to accurately capture multi-scale, high-frequency, or chaotic dynamics, including turbulent flows and long-time integration of nonlinear systems. This limitation is often attributed to the spectral bias of neural networks, whereby low-frequency components are learned preferentially over fine-scale features [240, 249]. Studies have shown that vanilla PINN may fail for high Reynolds number Navier–Stokes equations or convection-dominated diffusion problems due to increasingly complex loss landscapes [241]. Proposed remedies include domain decomposition approaches such as conservative or composite PINN [245], as well as causality-aware and time-marching formulations that respect the physical propagation of information from initial and boundary conditions [250].

From a theoretical standpoint, PINN currently lacks the rigorous convergence, stability, and error estimation guarantees that underpin classical numerical methods. PINN error can be decomposed into approximation, optimization, and generalization components, each of which is difficult to control in practice [251]. While recent work has established approximation bounds for wide neural networks in Sobolev spaces [252], these results do not ensure that the trained PINN converges to the true solution. A posteriori error estimation and reliability analysis remain largely open problems, particularly for high-dimensional PDEs and safety-critical applications [238].

Currently, PINN is best viewed not as direct replacements for classical solvers, but as complementary tools that thrive in inverse problems, data assimilation, surrogate modeling, and hybrid physics-data settings. Ongoing research focuses on improving optimization robustness, constraint enforcement,

multi-scale representation, and theoretical understanding, with the goal of bridging the gap between PINN and established numerical methods. Advances along these directions are essential for the broader adoption of PINN in real-world scientific and engineering applications.

2.2.3 Applications in cardiovascular flow

PINN has been increasingly applied to cardiovascular flow modeling as a physics-constrained framework that integrates sparse clinical measurements with governing hemodynamic equations. By embedding the incompressible Navier-Stokes equations directly into the learning process, PINN enables the reconstruction of velocity, pressure, and other fields of interest without requiring fully specified boundary or initial conditions, which are often unavailable in vivo. This capability makes PINN particularly suitable for patient-specific cardiovascular flow applications, where measurements are limited by imaging resolution, noise, and acquisition constraints [237, 253].

A key application of PINN in cardiovascular flows is the non-invasive reconstruction of hemodynamic quantities from medical imaging data. PINN has been used to infer arterial blood pressure fields directly from 4D flow MRI velocity measurements by enforcing momentum and mass conservation, thereby eliminating the need for invasive catheter-based pressure measurements [243]. In this setting, PINN acts as a physics-guided data assimilation tool that compensates for limited spatial and temporal resolution in clinical imaging while preserving physical consistency. In ventricular flows, recent extensions of left ventricular vector flow mapping incorporate PINN that utilizes color-Doppler ultrasound data together with continuity and Navier–Stokes constraints to reconstruct either velocity fields alone [149] or both velocity and pressure fields [28].

PINN has also been applied to cerebrovascular hemodynamics, where anatomical complexity and sparse measurements pose significant challenges for conventional CFD. Full velocity and pressure fields have been reconstructed from flow visualization or limited velocity data in intracranial aneurysm and cerebral artery configurations, enabling the estimation of clinically relevant quantities such as pressure gradients and wall shear stress [254, 255]. These approaches demonstrate that PINN can recover physically consistent hemodynamics in complex vascular geometries without the need for extensive meshing or precisely defined boundary conditions.

Another important application area is the estimation of wall shear stress, which plays a critical

role in vascular remodeling and disease progression. PINN-based formulations have shown improved robustness in recovering near-wall velocity gradients from sparse or noisy measurements by leveraging the governing equations as regularization, even when boundary conditions are uncertain [256]. This property addresses a key limitation of purely data-driven approaches and traditional post-processed CFD reconstructions in clinical settings.

Beyond flow field reconstruction, PINN has been extended to cardiovascular inverse problems involving flow-tissue interactions. In arterial thrombus modeling, PINN has been employed to infer thrombus material properties, such as permeability, by coupling the Navier-Stokes equations with phase-field transport models. This approach enables simultaneous estimation of velocity, pressure, and biomechanical parameters from limited observational data, providing a non-invasive pathway for characterizing pathological cardiovascular conditions [257].

Chapter 3

SUPER-RESOLUTION LEFT VENTRICULAR FLOW AND PRESSURE MAPPING BY NAVIER-STOKES-INFORMED NEURAL NETWORK (AI-VFM)

3.1 Introduction

Echocardiography is the most used imaging modality in the clinical setting to measure intracardiac flow, as it is non-invasive, portable, and widely available [121]. Techniques such as color-, pulsed-, and continuous-wave Doppler facilitate the bedside assessment of intracardiac blood flow by providing time-resolved velocities in the direction of the ultrasound beam. However, despite its well-known accuracy and accessibility, Doppler velocity measurements are unidirectional, often missing critical information about flow patterns.

Two- and three-directional vectorial maps of intracardiac flow can be acquired using phase-contrast cardiac magnetic resonance imaging (PC-MRI) [105, 106, 258, 259]. However, compared with echocardiography, PC-MRI has limited temporal and spatial resolutions, involves long acquisition times, and is less accessible, which can hamper its application in clinical practice [260–262]. Computed tomography (CT) is more widely available, can assess cardiovascular anatomy with high spatial resolution in the imaging plane, and offers the possibility to obtain volumetric data. However, no analysis tools quantify flow directly from CT images. While computational fluid dynamics (CFD) analysis can be performed using CT-derived segmentations [79], CFD is critically sensitive to modeling choices and boundary conditions.

Given its potential value in the diagnosis and prognosis of heart failure, ultrasound flow mapping inside the left ventricle (LV) has received significant attention. Several ultrasound methods have been developed to capture LV flow, balancing temporal and spatial resolutions with clinical accessibility. Echo particle image velocimetry (Echo-PIV) can obtain time-resolved 2D intraventricular velocity fields with good spatial and temporal resolution from contrast-infused images [263, 264], and 3D extensions have been proposed [265]. However, its sensitivity to operational parameters like image quality, frame rate, insonation angle, and absolute flow velocity [266], and the consequent need for a finely tuned

contrast infusion, have hindered the clinical application of Echo-PIV. Blood Speckle Imaging (BSI) is an emerging ultrasound modality that calculates two-dimensional blood flow by tracking speckles generated by moving blood cells in consecutive frames without the need for injecting contrast agent, simplifying LV flow mapping. However, this technique has been limited so far by the shallow penetration of high-frequency ultrasound. Although its feasibility has been demonstrated in adults [267], BSI is most often applied in the pediatric setting [268].

In the past decade, vector flow mapping (VFM) has gained acceptance as a modality to obtain time-resolved 2D velocity maps from color-Doppler echocardiographic acquisitions in the LV [31–37], as well as 3D velocity maps from triplane color-Doppler [147]. Early implementations of VFM recovered the cross-beam velocity component by enforcing planar mass conservation along circular arcs of the color-Doppler sector [31–33]. Subsequent extensions reformulated VFM as a maximum-likelihood estimation problem with priors based on mass conservation and noise reduction [34, 35]. Other VFM variants use the stream function-vorticity formulation [37, 144]. Despite these advances, existing VFM methods do not incorporate physical information about momentum balance (i.e., the Navier-Stokes equations) and remain sensitive to imaging artifacts that lead to spatial and temporal gaps in the color-Doppler input images. Moreover, current VFM implementations focus exclusively on recovering the cross-beam velocity, overlooking significant flow variables like pressure, which has to be recovered by secondary analyses [146]. The dynamical interplay between pressure and flow acceleration is not used to inform current VFM methods.

In parallel with these efforts, artificial intelligence (AI) models based on deep learning (DL) have swiftly transformed cardiovascular image analysis [50]. Initially focused on automating the interpretation of echocardiograms and electrocardiograms or the segmentation of cardiovascular structures [269–271], these models can enhance flow imaging data by removing phase-wrapping artifacts, reconstructing complete fields from scattered measurements, and increasing the resolution of flow measurements [256, 272, 273]. Physics-informed neural networks (PINNs) harness recent advancements in DL to infer hidden information in flow measurements constrained by the governing equations of fluid mechanics [253, 274]. A notable feature of PINNs is that they can be trained on sparse data and in a patient-specific manner, making them well-suited for VFM. The application of DL models or, more specifically, PINNs to VFM remains underexplored despite noteworthy ongoing work [149].

This chapter introduces AI-VFM, a vector flow mapping method that uses PINNs to reconstruct the

cross-beam flow velocity and fluctuating pressure in the apical long-axis view of the left ventricle. Like other VFM implementations, AI-VFM uses a color-Doppler sequence and a time-resolved delineation of the LV endocardial wall as inputs. Its outputs are a phase-unwrapped color-Doppler map together with reconstructed cross-beam velocity and fluctuating pressure maps. The PINN is trained for each patient-specific acquisition to minimize a loss function that aggregates the difference between the input and output Doppler maps, as well as the residuals of the continuity and Navier-Stokes equations in the imaging plane, including boundary conditions at the endocardial border. Most notably, AI-VFM leverages momentum balance to recover hidden information about flow dynamics which was inaccessible to previous methods.

The chapter is organized as follows. Section §3.2 describes the AI-VFM implementation and the methodology for validation and sensitivity analyses using ground-truth data from CFD simulations. Section §3.3 reports results from those analyses and illustrates the application of AI-VFM to echocardiographic acquisitions, including cases with evident missing information (e.g., truncated ultrasound sectors) or artifacts in the input Doppler images. Section §3.4 discusses the strengths and weaknesses of AI-VFM, suggests applications of this modality, and outlines the potential extensions of AI-VFM.

3.2 *Materials and Methods*

3.2.1 *Governing Equations of LV flow*

This work relies on mathematical models of LV flow to define loss function terms in the physics-informed neural network (PINN) and to generate synthetic ground-truth data for validation via computational fluid dynamics (CFD) analysis. These models use the same governing equations of fluid flow but differ in some aspects, like their physical domains and coordinate systems (see §3.2.1 and §3.2.7 below). In all cases, blood was modeled as an incompressible, Newtonian fluid of constant kinematic viscosity $\nu = 3 \times 10^{-6}$, m²/s. Consequently, the velocity ($\vec{v}$) and pressure (p) fields were assumed to obey the continuity and Navier-Stokes equations, i.e.,

$$\nabla \cdot \vec{v} = 0 \quad (3.1)$$

$$\frac{\partial \vec{v}}{\partial t} + \vec{v} \cdot \nabla \vec{v} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \vec{v} \quad (3.2)$$

where ρ is the density of blood. The boundary conditions for these equations are no-slip at the endocardial wall, i.e., $\vec{v} = \vec{v}_{wall}$.

3.2.2 LV geometrical domain for AI-VFM

In AI-VFM, we adopted a planar polar coordinate system with its origin at the ultrasound transducer (Figure 3.1A), as customary in the VFM literature [33]. In this reference frame, the radial velocity component v_r is parallel to the ultrasound beam and, thus, measurable by color-Doppler. In contrast, the polar velocity component, v_θ , and the pressure field p are unknown. The equations of fluid motion (3.1–3.2) were enforced on a domain Ω corresponding to an LV segmentation obtained from echocardiographic imaging as described in §3.2.10 below. Boundary conditions were imposed on the endocardial border of the segmentation, $\partial\Omega$.

3.2.3 Color-Doppler Phase Unwrapping in AI-VFM

During Doppler acquisitions, phase wrapping (i.e., aliasing) artifacts occur when blood velocity exceeds the maximum detectable velocity (e.g., Nyquist velocity), causing the detected Doppler phase to wrap around 2π and altering the sign and magnitude of the measured velocity, i.e.,

$$\frac{v_r}{v_N} \pi = \frac{v_{Dop}}{v_N} \pi + 2n\pi, \quad (3.3)$$

where v_r is the actual velocity in the direction parallel to the ultrasound beam, v_{Dop} is the measured Doppler velocity, v_N is the Nyquist velocity, and n is the number of phase wraps. Since aliased velocities do not obey the equations of fluid motion (3.1–3.2), Doppler measurements cannot be used directly as training data in AI-VFM. However, trigonometric functions (e.g., sines and cosines) of both sides of equation 3.3 are equal. Thus, we used $\sin(\pi v_r/v_N)$ and $\cos(\pi v_r/v_N)$ as training data fields in our neural network. A similar approach was previously implemented in PINN to perform phase-unwrapping in 4D flow MRI velocities [275]. We found this method yielded flow and pressure fields free of aliasing artifacts with no need for additional constraints and without considering n into the PINN.




Figure 3.1: AI-VFM flowchart. Raw color Doppler radial velocity field is used for training. B-mode wall segmentation to obtain LV boundary points to impose the boundary condition. A fully-connected neural network is employed to infer radial velocity v_r , angular velocity v_θ , and pressure field p . The neural network is penalized by a loss function constructed from training data L_{data} and underlying physical laws L_{BC} and L_{cons} . The total loss L_{total} is minimized using the *Adam* optimizer. Neural network parameters (e.g., weights ϑ^*) are updated via backpropagation during each iteration.

3.2.4 Neural Network Architecture and Super-Resolution Loss Function

AI-VFM (Figure 3.1A) uses physics-informed neural networks (PINNs) to reconstruct the cross-beam flow velocity and the fluctuating pressure from a color-Doppler sequence and a time-resolved delineation of the LV endocardial wall. A PINN is a neural network trained to minimize a residual including physical laws. Therefore, it learns the training data and the underlying (i.e., hidden) dynamics of the system. In AI-VFM, we use a fully-connected neural network to approximate the following mapping function: $[t, r, \theta] \mapsto [v_r, v_\theta, p]$, where the inputs are time and the spatial coordinates of the points inside the LV, and the outputs are the corresponding velocities (radial and cross-beam components) and pressure fluctuations. Input and output variables are defined on N_{out} frames along the cardiac cycle. In all the cases we studied, we adopted a neural network configuration with 4 layers, each consisting of 150 neurons. This configuration aimed to balance low expressivity (shallow layer) with overfitting and computational cost (deep layer). Given that the underlying physical constraints involve nonlinear equations, we employed nonlinear *swish* activation functions [276]. The neural network weights were initialized using Xavier's scheme [277], and weight normalization was employed to accelerate training [278].

The loss function employed to train the AI-VFM PINN comprises three terms representing the differences with respect to the training data ($\mathcal{L}_{data}$), the wall boundary conditions ($\mathcal{L}_{BC}$), and the governing mass and momentum conservation equations ($\mathcal{L}_{cons}$),

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \mathcal{L}_{BC} + \mathcal{L}_{cons}.$$

Color-Doppler v_r fields usually have different temporal resolutions than the B-mode acquisitions used to delineate the LV wall, where boundary conditions are imposed. To account for this disparity, and to achieve temporal super-resolution in the output flow and pressure fields, we devised the strategy outlined in Figure 3.1B. The data loss, $\mathcal{L}_{data}$, was sampled in the available N_{Dop} color-Doppler frames, whereas the boundary condition loss, $\mathcal{L}_{BC}$, was sampled in the available N_{BC} B-mode frames, and the governing equations loss, $\mathcal{L}_{cons}$, was sampled at the N_{cons} frames. We then set the number of output temporal frames, N_{out} , equal to $N_{cons} > N_{BC} > N_{Dop}$ to achieve temporal super-resolution. The data loss term was defined as

$$\mathcal{L}_{data} = \lambda_{data} \left(\frac{1}{M_d} \sum_{i=1}^{M_d} \left\| \cos \left(\frac{\pi v_{r,i}}{v_N} \right) - \cos \left(\frac{\pi v_{Dop,i}}{v_N} \right) \right\|^2 + \frac{1}{M_d} \sum_{i=1}^{M_d} \left\| \sin \left(\frac{\pi v_{r,i}}{v_N} \right) - \sin \left(\frac{\pi v_{Dop,i}}{v_N} \right) \right\|^2 \right), \quad (3.4)$$

where λ_{data} is a hyperparameter and i indicates members within a list of M_d sampled points in the segmented LV mask (Ω) across the available N_{Dop} color-Doppler frames. We accounted for possible holes in the color-Doppler measurements caused by, e.g., low signal power or imaging artifacts, by enabling the selective removal of subdomains of Ω in the calculation of $\mathcal{L}_{data}$. Therefore, the sampled points used to compute $\mathcal{L}_{data}$ were randomly redrawn each training iteration, excluding points within data holes. The boundary condition loss, weighed by the hyperparameter λ_{BC} , included the residual of the no-slip boundary condition sampled at M_{bc} points of the LV boundary ($\partial\Omega$) across the available N_{BC} bright-mode frames,

$$\mathcal{L}_{BC} = \lambda_{BC} \frac{1}{M_{bc}} \sum_{k=1}^{M_{bc}} \|\vec{v}_k - \vec{v}_{wall,k}\|^2. \quad (3.5)$$

Finally, the mass and momentum conservation loss was defined as

$$\mathcal{L}_{cons} = \sum_{j=1}^3 \lambda_{J_j} \frac{1}{M_p} \sum_{i=1}^{M_p} \|\vec{J}_{j,i}\|^2, \quad (3.6)$$

where the vector $\vec{J}_{j,i} = J_j(r_i, \theta_i, t_i)$ is formed by the residuals of the continuity equation and the radial and polar components of the Navier-Stokes equations in polar coordinates, i.e.,

$$\begin{aligned} J_1 &= \frac{1}{r} \frac{\partial}{\partial r} (rv_r) + \frac{1}{r} \frac{\partial v_\theta}{\partial \theta}, \\ J_2 &= \frac{\partial v_r}{\partial t} + v_r \frac{\partial v_r}{\partial r} + \frac{v_\theta}{r} \frac{\partial v_r}{\partial \theta} - \frac{v_\theta^2}{r} + \frac{1}{\rho} \frac{\partial p}{\partial r} - \nu \left[\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} (rv_r) \right) + \frac{1}{r^2} \frac{\partial^2 v_r}{\partial \theta^2} - \frac{2}{r^2} \frac{\partial v_\theta}{\partial \theta} \right], \\ J_3 &= \frac{\partial v_\theta}{\partial t} + v_r \frac{\partial v_\theta}{\partial r} + \frac{v_\theta}{r} \frac{\partial v_\theta}{\partial \theta} + \frac{v_r v_\theta}{r} + \frac{1}{\rho} \frac{1}{r} \frac{\partial p}{\partial \theta} - \nu \left[\frac{\partial}{\partial r} \left(\frac{1}{r} \frac{\partial}{\partial r} (rv_\theta) \right) + \frac{1}{r^2} \frac{\partial^2 v_\theta}{\partial \theta^2} + \frac{2}{r^2} \frac{\partial v_r}{\partial \theta} \right], \end{aligned}$$

To non-dimensionalize the above equations, the training Doppler data v_r and other flow variables were scaled with their characteristic magnitudes, i.e.,

$$\hat{r} = \frac{r}{D}, \quad \hat{\theta} = \theta, \quad \hat{t} = \frac{t}{D/U}, \quad \hat{v}_r = \frac{v_r}{U}, \quad \hat{v}_\theta = \frac{v_\theta}{U}, \quad \hat{p} = \frac{p}{\rho U^2},$$

Example		U [cm/s]	D [cm]	Re	Δr [mm]	$\Delta \theta$ [°]	N_{Dop}	N_{BC}	N_{cons}	Iterations
CFD		50.00	2.73	3413	0.50	1.00	15	100	200	200K
	#1	73.60	3.57	8758	0.54	1.34	17	64	200	45K
ECHO	#2	57.24	2.31	4407	0.54	1.34	19	74	200	45K
	#3	57.24	2.40	4579	0.54	1.35	19	83	200	45K

Table 3.1: Characteristic scales used to non-dimensionalize inputs. Reynolds number defined by $Re = UD/\nu$, where $\nu = 3 \times 10^{-6} \text{m}^2/\text{s}$ is the kinematic viscosity, and U and D are the peak diastolic jet velocity and the diameter of the mitral valve, respectively. Spatial resolution for CFD benchmark and echocardiographic data. Case-specific N_{Dop} , N_{BC} , and $N_{cons} = 200$. Iterations of AI-VFM for each example.

where U is the typical diastolic jet velocity, D is the mitral valve diameter, and ρ is the blood density. The values of these magnitudes are specified in **Table 3.1**. The reconstructed flow fields were then back-dimensionalized during the inference stage after training is complete. The resulting non-dimensional residuals were

$$\begin{aligned}
J_1 &= \frac{1}{\hat{r}} \frac{\partial}{\partial \hat{r}} (\hat{r} \hat{v}_r) + \frac{1}{\hat{r}} \frac{\partial \hat{v}_\theta}{\partial \hat{\theta}}, \\
J_2 &= \frac{\partial \hat{v}_r}{\partial \hat{t}} + \hat{v}_r \frac{\partial \hat{v}_r}{\partial \hat{r}} + \frac{\hat{v}_\theta}{\hat{r}} \frac{\partial \hat{v}_r}{\partial \hat{\theta}} - \frac{\hat{v}_\theta^2}{\hat{r}} + \frac{\partial \hat{p}}{\partial \hat{r}} - \frac{1}{Re} \left[\frac{\partial}{\partial \hat{r}} \left(\frac{1}{\hat{r}} \frac{\partial}{\partial \hat{r}} (\hat{r} \hat{v}_r) \right) + \frac{1}{\hat{r}^2} \frac{\partial^2 \hat{v}_r}{\partial \hat{\theta}^2} - \frac{2}{\hat{r}^2} \frac{\partial \hat{v}_\theta}{\partial \hat{\theta}} \right], \\
J_3 &= \frac{\partial \hat{v}_\theta}{\partial \hat{t}} + \hat{v}_r \frac{\partial \hat{v}_\theta}{\partial \hat{r}} + \frac{\hat{v}_\theta}{\hat{r}} \frac{\partial \hat{v}_\theta}{\partial \hat{\theta}} + \frac{\hat{v}_r \hat{v}_\theta}{\hat{r}} + \frac{1}{\hat{r}} \frac{\partial \hat{p}}{\partial \hat{\theta}} - \frac{1}{Re} \left[\frac{\partial}{\partial \hat{r}} \left(\frac{1}{\hat{r}} \frac{\partial}{\partial \hat{r}} (\hat{r} \hat{v}_\theta) \right) + \frac{1}{\hat{r}^2} \frac{\partial^2 \hat{v}_\theta}{\partial \hat{\theta}^2} + \frac{2}{\hat{r}^2} \frac{\partial \hat{v}_r}{\partial \hat{\theta}} \right],
\end{aligned}$$

where $Re = UD/\nu$ is the Reynolds number. Additionally, we employed a normalization scheme to ensure all training input variables (e.g., r , θ , t) had zero mean and unit variance. This step promotes training robustness and mitigates the issue of vanishing gradients during backpropagation [277].

These loss terms were sampled at M_p points randomly selected across N_{cons} frames. The points were chosen inside the time-dependent Ω domain obtained by interpolating the LV mask segmentation from N_{BC} frames to N_{cons} frames. The temporal and spatial derivatives appearing in the definition of $\vec{J}_{j,i}$ were computed using automatic differentiation [279] on the neural network.

3.2.5 Adaptive Loss Function Weight Adjustment

The loss function in AI-VFM has several terms weighed by hyperparameters that affect the convergence rate of the PINN. Even if we found AI-VFM to converge without tuning these hyperparameters, faster convergence can be achieved by careful tuning. However, manual tuning is tedious and not generalizable for different flows in different patients. Thus, we employed the automatic adaptive strategy proposed by Wang *et al* [280]. In this approach, derived from Neural Tangent Kernel (NTK) theory, the hyperparameter values (e.g., λ_{data} , λ_{BC} , $\lambda_{J_{1-3}}$) are updated every few iterations (100 in our case) starting from some initial values (one in our case). The updated values are computed as

$$\lambda_{data} = \frac{Tr(\mathbf{K})}{Tr(\mathbf{K}_{data})}, \quad \lambda_{BC} = \frac{Tr(\mathbf{K})}{Tr(\mathbf{K}_{BC})}, \quad \lambda_{J_1} = \frac{Tr(\mathbf{K})}{Tr(\mathbf{K}_{J_1})}, \quad \lambda_{J_2} = \frac{Tr(\mathbf{K})}{Tr(\mathbf{K}_{J_2})}, \quad \lambda_{J_3} = \frac{Tr(\mathbf{K})}{Tr(\mathbf{K}_{J_3})}, \quad (3.7)$$

where $Tr(*)$ indicates the trace of a matrix, and $\mathbf{K}_{data}$, $\mathbf{K}_{BC}$, $\mathbf{K}_{J_1}$, $\mathbf{K}_{J_2}$, and $\mathbf{K}_{J_3}$ are the NTK matrices associated with each loss term. These matrices are computed from the Jacobian matrices with respect to the neural network parameter ϑ^* (e.g., $\mathbf{K}_{data} = \mathbf{J}_{data}\mathbf{J}_{data}^T$). And the overall NTK matrix $\mathbf{K}$ is obtained by

$$\mathbf{K} = \begin{bmatrix} \mathbf{J}_{data} \\ \mathbf{J}_{BC} \\ \mathbf{J}_{J_1} \\ \mathbf{J}_{J_2} \\ \mathbf{J}_{J_3} \end{bmatrix} \begin{bmatrix} \mathbf{J}_{data}^T & \mathbf{J}_{BC}^T & \mathbf{J}_{J_1}^T & \mathbf{J}_{J_2}^T & \mathbf{J}_{J_3}^T \end{bmatrix}.$$

Since the eigenvalues of each loss term's NTK matrix measure that term's rate of convergence [280], equation 3.7 effectively increases the weight of the constraints whose residuals are not decreasing as fast as the other constraints in the loss.




Figure 3.2: Loss curves (panels A & D), loss weight evolution (panels B & E), and eigenvalues of neural kernel (panels C & F) associated with each loss component for the AI-VFM with NTK adaptive weight adjustment [280], calculated at the last iteration of training. The top and bottom rows respectively correspond to the CFD validation and clinical application patient #1 reported in Table 3.1.

3.2.6 PINN runs

In all the PINN runs reported in this manuscript, a batch size of 10K sampling points (e.g., $M_d = M_p = 10^4$) was used during each iteration to evaluate the losses with respect to training data and physical constraints, and the whole batch (i.e., all available points) was used to impose boundary conditions. The entire batch size for boundary conditions varied for CFD and ECHO depending on spatiotemporal resolution, with typical values being $M_{bc} \approx 59K$ and $23K$ for CFD and ECHO, respectively.

The PINNs were run until $\mathcal{L}_{data}$, $\mathcal{L}_{BC}$, and $\mathcal{L}_{cons}$ reached plateaus, as shown in Figure 3.2A and 3.2D (patient #1). The number of iterations run for each case is indicated in **Table 3.1**. The total loss function was minimized using the *Adam* adaptive optimization algorithm [281] with an initial learning rate of 0.001 that was lowered exponentially every 10^4 iterations with a decay rate of 0.9. The hyperparameters λ_{data} , λ_{BC} , and $\lambda_{J_{1-3}}$ associated with each loss term were updated every 100 iterations as discussed above. Figures 3.2B & E show that these hyperparameters were rapidly adjusted from their initial guesses during initial training, followed by a more gradual variation for later stages of training. We also examined the eigenvalues of each NTK matrices at the last iteration of training, with and without multiplying by their respective λ (Fig. 3.2C & F). These data showed that hyperparameter adjustment effectively equalized the convergence rate of all individual loss terms, since the eigenvalues of $\lambda\mathbf{K}$ were approximately the same for all terms.

AI-VFM was implemented in Python using TensorFlow [282] library as described in Raissi et al. 2020 [254]. The training was performed using an NVIDIA A40 GPU with 48GB RAM, providing favorable accuracies in the recovered pressure and velocity fields within 10^4 iterations (≈ 15 mins with one A40, Figure 3.16). All other pre- and postprocessing was done in MATLAB. The CFD validation data and codes are publicly available at <https://github.com/bmaidu/AIVFM>.

3.2.7 LV geometrical domain for CFD

The LV cavity of our CFD simulation data was modeled as a half-ellipsoid (Figure 3.3A) with the endocardial surface given by

$$\left(\frac{2x}{D_0}\right)^2 + \left(\frac{2y}{D_0}\right)^2 + \left(\frac{z}{L(t)}\right)^2 = 1, \quad (3.8)$$

where $D_0 = 5.57$ cm is the diameter of the LV base and $L(t)$ is the time-dependent LV long-axis




Figure 3.3: A) Sketch of the computational domain and idealized LV geometry, showing the mitral valve orifice in orange and the aortic valve orifice in blue; B) Time evolution of the LV volume prescribed in CFD simulations.

dimension (i.e., the distance between the LV base and its apex). LV beating in sinus rhythm at 70 b.p.m. was modeled by varying $L(t)$ between $L_{\min} = 3.62$ cm and $L_{\max} = 7.96$ cm, yielding the LV volume vs. time evolution shown in Figure 3.3B. The two increments of LV volume correspond to the E and A waves mediated by LV relaxation and atrial contraction, with a physiological E/A wave ratio of 1.1. LV ejection fraction ($EF = 0.54$) and sphericity index ($SpI = 0.49$) were also set within physiological ranges [283, 284].

The openings at the LV base shown in Figure 3.3A represent the aortic and mitral orifices. They were modeled without leaflets and with two states: fully open or fully closed. The aortic orifice was modeled as a circle of diameter $D_{AO} = 2.52$ cm centered at $x_{AO} = 0$ and $y_{AO} = -1.48$ cm (i.e., separated by a distance 0.05 cm to the edge of the LV base), yielding an area $A_{AO} = 4.97$ cm^2 . A half-moon-shaped mitral orifice was defined by the boundary line:

$$\left(\left(\frac{2x}{D_0} \right)^2 + \left(\frac{2y}{D_0} \right)^2 - 1 \right) \cdot \left(\left(\frac{2x}{D_{AO}} \right)^2 + \left(\frac{2(y - y_{AO})}{D_{AO}} \right)^2 - 1 \right) = \varphi \quad (3.9)$$

where the parameter φ was chosen to yield an orifice area $A_{MO} = 7.45$ cm^2 .

3.2.8 CFD runs

Following previous works [87], the blood flow simulations of the idealized LV configuration described above were performed using TUCAN, an in-house GPU flow solver [88, 142, 143, 285]. TUCAN uses a fractional step method to solve the Navier-Stokes equations for an incompressible fluid. The spatial discretization is based on second-order finite differences in a Cartesian staggered grid with constant grid spacing in all directions, Δx . Time integration is performed using a three-step, semi-implicit Runge-Kutta. The presence and motion of the LV walls are modeled in TUCAN using the immersed boundary method [286].

The simulations were initialized with zero velocity inside the LV and run for 5 cardiac cycles to ensure a quasi-periodic and converged flow. All the data presented in this paper correspond to the 5th cycle, for which 200 uniformly-spaced velocity and pressure fields were stored. For the present work, we chose $\Delta t = 1.47 \times 10^{-4}$ s and $\Delta x = 0.5$ mm, yielding a Courant-Friedrichs-Lewy number (CFL) < 0.6 throughout the cardiac cycle. These temporal and spatial resolution values agree with recent convergence studies for intracardiac flows [287].

3.2.9 Validation and Sensitivity Analysis on CFD-generated LV flow

Velocity and pressure fields from CFD simulations were interpolated into an $11 \text{ cm} \times 60$ degree fan-shaped sector mimicking the long-axis apical echocardiographic view (Figure 3.3A). In this scheme, the ultrasound transducer would be located at the sector's vertex, which also coincides with the origin of the polar coordinate system used in the VFM model (Figure 3.1A). To match the spatial resolution of standard color-Doppler acquisitions (see §3.2.10), the CFD data were interpolated using radial and polar step sizes $\Delta r = 0.5$ mm and $\Delta \theta = 1.0^\circ$, yielding a mesh with 220×50 points. Also, 200 temporal frames of flow fields evenly distributed over one cardiac cycle were obtained for training and validation.

To test the accuracies of *dynamic* and *kinematic* AI-VFMs' when varying the temporal resolution of the input color-Doppler sequence, we first considered using all available training data (e.g., $N_{Dop} = N_{BC} = N_{cons} = 200$). We then tested *dynamic* AI-VFM's ability to provide super-resolution against more realistic temporal resolution of wide-sector color-Doppler and bright-mode acquisitions (e.g., $N_{Dop} = 15, N_{BC} = 100, N_{cons} = 200$). To evaluate AI-VFM's ability to recover hidden information in untrained frames, we ran another *dynamic* AI-VFM using $N_{Dop} = N_{BC} = N_{cons} = 15$ and linearly

interpolated the results to populate the velocity field up to 200 frames. The sensitivity of AI-VFM with respect to imaging resolution was tested by considering down-sampled cases with $\Delta r = 1\text{mm}$, $\Delta\theta = 1.3, 1.8^\circ$ and one up-sampled case $\Delta r = 0.5\text{mm}$, $\Delta\theta = 0.3$. We also tested whether increasing N_{Dop} from 15 to 50, 100, or 200 improved the accuracy of the recovered fields in both *dynamic* and *kinematic* models.

Aliasing artifacts were simulated by introducing a Nyquist limit of $v_N = 50\text{ cm/s}$ in v_r unless stated otherwise. In addition, we investigated the effect of transducer misalignment with the LV and outflow jets by rotating the virtual Doppler sector as described in Figure 3.6(E & F). When varying one parameter, all other parameters are kept fixed at nominal values (i.e., $\Delta r = 0.5\text{mm}$, $\Delta\theta = 1.0^\circ$, $N_{Dop} = 200$, and $\alpha = 0$).

To compare the recovered flow fields vs. the ground-truth CFD data, we computed the coefficient of determination, r^2 , and the normalized root-mean-squared error (NRMSE) as follows:

$$NRMSE = \frac{1}{\max\|\vec{v}_{CFD}\|} \sqrt{\frac{1}{N_t} \sum_{i=1}^{N_t} \|\vec{v}_{VFM_i} - \vec{v}_{CFD_i}\|^2}. \quad (3.10)$$

An interesting feature of AI-VFM is that, by incorporating momentum equations in its loss functions, it can also infer flow pressure fields, $p(r, \theta, t)$. Pressure in incompressible flows is undetermined by a constant (i.e., pressure enters the governing equations 3.2 under a spatial gradient). Thus, to compare the inferred pressure maps with ground-truth maps from CFD, we examined the pressure fluctuations $p' = p - \bar{p}$ by subtracting each snapshot's mean pressure, $\bar{p}(t) = \Omega^{-1} \int_{\Omega} p(r, \theta, t) d\Omega$.

3.2.10 Echocardiographic Imaging

To assess the clinical feasibility of AI-VFM, color-Doppler image sequences were acquired using a Vivid 7 scanner with broadband transducers (GE Healthcare) from three patients in sinus rhythm, one with dilated cardiomyopathy (patient #1) and two with acute myocardial infarction (patient #2 & #3). Standard B-mode and Doppler data were obtained as recommended in established guidelines [21]. Pulsed-wave Doppler spectrograms were obtained at the mitral valve tips and LV outflow tract, from which key cardiac cycle events were identified using EchoPac software (version 214, GE Healthcare) to temporally align data for postprocessing. Color-Doppler images covering the entire LV chamber were obtained in the apical three-chamber view (~ 10 cycles, frame rate $\approx 15\text{ Hz}$), followed by the acquisition of 2D

cine-loops (~ 3 cycles, frame rate > 60 Hz). The Ethics Committee of Hospital Gregorio Maranon approved the study and all patients provided written informed consent. The LV myocardial wall was segmented from the LV B-mode series using speckle-tracking software to delineate the endocardium (EchoPac, version 214, GE Healthcare). The spatial and temporal resolutions of the acquired Doppler images (e.g., v_r) and flow parameters used in AI-VFM for training (N_{Dop}), imposing boundary condition (N_{BC}), and mass and momentum conservation equations (N_{cons}) are summarized in **Table 3.1**.

3.2.11 Reference (Vanilla) VFM and Phase-Unwrapping Methods

For reference in our validation study, we also derived time-resolved 2D blood velocity fields from color-Doppler data inside the LV using Garcia *et al.*'s original method (as described in references [31, 102], and hereon referred to as vanilla VFM). The vanilla VFM algorithm is fed by a color-Doppler acquisition and integrates the planar continuity equation, imposing no-penetration boundary conditions at the LV endocardium. This method used as input a color-Doppler sequence acquired over ~ 10 heartbeats (under breathhold) at ~ 15 frames per second. The input sequence is assumed to be periodic in time, and the frames from the ~ 10 different heartbeats are gated retrospectively to generate velocity fields with higher temporal resolution. This procedure uses a Fourier interpolant with a number of modes chosen by an Akaike information criterion, which usually yields 10-20 modes. Data from heartbeats with $> 5\%$ heart rate variation are rejected. In our experience with vanilla VFM, having patients hold their breath for ~ 10 seconds can introduce challenges. However, this is no longer an issue in AI-VFM, which only requires one heartbeat of color-Doppler input data given its temporal superresolution capabilities.

As reference dealiasing method, we used Loecher *et al.*'s one-step Laplacian correction [288] as the reference method for phase unwrapping of aliased color-Doppler values. In this method, a Laplacian equation is formulated for the number of phase wraps,

$$\nabla^2 n = (2\pi)^{-1} [\cos(\varphi_w) \nabla^2 \sin(\varphi_w) - \sin(\varphi_w) \nabla^2 \cos(\varphi_w) - \nabla^2 \varphi_w],$$

where $\varphi_w = \pi v_w / v_N$ is the wrapped phase and v_w the aliased velocity. The solution to this equation is rounded off to the nearest integer number and plugged into equation 3.3 to correct the velocity field.

3.3 Results

This section presents validation and sensitivity analyses of AI-VFM using CFD-generated data. We also illustrate the clinical application of AI-VFM to improve ultrasound flow data. Overall, AI-VFM accurately recovers the cross-beam velocity component and fluctuating pressure inside the LV at temporal resolution higher than that of the training Doppler data. It also deals with aliasing artifacts and recovers missing velocity values.

3.3.1 AI-VFM Validation on CFD-generated LV flow

Figure 3.4 displays snapshots of LV flow velocity and fluctuating pressure fields during the early filling, late filling, and ejection phases of the cardiac cycle. This figure includes the CFD-generated ground-truth data and the results from two AI-VFM runs. In the first run, AI-VFM was trained using $N_{cons} = N_{BC} = N_{Dop} = 200$ temporal frames, including the three timepoints represented in Figure 3.4. In the second run, AI-VFM was trained using a more realistic setting with $N_{cons} = 200$, $N_{BC} = 100$, $N_{Dop} = 15$. Of note, in this second AI-VFM run, none of the color-Doppler data from the three time points represented in Figure 3.4 was used for training. Still, AI-VFM inferred the entire velocity field at these time points, thereby increasing the effective temporal resolution of the original data.

The CFD-generated data, used as ground truth for evaluating AI-VFM, showed LV flow features in good agreement with previous simulations [67, 289] and clinical measurements [121]. As in the human LV, the CFD-simulated flows exhibited strong filling jets flanked by vortex rings and an emptying jet with prograde swirl. These patterns are the hallmarks of systolic and diastolic LV flow. Thus, although idealized, our CFD model provided an effective, physiologically representative benchmark flow for validating AI-VFM.

At time points with color-Doppler training data, AI-VFM inferred v_r almost exactly (Figure 3.4A), which was expected since the $\mathcal{L}_{data}$ loss directly penalized differences between the original and the reconstructed v_r component. Moreover, outside of the timepoints with color-Doppler training data, AI-VFM also provided a faithful estimation of v_r . The cross-beam velocities from AI-VFM closely matched the ground-truth CFD maps (Figure 3.4B). Because AI-VFM did not use any v_θ data for training, this flow component estimation was not as precise as v_r . Nevertheless, AI-VFM still recovered the dominant patterns of v_θ even if some fine-scale features were not captured. The availability of color-Doppler




Figure 3.4: Velocity and pressure maps inferred by AI-VFM compared to ground-truth CFD data at early filling, late filling, and ejection (i.e., $t = 0.85, 0.45, 0.21T$, where T is the cardiac cycle period). Data from two AI-VFM implementations, using PINNs trained with $N_{Dop} = 200$ and $N_{Dop} = 15$, are shown. As indicated by the dials symbolizing the cardiac cycle, the $N_{Dop} = 200$ AI-VFM is shown at frames (•) within the training set (○). In contrast, the $N_{Dop} = 15$ AI-VFM is shown at frames outside of the training set. A) Radial velocity v_r ; B) polar velocity v_θ ; C) Flow vectors and color map of velocity magnitude; D) Flow vectors and color map of fluctuating pressure p' . Units for velocity and pressure are cm/s and $mmHg$, respectively.

training data at the interrogated temporal frame did not significantly influence the reconstruction of v_θ . This result indicates that AI-VFM learned the dynamics of the flow to incorporate data across different time points despite the low temporal resolution in the training data.

The p' maps produced by AI-VFM agreed well with the ground-truth fields. During early filling, AI-VFM recovered the pressure fluctuations caused by the vortex ring flanking the mitral jet, and the vortex ring from the preceding beat as it impinged the LV apex. During late filling, AI-VFM reproduced pressure fluctuations induced by vortices as well as the adverse pressure gradient that decelerates the LV filling jet during this phase. Likewise, the pressure inferred by AI-VFM during ejection agreed with the CFD data, showing a favorable pressure gradient that accelerates the flow toward the aortic orifice.

The only noticeable difference between the pressure maps in AI-VFM and CFD was that AI-VFM produced smoother results, reaching lower p' peak values than CFD. This effect was more pronounced when AI-VFM was interrogated at temporal frames with no color-Doppler training data. PINNs are known to implicitly introduce smoothing by prioritizing learning the low-frequency features in the training data, a phenomenon known as spectral bias [280, 290]. To investigate whether this effect could be palliated by reducing the importance of the viscous terms in the PINN, we ran *dynamic* AI-VFM setting $Re = \infty$ in the Navier-Stokes loss terms (i.e., J_2 and J_3). The resulting flow maps were almost indistinguishable from those obtained using the Reynolds number value used in the CFD, $Re = 3,413$ (Figure 3.14).

Point-wise Comparison of Velocity and Pressure Values

Scatter plots comparing v_r , v_θ , and p' from AI-VFM (y coordinate) and CFD (x coordinate) were built by pooling all the spatial points within the interrogated temporal frames, as indicated below. We overlaid contours of the joint probability density functions (pdfs) of VFM and CFD values and computed their r^2 values. For reference, the outermost contour of these pdfs was chosen so that it left out 1% of the sampled points.

We tested the *dynamic* and *kinematic* versions of AI-VFM, along with the vanilla VFM method of Garcia *et al* [31], using all the CFD validation frames as input frames (e.g., $N_{Dop} = N_{BC} = N_{cons} = 200$ for the PINNs). The results are displayed in the first row of Figure 3.5. The v_r and v_θ scatter plots from *kinematic* AI-VFM and vanilla VFM are shown as insets in the corners of each panel (top left and




Figure 3.5: Point-wise comparison of velocity and fluctuating pressure inferred by AI-VFM with ground-truth CFD data. Each panel shows scatter plots where the x and y coordinates are, respectively, the ground-truth CFD and AI-VFM values of flow variables sampled at uniformly spaced points inside the LV domain. The contour maps of joint probability densities of CFD and AI-VFM values are superimposed on each scatter plot. Units for velocity and pressure are cm/s and $mmHg$, respectively. A-C) AI-VFMs vs. CFD, $N_{Dop} = N_{BC} = N_{cons} = 200$; D-F) AI-VFMs vs. CFD, $N_{Dop} = 15$, $N_{BC} = 100$, $N_{cons} = 200$ pooling frames within the $N_{Dop} = 15$ color-Doppler dataset; G-I) AI-VFMs vs. CFD, $N_{Dop} = 15$, $N_{BC} = 100$, $N_{cons} = 200$ pooling frames outside of the $N_{Dop} = 15$ color-Doppler dataset

bottom right corners of Figures 3.5A-B, respectively). Consistent with the results shown in Figure 3.4, the scatter plots for v_r aligned tightly along the identity line $y = x$ for the three methods considered (Figure 3.5A), with r^2 hovering around 0.99 in all the cases. Although the scatter plots of v_θ experienced more dispersion than those of v_r , *dynamic* AI-VFM still showed favorable agreement with the CFD ground-truth data, leading to $r^2 = 0.62$. We found r^2 to vary little with the number of training iterations for v_r , implying that a high accuracy was already achieved with few iterations ($\sim 10K$, Figure 3.16). On the other hand, r^2 increased appreciably with the number of iterations for v_θ (Figure 3.16). In comparison to *dynamic* AI-VFM, *kinematic* AI-VFM and vanilla VFM reached $r^2 = 0.36$ and $r^2 = 0.25$ respectively. Correspondence with ground-truth data was better for large v_θ values, while more significant differences were observed around $v_\theta = 0$. Of note, *kinematic* AI-VFM and vanilla VFM showed larger errors than *dynamic* AI-VFM when reconstructing negative values of v_θ . Pressure fluctuations from *dynamic* AI-VFM showed good agreement with CFD data (Figure 3.5C), with an overall $r^2 = 0.94$. The agreement was particularly good for low and intermediate $|p'|$ values, whereas AI-VFM tended to underestimate extreme values, consistent with the visualizations of Figure 3.4D. Similar to v_r , high r^2 values were achieved for p' with as little as $10K$ training iterations, and the accuracy increased modestly beyond that (Figure 3.16).

Next, we tested the ability of AI-VFM to learn hidden information about flow physics and recover velocity and pressure fields with finer temporal resolution than the training data. Thus, we ran AI-VFM using only $N_{Dop} = 15$ frames with training v_r data, but we imposed the boundary conditions and governing equations on more frames along the cardiac cycle ($N_{BC} = 100$, $N_{cons} = 200$). We then compared AI-VFM with ground-truth CFD data both at the $N_{Dop} = 15$ within the training set (Figure 3.5D-F) and at the 185 frames outside of the training set (Figure 3.5G-I). Overall, the accuracy of *dynamic* AI-VFM decreased modestly when N_{Dop} was lowered from 200 to 15, with r^2 values for v_r , v_θ and p' varying from 0.99, 0.62, and 0.94 to 0.97-1.00, 0.56-0.57, and 0.86-0.87, respectively. Remarkably, *dynamic* AI-VFM had similar accuracy when interrogated within the 15 frames containing training data (Figure 3.5D-F) and the 185 frames not containing training data (Figure 3.5G-I).

Kinematic AI-VFM conserved its performance when trained with $N_{Dop} = 15$ instead of $N_{Dop} = 200$, both within the training range (top left corner insets of Figures 3.5D, E) and outside of it (top left corner insets of Figures 3.5G, H). However, its accuracy decreased when tested outside of the training dataset, displaying drops in r^2 from 1.00 to 0.92 and from 0.37 to 0.33 respectively for v_r and v_θ . Moreover, we

ran *dynamic* AI-VFM using $N_{out} = N_{cons} = N_{BC} = N_{Dop} = 15$, and linearly interpolated its results in time at the 185 outside of the training dataset. This run aimed to evaluate the level of hidden flow information recovered by the PINN for different values of N_{BC} and N_{cons} . Their scatter plots are shown in the lower right corner insets of Figures 3.5G-I. The accuracy of the linearly interpolated AI-VFM maps dropped significantly, reaching values very similar to those obtained by *kinematic* AI-VFM.

We also tested the accuracy of imposing the no-penetration boundary condition routinely used in most VFM methods [33] since AI-VFM provides flexibility to enforce either free slip or no slip at the LV walls. The results yielded (scatter plots not shown) $r^2 = 0.99, 0.50, 0.76$ for v_r , v_θ , and p' within the training dataset and $r^2 = 0.96, 0.51, 0.80$ outside of it. Therefore, the accuracy of the no-penetration boundary condition only decreased slightly with respect to the no-slip boundary condition.

Sensitivity Analyses

To assess the robustness of AI-VFM, we repeated the CFD-based validation for different spatial and temporal resolutions and transducer locations, as outlined in §3.2.9. To evaluate the distinct contributions of the Navier-Stokes and mass conservation constraints, we performed sensitivity analyses of the *dynamic* and *kinematic* implementations of AI-VFM. Figure 3.6 summarizes the results of these analyses, reporting NRMSEs of v_r , v_θ , and p' for *dynamic* AI-VFM, and NRMSEs of v_r and v_θ for *kinematic* AI-VFM.

First, we varied the angular resolution while keeping nominal values for transducer positioning and radial and temporal resolutions. Since a standard color-Doppler acquisition has an angular step size $\Delta\theta \approx 1.0^\circ$, we considered step sizes $\Delta\theta \approx 0.3, 1.0, 1.3$, and 1.8° in our sensitivity analysis. The NRMSEs of the reconstructed flow variables varied little with $\Delta\theta$ in the studied range for both *dynamic* AI-VFM (Figure 3.6A) and *kinematic* AI-VFM (Figure 3.6B). However, *dynamic* AI-VFM produced lower errors consistent with the results reported in the previous section. We also addressed the effect of radial resolution, varying the standard echocardiographic value, $\Delta r \approx 0.5mm$, to $\Delta r \approx 1mm$. Similar to the case of $\Delta\theta$, this perturbation did not significantly modify the NRSME of the reconstructed flow variables in any of the two AI-VFM models (Figure 3.6C-D). Once again, *dynamic* AI-VFM was more accurate than *kinematic* AI-VFM over the range of studied resolutions.

We also examined the effect of transducer positioning by varying the angle α formed by the ultrasound sector axis and the CFD-LV long axis. For *dynamic* AI-VFM (Figure 3.6E), the NRMSEs of v_r and




Figure 3.6: Normalized root mean squared error (NRMSE) of the radial velocity v_r (blue), polar velocity v_θ (red), and pressure fluctuation p' (yellow) maps recovered by AI-VFM using CFD analysis as ground-truth data. The top row (panels A, C, E, and G) corresponds to *dynamic* AI-VFM whereas the bottom row (panels B, D, F, and H) corresponds to *kinematic* AI-VFM. A-B) NRMSE vs. spatial resolution in the polar direction (i.e., angular step size $\Delta\theta$); C-D) NRMSE vs. spatial resolution in the radial direction (i.e., radial step size Δr); E-F) NRMSE vs. ultrasound sector orientation (i.e., angle α between sector axis and LV long axis); G-H) NRMSE vs. temporal resolution of color-Doppler training sequence (i.e., N_{Dop} or $\Delta t = 1/N_{Dop}$). When varying one parameter, all other parameters are kept fixed at nominal values (i.e., $\Delta r = 0.5\text{mm}$, $\Delta\theta = 1.0^\circ$, $N_{dop} = 200$, and $\alpha = 0$).

p' were almost independent of α in the range $[-45^\circ, 45^\circ]$ while the NRMSE of v_θ experienced a shallow increase with $|\alpha|$. In the case of *kinematic* AI-VFM (Figure 3.6F), the NRMSEs of v_r were also insensitive to α . However, the reconstruction error of v_θ increased more sharply with $|\alpha|$ in *kinematic* AI-VFM than in *dynamic* AI-VFM. In particular, the NRMSEs of *dynamic* AI-VFM at $\alpha = \pm 45^\circ$ (0.21 and 0.30) were comparable to that of *kinematic* AI-VFM at $\alpha = 0$ (0.25).

Finally, we studied the AI-VFM performance while varying N_{Dop} between 15 and 200. We found that both *dynamic* and *kinematic* AI-VFM were relatively insensitive to N_{Dop} . However, a modest increase in the NRMSEs of v_r was observed between $N_{Dop} = 50$ and $N_{Dop} = 15$ for *kinematic* AI-VFM (Figure 3.6H).

3.3.2 Application to Clinical Acquisitions

This section demonstrates the feasibility of performing AI-VFM on clinical echocardiographic acquisitions. We illustrate this technique's ability to correct aliasing artifacts in the color-Doppler data and reconstruct 2D maps of flow velocity vectors and pressure fluctuations. We also demonstrate that AI-VFM can restore missing data in clinical acquisitions, including entire color-Doppler frames. Since the validation and sensitivity analyses reported in previous sections suggest *dynamic* AI-VFM is more accurate and robust than *kinematic* AI-VFM, our clinical demonstration of AI-VFM focused on the *dynamic* implementation. Unless otherwise stated, we ran the PINNs using patient-specific values for N_{Dop} and N_{BC} in **Table 3.1** given by the temporal resolution of each echocardiographic acquisition. In contrast, $N_{cons} = 200$ frames were used to enforce mass conservation and momentum balance.

Figure 3.7 displays the velocity and pressure maps obtained by AI-VFM for patient # 1. These maps are represented during early filling, late filling, and ejection. For completion, the raw color-Doppler maps at those time instants, some of which contain aliasing artifacts, are also represented (Figure 3.7A-C). Similar plots for two additional patients are shown in Supplementary Figures 3.11 and 3.12. The velocity vectors captured the early filling jet and its associated vortex ring (Figure 3.7D), which evolves into a single prograde swirling structure that occupies most of the LV chamber by late diastole (Figure 3.7E) and persists through ejection (Figure 3.7F), routing incoming blood toward the outflow tract. This spatiotemporal pattern is well-established as the representative LV flow pattern in the apical long-axis view [121]. In addition to the velocity field, AI-VFM calculated the pressure fluctuations inside the LV




Figure 3.7: Example of clinical application of AI-VFM in patient #1. The first (panels A, D, and G), second (panels B, E, and H), and third (panels C, F, and I) columns display respectively early filling, late filling, and ejection. A-C) Raw color-Doppler training frames at indicated time points. D-F) Flow vectors and color map of velocity magnitude from AI-VFM at indicated time points. G-I) Flow vectors and pressure fluctuation at indicated time points. Velocity units are $[m/s]$ and pressure units are $[mmHg]$. $N_{Dop} = 17$, $N_{BC} = 64$, $N_{cons} = 200$ for *dynamic* AI-VFM.




Figure 3.8: Example of clinical application of AI-VFM in patient #1 to color-Doppler sequences with scan sector truncation near the LV boundary. The first (panels A, D, and G), second (panels B, E, and H), and third (panels C, F, and I) columns display respectively early filling, late filling, and ejection. A-C) Raw color-Doppler training frames at indicated time points displaying narrowed scan sector near the LV boundary. D-F) Flow vectors and color map of velocity magnitude from AI-VFM at indicated time points. Note that AI-VFM reconstructs the velocity field inside the truncated regions. G-I) Flow vectors and pressure fluctuation at indicated time points. Note that AI-VFM reconstructs the velocity and pressure fields inside the truncated regions. Velocity units are $[m/s]$ and pressure units are $[mmHg]$. $N_{Dop} = 17$, $N_{BC} = 64$, $N_{cons} = 200$ for *dynamic* AI-VFM.

from clinical color-Doppler acquisitions (Figures 3.7G-I). These pressure maps identified vortex cores since those are associated with local pressure minima, and agreed well with previously reported maps by secondary analyses of VFM maps [146].

To illustrate the capacity of AI-VFM to handle images with holes and regenerate the missing information, we reran the algorithm on the same patient, i.e., patient #1. However, this time, we truncated the color-Doppler sector in all the frames (N_{Dop}) of the input sequence (Figure 3.8A-C). Sector truncation is a frequent manipulation that increases the acquisition frame rate of color-Doppler sequences. Despite the truncated sector, the flow and pressure maps reconstructed by AI-VFM, shown in Figures 3.8D-I, agreed favorably with those obtained when running AI-VFM on intact input images (Figures 3.7D-I), yielding $r^2 = 0.87, 0.64, 0.88$ for v_r, v_θ , and p' , respectively. We also tested the case where the color-Doppler images had a hole in the center of the LV cavity, as in Figure 3.13. The flow and pressure maps reconstructed by AI-VFM, shown in Figure 3.13D-I, were almost indistinguishable from Figure 3.7D-I, yielding $r^2 = 0.97, 0.91, 0.96$ for v_r, v_θ , and p' .

Finally, we followed a leave-one-out approach to show that AI-VFM can also recover complete missing frames in clinical acquisition sequences. We created decimated versions of patient #1's color-Doppler acquisition by sequentially removing one of its N_{Dop} color-Doppler frames. Then, we trained the PINN in *dynamic* AI-VFM using the remaining $N_{Dop} - 1$ frames, and used the color-Doppler frame excluded from training as ground-truth data for validation. Figure 3.9 compares ground-truth frames with the corresponding AI-VFM-reconstructed frames during early diastole, late diastole, and ejection, showing favorable visual agreement, and yielding $r^2 = 0.84$ with respect to training with all available (N_{Dop}) frames. In comparison, linear interpolation between phase-unwrapped frames yielded $r^2 = 0.73$.

3.3.3 Phase unwrapping in AI-VFM

Phase wrapping artifacts are common in clinical color-Doppler images and can be particularly problematic during early diastole. During this phase, the flow velocity reaches large negative values inside the filling jet. If these velocity values are aliased, a large region of positive Doppler velocity, including the LV diastolic vortices, can be created, as shown in Figure 3.7A. This pattern can confound physics-unaware phase unwrapping algorithms, leading to jagged velocity contours that require artificial smoothing. We tested the accuracy and sensitivity of phase-unwrapping in AI-VFM using the CFD-generated flow as




Figure 3.9: Leave-one-out validation of AI-VFM on the clinical color-Doppler sequence of patient #1. The three columns show results from three runs, where *dynamic* AI-VFM was trained on $N_{Dop} - 1$ color-Doppler frames, and one frame was dropped. The top image in each column is the dropped color-Doppler frame. The bottom image is the v_r map recovered by AI-VFM for the instant of time corresponding to the dropped frame. Velocity units are $[m/s]$. $N_{Dop} = 17$, $N_{BC} = 64$, $N_{cons} = 200$.

benchmark. Aliased regions were created (Figure 3.10A) by prescribing a Nyquist velocity of 40 cm/s. The one-step Laplacian method of Loecher *et al* [288] was used as the reference algorithm. We found that both the Laplacian method (Figure 3.10B) and AI-VFM (Figure 3.10C) correctly unwrapped most of the aliased regions. However, the Laplacian method failed in small regions close to the boundary and the inlet (see arrows in Figure 3.10B). Consistent with these data, the scatter plots of phase unwrapped v_r vs. the original CFD v_r (Figure 3.15) showed a small but noticeable amount of pixels along the lines $y = x \pm v_N$ for the Laplacian method when $v_N = 40$ cm/s. We also examined the sensitivity of both methods to the Nyquist velocity (Figure 3.10D), finding that AI-VFM performed well for Nyquist velocities down to $v_N = 20$ cm/s. On the other hand, the Laplacian method deteriorated for $v_N \leq 40$ cm/s.




Figure 3.10: Color-Doppler dealiasing of radial velocity v_r analysis in CFD model (top row) and one example in a clinical acquisition (bottom row). A) CFD frame simulated during early diastole with an aliased region (orange) manually created in the center part of inflow jet (blue) with the Nyquist velocity $v_N = 40$ cm/s. B) CFD frame dealiased using the one-step Laplacian method of Loecher *et al.*'s [288]. C) CFD dealiased output frame from AI-VFM. D) NRMSE vs. Nyquist velocity in the idealized CFD LV model for both methods.

3.4 Discussion

Machine learning (ML) and AI are shaking the landscape of fluid mechanics research, offering unprecedented ability to study complex flows. In cardiovascular research, ML models are powerful tools to discover hidden relationships between clinical variables and diseases, offering valuable metrics for disease risk stratification [236,291,292]. The impact of deep learning (DL) models has been particularly

significant [293]. Among these models, deep convolutional networks (CNNs) have proven effective at pattern recognition and feature extraction in complex flow data. Applied to cardiovascular flows, CNNs have proven helpful in resolving Doppler images aliasing artifacts [272], aortic wall shear stress estimation from 4D flow MRI [273], and real-time estimation of thrombotic risk in left atrial appendage [294]. Among the growing body of DL models, PINNs are well suited to study flow phenomena since they effectively incorporate the laws of physics as constraints in the learning process [237, 253]. PINNs have been used to enhance flow measurements from laboratory experiments [295] and medical imaging acquisitions [256, 275]. However, its applications to recover and visualize intracardiac flows from clinical data are scarce, and only preliminary approaches to the cardiovascular system have been recently conducted [243, 256, 257]. Some of these applications aim to use PINNs for improving cardiac simulations or modeling [271, 296, 297].

The application of DL models to reconstruct LV flow from color-Doppler data is just beginning to be explored but offers significant promise of improvement over traditional VFM methods [149, 298]. Specifically, Ling *et al* [149] proposed physics-guided neural networks by PINNs and nnU-Net [299] to reconstruct intraventricular flow from Doppler data based solely on mass conservation. Their comparative study has demonstrated the generalizability and robustness of nnU-Net method trained on a diverse range of simulated and clinical Doppler data. While near real-time inference of flow fields using pre-trained model shows promise in clinical settings, the requirement of a alias-free Doppler as training data complicates the applicability of nnU-Net method, and no constraints on flow dynamics require secondary analysis to obtain other flow fields such as pressure. On the other hand, our study centered around PINNs as the means of flow reconstruction, focusing on the performance gains obtained by incorporating physical constraints related to flow kinematics (i.e., mass conservation) and dynamics (i.e., momentum balance). Unprecedented in DL-based VFM approaches, we included momentum balance in the AI-VFM via loss terms that measure the residuals of the planar Navier-Stokes equations. Informing AI-VFM with flow dynamics allowed this modality to provide temporal super-resolution and recover entire missing frames from an acquisition sequence. Furthermore, since the Navier-Stokes equations include pressure gradients, the pressure was an output variable of the PINNs, allowing AI-VFM to recover pressure fluctuations inside the LV. However, since blood is an incompressible fluid, these fluctuations do not account for the time-periodic mean chamber pressure or hydrostatic variations due to gravity.

Our comparative analyses of *kinematics*-informed vs. *dynamics*-informed AI-VFM suggested that enforcing momentum balance in the PINN increased the accuracy and robustness of VFM. Overall, *dynamic* AI-VFM was more accurate than *kinematic* AI-VFM, which in turn was more accurate than the vanilla VFM implementation that integrates the continuity equation separately along each arc of the Doppler sector [31]. Moreover, our sensitivity analyses showed *dynamic* AI-VFM to be less sensitive to the input data's spatial and temporal resolutions and the ultrasound sector's orientation. The sensitivity to ultrasound sector alignment is particularly relevant since VFM relies on the inflow and outflow jets of the LV being captured by the color-Doppler velocity component. Because those jets are primarily aligned with the LV long axis, the performance of VFM deteriorates as the angle α between the LV long axis and the ultrasound sector increases [31]. Still, we found *dynamic* AI-VFM to perform as accurately for $\alpha = 45^\circ$ as *kinematic* AI-VFM for $\alpha = 0$.

The Navier-Stokes equations include acceleration and viscous terms representing physical mechanisms distinct from mass conservation. Informing AI-VFM with flow acceleration allowed this method to connect information from different temporal frames within each training set. Consequently, *dynamic* AI-VFM produced temporally super-resolved outputs with velocity and pressure maps at time points without color-Doppler training data. When forcing *dynamic* AI-VFM to linearly interpolate between training frames without enforcing momentum balance inside the intervals defined by these frames, its performance degenerated to that of *kinematic* AI-VFM. This experiment supports the idea that enforcing momentum balance is crucial to achieving temporal super-resolution. From a practical point of view, handling color-Doppler input sequences with few more than a dozen frames spanning one single heart-beat facilitates the clinical adoption of VFM because it reduces the time patients need to hold their breath while being scanned from ~ 10 s to ~ 1 s. In our experience, breath-holding has been a limiting factor of VFM, especially in patients with severely impaired cardiac function. This feature also opens the door to applying AI-VFM to patients with heart rhythm disorders.

In color-Doppler acquisitions, the Nyquist velocity limit beyond which the Doppler phase wraps around 2π is dictated by the ultrasound pulse repetition frequency (PRF). Increasing the PRF allows for detecting higher flow velocities but compromises the scanning depth of the acquisition. Consequently, color-Doppler sequences for VFM are often acquired using Nyquist velocities lower than the peak velocities found in the inflow and outflow jets of the LV. Thus, these sequences can contain significant aliasing artifacts. Since aliasing artifacts are common to other flow imaging modalities like phase-

contrast MRI, they have received significant attention, and standalone algorithms have been developed for their removal [272, 288, 300, 301]. These algorithms use knowledge about phase-wrapping, training data, and/or smoothness constraints to decide whether jumps in velocity between neighboring pixels are caused or not by phase wrapping. Using physical constraints should help make this decision but had not been attempted in color-Doppler echocardiography so far. In AI-VFM, we followed an approach recently proposed for 4D-flow MRI [275]. We made the data loss term in the PINN insensitive to 2π jumps in the phase of the color-Doppler signal and let the physics constraints in the PINN ensure the regularity of the output flow maps. We found this approach was sufficient to correct aliasing artifacts in several clinical color-Doppler acquisitions without introducing any explicit dealiasing step in the algorithm. We also found that the dealiased color-Doppler maps generated by AI-VFM compared well with those obtained with standalone methods informed by phase wrapping rules [288].

In addition to aliasing, clinical color-Doppler images sometimes include sizeable areas of unreliable data caused by ultrasound sector truncation, insufficient signal power, device-related ultrasound reflections [302], etc. These artifacts are more challenging to remedy than aliasing because they are not governed by a simple phase-wrapping rule. We have shown that AI-VFM seamlessly tolerates large holes in the input data, using the hidden information about flow physics to restore the flow inside these holes. Therefore, AI-VFM offers a new method to analyze artifact-ridden images as long as the unreliable pixels in each image can be labeled and excluded from training. Future efforts should focus on the automation of unreliable pixel labeling and the design of data loss terms that tolerate such pixels, similar to the 2π -insensitive data loss term used in this work.

The inclusion of viscous terms in the PINN loss should contribute to regularizing the velocity fields in ways consistent with flow physics, even if flow inside the human LV is dominated by pressure and inertial forces. The flexibility of the PINN framework allows for blood viscosity to be an output parameter [237], so no prior knowledge of this parameter is needed. Nevertheless, our results suggest that spectral bias in the PINNs [280, 290] created at least as much smoothing as the viscous terms for physiologically representative Reynolds number values. Therefore, the *effective* Re of the AI-VFM-reconstructed flow maps was lower than the actual Re of the flow and, consequently, the peak values of, e.g., pressure fluctuations, were underestimated by AI-VFM. On the other hand, except for the presence of imaging noise, the clinical color-Doppler maps do not exhibit the fine-scale flow patterns expected to be found in LV flow at physiologically representative Reynolds number values of $Re \approx 4,000$. This averaging effect

in color-Doppler maps was previously noted by Seo *et al* [289] by comparing with their CFD simulations. Therefore, while further efforts are warranted to mitigate spectral bias in AI-VFM [247,280], the clinical application of this technique should not be hindered by this limitation. Furthermore, our leave-one-out validation of AI-VFM on clinical acquisitions suggests that the PINN captured most if not all the flow features in the color-Doppler maps.

The PINN approach followed in AI-VFM ensures this technique is trained using data from one patient at a time, bypassing the need for training on large pre-existing data sets. PINNs are generally less sensitive to common issues of DL models like the misrepresentation of features underrepresented in the training set [50, 291, 292]. While the need for re-training the PINN for every single patient may seem computationally expensive, this demand is manageable in VFM given the relatively light weight of the input 2D image sequences. AI-VFM provides high accuracy for v_r and p and reasonable accuracy for v_θ with < 10 minutes of training on an off-the-shelf contemporary graphics processing unit (GPU), and the accuracy of v_θ increases significantly by increasing the number of training iterations. AI-VFM works robustly without the need for carefully fine-tuning hyperparameters to the specifics of each acquisition, and improved convergence and accuracy are achieved with adaptive loss function weight adjustment [247, 250, 280]. In particular, adjusting the PINN loss weights using the eigenvalues of the neural tangent kernel [280] is an effective approach to equalize the rate of convergence of the different loss terms in AI-VFM. Considering the rapid progress in GPU architecture, and the Tensorflow software engine used by the PINNs [303], we expect that further reductions in compute time that could bring AI-VFM closer to real-time execution are well within reach. On the other hand, it is important to keep in mind that VFM uses a planar flow simplification, i.e., the physical loss terms are formulated by restricting the 3D mass conservation and momentum balance equations within the 2D plane formed by the long-axis apical view of the LV. Therefore, one should expect the loss function in AI-VFM to have a floor that is related to the error of the planar flow approximation. Future efforts should leverage the flexibility and modest computational cost of PINNs to extend AI-VFM to three dimensions.

The fully-connected neural network configuration in AI-VFM provides a flexible framework that can be expanded without great difficulty to include additional input and output variables, as well as physical constraints. In the field of echocardiography, this flexibility could be exploited to assimilate flow fields from blood speckle imaging [267, 268], a technique based on high-frequency ultrasound that measures 2D blood flow vectors but often produces gappy data. In addition, future directions will explore including

additional physical constraints in the PINNs of AI-VFM to map LV blood residence time [304] or the concentration of coagulation species [305].

3.5 Conclusion

This chapter describes AI-VFM, a new vector flow mapping (VFM) method enabled by recent advances in artificial intelligence (AI). AI-VFM uses physics-informed neural networks (PINNs) encoding mass conservation, momentum balance, and boundary conditions to recover intraventricular flow and pressure fields from standard echocardiographic scans. AI-VFM performs phase unwrapping and recovers missing data in the form of spatial and temporal gaps in the input color-Doppler data, thereby producing super-resolution flow maps. AI-VFM is solely informed by each patient's flow physics; it does not utilize explicit smoothness constraints or incorporate data from other patients or flow models. AI-VFM shows good validation against ground-truth data from CFD, outperforming traditional VFM methods as well as similar PINN-based VFM formulations relying exclusively on mass conservation.

3.6 Acknowledgements

This work was supported by the National Institutes of Health under grants 1R01HL160024 and 1R01HL158667, the National Science Foundation Graduate Research Fellowship under grant No.DGE-2140004, the Spanish Research Agency and the European Regional Development Fund under grant PID2019-107279RB-I00, the Comunidad de Madrid and the European Regional Development Fund under grant Y2018/BIO-4858 PREFI-CM, and the Instituto de Salud Carlos III under grants PI20/00587-AI4NHEM, PI21/00274-PACER, and DTS/1900063-ISBIFLOW.

The results presented in this chapter are a published manuscript [28]: B. Maidu, P. Martinez-Legazpi, M. Guerrero-Hurtado, C. M. Nguyen, A. Gonzalo, A. M. Kahn, J. Bermejo, O. Flores, and J. C. del Alamo. Super-resolution left ventricular flow and pressure mapping by Navier–Stokes-informed neural networks. *Computers in Biology and Medicine*, 185:109476, 2025. The dissertation author is the first-author in this paper.

3.7 *Supporting Information*

Figures 3.11 and 3.12 in this section demonstrate the clinical application of AI-VFM to two additional patients (#2 and #3, see Table 3.1 in the main text).

Figure 3.13 demonstrates the clinical application of AI-VFM to patient with missing data (holes) at the center of LV chamber due to acquisition limitations.

Figure 3.14 compares CFD and AI-VFM data when the PINN is informed with and without viscous terms.

Figure 3.15 shows the point-wise comparison of dealiasing results by One-step Laplacian method and AI-VFM in CFD LV model.

Figure 3.16 shows the point-wise comparison of velocity and pressure inferred by AI-VFM with CFD ground-truth data for different number of training iterations and point-wise comparison of velocity and fluctuating pressure values inferred by AI-VFM with CFD ground-truth data for different number of batch sizes (M_d, M_p) used to sample training data. Both plots are normalized by their common origins.




Figure 3.11: Example of clinical application of AI-VFM in patient #2. The first (panels A, D, and G), second (panels B, E, and H), and third (panels C, F, and I) columns display respectively early filling, late filling, and ejection. A-C) Raw color-Doppler training frames at indicated time points. D-F) Flow vectors and color map of velocity magnitude from AI-VFM at indicated time points. G-I) Flow vectors and pressure fluctuation at indicated time points. Velocity units are $[m/s]$ and pressure units are $[mmHg]$. $N_{Dop} = 19$, $N_{BC} = 74$, $N_{cons} = 200$ for *dynamic* AI-VFM.




Figure 3.12: Example of clinical application of AI-VFM in patient #3. The first (panels A, D, and G), second (panels B, E, and H), and third (panels C, F, and I) columns display respectively early filling, late filling, and ejection. A-C) Raw color-Doppler training frames at indicated time points. D-F) Flow vectors and color map of velocity magnitude from AI-VFM at indicated time points. G-I) Flow vectors and pressure fluctuation at indicated time points. Velocity units are $[m/s]$ and pressure units are $[mmHg]$. $N_{Dop} = 19$, $N_{BC} = 83$, $N_{cons} = 200$ for *dynamic* AI-VFM.




Figure 3.13: Example of clinical application of AI-VFM in patient #1 to color-Doppler sequences with holes inside the LV domain. The first (panels A, D, and G), second (panels B, E, and H), and third (panels C, F, and I) columns display respectively early filling, late filling, and ejection. A-C) Raw color-Doppler training frames at indicated time points displaying a large hole near the center of the LV chamber. D-F) Flow vectors and color map of velocity magnitude from AI-VFM at indicated time points. Note that AI-VFM reconstructs the velocity field inside the holes. G-I) Flow vectors and pressure fluctuation at indicated time points. Note that AI-VFM reconstructs the velocity and pressure fields inside the holes. Velocity units are $[m/s]$ and pressure units are $[mmHg]$. $N_{Dop} = 17$, $N_{BC} = 64$, $N_{cons} = 200$ for *dynamic* AI-VFM.




Figure 3.14: Velocity and pressure maps inferred by AI-VFM compared to ground-truth CFD data at early filling, late filling, and ejection (i.e., $t = 0.85, 0.45, 0.21T$, where T is the cardiac cycle period). Data from two AI-VFM implementations, using PINNs trained with $Re = 4,550$ and $Re = \infty$, are shown. Both are trained with $N_{Dop} = 15$, $N_{BC} = 100$, and $N_{cons} = 200$. As indicated by the dials symbolizing the cardiac cycle, bot are shown at frames ($\bullet$) outside of the training set ($\circ$). A) Radial velocity v_r ; B) polar velocity v_θ ; C) Flow vectors and color map of velocity magnitude; D) Flow vectors and color map of fluctuating pressure p' . Units for velocity and pressure are cm/s and $mmHg$, respectively.




Figure 3.15: Point-wise comparison of de-aliased radial velocity from One-step Laplacian method (left) and AI-VFM (right) with ground-truth CFD data. Each panel shows scatter plots where the x and y coordinates are, respectively, the ground-truth CFD, radial velocity de-aliased by Loecher *et al* and AI-VFM sampled at uniformly spaced points inside the LV domain. The contour maps of joint probability densities superimposed on each scatter plot. Units for velocity is cm/s .




Figure 3.16: Left: point-wise comparison of velocity and fluctuating pressure values inferred by AI-VFM with CFD ground-truth data for different number of training iterations normalized by r_{10K}^2 ; Right: point-wise comparison of velocity and fluctuating pressure values inferred by AI-VFM with CFD ground-truth data for different number of batch sizes (M_d, M_p) used to sample training data normalized by r_{2500}^2 .

Chapter 4

TIME-RESOLVED THREE-DIMENSIONAL AI-VFM BY TRIPLANE COLOR-DOPPLER IMAGING

4.1 Introduction

In Chapter 3, we developed a novel left ventricular (LV) vector flow mapping (VFM) framework that uses physics-informed neural networks [254], referred to as AI-VFM [28], to reconstruct two-dimensional flow and pressure fields from standard color-Doppler imaging. We validated AI-VFM using an idealized LV dataset generated from computational fluid dynamics (CFD) simulations and successfully applied the framework to clinical electrocardiographic imaging data, obtaining physiologically meaningful flow and pressure fields.

In the domain of color-Doppler ultrasound, several approaches have been proposed to reconstruct three-dimensional flow and pressure fields. For example, prior studies have demonstrated the feasibility of using echo-particle image velocimetry to recover full-volume intraventricular flow fields [306–308]. The classical *i*VFM formulation has also been extended to an optimization framework for reconstructing three-dimensional LV flow fields from triplane Doppler imaging [147]. In addition, volumetric Doppler measurements have been used to recover complete three-dimensional flow fields through image-registration-based approaches [309, 310] and optimization-based regularization methods [311]. More recently, Wong *et al.* [150] proposed a coupled sequential-frame PINN framework to reconstruct three-dimensional flow and pressure fields in a CFD dataset, although the underlying data assumption is more closely aligned with volumetric Doppler imaging than with standard two-dimensional color-Doppler acquisition.

Given that our 2D AI-VFM can be trained using incomplete, sparse, or partially missing spatiotemporal color-Doppler data, we now extend the 2D AI-VFM framework to three dimensions. In this chapter, we present 3D AI-VFM framework and the validation studies demonstrating the use of triplane color-Doppler imaging as training data to reconstruct time-resolved, three-dimensional LV flow and pressure fields.

4.2 Materials and Methods




Figure 4.1: Synthetic triplane Doppler data acquisition on an idealized CFD data. Left panel depicts the triplane separated by 60° . Right panel shows a representative radial velocity component v_r snapshot during early filling for the triplane.

4.2.1 Synthetic triplane Doppler generation

We used the idealized LV model described in Section 3.2.7, together with the CFD-generated ground-truth flow and pressure fields, to validate the 3D AI-VFM framework introduced in this chapter. To mimic clinical triplane Doppler acquisition, we placed a virtual transducer away from the apex and interpolated the CFD flow field onto a fan-shaped polar grid to obtain the radial Doppler velocity field, as shown in the left panel of Fig. 4.1. The polar grid was then rotated twice, with an angular increment of 60° , to generate the other two triplane Doppler velocity fields. The polar grid has a spatial resolution of $\Delta r = 0.5$ mm

and $\Delta\theta = 1.0^\circ$, with a radial length of 11 cm and an angular width of 60° . The spatial coordinates of the triplane Doppler velocity fields were represented in both polar and Cartesian coordinates to facilitate coordinate transformation in the neural network loss formulation. The full volumetric LV mesh is represented in Cartesian coordinates with a spatial resolution of $\Delta x = \Delta y = \Delta z = 0.5$ mm. A representative triplane Doppler snapshot during early filling is shown in the right panel of Fig. 4.1.




Figure 4.2: Flowchart of the proposed framework in this chapter. Spatiotemporal coordinates of the triplane Doppler data (top left) are used as inputs to the 3D AI-VFM (top right). After training, it reconstructs full volumetric flow (bottom left) and pressure fields (bottom right).

4.2.2 3D AI-VFM architecture & training strategy

The Cartesian spatiotemporal coordinates of the LV mesh are used as the input to our 3D AI-VFM framework (Fig. 4.2, top row). The input coordinates $[t, \mathbf{x}]$ are first mapped into a high-dimensional feature space using a random Fourier embedding layer, defined as $\Phi(\mathbf{x}) = [\cos(\mathbf{B}\mathbf{x}), \sin(\mathbf{B}\mathbf{x})]$, where the entries of $\mathbf{B}$ are sampled from $\mathcal{N}(0, s^2)$ with $s = 1.0$. Random Fourier feature embedding helps mitigate the spectral bias of PINNs by introducing high-frequency and multiscale components into the neural network representation [247, 250, 280]. A second structured Fourier layer is applied to the temporal coordinate to enforce exact temporal periodicity [250, 312], given by $\Psi(t) = [\cos(2\pi t/T), \sin(2\pi t/T)]$, where $T = 0.86$ s is the cardiac cycle duration. The low-frequency inputs $[t, x, y, z]$ are also normalized using a z-score scheme and concatenated with the high-frequency Fourier features before being propagated through the network. The resulting transformed input is passed through four adaptive residual blocks [313], each consisting of three MLPs, and the network outputs the four variables $[u, v, w, p]$.

The predicted flow and pressure variables are then used to construct the training loss functions. Specifically, the velocity field $\mathbf{u} = [u, v, w]$ is used to penalize the mismatch between the network predictions and the triplane Doppler measurements, defined as

$$L_{\text{data}} = \lambda_{\text{data}} \left(\frac{1}{M} \sum_{i=1}^M \|\vec{\mathbf{u}} \cdot \vec{\mathbf{e}}_r - v_{r,i}\|^2 \right), \quad (4.1)$$

where λ_{data} is the data-loss weighting coefficient, and i indexes the $M = 10^4$ sampled points from the triplane Doppler planes. We also impose the no-slip boundary condition through a boundary loss weighted by λ_{bc} , evaluated at $M_{\text{bc}} = 10^4$ sampled points on the full volumetric LV boundary,

$$L_{\text{bc}} = \lambda_{\text{bc}} \frac{1}{M_{\text{bc}}} \sum_{i=1}^{M_{\text{bc}}} \|\vec{\mathbf{u}}_i - \vec{\mathbf{u}}_{\text{wall},i}\|^2. \quad (4.2)$$

The mass and momentum balance partial differential equations (PDEs) losses were defined as

$$\mathcal{L}_{\text{pde}} = \sum_{j=1}^4 \lambda_{J_j} \frac{1}{M} \sum_{i=1}^M \|\vec{\mathbf{J}}_{j,i}\|^2, \quad (4.3)$$

where $\vec{\mathbf{J}}_{j,i} = J_j(t_i, x_i, y_i, z_i)$ denotes the j th residual evaluated at the i th sampled point, with $\mathbf{J}$ constructed from the continuity equation and the three Cartesian components of the Navier–Stokes equations. To non-dimensionalize these equations, the training Doppler data v_r and other flow variables

are scaled by their characteristic magnitudes, i.e.,

$$\hat{x} = \frac{x}{D}, \hat{y} = \frac{y}{D}, \hat{z} = \frac{z}{D}, \hat{t} = \frac{t}{D/U}, \hat{v}_r = \frac{v_r}{U}, \hat{u} = \frac{u}{U}, \hat{v} = \frac{v}{U}, \hat{w} = \frac{w}{U}, \hat{p} = \frac{p}{\rho U^2},$$

where $U = 50$ cm/s and $D = 2.73$ cm, and ρ is the blood density. The resulting non-dimensional residuals were,

$$\begin{aligned} J_1 &= \frac{\partial \hat{u}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{u}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{u}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{u}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{x}} - \frac{1}{Re} \left[\frac{\partial^2 \hat{u}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{z}^2} \right], \\ J_2 &= \frac{\partial \hat{v}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{v}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{v}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{v}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{y}} - \frac{1}{Re} \left[\frac{\partial^2 \hat{v}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{v}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{v}}{\partial \hat{z}^2} \right], \\ J_3 &= \frac{\partial \hat{w}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{w}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{w}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{w}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{z}} - \frac{1}{Re} \left[\frac{\partial^2 \hat{w}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{w}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{w}}{\partial \hat{z}^2} \right], \\ J_4 &= \frac{\partial \hat{u}}{\partial \hat{x}} + \frac{\partial \hat{v}}{\partial \hat{y}} + \frac{\partial \hat{w}}{\partial \hat{z}}, \end{aligned}$$

where $Re = UD/\nu = 3412.5$ is the Reynolds number. The reconstructed flow fields were then back-dimensionalized during the inference stage after training is complete (Fig. 4.2 bottom row).

The total loss function consists of multiple terms with different physical scales and numerical magnitudes. To mitigate this imbalance and stabilize the training process, we employ adaptive loss term weighting based on the GradNorm scheme proposed in prior work [247].

$$\lambda_i = \frac{||\nabla_{\theta} L_{ic}(\theta)|| + ||\nabla_{\theta} L_{bc}(\theta)|| + ||\nabla_{\theta} L_{pde}(\theta)||}{||\nabla_{\theta} L_i(\theta)||}$$

where $i \in \{\text{data}, \text{bc}, \text{pde}\}$, and $||\cdot||$ denotes the L^2 norm. The GradNorm scheme dynamically adjusts the loss term weights by equalizing the gradient norms of the weighted losses, thereby preventing any single loss term from dominating the training process. The computed weights are updated using a running average with a momentum coefficient of 0.95, and the GradNorm weights are recalculated every 1000 training iterations.

The neural network weights were initialized using the Glorot scheme [277]. We used the *swish* activation function, and random weight factorization was employed to accelerate training [313]. The network was optimized using the *Adam* optimizer with an initial learning rate of 10^{-3} , which was exponentially decayed by a factor of 0.8 every 25,000 training iterations. All 200 triplane Doppler frames, synthetically acquired from the CFD simulations, were used to train the network and reconstruct the full volumetric flow and pressure fields. The reconstructed fields were then validated against the

CFD-generated ground truth. The results presented in this chapter were obtained after 500,000 training iterations. The 3D AI-VFM framework was implemented in TensorFlow, with training accelerated using JIT and XLA compilation, which have been shown to provide substantial speedups for PINN training [303].

4.3 Results

This section presents qualitative 3D LV flow visualization and quantitative validation of 3D AI-VFM against CFD ground-truth data. Overall, 3D AI-VFM can accurately reconstruct full volumetric flow and pressure fields by training only on triplane Doppler data.




Figure 4.3: Velocity and vorticity maps inferred by 3D AI-VFM (bottom row) compared to CFD ground-truth data (top row) at early filling (left column), late filling (mid column), and ejection (right column). The blue volumes are the vortex regions computed by the Q-criterion ($3000/s^{-2}$ threshold).

4.3.1 Flow and pressure reconstruction visualization and comparison

Fig. 4.3 shows snapshots of the LV flow velocity and vorticity, while Fig. 4.4 shows the corresponding fluctuating pressure maps during early filling (panels *A, D*), late filling (panels *B, E*), and ejection (panels *C, F*) within one cardiac cycle. These figures compare the 3D CFD ground-truth flow fields (top rows) with the reconstructed results from the 3D AI-VFM framework (bottom rows) described in §4.2.2.

Overall, the velocity field inferred by 3D AI-VFM shows very good agreement with the CFD ground truth. During early filling (Fig. 4.3A,D), the inferred results capture the strong mitral jet entering the LV and forming a large vortex ring. During late filling (Fig. 4.3B,E), 3D AI-VFM continues to recover the overall velocity field well, although some smaller vortex structures are not captured and the velocity magnitude is slightly underestimated, as indicated by the shorter velocity arrows. During ejection (Fig. 4.3C,F), the agreement remains strong, despite the loss of fine-scale vortex details and continued underestimation of the relatively small velocities near the inlet region.




Figure 4.4: Pressure fluctuation (iso-surface) maps overlay with velocity vectors (green arrows) inferred by 3D AI-VFM (bottom row) compared to CFD ground-truth data (top row) at early filling (left column), late filling (mid column), and ejection (right column).

Because pressure in incompressible flow is determined only up to an additive constant, we subtracted the spatial mean pressure at each time instant from both the CFD ground-truth and the inferred results, i.e., $p' = p - \bar{p}$. The pressure fluctuation maps p' inferred by 3D AI-VFM (Fig. 4.4) agree well with the CFD ground-truth. During early filling (Fig. 4.4A & D), 3D AI-VFM recovers the pressure distribution associated with the vortex ring, including its overall shape and magnitude. During late filling (Fig. 4.4B & E), the inferred pressure field captures the adverse pressure gradient that decelerates LV filling. Similarly, during ejection (Fig. 4.4C & F), the inferred pressure fluctuation maps agree well with the CFD ground-truth and recover the pressure gradient that drives the flow toward the outlet.

4.3.2 Point-wise and global comparisons against CFD ground-truth data




Figure 4.5: Point-wise comparison of velocity and fluctuating pressure inferred by 3D AI-VFM with CFD ground-truth data.

Fig. 4.5 shows scatter plots comparing the inferred velocity components u , v , and w , as well as the pressure fluctuation p' , against the CFD ground-truth, together with the corresponding coefficients of determination (R^2). Consistent with the flow and pressure fields shown in Fig. 4.3 and Fig. 4.4, the scatter plot for the velocity component u , corresponding to the LV long-axis direction, aligns closely with the identity line $y = x$, with $R^2 = 0.90$. The inferred velocity components v and w also show good overall agreement, although moderate underestimation is observed, particularly near the peak values. The corresponding R^2 values for v and w are 0.67 and 0.42, respectively. In contrast, the pressure fluctuation p' shows very good agreement with the CFD ground-truth, with the scattered points clustered closely around the identity line and only minor underestimation.




Figure 4.6: Volume-averaged kinetic energy at each time instant comparison inferred by 3D AI-VFM with CFD ground-truth data.

Fig. 4.6 shows the volume-averaged kinetic energy computed at each time instant, defined as $\overline{KE} = \frac{1}{2} (\overline{u^2} + \overline{v^2} + \overline{w^2})$. Consistent with the previous velocity field comparisons, the overall mean kinetic energy is moderately underestimated by 3D AI-VFM, reflecting the underestimation of the reconstructed velocity magnitudes.

Frame-wise mean vorticity was also computed and is shown in Fig. 4.7. Despite the minor-to-moderate underestimation of velocity magnitudes described above, the overall vorticity is largely captured




Figure 4.7: Mean vorticity ($[s^{-1}]$) at each time instant comparison inferred by 3D AI-VFM with CFD ground-truth data.

well by 3D AI-VFM. The most noticeable differences are observed in the ω_y component, while the overall trends remain consistent with the vortex structures identified using the Q-criterion in Fig. 4.3.




Figure 4.8: Loss terms and weights vs. training iterations.

The training losses and associated loss-term weights are shown in Fig. 4.8. Overall, the training process remained stable: the loss values decreased rapidly during the initial stage of training and then gradually plateaued. The loss weights were dynamically updated during training, helping to balance the contributions of the individual loss terms and stabilize the overall training dynamics.

4.4 Discussion

This chapter extends our previous 2D AI-VFM framework to 3D by using triplane Doppler imaging to infer time-resolved volumetric LV flow and pressure fields. Motivated by prior 2D *i*VFM work [147], we used triplane Doppler measurements to train the AI-VFM framework while imposing physical constraints, including the incompressible Navier–Stokes equations, on full-volume spatiotemporal points. Triplane Doppler imaging provides sparse velocity measurements with good temporal resolution, enabling AI-VFM to infer flow fields in regions of the LV that are not directly observed. Unlike conventional 3D *i*VFM, 3D AI-VFM can also recover the pressure field because the momentum balance is included in the loss function. In spirit, this approach is similar to recent PINN-based VFM work [150]; however, it differs in imaging modality and formulation. Specifically, the PINN-based VFM framework of Wong *et al.* relies on prior knowledge of volumetric Doppler acquisitions, which are not readily available in most clinical ultrasound settings, whereas 3D AI-VFM uses clinically more accessible triplane Doppler measurements.

We validated the 3D AI-VFM framework using an idealized LV model, for which the ground-truth flow and pressure fields were generated from CFD simulations. The inferred velocity fields and pressure fluctuations showed good overall agreement with the CFD ground truth, although minor to moderate underestimation was observed near peak values. This behavior is commonly observed in PINN-inferred flow fields and is often associated with spectral bias. Although several techniques, such as Fourier feature embeddings, were used to mitigate spectral bias, its effects remain noticeable and limit the ability of 3D AI-VFM to capture fine-scale hemodynamic structures. In the present study, we did not perform sensitivity analyses with respect to spatial resolution, temporal resolution, Doppler aliasing, or measurement noise, which are especially important in clinical settings. Likewise, clinical application of 3D AI-VFM was not evaluated in this study.

Future work will include a comprehensive sensitivity analysis using CFD-generated flow fields

to assess the effects of imaging parameters, Doppler aliasing, and measurement noise, followed by validation using clinical triplane Doppler imaging. Although the 3D AI-VFM framework is flexible and can accommodate different neural network architectures and loss function formulations, the combination of these techniques increases training complexity and computational cost, with the results presented here requiring approximately 25 hours of training. Future work should therefore also focus on accelerating 3D AI-VFM through multi-GPU training, improved optimization strategies, or related computational advances [314].

Chapter 5

LEFT ATRIAL FLOW AND THROMBOSIS RISK FROM 4D CT CONTRAST DYNAMICS BY PHYSICS-INFORMED NEURAL NETWORKS AND INDICATOR DILUTION THEORY

5.1 Introduction

During atrial fibrillation (AF), the left atrium (LA) contracts weakly and irregularly, slowing down blood flow in regions like the left atrial appendage (LAA). These conditions promote the formation of thrombi, some of which may travel to the brain and cause ischemic strokes. Consequently, 15% of all strokes are thought to be caused by AF [315]. These strokes are associated with worse outcomes, more disability, and higher mortality than ischemic strokes in the absence of AF [316]. Anticoagulation therapies for stroke prevention are based on clinical evidence that their benefits statistically outweigh their risks in large groups of AF patients with certain demographic and comorbid factors [317]. These factors are usually aggregated into numerical scores (e.g., CHA₂DS₂-VASc) that, when applied to individual patients, have modest accuracy [318]. Currently, there are no quantitative tools to personalize stroke risk prediction in AF patients.

Given that the vast majority of LA thrombi form inside the LAA [15], estimating LAA thrombosis risk is considered comparable to estimating stroke risk during AF. From a mechanistic standpoint, prolonged blood residence time in the LAA is an appealing biomarker because it directly quantifies the slow flow conditions that favor thrombogenesis. Slow or reduced swirling blood flow inside the LA, measured by 4D flow magnetic resonance imaging (MRI) [319–321], has been associated with AF and ischemic stroke. However, current 4D flow MRI methods cannot resolve flow inside the small-size LAA, especially its distal part. Echocardiographic vector flow mapping (VFM) detects slow flow regions to predict the risk of left ventricular (LV) thrombosis and associated brain embolisms [322–324]. Applying VFM to the LAA would require transesophageal echocardiography, which is unlikely to be adopted to risk-stratify AF patients due to its invasive nature.

Computed tomography (CT) produces high-resolution images with quick acquisitions. While no

methods are yet available to measure blood flow from CT, it is recognized that slow flow can create noticeable contrast gradients in the LAA. In particular, poor CT enhancement in the distal LAA has been associated with stroke risk in AF patients [325, 326]. However, the temporal manifestation and spatial steepness of these gradients are likely to depend on each individual's LA and LAA filling dynamics. Consequently, CT studies using different imaging and analysis protocols have not reached consensus about the relationship between LAA contrast gradient and stroke risk [327]. This limitation suggests that qualitative assessment of LAA opacification alone may not provide a robust patient-specific marker of stasis. However, because iodinated contrast behaves approximately as a passively advected scalar, its spatiotemporal dynamics still encode information about intracardiac transport, motivating quantitative approaches that infer residence time from contrast time curves while accounting for the governing flow physics.

Computational fluid dynamics (CFD) analysis is garnering attention since it can finely resolve stagnant regions inside each patient's LAA from CT image segmentation [87, 89, 90, 135, 137, 328–334]. When parameterized correctly, CFD models can compute essentially exact flow fields by representing the governing equations of flow physics. However, CFD parameterization is sensitive to assumptions about LA wall motion [335], blood rheology [88], and inflow boundary conditions [336]. Furthermore, creating CFD meshes, selecting numerical discretization schemes, and setting CFD resolution to balance accuracy with compute time all require significant technical expertise [337]. Machine learning and, in particular, deep learning (DL) offers a versatile alternative to CFD for finding complex relationships between cardiac chamber anatomy, blood flow, and clotting risk [236, 294, 338]. However, these approaches are often constrained by the requirement for extensive, high-quality labeled biomedical imaging training data, which limits their clinical applicability [292].

Physics-informed neural networks (PINNs) combine patient-specific data with governing equations to infer hidden flow information, with a paradigmatic example being flow reconstruction in a cerebral aneurysm from contrast dynamics [253, 254]. Recent developments in low-dose multi-cycle 4D contrast CT enable imaging of LA/LAA filling and emptying over multiple heartbeats [172], creating an opportunity to apply PINNs to intracardiac transport. In this context, PINNs provide a framework to infer contrast concentration, flow and residence time fields from spatiotemporal contrast data. Patient-specific training minimizes a loss function that includes a data-fidelity term for contrast concentration, the residuals of the Navier–Stokes equations, and transport equations for both contrast concentration

and residence time. Notably, unlike conventional CFD workflows, such a framework does not require an explicit computational mesh of the LA or complete explicit prescription of flow boundary conditions during inference.

Indicator dilution theory (IDT) has a long clinical history for estimating blood flow parameters, vascular volume, and cardiac output following bolus injection of a tracer [339]. Its simplicity and rapid inference capabilities have motivated extensive applications across medical imaging modalities, including ultrasound [340], MRI [341], and CT [342]. In this setting, IDT can be used to infer residence time from contrast time curves by modeling chamber transport as a linear response system with a parameterized transport function (e.g., a gamma-variate function) [343].

In this study, we use CFD-generated contrast dynamics as ground truth to investigate two complementary approaches for inferring LA/LAA residence time from contrast CT: (i) a PINN that reconstructs velocity, contrast concentration, and residence time fields by enforcing the incompressible Navier–Stokes equations together with transport equations for contrast and residence time, and (ii) an IDT framework that computes voxelwise residence time directly from temporal contrast curves under a tank-in-series transport model along flow pathlines. These two approaches are complementary: the PINN targets richer hemodynamic inference, whereas IDT targets rapid residence-time mapping.

This chapter is organized as follows. Section §5.2 describes the dataset, the CFD ground-truth simulations, the PINN and IDT implementations, and the validation procedure. Section §5.3 presents flow visualizations and quantitative comparisons of both methods against CFD ground truth. Section §5.4 discusses the strengths and limitations of the proposed approaches and outlines potential extensions and clinical applications.

5.2 *Materials and Methods*

5.2.1 *Image acquisition and segmentation*

We retrospectively examined $N = 6$ patients from a previous study [87]. Each patient’s left atrial (LA) anatomy was segmented from time-resolved, three-dimensional computed tomography (i.e., 4D CT) imaging following the standard protocols at each participating location (National Institute of Health, Bethesda, MD and University of California San Diego, La Jolla, CA). DICOM files, with spatial resolution between 0.32 and 0.48 mm in the $x - y$ axial plane and 0.5 – 1 mm in the z -direction, were

generated by reconstructing the images following the CT scanner manufacturers' standard algorithms. Time-resolved images were acquired at regular intervals throughout the cardiac cycle, corresponding to 5%–10% of the R–R interval. Three of the patients did not have left atrial appendage thrombus (LAAT) or a history of transient ischemic attacks (TIAs), and had normal LA function. The other three had impaired LA global function and two of them had an LAAT at the time of imaging while the other one had a history of TIAs. These patients are hereafter referred to as LAAT/TIA-negative and LAAT/TIA-positive, respectively. In the LAAT/TIA-positive patients thrombi were digitally removed prior to generating meshes for CFD simulations.

Each patient's LA, left ventricle (LV), pulmonary vein (PVs) inlets, mitral valve (MV) outlet, and LAA were identified and segmented. These data were used to create inlets, outlets, and immersed-boundary Lagrangian meshes of the LA for the CFD simulations. LV and LA volumes were computed to determine the flow rates through the PV inlets and MV outlet as described below. Segmentation and meshing were performed using ITK-SNAP [344] and custom MATLAB scripts similar to our previous works [87, 88, 336]. Table 5.1 summarizes the anatomical and functional features parameters of the patients studied in our work and Fig. 5.1 shows the final segmented left atrial geometries from 4D CT images used for CFD simulations, PINN, and IDT framework.

Case Number	1	2	3	4	5	6
LAA thrombus (or history of TIA)	No	No	No	TIAs	Yes	Yes
Atrial fibrillation	No	No	No	Persistent	Persistent	Persistent
Sinus rhythm	Yes	Yes	Yes	No	No	No
D_{MV} (cm)	2.14	1.78	2.40	2.42	2.43	2.33
U_{MV} (cm/s)	80.44	140.77	105.37	72.15	67.71	138.40
Reynolds number (Re)	4297	6266	6315	4361	4111	8058

Table 5.1: Anatomical and functional parameters of the cases reported in this study. D_{MV} is the mitral valve diameter $D_{MV} = \sqrt{4A_{MV}/\pi}$ used as characteristic length, where A_{MV} is the area of the MV. U_{MV} is the peak velocity during the E-wave $U_{MV} = Q_{MV}^{E-wave}/A_{MV}$ used as characteristic velocity, where Q_{MV} is the flow rate through the MV during E-wave. Reynolds number at the mitral valve annulus is calculated using the characteristic scales and the blood viscosity i.e., $Re = D_{MV}U_{MV}/\nu$.




Figure 5.1: Patient-specific left atrial meshes segmented from computed tomography (CT) images depicting left atrium wall (in gray), left atrial appendage (in red), mitral valve plane (in green), and the pulmonary veins inlets (in blue). Top row shows the LAAT/TIA-negative cases (1-3) whereas the bottom row shows the LAAT/TIA-positive cases (4-6).

5.2.2 Computational fluid dynamics analysis

We used patient-specific CFD simulations to generate reference data for training and testing the PINN models and the IDT framework. Specifically, blood flow, contrast transport, and residence time within the LA were computed using TUCAN [87, 88, 336], an in-house CFD solver for the partial differential equations (PDEs) governing mass conservation, momentum balance (Navier–Stokes), transport of contrast concentration (c), and residence time (T_R),

$$\nabla \cdot \vec{v} = 0, \quad (5.1)$$

$$\frac{\partial \vec{v}}{\partial t} + \vec{v} \cdot \nabla \vec{v} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \vec{v}, \quad (5.2)$$

$$\frac{\partial c}{\partial t} + \vec{v} \cdot \nabla c = 0, \quad (5.3)$$

$$\frac{\partial T_R}{\partial t} + \vec{v} \cdot \nabla T_R = 1, \quad (5.4)$$

where $\vec{v}$ and p are the velocity and pressure fields. Blood was modeled as a Newtonian fluid of constant density ρ and kinematic viscosity $\nu = 0.04 \text{ cm}^2/\text{s}$. The cardiac cycle was set to $t_c = 1 \text{ s}$ in all cases. The contrast agent was modeled as a passive scalar, assuming it was too diluted to significantly affect blood flow via changes in density or viscosity. Residence time was computed by solving a transport equation with unit forcing, as previously described [304].

The simulations were initialized with $c(\vec{x}, 0) = 0$, $T_R(\vec{x}, 0) = 0$ together with fully converged $\vec{v}$ and p fields obtained after running our patient-specific CFD solver for $t \geq 20t_c$ in previous studies [87, 88, 336]. Starting from these conditions, both the Navier–Stokes and transport equations Eqs. (5.2)–(5.4) were integrated for an additional $60 t_c$ to simulate contrast bolus injection.

The bolus was prescribed through time-dependent contrast concentration profiles at the PV inlets. These profiles were derived from a gamma-variate model fitted to contrast-enhanced CT data from a separate patient, distinct from the six cases retrospectively analyzed here [172]. The fitted inlet profiles for the four PVs are shown in Fig. 5.2A. At the PV inlets, the residence time was set to zero and inflow velocities were set to $|\vec{v}_i| = Q/(4A_i)$ (i.e., even split), where A_i represents each PV's cross-sectional area and Q is the total flow rate through the PVs, calculated from the patient-specific LA and LV volumes obtained by 4D CT. The motion of the LA wall obtained from the patient-specific 4D CT images was imposed via no-slip boundary conditions on $\vec{v}$. This no-slip condition included the mitral valve (MV) when it was closed. No particular spatial or temporal profile was considered for transmitral flow when the MV was open.

The patient-specific LA geometry was embedded in a cubic computational box of side length $L = 13 \text{ cm}$. At the outer boundaries of this box, free-slip boundary conditions were imposed on $\vec{v}$, homogeneous Neumann boundary conditions were imposed on c , and homogeneous Dirichlet boundary conditions were imposed on residence time. Our formulation does not require boundary conditions for

the pressure other than fixing its value at one point of the computational mesh. This value is immaterial since the pressure is undetermined by a constant in an incompressible flow.

Boundary conditions at the LA wall, mitral valve (MV), and pulmonary vein (PV) inlets were enforced with an immersed boundary method (IBM) [286], implemented as volumetric forcing terms in the corresponding evolution equations. At the LA wall, and at the MV plane when the valve was closed, the IBM forcing enforced no-slip conditions on $\vec{v}$. When the MV opened, this forcing was removed, yielding an idealized instantaneous opening/closing model. At each PV inlet, the prescribed inflow was imposed through a cylindrical buffer region composed of 12 cross-sectional layers extending upstream along the local inlet normal. A smoothed forcing term applied over this buffer region drove the flow toward the target inlet velocity, following [87].

We integrated the evolution equations (5.2)–(5.4) in time using a low-storage, three-step, semi-implicit Runge-Kutta. The time step was set constant to the Courant-Friedrichs-Lewy (CFL) number below 0.2. Continuity was enforced by splitting the temporal integration of the Navier-Stokes equations using a fractional step method with pressure correction. The spatial derivatives in the Navier-Stokes equations were approximated using second-order finite differences on a uniform Cartesian staggered grid. A third-order WENO scheme was used for the spatial discretization of convective terms in the transport equations (5.3)–(5.4) to mitigate spurious oscillations near sharp gradients [345]. The spatial resolution was $\Delta x = 0.050$ cm, matching that used in previous simulations [87, 88, 336]. To reduce the computational cost of PINN training, the CFD results were subsequently downsampled to $\Delta x = 0.090$ cm and used as the reference dataset for PINN training and validation.

5.2.3 *Physics-informed neural network*

Architecture

We built a physics-informed neural network (PINN) with PirateNet [313] as the underlying architecture to infer flow and residence time from the 4D contrast dynamics obtained using CFD. PirateNet is a residual architecture with adaptive residual connections that improves stability and mitigates the deep multilayer perceptrons (MLP) derivatives' initialization difficulties in training deep PINNs. Through its trainable residual connections, PirateNet initially learns toward a stable linear combination representation and progressively recovers richer nonlinear expressiveness during training. In our study, the PINN




Figure 5.2: Overview of the study. (A) Patient-specific pulmonary vein (PV) contrast inflow profiles from 4D contrast CT are modeled with gamma-variate fits. (B) Contrast concentration maps obtained from CFD simulations. (C) Physics-informed residual adaptive neural network (PirateNet) framework: spatiotemporal inputs are encoded with Fourier features and passed through four residual blocks to infer contrast concentration, velocity, pressure, and residence time. Each block consists of three multilayer perceptrons (MLPs) and an adaptive skip connection step with a trainable parameter α initialized at 0. Training minimizes a composite loss including data, residence time inflow conditions, Navier-Stokes, and transport equations. (D) Indicator dilution theory (IDT) framework: chamber transport is modeled through an impulse response inferred from the PV inflow profiles and simulated contrast maps. Both inflow and response curves are parameterized with gamma-variate models, and residence time is obtained from the closed-form integral of the corresponding survival function (a hypergeometric function).

approximates the mapping function $[t, \vec{x}] \rightarrow [c, \vec{v}, p, T_R]$, and the underlying PirateNet is composed of 4 residual blocks, each containing 3 MLPs with 200 neurons per layer and the nonlinear *swish*

activation function [276]. This deep and dense configuration was chosen because of the complexity of the geometries and the underlying flow dynamics. The neural network parameters (i.e., weights) were initialized using Xavier’s scheme [277], and weight normalization [278] was employed to accelerate training.

We employed Fourier feature embeddings to enforce exact temporal periodic condition and to enrich input spatial features proposed in previous literature to circumvent spectral bias in the PINN [247]. Specifically, the spatiotemporal contrast inputs were passed through a structured Fourier embedding layer comprising sine and cosine functions,

$$\left[\sin(k_t \frac{2\pi}{T} t), \quad \cos(k_t \frac{2\pi}{T} t), \quad \sin(k_{\vec{x}} \frac{2\pi}{L} \vec{x}), \quad \cos(k_{\vec{x}} \frac{2\pi}{L} \vec{x}) \right]$$

where $[t, \vec{x}]$ are spatiotemporal inputs, T is the period of the cardiac cycle, and L is the domain size. k_t and $k_{\vec{x}}$ are array of mode indices (i.e., $k_t = [1, 2, 3, \dots, M_t]$ and $k_{\vec{x}} = [1, 2, 3, \dots, M_{\vec{x}}]$). We chose the number of modes in the Fourier layer $M_t = 2$ and $M_{\vec{x}} = 4$. This configuration was selected empirically to balance computational cost and high frequency Fourier feature expressibility in the MLP [247]. Higher Fourier modes were also evaluated for $M_{\vec{x}} = 8$ and $M_{\vec{x}} = 16$ to encode finer spatial features but they did not lead to improved accuracy or performance.

To impose residence time initial condition in the PINN, we chose to impose it as a hard constraint [346] by reparameterizing the network output as,

$$T_R^* = T_{R_0} + (t - t_0) \cdot T_R \quad (5.5)$$

where T_R denotes the raw output of the MLP, t_0 is the initial time, and $T_{R,0}$ is the desired residence time map at t_0 ; here, $T_{R,0} = 0$ everywhere. The reparameterized variable T_R^* was then used as the residence time field in Eq. 5.4. This construction enforces the initial condition exactly, rather than penalizing deviations from it through an additional loss term. For simplicity, we continue to denote the residence time field as T_R throughout the manuscript.

The model was trained by minimizing the composite loss function

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \mathcal{L}_{T_R|inlet} + \mathcal{L}_{PDEs},$$

which penalizes mismatches between the predicted contrast concentration c and the training data ($\mathcal{L}_{data}$), departures from zero residence time at the PV inlets ($\mathcal{L}_{T_R|inlet}$), and residuals of the PDEs governing blood flow and transport of c and T_R ($\mathcal{L}_{PDEs}$).

At each training iteration, the optimizer sampled the CFD-derived 4D contrast dataset c_{data} to evaluate the data-fidelity term,

$$\mathcal{L}_{\text{data}} = \lambda_{\text{data}} \left(\frac{1}{M_c} \sum_{i=1}^{M_c} \|c_{\text{data}} - c\|^2 \right), \quad (5.6)$$

where λ_{data} is a weight coefficient and c is the MLP contrast output. The training batch comprised M_c randomly drawn points from the CFD mesh across N_{Cca} frames.

The loss term enforcing the inlet condition on residence time was defined as

$$\mathcal{L}_{T_R|\text{inlet}} = \lambda_{\text{inlet}} \left(\frac{1}{M_{\text{inlet}}} \sum_{i=1}^{M_{\text{inlet}}} \|T_R(t, \vec{x}_{\text{inlet}})\|^2 \right), \quad (5.7)$$

where λ_{inlet} is the corresponding weight. For this term, all available points at the four PV inlets were sampled across all frames.

The loss terms corresponding to the residuals of the governing equations were defined as

$$\mathcal{L}_{PDEs} = \sum_{j=1}^6 \frac{1}{M_{e,j}} \sum_{i=1}^{M_{e,j}} \lambda_{J_j} \|J_{j,i}\|^2, \quad (5.8)$$

where $J_{j,i} = J_j(t_i, \vec{x}_i)$ represents the residuals of the j -th PDE,

$$\begin{aligned} J_1 &= \frac{\partial c}{\partial \hat{t}} + \hat{u} \frac{\partial c}{\partial \hat{x}} + \hat{v} \frac{\partial c}{\partial \hat{y}} + \hat{w} \frac{\partial c}{\partial \hat{z}}, \\ J_2 &= \frac{\partial \hat{u}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{u}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{u}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{u}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{x}} - \frac{1}{Re} \left(\frac{\partial^2 \hat{u}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{z}^2} \right), \\ J_3 &= \frac{\partial \hat{v}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{v}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{v}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{v}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{y}} - \frac{1}{Re} \left(\frac{\partial^2 \hat{v}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{v}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{v}}{\partial \hat{z}^2} \right), \\ J_4 &= \frac{\partial \hat{w}}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{w}}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{w}}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{w}}{\partial \hat{z}} + \frac{\partial \hat{p}}{\partial \hat{z}} - \frac{1}{Re} \left(\frac{\partial^2 \hat{w}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{w}}{\partial \hat{y}^2} + \frac{\partial^2 \hat{w}}{\partial \hat{z}^2} \right), \\ J_5 &= \frac{\partial \hat{u}}{\partial \hat{x}} + \frac{\partial \hat{v}}{\partial \hat{y}} + \frac{\partial \hat{w}}{\partial \hat{z}}, \\ J_6 &= \frac{\partial \hat{T}_R}{\partial \hat{t}} + \hat{u} \frac{\partial \hat{T}_R}{\partial \hat{x}} + \hat{v} \frac{\partial \hat{T}_R}{\partial \hat{y}} + \hat{w} \frac{\partial \hat{T}_R}{\partial \hat{z}} - 1, \end{aligned}$$

obtained at i -th the sampling point, $(t_i, \vec{x}_i)$. The partial derivatives were computed using automatic differentiation [279]. The residuals were computed in non-dimensional form to improve training convergence [238, 253, 254, 274]. Specifically, we used the change of variables,

$$\hat{x} = \frac{x}{D}, \quad \hat{y} = \frac{y}{D}, \quad \hat{z} = \frac{z}{D}, \quad \hat{t} = \frac{t}{D/U}, \quad \hat{u} = \frac{u}{U}, \quad \hat{v} = \frac{v}{U}, \quad \hat{w} = \frac{w}{U}, \quad \hat{p} = \frac{p}{\rho U^2}, \quad \hat{T}_R = \frac{T_R}{D/U},$$

where U is the peak mitral valve velocity during rapid LV filling (E-wave), D is the mitral valve diameter, ρ is the blood density, and $Re = DU/\nu$ is the Reynolds number. These characteristic scales are listed in Table 5.1. In addition, all variables, including the contrast training data c , were standardized to zero mean and unit variance to improve training robustness and reduce vanishing-gradient issues [277]. After training, the inferred fields were transformed back to physical units. The same nondimensionalization and normalization procedures were applied to all loss terms. Equation residuals were evaluated at M_e collocation points, randomly resampled across space and time from the available N_{Cca} and N_{eqns} frames at each training iteration.

Patient-specific training and validation

Contrast released at the PV inlets required approximately 10 cardiac cycles to fill a substantial portion of the LA. Accordingly, frames with $t \leq 10t_c$ were excluded, and for each patient-specific case the loss terms were sampled from cycles 10 through 25 using $N_{Cca} = N_{eqns} = 300$ frames. For each iteration, we used batches of size $M_c = M_e = 10^4$ randomly sampled points for L_{data} and L_{PDEs} and full-batch for $L_{TR|inlet}$, and the composite loss function was minimized with the *Adam* optimizer [281]. Training was initialized with a learning rate of 10^{-3} for 10^5 iterations, followed by a fine-tune phase using a learning rate of 10^{-4} for additional 5×10^4 iterations as illustrated for a representative case in Fig. 5.7. The weighting coefficients associated with all loss terms were set equal to one ($\lambda_{data} = 1$, etc.) in all cases. The PINN framework was implemented in TensorFlow [282], building on our previous work [28], and training was performed on an NVIDIA A100 GPU with 80 GB of memory. Total runtime ranged from 10 to 15 hours per case. All other processing and visualization steps were carried out using in-house MATLAB and ParaView scripts.

Correlation coefficients r were computed for the velocity fields and residence-time maps, and the corresponding statistical distributions of residence time was also compared. In addition, velocity-field discrepancies between the contrast-trained PINN and the CFD reference were quantified using the normalized magnitude (ϵ_M) and orientation (ϵ_O) errors:

$$\begin{aligned}\epsilon_M &= 2 \frac{||\vec{v}_{PINN}|| - ||\vec{v}_{CFD}||}{||\vec{v}_{PINN}|| + ||\vec{v}_{CFD}||}, \\ \epsilon_O &= \cos^{-1} \left(\frac{\vec{v}_{PINN} \cdot \vec{v}_{CFD}}{||\vec{v}_{PINN}|| \cdot ||\vec{v}_{CFD}||} \right),\end{aligned}\tag{5.9}$$

where $\vec{v}_{\text{PINN}}$ and $\vec{v}_{\text{CFD}}$ denote the velocity fields predicted by the contrast-trained PINN and computed by CFD, respectively. The orientation error measures the angle between $\vec{v}_{\text{PINN}}$ and $\vec{v}_{\text{CFD}}$. Positive values of ϵ_M indicate larger velocity magnitudes in the PINN than in CFD, and *vice versa*.

5.2.4 Indicator dilution theory

Chamber response and residence time derivations

Indicator dilution theory (IDT) [347] models tracer transport between an injection site $\mathbf{x}_i$ and a downstream location $\mathbf{x}$ as a statistically stationary linear process, fully characterized by a transport function $h_{\mathbf{x}_i \rightarrow \mathbf{x}}(t)$ [348]. Accordingly, the tracer concentration at $\mathbf{x}$ is given by the convolution of the injection concentration profile $c_{\mathbf{x}_i}(t)$ with the transport function,

$$c_{\mathbf{x}}(t) = \int_0^\infty c_{\mathbf{x}_i}(t - \tau) h_{\mathbf{x}_i \rightarrow \mathbf{x}}(\tau) d\tau = c_{\mathbf{x}_i} * h_{\mathbf{x}_i \rightarrow \mathbf{x}}. \quad (5.10)$$

The transport function represents the system's impulse response to a Dirac delta bolus injection. It has been commonly represented by a gamma-variate function following a tank-in-series model [349],

$$h^{q,l}(\tau) = \frac{q^l \tau^{l-1} e^{-q\tau}}{\Gamma(l)}, \quad (5.11)$$

where l and q denote the shape and rate parameters, respectively. These parameters are often interpreted phenomenologically as reflecting the effective volume and flow rate of the compartment, and $\Gamma(\cdot)$ denotes the gamma function. Assuming a bolus injection that delivers a finite dose over a time interval that is short relative to the transport dynamics, the injection profile $c_{\mathbf{x}_i}(t)$ may be idealized as a Dirac delta function. In this case, the observed concentration $c_{\mathbf{x}}(t)$ coincides with the transport function itself. Accordingly, empirical concentration–time curves are often fitted with gamma-variate models to estimate the underlying transport function [343].

The associated survival function

$$R_{\mathbf{x}_i \rightarrow \mathbf{x}}(t) = \int_t^\infty h_{\mathbf{x}_i \rightarrow \mathbf{x}}(\tau) d\tau, \quad (5.12)$$

quantifies the fraction of tracer remaining within the compartment bounded by $\mathbf{x}_i$ and $\mathbf{x}$ at time t . The corresponding residence time is defined as

$$T_{R,\mathbf{x}_i \rightarrow \mathbf{x}}(t) = \int_0^t R_{\mathbf{x}_i \rightarrow \mathbf{x}}(s) ds, \quad (5.13)$$

and its limit when $t \rightarrow \infty$ is the expectation of the transport function,

$$\mathbb{E}[h_{\mathbf{x}_i \rightarrow \mathbf{x}}] = T_{R, \mathbf{x}_i \rightarrow \mathbf{x}}^\infty = \int_0^\infty R_{\mathbf{x}_i \rightarrow \mathbf{x}}(s) ds = \int_0^\infty \tau h_{\mathbf{x}_i \rightarrow \mathbf{x}}(\tau) d\tau = \frac{l}{q}. \quad (5.14)$$

The expectation and variance of the transport function are additive, i.e.,

$$\mathbb{E}[h_{\mathbf{x}_i \rightarrow \mathbf{x}}] = T_{R, \mathbf{x}_i \rightarrow \mathbf{x}_2}^\infty = \mathbb{E}[h_{\mathbf{x}_i \rightarrow \mathbf{x}_1}] + \mathbb{E}[h_{\mathbf{x}_1 \rightarrow \mathbf{x}_2}] = T_{R, \mathbf{x}_i \rightarrow \mathbf{x}_1}^\infty + T_{R, \mathbf{x}_1 \rightarrow \mathbf{x}_2}^\infty = \frac{l_1}{q_1} + \frac{l_2}{q_2}, \quad (5.15)$$

and

$$\text{Var}[h_{\mathbf{x}_i \rightarrow \mathbf{x}}] = \text{Var}[h_{\mathbf{x}_i \rightarrow \mathbf{x}_1}] + \text{Var}[h_{\mathbf{x}_1 \rightarrow \mathbf{x}_2}] = \frac{l_1^2}{q_1^2} + \frac{l_2^2}{q_2^2}. \quad (5.16)$$

In contrast, the convolution of two gamma-variate transport functions only yields another gamma-variate function when they have the same rate parameter, in which case the shape parameter becomes additive, i.e.,

$$h_{\mathbf{x}_i \rightarrow \mathbf{x}_1}^{q, l_1} * h_{\mathbf{x}_1 \rightarrow \mathbf{x}_2}^{q, l_2} = h_{\mathbf{x}_i \rightarrow \mathbf{x}_2}^{q, l_1 + l_2}.$$

We seek to apply IDT to model contrast dynamics within the LA voxel by voxel, using the contrast profile measured at the pulmonary veins (PVs) as the inlet reference to determine the transport function $h_{\text{PV} \rightarrow \text{LAA}}(\tau)$. This approach introduces a critical challenge: the PV concentration profile is not a Dirac delta input, and $h_{\text{PV} \rightarrow \text{LAA}}$ may not share the same rate parameter as the upstream transport function $h_{\mathbf{x}_i \rightarrow \text{PV}}$, which reflects the branched transport from the bolus injection site, typically a peripheral vein in the arm, to the PVs. Therefore, modeling $c_{\text{LAA}}(t)$ as a gamma-variate function would not directly yield $h_{\text{PV} \rightarrow \text{LAA}}(\tau)$ nor the corresponding residence time $T_{R, \mathbf{x}_i \rightarrow \mathbf{x}}(t)$. Accordingly, we assume that both the PV inlet contrast profile and the PV $\rightarrow$ LAA transport function follow gamma-variate forms,

$$c_{\text{PV}}(t) = h^{\mu, n}(t), \quad h_{\text{PV} \rightarrow \text{LAA}}(\tau) = h^{\lambda, \alpha}(\tau), \quad c_{\text{LAA}}(t) = c_{\text{PV}} * h_{\text{PV} \rightarrow \text{LAA}}. \quad (5.17)$$

To determine the parameters α and λ of the atrial transport function, we employ an integral-moment-based method that enables fast and robust estimation without explicit fitting of temporal concentration curves. Leveraging the additive properties of the mean and variance (Eqs. 5.15–5.16), we write

$$\mathbb{E}[c_{\text{LAA}}] = \mathbb{E}[c_{\text{PV}}] + \mathbb{E}[h_{\text{PV} \rightarrow \text{LAA}}], \quad \text{Var}[c_{\text{LAA}}] = \text{Var}[c_{\text{PV}}] + \text{Var}[h_{\text{PV} \rightarrow \text{LAA}}]. \quad (5.18)$$

Assuming that both $c_{\text{PV}}(t)$ and $c_{\text{LAA}}(t)$ can be measured, Eqs. (5.17)–(5.18) allow the transport parameters to be computed as

$$\lambda = \frac{\mathbb{E}[c_{\text{LAA}}] - \mathbb{E}[c_{\text{PV}}]}{\text{Var}[c_{\text{LAA}}] - \text{Var}[c_{\text{PV}}]}, \quad \alpha = \frac{(\mathbb{E}[c_{\text{LAA}}] - \mathbb{E}[c_{\text{PV}}])^2}{\text{Var}[c_{\text{LAA}}] - \text{Var}[c_{\text{PV}}]}. \quad (5.19)$$

To derive the LAA residence time, we define the atrial survival function

$$R_{PV \rightarrow LAA}(t) = \int_t^\infty h_{PV \rightarrow LAA}(\tau) d\tau = \frac{\Gamma(\alpha, \lambda t)}{\Gamma(\alpha)} = Q(\alpha, \lambda t), \quad (5.20)$$

where $Q(\alpha, \lambda t)$ denotes the upper incomplete gamma function, which is fully determined by the transport parameters α and λ .

The residence time is then obtained by integrating the survival function,

$$T_{R,PV \rightarrow LAA}(t) = \int_0^t R_{PV \rightarrow LAA}(\tau) d\tau = \frac{1}{\lambda} [\alpha + \lambda t Q(\alpha, \lambda t) - \alpha Q(\alpha + 1, \lambda t)]. \quad (5.21)$$

This closed-form expression satisfies the initial condition $T_{R,PV \rightarrow LAA}(0) = 0$, since $Q(\alpha + 1, 0) = 1$ provided that $h_{PV \rightarrow LAA}$ is a properly normalized probability density function. In the limit $t \rightarrow \infty$, the last two terms in Eq. (5.21) cancel each other out, yielding the expected result

$$T_{R,PV \rightarrow LAA}^\infty = \mathbb{E}[h^{\lambda, \alpha}] = \frac{\alpha}{\lambda}.$$

An alternate working solution to the challenge created by the PV concentrations not being dirac Deltas is to use gamma-variate fits to estimate the transport functions $h_{\mathbf{x}_i \rightarrow PV}$ and $h_{\mathbf{x}_i \rightarrow LAA}$, from which $T_{R,\mathbf{x}_i \rightarrow PV}^\infty$ and $T_{R,\mathbf{x}_i \rightarrow LAA}^\infty$ can be inferred. The residence time between the PVs and the LAA is then obtained as $T_{R,PV \rightarrow LAA}^\infty = T_{R,\mathbf{x}_i \rightarrow LAA}^\infty - T_{R,\mathbf{x}_i \rightarrow PV}^\infty$ [172]. However, this approach defines the origin of residence time at the instant of bolus injection rather than at arrival at the PVs, and is therefore valid only in the limit $t \rightarrow \infty$. This estimate can overestimate residence time at finite times when comparing to CFD data from finite-time simulations.

Assumptions and rationale for applying IDT in periodic 3D atrial transport

Our IDT formulation approximates LA blood transport by an effective linear tank-in-series process when analyzed at a fixed cardiac phase. Its voxel-wise application to 3D LA flow departs from traditional applications to 0D representations of the circulatory system, for which the assumptions of linearity and stationarity have been empirically validated [350]. Here, we reason that the 0D approximation can be extended to 3D under two assumptions: (i) the transport process is stationary in the mean across cycles, and (ii) contrast behaves as a passive tracer with negligible molecular diffusion (high Péclet number). Neglecting diffusion is consistent with the low diffusivity of iodine contrast agents in blood, and is important because it allows residence time to be inferred from temporal moments of advective

transport alone, without introducing variables that could confound the identification and interpretation of the transport parameters (q, l) .

Under these assumptions, we write $t = t^*T + \tau$ with $\tau \in [0, T)$, where t^* is a slow time variable (cycle count) and τ is within-cycle time. Fixing a phase τ_0 and sampling once per cycle at that phase defines the phase-locked field $c(\vec{x}, t^*) = c(\vec{x}, t^*T + \tau_0)$. By flow periodicity, $\vec{v}(\vec{x}, t^*T + \tau_0) = \vec{v}(\vec{x}, \tau_0) = \vec{v}^*(\vec{x})$, so $\vec{v}^*$ is independent of t^* . The phase-locked contrast dynamics across cycles is governed by

$$\frac{\partial c}{\partial t^*} + \vec{v}^*(\vec{x}) \cdot \nabla c = 0. \quad (5.22)$$

To motivate the tank-in-series model in this setting, consider transport restricted to a phase-locked pathway coordinate s with local advective speed $u(s)$, so that

$$\frac{\partial c}{\partial t^*} + u \frac{\partial c}{\partial s} = 0. \quad (5.23)$$

Discretizing the advective term in Eq. (5.23) with a first-order upwind scheme on a grid $s_j = j \Delta s$ yields the semi-discrete system

$$\frac{dc_j}{dt^*} = -u \frac{c_j - c_{j-1}}{\Delta s} = -q (c_j - c_{j-1}), \quad (5.24)$$

i.e., a cascade of l first-order relaxation (mixing) elements in series with per-stage rate $q = u/\Delta s$ (or locally $q_j = u(s_j)/\Delta s$), which motivates the gamma-variate (tank-in-series) impulse response used in §5.2.4. The associated numerical diffusivity of the upwind discretization, $D_{\text{num}} = u \Delta s/2 = u^2/(2q)$, quantifies the dissipative coarse-graining implicit in the upwind/tank analogy rather than physical dispersion. Moreover, it produces no transverse mixing between trajectories consistent with the high Péclet number approximation, since it acts only along the pathway coordinate s . Finally, for a pathway section length L_s the l -stage tank-in-series response has a transit time given by $\mathbb{E}[h] = l/q = L_s/u$, which is independent of Δs and therefore independent of D_{num} .

Voxelwise IDT Implementation

To obtain voxel-wise residence times, we apply the tank-in-series model along the flow pathlines, effectively extending it to 3D unsteady flow. Accordingly, we define

$$T_R(\mathbf{x}, t) = T_R^{\lambda(\mathbf{x}), \alpha(\mathbf{x})}(t),$$

Algorithm 1 Residence time computation via indicator dilution theory

Input:

- 1) Contrast concentration field $c_p(\mathbf{x}, t_j)$, $p = 1, \dots, N_p$, $t_j \in \{t_{p+jN_p}\}_{j=0}^{N_{\text{cyc}}-1}$, $\mathbf{x} \in \Omega_{\text{LA}}$ or $\mathbf{x} \in \Omega_{\text{LAA}}$, where $N_p = 20$ is the number of phases in one cardiac cycle, and $N_{\text{cyc}} = 60$ is the number of cardiac cycles;
- 2) Pulmonary vein inlet profiles $\{g_i(t)\}_i^4$.

Output: voxelwise chamber response $h(\mathbf{x}, t)$ shape parameter $\alpha(\mathbf{x}, t)$, rate parameter $\lambda(\mathbf{x}, t)$, and residence time field $T_R(\mathbf{x}, t)$.

Step 1 - PVs profile gamma-variate fit:

for each PV $i = 1 : 4$

- 1) Compute PV inflow contrast total area: $A_i(t) = \int_0^t g_i(t)dt$, and normalize contrast profile as $g_{\text{pv},i} = g_i(t)/A_i$.

- 2) Define PVs gamma-variate model:

$$g_{\text{PV}}(t; n, \mu) = \frac{\mu^n t^{n-1} e^{-\mu t}}{\Gamma(n)}.$$

- 3) Compute temporal moments for gamma-variate parameter initial guesses:

$$m_i = \int_0^t t g_{\text{pv},i}(t) dt, \quad v_i = \int_0^t (t - m_i)^2 g_{\text{pv},i}(t) dt.$$

- 4) Initialize gamma-variate parameters $n_{0,i} = m_i^2/v_i$ and $\mu_{0,i} = n_{0,i}/m_i$.

- 5) Define parameter lower and upper bounds as $l_b = [1, 10^{-6}]$ and $u_b = [500, 10^3]$.

- 6) Perform nonlinear least-squares fit:

$$[n_i, \mu_i] = \text{lsqcurvefit}(g_{\text{PV}}, [n_{0,i}, \mu_{0,i}], t, g_{\text{pv},i}, l_b, u_b);$$

end

Compute mean of the four fits for downstream calculations: $\hat{n} = \frac{1}{4} \sum_{i=1}^4 n_i$, $\hat{\mu} = \frac{1}{4} \sum_{i=1}^4 \mu_i$.

Compute PV inlet mean m_g and variance v_g : $m_g = \hat{n}/\hat{\mu}$, $v_g = \hat{n}/\hat{\mu}^2$.

Step 2 - Chamber response parameters and residence time calculations:

for each phase $k = 1 : P$

- 1) Extract contrast concentration data: $c_k = [c_{ij}] \in R^{N_{\text{voxel},k} \times N_{\text{cyc}}}$ and $t_k = [t_1, t_2, t_3, \dots, t_{\text{end}}]$

- 2) Normalize each voxel time series contrast concentration: $c_k(\mathbf{x}, t) = c_k(\mathbf{x}, t) / \int_0^{t_k} c_k(\mathbf{x}, t) dt$.

- 3) Compute voxelwise mean and variance based on temporal moments:

$$m_c(\mathbf{x}, t) = \int_0^{t_k} c_k(\mathbf{x}, t) dt, \quad v_c(\mathbf{x}, t) = \int_0^{t_k} (t_k - m_c(\mathbf{x}, t))^2 c_k(\mathbf{x}, t) dt.$$

- 4) Subtract inlet moments to obtain chamber response moments:

$$m_h(\mathbf{x}, t) = m_c(\mathbf{x}, t) - m_g, \quad v_h(\mathbf{x}, t) = v_c(\mathbf{x}, t) - v_g.$$

- 5) Compute per voxel moment-based chamber response parameters:

$$\alpha_k(\mathbf{x}, t) = m_h^2(\mathbf{x}, t)/v_h(\mathbf{x}, t), \quad \lambda_k(\mathbf{x}, t) = m_h(\mathbf{x}, t)/v_h(\mathbf{x}, t).$$

end

Define atrial survival function:

$$R(\mathbf{x}, t) = \int_t^\infty h(\mathbf{x}, \tau) d\tau = \frac{\Gamma(\alpha(\mathbf{x}, t), \lambda(\mathbf{x}, t)t)}{\Gamma(\alpha(\mathbf{x}, t))} = Q(\alpha(\mathbf{x}, t), \lambda(\mathbf{x}, t)t)$$

Compute residence time per voxel:

$$\begin{aligned} T_R(\mathbf{x}, t) &= \int_0^t R(\mathbf{x}, \tau) d\tau \\ &= \frac{1}{\lambda} [\alpha(\mathbf{x}, t) + \lambda(\mathbf{x}, t)t Q(\alpha(\mathbf{x}, t), \lambda(\mathbf{x}, t)t) - \alpha(\mathbf{x}, t) Q(\alpha(\mathbf{x}, t) + 1, \lambda(\mathbf{x}, t)t)]. \end{aligned}$$

where $\alpha(\mathbf{x})$ and $\lambda(\mathbf{x})$ are obtained by applying the integral-moment relations in Eq. (5.19) to each voxel's contrast profile $c(\mathbf{x}, t)$. For each patient's PV, we fit the inlet contrast time curve (Fig. 5.2A) with a gamma-variate model using nonlinear least squares to estimate the shape and rate parameters (n, μ) . Because each left atrium has multiple (typically four) PV inlets, we used the mean values of n and μ across PVs as the inlet reference for the downstream moment calculations.

The full CFD contrast dataset spanning 60 cardiac cycles, ensuring that both the wash-in and wash-out phases were included. To compute voxelwise mean and variance, we grouped the contrast concentration fields by cardiac phase (i.e., the ℓ -th frame within each cardiac cycle) and stacked the resulting time series. This procedure emulates clinical low-dose multi-cycle 4D contrast CT acquisitions, where only one 3D volume is reconstructed at a given cardiac phase per heartbeat (or every $\sim k$ heartbeats), due to temporal-resolution and radiation-dose constraints [172]. For each voxel and phase, the stacked concentration curve was normalized by its temporal integral and then integrated in time to compute the mean m_c and variance v_c . These moments, together with the inlet reference moments, were used to compute the voxelwise transport parameters α and λ via Eq. (5.19). Finally, the residence time evolution at each voxel was obtained from the closed-form expression in Eq. (5.21). All fitting and computations were performed using in-house MATLAB scripts that run on CPUs in a personal laptop; additional details are provided in Algorithm 1.

5.3 Results

This section presents flow visualizations and quantitative validation of the PINN and IDT methods, using CFD-derived contrast concentration fields for training, and residence time and flow fields for validation. For brevity, we highlight flow visualizations from two representative cases: one LAAT/TIA-negative patient (case 3) and one LAAT/TIA-positive patient (case 5), while reporting quantitative statistics across all six cases. Additional visualizations for the remaining cases are provided in the Supporting Information.




Figure 5.3: Contrast concentration (c) and residence time (T_R) maps comparison of case 3 (top) and case 5 (bottom) between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.

Overall, the PINN accurately reconstructs the dominant LA flow structures across the cardiac cycle and yields residence time maps T_R that show reasonable agreement with CFD, with a tendency to underestimate the largest T_R values, and requiring hours of patient-specific training. IDT does not recover velocity fields, but infers residence time maps from contrast curves within seconds with excellent agreement with CFD. Despite these differences, both methods discriminate between LAAT/TIA-negative and LAAT/TIA-positive cases based on the inferred LAA T_R .

5.3.1 Spatiotemporal Comparison of T_R Fields: CFD vs. PINN and IDT

Figure 5.3 shows snapshots of the contrast concentration and residence time fields, $c(\vec{x}, t)$ and $T_R(\vec{x}, t)$, from the CFD simulations for representative cases 3 (LAAT/TIA-negative) and 5 (LAAT/TIA-positive). The corresponding PINN-inferred fields are also shown, together with IDT-derived $T_R(\vec{x}, t)$. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Each snapshot is selected at the same phase of the cardiac cycle (early atrial diastole) but spanning different cycles. Analogous maps for the remaining four cases are provided in the Supporting Information (Figs. 5.11–5.13). Because the contrast data-misfit term directly constrains $c(\vec{x}, t)$ during training, the PINN reproduces the contrast concentration with minimal deviation from the CFD reference across cases and time points.

In the LAAT/TIA-negative case (top panels of Fig. 5.3), during the early stages of bolus transport ($t_c = 12T$) contrast enters the LA through the PVs but does not yet penetrate deeply into the LAA, so $c(\vec{x}, t)$ remains low and largely confined to the LA body. At this time, the PINN recovered the overall spatial organization of T_R ($t_{T_R} = 2T$), but it smoothed fine-scale features and tended to underestimate T_R within the LAA. By comparison, IDT yielded a more detailed T_R map that more closely matched the CFD reference.

As the bolus continues to enter through the PVs ($t_c = 18T$), $c(\vec{x}, t)$ increases throughout the LA body and begins to rise within the LAA, with a delayed wash-in relative to the LA. This lag is consistent with longer residence times in the appendage, as reflected by the CFD results. At this stage ($t_{T_R} = 8T$), the PINN reproduced the overall $T_R(\vec{x}, t)$ distribution but underestimated T_R within the LAA. Notably, the IDT reconstruction resolved the sharp T_R gradient at the LAA ostium, in close agreement with the

CFD reference, while tending to overestimate T_R in distal regions. At later cycles ($t_{T_R} = 15T$), the PINN still underestimated T_R in case 3, particularly in high-residence-time regions of the LAA. On the other hand, IDT remained closer to the CFD reference and continued to capture the ostial T_R gradient, although it slightly overestimated distal T_R values.

The performance of both inference methods was more favorable in the LAAT/TIA-positive case (bottom panels of Fig. 5.3). Across all reported time points during bolus wash-in and wash-out, both PINN and IDT yielded $T_R(\vec{x}, t)$ fields in good visual agreement with the CFD reference, although the PINN again appeared smoother and missed some fine-scale structure. The agreement between IDT-derived and CFD-derived residence time was remarkable.

5.3.2 Statistical Validation of Inferred LAA T_R

Given the good agreement of PINN- and IDT-derived T_R patterns in the LA body, we focus our quantitative evaluation on the LAA, the predominant site of thrombus formation and the region most sensitive to prolonged residence time. We therefore estimate the joint probability density functions relating inferred and CFD residence times within the LAA, i.e., $p(T_{R,\text{CFD}}, T_{R,\text{PirateNet}})$ and $p(T_{R,\text{CFD}}, T_{R,\text{IDT}})$, and the corresponding correlation coefficients $R_{\text{CFD},\text{PirateNet}}$ and $R_{\text{CFD},\text{IDT}}$. These statistics are calculated by pooling all LAA voxels over 15 cardiac cycles. Fig. 5.4 displays contours of $p(T_{R,\text{CFD}}, T_{R,\text{PirateNet}})$ and $p(T_{R,\text{CFD}}, T_{R,\text{IDT}})$ for the six patient-specific cases analyzed. The outermost PDF contour contains 99% of the sampled points and the innermost one contains 75%. In line with the map-level comparisons in Fig. 5.3, IDT exhibited near-uniform accuracy across the full range of T_R values, with deviations confined to an overestimation of the upper tail. The PINN, in turn, showed good agreement with CFD for low-to-intermediate residence times but progressively underestimated T_R as values increase, leading to a compression of the high- T_R range. This underestimation was more severe in the LAAT/TIA-negative cases (Fig. 5.3A–C), where the PINN rarely predicted a T_R value higher than 4 cycles. As a result, the PINN provided comparatively better residence-time reconstructions in the LAAT/TIA-positive cohort. Violin plots of the univariate distributions $p(T_R)$ from CFD, PINN, and IDT (Supplementary Fig. 5.14) confirmed these trends. Likewise, the computed correlation coefficients (Table 5.2) indicated that IDT achieved higher correlations with CFD ($r = 0.81\text{--}0.94$) than the PINN ($r = 0.51\text{--}0.82$), and the PINN correlations were obtained highest in the LAAT/TIA-positive cases than




Figure 5.4: Comparison of LAA T_R by PirateNet and IDT with CFD ground-truth data. The top panels (A–F) display joint probability density functions of CFD residence times (x -axis) and IDT-inferred residence time (y -axis, filled black contours), $p(T_{R,\text{CFD}}, T_{R,\text{IDT}})$, or PirateNet-inferred residence time (y -axis, blue line contours), $p(T_{R,\text{CFD}}, T_{R,\text{PirateNet}})$. 15 cycles of residence time are pooled to display the joint pdfs. Panels (A) through (F) respectively represent cases 1 through 6. Each contour, from dark to light, contains 75%, 90%, 95%, and 99% of the data, respectively. The bottom right corner of each panel shows the correlation coefficients $R_{\text{PirateNet}}$ and R_{IDT} . The bottom panels (G–I) are scatter plots of mean LAA T_R from PirateNet and IDT vs. with mean T_R from CFD, averaged within the LAA over three different cardiac cycles: G), $4T \leq t \leq 5T$; H), $9T \leq t \leq 10T$; I) $14T \leq t \leq 15T$. Hollow and solid symbols respectively represent LAAT/TIA-negative and cases LAAT/TIA-positive cases.

in the LAAT/TIA-negative cases.

To assess whether the PINN and IDT methods can discriminate LAAT/TIA-positive from LAAT/TIA-negative cases despite their biases, we compared mean LAA residence times against CFD using scatter plots (Fig. 5.4G–I). These data confirmed the trends observed above: IDT tended to overestimate mean T_R , whereas the PINN tended to underestimate it. Notably, case 3—used for the detailed flow and T_R visualizations—showed one of the strongest manifestations of these biases in panels G–I. Despite these offsets, the inference was favorable relative to the intrinsic inter-patient variability in LAA residence time. Importantly, both methods preserved the clinically relevant ordering across patients: PINN- and IDT-inferred mean LAA T_R values are higher in LAAT/TIA-positive cases than in LAAT/TIA-negative cases. This behavior was conserved across a different cardiac cycles as the bolus washes in and out the LA and LAA, with a trend for biases to increase as T_R increases (compare panels G–I in Fig. 5.4).

Case	$T_{R,LAA}$ correlation coefficient			Velocity errors							
	PINN vs. CFD		IDT vs. CFD	ϵ_M^{LA} (-)		ϵ_O^{LA} (°)		ϵ_M^{LAA} (-)		ϵ_O^{LAA} (°)	
	PirateNet	MLP	—	PirateNet	MLP	PirateNet	MLP	PirateNet	MLP	PirateNet	MLP
1	0.67	0.65	0.81	0.98	0.99	32.04	32.64	0.88	0.93	38.81	38.02
2	0.53	0.37	0.86	0.90	0.81	29.26	28.15	0.61	0.56	35.90	36.35
3	0.67	0.63	0.92	0.82	0.77	28.43	28.19	0.74	0.73	30.86	29.39
4	0.53	0.59	0.84	0.56	0.48	27.37	27.58	0.47	0.48	25.34	30.62
5	0.82	0.84	0.94	0.35	0.35	19.92	19.27	0.35	0.35	23.64	23.27
6	0.51	0.51	0.94	0.76	0.79	30.68	30.49	0.38	0.40	27.96	28.31

Table 5.2: Global velocity comparison metrics between PINN and CFD for the six studied cases. PINN results are reported for both PirateNet and vanilla MLP architectures. ϵ_M represents the normalized magnitude discrepancy and ϵ_O is the global orientation error (in degrees). LAA residence time T_R correlations are shown for PINN vs. CFD and IDT vs. CFD.




Figure 5.5: Comparison of volume-averaged LAA residence time T_R inferred by the PirateNet (blue) and IDT (black) against CFD ground-truth (red). Panels (A)–(F) correspond to cases 1–6. Insets show representative training contrast concentration time curves at inlet and chamber locations (lighter/darker shades indicate PVs/LAA). For the PINN, the LAA volume-averaged contrast and the cross-section-averaged contrast at each PV inlet are shown. For IDT, the LAA volume-averaged contrast and the cross-section-averaged contrast pooled across the four PV inlets are shown. In each inset, only the time interval used to train the corresponding method is shown.

5.3.3 Temporal Dynamics of Inferred LAA T_R

The temporal evolution of LAA residence time in CFD typically shows two regimes [87]. Starting from $T_R = 0$, residence time initially grows approximately linearly over a period set by the advective delay for tracer to reach the LAA from the PVs (see insets in Fig. 5.5). Once the LAA begins to exchange fluid significantly with the atrial body, T_R grows more slowly and is modulated by cardiac-cycle oscillations, consistent with dispersive mixing. Over longer times, the mean T_R approaches a quasi-steady plateau, whose level reflects the characteristic time required to renew the LAA blood pool. This behavior is

illustrated in Fig. 5.5, which shows LAA volume-averaged residence time, $T_R^{\text{LAA}}(t)$ for each case over 15 cardiac cycles. Shown in this way, the data tests whether the PINN and IDT frameworks trained on contrast can reproduce the temporal dynamics of T_R^{LAA} .

Overall, both PINN- and IDT-inferred $T_R^{\text{LAA}}(t)$ captured the initial linear-growth regime and its duration. In the subsequent dispersive regime, both methods also reproduced the pulsatile modulation of $T_R^{\text{LAA}}(t)$, including the correct cardiac frequency and the larger oscillation amplitude in LAAT/TIA-negative than in LAAT/TIA-positive cases. However, the two methods differed in the slow-time evolution of the mean $T_R^{\text{LAA}}(t)$ in the dispersive regime. The PINN tended to flatten earlier, yielding a lower long-time envelope, whereas IDT relaxed more gradually and remained higher over the same window. These differences translated into a net underestimation of LAA T_R by the PINN and a net overestimation by IDT, consistent with the biases identified in the previous section

5.3.4 Left Atrial Flow Reconstruction Using PINN

We next assessed the ability of the PINN and IDT frameworks to resolve flow within the LA and LAA. The PINN directly inferred flow velocities while the flow rate parameter λ inferred by IDT method could be interpreted as a measure of local flow velocity magnitude. However, spatial maps of voxel-wise inferred rate parameter, $\lambda(\mathbf{x}, t)$, not shown, were not informative of flow patterns. Therefore, this section focuses on PINN-inferred velocities.

Figure 5.6 compares CFD and PINN instantaneous LA velocity fields at three representative phases of the cardiac cycle (atrial diastole, LV early filling, and LV late filling) for case 3 and case 5. Vector fields from both sources as well as relative differences in magnitude (ϵ_M) and orientation (ϵ_O), defined in equations (5.9), are presented. The accompanying pulmonary vein (PV) and mitral valve (MV) flow rate waveforms, and LA and LAA volume traces, provide hemodynamic context for the selected time instants.

During atrial diastole (Fig. 5.6A–B), the PINN-inferred velocity field agreed well with the CFD reference in both patients, reproducing the PV inflow jets, their collision in the LA body, and the subsequent redirection of flow toward the MV and the LAA. The associated LAA filling pattern was also captured. The ϵ_M map indicated a mild underestimation of velocity magnitude in the LA body ($\epsilon_M < 0$), with a slight overestimation in parts of the LAA ($\epsilon_M > 0$). Additionally, the ϵ_O map showed




Figure 5.6: Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 3 (LAAT/TIA-negative, normal LA function), whereas right panels (B,D,F) show case 5 (LAAT/TIA-positive, impaired LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.0 in case 3 and 1.8 in case 5, respectively.

small angular differences between CFD and PINN velocity vectors across most of the LA, indicating good agreement in flow orientation.

Blood continues to enter through the PVs during early LV filling, and LV suction together with passive LA recoil cause the MV to open and drive a strong transmitral jet. During this phase, the LAA also empties, generating outward flow that is partially redirected toward the MV near the LAA ostium. These dominant flow patterns were reproduced by the PINN (Fig. 5.6C–D). However, the ϵ_M map revealed a bias toward underestimating velocity magnitudes, particularly in case 3. Nevertheless, flow orientations were largely preserved: elevated ϵ_O values were mainly confined to regions of near-zero velocity (e.g., near the left PVs in case 3 and near the superior LAA wall in case 5).

During LV late filling, the atrium in healthy case 3 contracts, pumping blood into the LV while also generating backflow into the PVs. And a much much weaker flow pattern is observed in case 5 due to absence of active contraction. Accordingly, the LAA continues to evacuate blood through the ostium, forming a jet that is partially deflected toward the MV (Fig. 5.6E–F). These dominant flow patterns were reproduced by the PINN, although the LAA emptying jet strength was underestimated in case 3 (consistent with the $\epsilon_M < 0$ region in Fig. 5.6E). As in atrial diastole and early LV filling, flow orientation was overall well captured, with elevated ϵ_O values mainly confined to low-speed regions where the velocity direction is effectively irrelevant.

Instantaneous 3D velocity vector fields for the remaining two LAAT/TIA-negative cases are shown in Fig. 5.8, while the other two LAAT/TIA-positive cases are shown in Fig. 5.9. Overall, these data confirm the trends reported for cases 3 and 5, suggesting that the PINN-inferred flow fields captured the orientation of the CFD ground-truth velocity, but tended to underestimate its magnitude. The globally averaged ϵ_M and ϵ_O for all the cases, shown in Table 5.2, indicate that the PINN reproduced the flow velocity field particularly well in the LAA, where $|\epsilon_M^{\text{LAA}}| < 0.50$ and $|\epsilon_O^{\text{LA}}| \lesssim 30$ deg for three out of the six cases studied, while ϵ_M increased in the LA body with mixed performances in terms of ϵ_O .

5.4 Discussion

This study investigates methods for inferring 4D (3D in space, time-resolved) residence time (T_R) fields in a flowing fluid from 4D distributions of a passively advected scalar (c) within the same flow. The work is motivated by the existence of medical imaging modalities which enable the quantification of

passive scalar dynamics, e.g., 4D CT contrast agent transport, but are incapable of directly measuring residence time distributions. We evaluate on two methods: (i) a physics-informed Neural Network (PINN) constrained by the Navier–Stokes and transport equations for c and T_R , and (ii) an indicator dilution theory (IDT) formulation that estimates T_R voxel-wise from contrast dynamics curves. Both methods are benchmarked against patient-specific CFD simulations in anatomically realistic, moving left atrial geometries.

Residence time is a relevant biomarker that reflects blood stasis, a central driver of thrombogenesis. We focus on the left atrium (LA), and particularly on the left atrial appendage (LAA), because the LAA is a major site of thrombus formation and embolic stroke. In non-valvular atrial fibrillation (AF), approximately 90% of left atrial thrombi originate in the LAA [15], and AF itself accounts for ~ 20 –30% of all ischemic strokes. The LAA also contributes to thromboembolic risk in the absence of documented AF: 15–38% of atrial thrombi in non-AF cardiomyopathy arise from the LAA.

Patient-specific CFD studies can quantify LAA residence time and related hemodynamic stasis metrics [87, 90, 329–331]. However, routine clinical translation of CFD remains limited by the need for detailed patient-specific boundary conditions, meticulous geometric meshing, and substantial computational resources [337]. Color-Doppler ultrasound remains the first-line modality for cardiac flow imaging, but transesophageal echocardiography—the technique used to visualize the LAA—is invasive, time-consuming, and provides only limited assessment of 3D flow. 4D flow MRI enables direct measurement of 3D flow patterns and stasis [319–321], but its spatiotemporal resolution in the LAA is generally insufficient. CT angiography provides high spatial resolution and has shown that impaired LAA opacification or delayed contrast washout is associated with stroke risk and LAA thrombus formation [325, 326]. Taken together, these constraints underscore the need for imaging-based approaches that can reliably infer residence time in the LAA.

Recent CT studies have linked impaired LAA opacification and delayed washout to LAA thrombus formation and stroke risk [326, 327]. Building on this premise, Severance *et al.* showed that IDT can be applied to clinically acquired multi-cycle 4D contrast CT to generate 3D maps of LAA transport parameters, including an asymptotic residence time $T_R(\mathbf{x}, t \rightarrow \infty)$, using voxelwise gamma-variate modeling [172]. In IDT, a downstream concentration–time curve is modeled as the convolution of an input bolus with a transport (impulse-response) function [349]. Historically, this framework has been used in reduced-order, compartment-based models to estimate global indices such as flow, mean transit

time, and effective volume [339, 351]. In modern contrast-enhanced CT and MRI, closely related IDT formulations are applied at the voxel level in tissue perfusion imaging, where each voxel is treated as an effective vascularized compartment and deconvolution or parametric modeling yields voxelwise perfusion metrics [352–354].

Here, we extend IDT to a moving intracardiac chamber by constructing a PV-referenced impulse response for each chamber voxel and combining it with a tank-in-series model along flow pathlines to obtain voxelwise, time-resolved residence-time maps $T_R(\mathbf{x}, t)$ from contrast dynamics. This PV-referenced formulation isolates the intrinsic chamber response from upstream transport history avoiding subtracting large, nearly equal upstream transit-time estimates (e.g., ostium- and LAA-based fits), which can be poorly conditioned when the chamber contribution is small compared with total venous-to-atrial transit. Additionally, by computing $T_R(\mathbf{x}, t)$ rather than the limit $T_R(\mathbf{x}, t \rightarrow \infty)$, our estimate is naturally tied to the observed number of cardiac cycles and reduces the risk of overestimation from extrapolating an unobserved washout tail. We expect this pathline-based method should translate to other convection-dominated vascular domains, e.g., intracranial aneurysms, carotid disease, or abdominal aortic aneurysms.

PINNs have emerged as a powerful framework to merge measurement data with physical laws in fluid mechanics and related areas of computational science [253], with applications ranging from aerodynamics to cardiovascular flows [256]. In cardiovascular biomechanics, PINNs have been used to infer latent hemodynamic variables—including pressure fields, wall shear stress, and constitutive parameters—from sparse or indirect measurements [355–358]. A canonical benchmark setting in which PINNs can infer latent velocity fields is advection of a passive scalar, because scalar dynamics provide dense spatiotemporal supervision while velocity remains unobserved. This has motivated tracer-/contrast-informed PINN formulations for flow and parameter inference in biological and medical-imaging contexts [359–361].

A substantial fraction of cardiovascular PINN studies has emphasized reduced-order hemodynamic models and/or vascular domains with relatively simple geometries and little or no wall motion. By comparison, applications targeting fully 3D intracardiac flow in atria or ventricles remain relatively scarce despite recent advances [28, 362]. Consistent with this gap, recovering intracardiac flow fields from contrast-enhanced CT remains underexplored in the PINN literature. Explicit inference of blood residence time is even less commonly considered in cardiovascular PINN studies.

Several aspects of our PINN formulation are tailored to the specific challenges of intracardiac flow reconstruction from contrast data. First, the use of spatiotemporal Fourier feature embeddings helps alleviate the spectral bias of standard MLPs and improves the representation of high-frequency flow features [247]. This is particularly important for capturing sharp inflow jets and complex LAA recirculation patterns. Second, we enforce the initial condition of residence time as a hard constraint by reparameterizing the T_R output [346], ensuring that T_R is exactly zero at $(t = t_0)$ without requiring an additional loss term. Third, we frame the learning problem in non-dimensional variables scaled by clinically meaningful quantities such as MV diameter and peak E-wave velocity, which improves numerical conditioning and may facilitate comparison across patients with different hemodynamic regimes. Lastly, we use PirateNet structure through adaptive residual connections, reported to help stabilize the training and mitigate initialization difficulties in training deep PINNs [313]. PirateNet yielded improved accuracy in T_R inference, especially when the vanilla MLP-based structure performed poorly. Together, these design choices enable the PINN to reconstruct volumetric fields in a moving LA, where the wall motion is implicitly encoded through the spatiotemporal distribution of contrast.

When compared with CFD, both PINN- and IDT-inferred T_R fields reproduce the expected trends: higher residence time in the LAA than in the LA body, especially in distal LAA regions, and systematically higher T_R in LAAT/TIA-positive cases than in negative ones. Beyond these qualitative trends, IDT shows the closest agreement with CFD across the full range of T_R values, capturing sharp spatial gradients (e.g., at the LAA ostium) with only a mild tendency to overestimate the upper tail. The PINN also recovers the main T_R patterns, but produces smoother maps and tends to underestimate high- T_R regions, most noticeably in distal LAA. Despite these biases, both methods preserve patient ranking and separate LAAT/TIA-positive from LAAT/TIA-negative cases based on mean LAA T_R .

The opposite biases of IDT and PINN reflect the different approximations each method makes. IDT models transport with no explicit diffusion, which can overestimate T_R . By contrast, the PINN balances data fit with PDE residuals, favoring smooth, low-frequency solutions in advection-dominated settings due to spectral bias [238, 253]. This phenomenon can smear sharp features and compress the high- T_R tail, leading to underestimation of the largest residence times.

In addition to residence time, the PINN reproduces the main atrial flow structures across the cardiac cycle, including PV inflow jets, their collision and redirection toward the MV, as well as LAA filling and recirculation. PINN also captures the phase-dependent flow reorganization associated with MV

opening and closing: during early LV filling it recovers the dominant transmitral jet and the concurrent LAA emptying flow through the ostium, and during LV late filling it reproduces the LAA evacuation jet and the associated flow redirection toward the MV. Across cases, the agreement is stronger in flow orientation than in magnitude, with the PINN underestimating peak speeds while preserving the overall vector field topology. This behavior is consistent with the global velocity statistics.

A key contribution of this study is showing that contrast dynamics can support two complementary inference pathways: a PINN that reconstructs full-field hemodynamics (including residence time) and an IDT-based approach that yields residence-time maps in seconds from contrast curves alone. The PINN provides velocity and pressure fields together with T_R , whereas IDT trades that detail for speed and robustness and is therefore appealing for clinical translation. This complementarity suggests a tiered workflow: IDT for rapid stasis screening, and the PINN for cases where detailed flow information is needed. Beyond this, rapid IDT inference could guide a second-stage (nested) PINN. For example, the IDT-inferred residence time $T_{R,IDT}(\mathbf{x}, t)$ could be treated as an additional “measurement” by adding a data-mismatch term, e.g., $\mathcal{L}_{IDT} = |T_{R,PINN} - T_{R,IDT}|^2$, and/or it could be included as an extra network input so the PINN learns fields conditioned on an IDT-based stasis estimate (with physics enforced using material derivatives parameterized by $T_{R,IDT}$). These hybrid IDT–PINN strategies could improve PINN accuracy in high- T_R regions and reduce underestimation of peak velocities, while preserving the PINN’s ability to recover velocity orientation.

Iodinated CT contrast can, in principle, change the local density and viscosity of blood [363], creating both a practical modeling issue and a more fundamental concern: the injection could perturb the very hemodynamics we aim to infer from the contrast dynamics. In our setting, however, pilot CFD simulations at clinically relevant iodine concentrations (Boussinesq approximation) did not show noticeable changes in left atrial velocity fields, suggesting that this effect is negligible for the protocols considered here. If such effects were non-negligible, they could be incorporated into the PINN by modifying the momentum/transport physics terms to allow property variations (e.g., variable density/viscosity), whereas the IDT approach does not require an explicit flow model and would not need to be modified beyond treating the measured curves as the effective transport response.

Training the PINN remains computationally demanding (on the order of 10–15 hours per patient on a high-end GPU), and retraining is currently required for each new patient, although this burden could be reduced with transfer learning (i.e., warm-start PINNs; [364]). Furthermore, our validation relies on

CFD simulations, which themselves depend on assumptions regarding blood rheology, wall motion, and boundary conditions [87, 88, 335, 336, 365], and are not equivalent to *in vivo* measurements. A natural next step is therefore to leverage emerging multi-cycle 4D contrast CT datasets [172] to build matched CFD and PINN models directly from clinically acquired contrast data, and to compare CFD-derived T_R against T_R inferred from the CT time curves using IDT and PINNs. Ultimately, direct validation remains challenging because no clinical modality can measure intracardiac residence time as a ground-truth field.

The IDT formulation, while computationally efficient, assumes that the LA behaves as a linear, time-invariant system with respect to contrast input. In our derivation, this approximation is justified by analyzing transport at a fixed cardiac phase and treating the phase-locked velocity field as “steady on the mean” across cycles, with negligible molecular diffusion (high Péclet number). Under these conditions, the phase-stacked voxel time curves can be interpreted through an effective impulse response (modeled here with a tank-in-series/gamma-variate form). In reality, LA contraction, PV inflow variability, and beat-to-beat fluctuations may violate these assumptions, particularly in AF. Moreover, IDT does not provide velocity or pressure fields and thus cannot capture mechanistic features such as swirling patterns, shear rate, or pressure gradients.

Despite our encouraging results, there are clear paths to improve both accuracy and practicality. On the modeling side, performance may benefit from more expressive architectures (e.g., operator-learning or multi-network decompositions), adaptive or curriculum-based loss weighting, and hard constraints for inlet residence-time conditions [346]. On the implementation side, migrating to modern platforms (e.g., TensorFlow 2.x or JAX with XLA compilation) could substantially reduce training time and ease deployment [303]. On the data side, an important next step is to train and test on clinically acquired 4D contrast CT with fewer cardiac phases, motion artifacts [366], and variable injection protocols.

5.5 Conclusion

In summary, we present and validate two complementary, physics-based approaches to infer left atrial stasis from contrast dynamics: a PINN that reconstructs full flow and residence time maps, and an IDT framework that produces residence time maps rapidly from contrast curves alone. Benchmarked against patient-specific CFD in moving atrial geometries, both methods recover key LAA stasis patterns and differentiate LAAT/TIA-positive from LAAT/TIA-negative cases based on mean LAA T_R . Together,


these results support contrast-enhanced 4D CT as a scalable route toward quantitative, patient-specific assessment of LAA thrombosis risk.

5.6 Acknowledgments

The results presented in this chapter are included in a preprint manuscript [312]: B. Maidu, A. Gonzalo, M. Guerrero-Hurtado, C. Bargellini, P. Martinez-Legazpi, J. Bermejo, F. Contijoch, O. Flores, M. García-Villalba, E. McVeigh, A. M. Kahn, and J. C. del Alamo. Left atrial flow and thrombosis risk from 4D CT contrast dynamics by physics-informed neural network and indicator dilution theory. *bioRxiv*, 2026. The dissertation author is the first-author of this manuscript.

This study was supported by the US National Institutes of Health (NIH) under award numbers R01HL160024 (to JCA, EMV, AK) and R01HL158667 (to JCA), the Spanish Research Agency under award number PID2023-146861OB-I00 (to PML) and the Instituto de Salud Carlos III, Spain, under grant number PI21/00274-PACER1 (to JB). The EU—European Regional Development Fund also supported this work.

5.7 Supporting Information



Training iterations

Figure 5.7: Loss curves for each loss component of a representative case 1 from PirateNet training.




Figure 5.8: Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 1 (LAAT/TIA-negative, normal LA function), whereas right panels (B,D,F) show case 2 (LAAT/TIA-negative, normal LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.5 in case 1 and 2.0 in case 2, respectively.




Figure 5.9: Instantaneous 3D velocity vector fields computed by CFD and inferred by contrast-trained PirateNet, together with normalized errors in PirateNet-inferred velocity magnitude (ϵ_M) and angular orientation (ϵ_O). Left panels (A,C,E) show case 4 (LAAT/TIA-positive, impaired LA function), whereas right panels (B,D,F) show case 6 (LAAT/TIA-positive, impaired LA function). Each row represents one phases of the cardiac cycle: Top panels (A, B) represent atrial diastole; center panels (C, D) represent early LV filling; bottom panels (E, F) represent and LV late filling. As temporal reference, flow-rate waveforms through the pulmonary veins and mitral valve, together with LA and LAA volume traces, are shown alongside the vector fields; magenta dashed lines indicate the corresponding time instants. Vectors are colored by proximity to anatomical regions: right PVs (blue), left PVs (cyan), LAA (green), and MV (red). To facilitate visualization, the length of PINN-inferred velocity vectors are scaled by a factor of 2.0 for both cases.




Figure 5.10: Contrast concentration (c) and residence time (T_R) maps comparison of case 1 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.




Figure 5.11: Contrast concentration (c) and residence time (T_R) maps comparison of case 2 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.




Figure 5.12: Contrast concentration (c) and residence time (T_R) maps comparison of case 4 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.




Figure 5.13: Contrast concentration (c) and residence time (T_R) maps comparison of case 6 between CFD, PirateNet, and IDT. Three 2D sections of the left atrium are shown at selected time instants. Contrast concentration snapshots are shown at $t_c/T = 12, 18$, and 25 , whereas residence time snapshots are shown at $t_{T_R}/T = 2, 8$, and 15 . Here t_c is referenced to the contrast-training window (cycles 10–25), and t_{T_R} is referenced to the start of residence time integration. Contrast concentration is dimensionless, whereas residence time is given in seconds.




Figure 5.14: Violin plots of the probability density function with mean and median of the residence time T_R from CFD, PirateNet, and IDT inside the left atrial appendage are shown for the $N = 6$ studied cases.

Chapter 6

PATIENT-SPECIFIC PARAMETER INFERENCE IN NORWOOD CIRCULATION BY LUMPED-PARAMETER-MODEL-INFORMED NEURAL NETWORK

6.1 Introduction

Hypoplastic left heart syndrome (HLHS) is a congenital heart defect in infants in which the left side of the heart is severely underdeveloped, and where the left ventricle is unable to support systemic blood flow [367]. Since HLHS is fatal if left untreated, a Norwood procedure [368] is required to support blood flow into the pulmonary and systemic circulation from the right ventricle due to the loss of function in the left ventricle. After the Norwood operation, the atria are permanently connected, the aorta is attached to the right ventricle, and a shunt is used to connect the systemic and pulmonary circulations.

Two other subsequent operations (i.e., Glenn and Fontan) [369] are needed in between four to six months and two to four years old, respectively. In order to reduce univentricle workload and balance oxygenated and deoxygenated blood flow in the pulmonary and systemic circulations, after the Glenn and Fontan procedures, the shunt is removed and the superior vena cava and inferior vena cava are connected directly to the pulmonary artery to let the blood passively flow into pulmonary circulation.

Shunt-related complications are a major cause of morbidity and mortality, as achieving appropriate pulmonary blood flow represents a “Goldilocks” dilemma. A systemic-to-pulmonary shunt creates a parallel circulation, leaving patients vulnerable to either insufficient or excessive pulmonary flow, both potentially fatal. Flow depends on ventricular function and the relative resistances of the systemic, pulmonary, and shunt pathways, with small changes in vessel or shunt diameter having large effects (per Poiseuille’s law). When shunt resistance is too high due to thrombosis, intimal thickening, or patient growth during the first six months, pulmonary flow becomes inadequate, causing hypoxia. Conversely, when resistance is too low, pulmonary overcirculation can compromise systemic perfusion, leading to acidosis and death. This can result from systemic hypertension (e.g., coarctation) or the normal neonatal drop in pulmonary vascular resistance over the first ~ 6 weeks of life [370].

In short, both pathological events and rapid physiological changes create dynamic hemodynamic

imbalances. Current static PTFE shunts cannot adapt to these changes, making sustained regulation of pulmonary blood flow extremely challenging and underscoring the need for new solutions.

Lumped-parameter models (LPMs) offer an efficient computational approach to simulate hemodynamics within the cardiovascular system. For blood circulation through cardiac chambers, LPMs relate flow rate, pressure, and compartment volumes via an electric circuit analogy. These variables are obtained by solving a system of ordinary differential equations coupled with patient-specific parameters [371]. LPMs have also been applied to single-ventricle circulation models in HLHS patients, showing promising results for parameter inference via iterative or optimization approaches [218].

Physics-informed neural network (PINN) approaches have gained traction within the deep learning domain. PINNs embed physical priors that can be expressed as differential equations to infer hidden variables from complete or partial measurement data [254]. Although PINNs have been applied mostly to infer detailed spatio-temporal physical fields [253], one of their features is that they can infer the scalar parameters of the differential equations while also inferring the physical fields represented by the equations with good accuracy [257].

In this chapter, we propose a PINN framework combined with LPMs to non-invasively infer patient-specific hemodynamic parameters in HLHS patients after the Norwood procedure. By leveraging the governing differential equations of cardiovascular circulation together with partial clinical measurements, our approach enables accurate estimation of parameters that are otherwise only obtainable through complex invasive procedures. This framework aims to address the critical clinical need for better monitoring of shunt-related complications and improved decision-making in the timing and eligibility of subsequent surgical interventions.

The rest of the chapter is organized as follows. Section §6.2 describes our in-house algorithm to generate LPM synthetic data and PINN framework. Section §6.3 shows the feasibility of our framework for patient-specific parameter inference and discusses the strengths and limitations of our approaches and future works.

6.2 Materials and Method

6.2.1 Lumped-Parameter Model (LPM)

This work centers around the mathematical formulation of an LPM, describing a post-Norwood operation HLHS patient, to generate synthetic ground-truth data for training and validation, and to define residual equations in our PINN loss functions. The LPM is inspired by previous work on Norwood circulation [188], consisting of four components: pulmonary loop, systemic loop, right heart, and left heart (Figure 6.1). A systemic-to-pulmonary (S-P) shunt is used to connect the systemic and pulmonary loops. Each compartment models blood flow analogous to that of an electric circuit which has a combination of resistors, capacitors, and inductors.




Figure 6.1: Post-Norwood circulation model used to generate training and validation data in LPM-PINN.

Physiologically, a resistor models pressure losses due to viscous friction (Poiseuille's law), a capacitor represents the compliance or elasticity of the vessel wall, and an inductor accounts for the inertia of the

blood flow moving through the vessel. A detailed list of compartment subscripts can be found in Table 6.1.

Subscript	Description	Subscript	Description
pa	pulmonary artery	svb	systemic venous bed
pab	pulmonary arterial bed	slv	systemic large veins
pvb	pulmonary venous bed	po	pulmonary valve
ao	aortic valve	rv	right ventricle
da	descending aorta	ti	tricuspid valve
sab	systemic arterial bed	ra	right atrium
la	left atrium	asd	atrial septal defect

Table 6.1: List of subscripts in the lumped-parameter model.

The governing equations of our post-Norwood circulation model are a system of ordinary differential equations (ODEs) and algebraic equations described in Eq. 6.1 and Eq. 6.2, encompassing a total of 25 physical variables and corresponding equations.

$$\begin{aligned}
\frac{dP_{pa}}{dt} &= \frac{Q_{shunt} - Q_{pa}}{C_{pa}} \\
\frac{dQ_{pa}}{dt} &= \frac{P_{pa} - P_{pab} - R_{pa}Q_{pa} - K_{pa}Q_{pa}^2}{L_{pa}} \\
\frac{dP_{pab}}{dt} &= \frac{Q_{pa} - Q_{pab}}{C_{pab}} \\
\frac{dP_{pvb}}{dt} &= \frac{Q_{pab} - Q_{pvb}}{C_{pvb}} \\
\frac{dP_{svb}}{dt} &= \frac{Q_{sab} - Q_{svb}}{C_{svb}} \\
\frac{dP_{slv}}{dt} &= \frac{Q_{svb} - Q_{slv}}{C_{slv}} \\
\frac{dQ_{da}}{dt} &= \frac{P_{da} - P_{sab} - R_{da}Q_{da}}{L_{da}} \\
\frac{dP_{da}}{dt} &= \frac{Q_{ao} - Q_{da}}{C_{da}} \\
\frac{dQ_{sab}}{dt} &= \frac{P_{sab} - P_{svb} - R_{sab}Q_{sab}}{L_{sab}} \\
\frac{dP_{sab}}{dt} &= \frac{Q_{da} - Q_{sab}}{C_{sab}} \\
\frac{dV_{rv}}{dt} &= Q_{ti} - Q_{po} \\
\frac{dV_{ra}}{dt} &= Q_{slv} + Q_{la} - Q_{ti} \\
\frac{dV_{la}}{dt} &= Q_{pvb} - Q_{la}
\end{aligned} \tag{6.1}$$

$$\begin{aligned}
P_{pa} &= P_{da} - R_{shunt} Q_{shunt} - K_{shunt} Q_{shunt}^2 & P_{rv} &= e_{rv}(t) \left(E_{1rv} + E_{2rv}(V_{rv} - V_{rv}^0) \right) (V_{rv} - V_{rv}^0) \\
Q_{pab} &= \frac{P_{pab} - P_{pvb}}{R_{pab}} & & + P_{rv}^0 \left(\exp(K_{rv}(V_{rv} - V_{rv}^0)) - 1 \right) \\
Q_{pvb} &= \frac{P_{pvb} - P_{la}}{R_{pvb}} & Q_{po} &= C_{po} A R_{po} \sqrt{P_{rv} - P_{da}} H(P_{rv} - P_{da}) \\
Q_{svb} &= \frac{P_{svb} - P_{slv}}{R_{svb}} & P_{la} &= E_{la}(t) \left(E_{la}(V_{la} - V_{la}^0) \right) \\
Q_{slv} &= \frac{P_{slv} - P_{ra}}{R_{slv}} & & + P_{la}^0 \left(\exp(K_{la}(V_{la} - V_{la}^0)) - 1 \right) \\
Q_{la} &= \frac{P_{la} - P_{ra}}{R_{asd}} & P_{ra} &= E_{ra}(t) \left(E_{ra}(V_{ra} - V_{ra}^0) \right) \\
Q_{ao} &= Q_{po} - Q_{shunt} & & + P_{ra}^0 \left(\exp(K_{ra}(V_{ra} - V_{ra}^0)) - 1 \right) \\
& & Q_{ti} &= C_{ti} A R_{ti} \sqrt{P_{ra} - P_{rv}} H(P_{ra} - P_{rv})
\end{aligned} \tag{6.2}$$

6.2.2 Synthetic Data Generation

Blood flow simulations of our LPM were performed using our in-house MATLAB solver using the *ode15i* time integration function, validated using an in-house time integration function based on an explicit fourth-order Runge Kutta time integrator. The parameters (e.g., resistance R , compliance C , inductance L) associated with each LPM compartment are defined using values from the literature [188], given in Table 6.2. The simulations were initialized with given initial conditions in Table 6.3 for each variable. For this work, we chose a uniform spacing of 5,000 points in time, and we set the heart rate to be 120 BPM (cardiac cycle $T = 0.5s$). We first let our solver run for 19 cardiac cycles to let it fully converge, and then used data from the 20th cycle for training and validation. Details of the code implementation are explained in Algorithm 2.

Parameter	Value	Parameter	Value
C_{pa}	0.27410	d	3.5
R_{pa}	0.02835	K_{pa}	0.00378
L_{pa}	0.00412	R_{pab}	0.41688
C_{pab}	0.04078	R_{pvb}	0.01097
C_{pvb}	0.88750	C_{sab}	0.44296
R_{sab}	3.51120	L_{sab}	0.00535
R_{svb}	0.32255	C_{svb}	0.31030
R_{slv}	0.08265	C_{slv}	4.07890
C_{da}	0.06118	R_{da}	0.19883
L_{da}	0.00287	R_{shunt}	$\frac{3.456}{d^4}$
K_{shunt}	$\frac{18.72}{d^4}$	R_{asd}	0.001

Table 6.2: List of model parameters. Units for C , R , L , and K are mL/mmHg, mmHg · min/mL, mmHg · min²/mL, and mmHg · min · mm⁴, respectively. The value d is the shunt diameter in mm.

Parameter	Value	Parameter	Value
Q_{ti}	50	Q_{po}	50
Q_{la}	50	Q_{shunt}	10
Q_{ao}	40	Q_{da}	40
Q_{sab}	40	Q_{svb}	40
Q_{slv}	40	Q_{pa}	10
Q_{pab}	10	Q_{pvb}	10
P_{da}	100	P_{sab}	100
P_{svb}	100	P_{slv}	100
P_{pa}	30	P_{pab}	30
P_{pvb}	30	P_{rv}	0.9
P_{ra}	0.17	P_{la}	0.17
V_{rv}	4	V_{ra}	1
V_{la}	1		

Table 6.3: List of initial conditions. Units for flow rate (Q), pressure (P), and volume (V) are ml/s, mmHg, and ml, respectively.

Algorithm 2 Synthetic Data Generation by Solving LPM DAEs Over Time

Input: Physical variables, parameters, and differential-algebraic equations (DAEs) to be solved.

Output: Solutions of flow rate (Q), pressure (P), and volume (V) over time.

1. Set initial conditions, parameters, and allometric scalars (see Table 6.2, Table 6.3).
2. Compute consistent initial conditions for the *ode15i* solver:
 - (a) Fix right-ventricle and right-atrium pressures (P_{rv} , P_{ra}). Provide start time t_0 and initial guesses for the state and its derivative, (y_0, y'_0) .
 - (b) Adjust unfixed variables to satisfy the implicit DAEs via a damped Newton method:
 - i. Enforce the consistency condition

$$F(t_0, y_n, y'_n) = \mathbf{0}.$$

- ii. Define

$$z_k = \begin{bmatrix} y_k \\ y'_k \end{bmatrix}, \quad \Delta z_k = \begin{bmatrix} \Delta y_k \\ \Delta y'_k \end{bmatrix},$$

where Δy_k and $\Delta y'_k$ are updates scaled by a step size α .

- iii. Linearize F at z_k using the Jacobian with respect to the unfixed variables:

$$F(t_0, z_k + \Delta z_k) \approx F(t_0, z_k) + J_k \Delta z_k,$$

where

$$J_k = \begin{bmatrix} \partial F / \partial y & \partial F / \partial y' \end{bmatrix}.$$

- iv. Solve the Newton step

$$J_k \Delta z_k = -F(t_0, z_k),$$

and iterate until $\|F(t_0, z_k)\|$ is sufficiently small.

3. Numerically solve the DAE system with *ode15i*:
 - (a) Advance to a new time point using a chosen step size and approximate $y'(t)$ using MATLAB's built-in backward differentiation formulas (BDFs).
 - (b) Apply a Newton method (as above) to solve for y given $y' = \sum_{k=1}^n \alpha_k y_k$.
-

6.2.3 Physics-Informed Neural Network

We use a physics-informed neural network (PINN) to infer flow variables and compartment parameters in our LPM from clinically measurable data (see Fig. 6.2). The PINN is a fully-connected neural network that does not require training on previously labeled data. Instead, training takes as inputs of the temporal information of the LPM and the outputs are the flow variables and compartment parameters. For all the analyses presented in this paper, we adopt a configuration of 4 layers with 100 neurons each, and we employ a sinusoidal activation function (i.e., $\sigma = \sin(\cdot)$). The neural network weights θ are initialized using the Xavier scheme [277] and a weight normalization method was used [278].




Figure 6.2: Physics-informed neural network training schematic. Note that θ refers to the neural network weights.

We train the PINN by penalizing a loss function ($\mathcal{L}_{total}$) consisting of two terms representing the L2 norms of the discrepancy with respect to training data ($\mathcal{L}_{data}$) and the residual of the LPM governing equations ($\mathcal{L}_{LPM}$). We assume that flow rate through the shunt (Q_{shunt}), maximum, and minimum of the pressure at the descending aorta (P_{da}) can be measured from clinical imaging as demonstrated in

previous works [372,373]. Thus these variables are used as training data (see 6.2). The PINN outputs are used to compute the residuals of the LPM equations. For example for the first equation in 6.1 we have,

$$e_1 = \frac{dP_{pa}}{dt} - \frac{Q_{shunt} - Q_{pa}}{C_{pa}}.$$

The discontinuous nature of the LPM equations representing cardiac valves presents a challenge because the PINN can only train on and infer continuous variables. Thus, we assume that the times of the valve opening and closure are known and add two additional loss terms to enforce that transvalvular flow rates (i.e., Q_{po} , Q_{ti}) are zero where the valves are closed. We also note that the valves are modeled with quadratic resistance terms in the LPM (based on Bernouilli's law). While this nonlinearity is seamlessly handled by the iterative Newton's solver in the ode15i solver, the presence of square root terms can complicate PINN convergence by creating vanishing gradients or NaNs in the loss terms. Thus, we rearrange the valve equations as,

$$\begin{aligned} Q_{po}^2 &= C_{po}^2 AR_{po}^2 (P_{rv} - P_{da}), \\ Q_{ti}^2 &= C_{ti}^2 AR_{ti}^2 (P_{ra} - P_{rv}), \end{aligned}$$

to compute the PINN loss. The final total loss function is defined as:

$$\begin{aligned} L_{total} &= L_{data} + L_{LPM} \\ &= \lambda_{Q_{shunt}} L_{Q_{shunt}} + \lambda_{P_{da}} L_{P_{da}} + \lambda_{Q_{po}} L_{Q_{po}} + \lambda_{Q_{ti}} L_{Q_{ti}} + \sum_{i=1}^{25} \lambda_{e_i} L_{e_i}, \end{aligned}$$

where the λ coefficients weigh the contribution of each term to the total loss. The loss terms for each

component without their weighing coefficients are defined as:

$$\begin{aligned}
L_{Q_{\text{shunt}}} &= \frac{1}{N} \sum_{i=1}^N \|Q_{\text{shunt}}^{\text{train}} - Q_{\text{shunt}}^{\text{pred}}\|^2 \\
L_{P_{\text{da}}} &= \frac{1}{N} \sum_{i=1}^N \|P_{\text{da,min}}^{\text{train}} - P_{\text{da,min}}^{\text{pred}}\|^2 + \frac{1}{N} \sum_{i=1}^N \|P_{\text{da,max}}^{\text{train}} - P_{\text{da,max}}^{\text{pred}}\|^2 \\
L_{Q_{\text{po}}} &= \frac{1}{N_{Q_{\text{po}}}} \sum_{i=1}^{N_{Q_{\text{po}}}} \|Q_{\text{po}}\|^2 \\
L_{Q_{\text{ti}}} &= \frac{1}{N_{Q_{\text{ti}}}} \sum_{i=1}^{N_{Q_{\text{ti}}}} \|Q_{\text{ti}}\|^2 \\
L_{e_i} &= \frac{1}{N} \sum_{j=1}^N \|e_{i,j}\|^2,
\end{aligned}$$

where N , $N_{Q_{\text{po}}}$, and $N_{Q_{\text{ti}}}$ are the numbers of sampling points used for each term during training. The temporal derivatives in the ODEs are computed using automatic differentiation [279] within the neural network during backpropagation. Additionally, we use a normalization scheme for the input temporal training data (e.g., t) so that it has zero mean and unit variance. This process is used to ensure training stability and promote robustness [277].

The loss function in our PINN has 29 terms in total, each multiplied by their weighing coefficients (e.g., λ 's). Since manual tuning of these weights is impractical, we employ an adaptive strategy derived from Neural Tangent Kernel (NTK) theory [280], which automatically adjust the λ 's to ensure that all the loss terms converge approximately at the same rate without any one term dominating the other. The weight adaptive adjustment initializes all the weights from unity and is updated in an user-defined interval. In this approach, the weights are computed as:

$$\lambda_i = \frac{\mathbf{K}}{\mathbf{K}_i},$$

where i , $\mathbf{K}_i$, and $\mathbf{K}$ represent each loss term, the NTK matrix associated with each loss term, and the total NTK matrix, respectively. And these NTK matrices are obtained by computing the Jacobian matrices of the loss term with respect to the neural network weights (e.g., θ), and the overall NTK matrix is obtained as:

$$\mathbf{K} = \sum_{i=1}^{29} \mathbf{J}_i \mathbf{J}_i^T.$$

In all the cases presented in this paper, we use a full-batch training approach. The PINN is trained until all loss terms reach plateaus or a maximum of 10^6 iterations. The total loss function is minimized using the *Adam* adaptive optimizer [281] with a fixed learning rate of 0.001, and the loss terms weights were updated every 1000 training iterations. Our framework is implemented in Python using TensorFlow derived from the original work [254]. We train all the cases using an NVIDIA A40 GPU with 48GB RAM, and all other processes and analysis were done in MATLAB.

6.2.4 Parameter Identifiability Analysis

The system parameters must be identifiable to be inferred by the model. An identifiable parameter is one that, given different values, independently changes the model outputs. To determine the subset of identifiable parameters, we performed local sensitivity analysis [374] and analyzed its Fisher information matrix.

The local sensitivity analysis was performed using a finite differences method in which each parameter was, one at a time, varied by a small step size, $\epsilon = 10^{-6}$, in the positive and negative directions. Letting f represent the model output as a function of the varied parameter θ_i , we get a $n \times m \times t$ matrix where n is the number of parameters, m is the number of output variables, and t is the number of timepoints in the cardiac cycle. The effect of each parameter θ_i can be observed with the below equation:

$$\frac{\partial f}{\partial \theta_i} \approx \frac{f(\theta_i + \epsilon) - f(\theta_i - \epsilon)}{2\epsilon}$$

For each parameter, we then take the 2-norm of the output variables over time as a metric for sensitivity to get an $n \times m$ matrix that we will call S . This gives a value for the sensitivity of each output variable with respect to the varied parameter θ_i . We can further simplify the effect of each parameter by taking the 2-norm of the variable-specific sensitivities to give a single number representing the sensitivity of each parameter.

However, since some parameters may only be jointly identifiable, it is not sufficient to observe the resulting sensitivities and draw conclusions about identifiability. As such, we calculated the Fisher information matrix, $\mathcal{F} = S^T S$. Furthermore, we compare $\|\mathcal{F}\|_F$ and $\|\mathcal{F} - \text{diag}(\mathcal{F})\|_F$ to assess how much parameters can be identified independently versus jointly.

6.3 Results and Discussions

This section presents the validation of our framework, parameter inference feasibility, and sensitivity analysis using simulations generated data. In general, the framework can accurately infer all the LPM variables and infer patient-specific parameters such as resistances. In addition, with the adaptive loss weights adjustment strategy, the PINN demonstrates stable and robust convergence without ad-hoc manual tuning of hyperparameters such as the loss-specific weight coefficients.

6.3.1 LPM Inference with Fixed Parameters

Figure. 6.3 presents the PINN-inferred flow rate, pressure, and volume for each compartment of the Norwood-shunt patient, compared against ground-truth data obtained by solving the LPM ODE system. In this experiment, all LPM parameters (resistance, compliance, and inductance) are fixed to their ground-truth values. Flow, volume, and pressure in all compartments are inferred from measurements of shunt flow rate vs. time and aortic systolic/diastolic (max/min) pressure values. Overall, the inferred variables closely match the ground truth, with moderate deviations ($\sim 20\%$) observed in P_{svb} , P_{slv} , and Q_{svb} , particularly during late diastole. These compartments are the most distal from the shunt and aorta, where the input data are prescribed to the PINN and immediately adjacent to the heart (see Fig. 6.1), which may explain why they display larger errors.




Figure 6.3: Flow rate (Q), pressure (P), and volume (V) inferred by PINN (---) compared to ground-truth (—) with fixed parameters in the LPM.

The loss terms for each constraint, each scaled by their weight coefficients λ , are shown in Fig. 6.4. As seen in the figure, losses decrease sharply within the first 3×10^5 training iterations, then adopt a much more gradual decay rate, indicating PINN convergence.




Figure 6.4: Loss term weights vs. training iterations for the fixed parameter analysis.

Figure. 6.5 displays the eigenvalues of each NTK matrix (solid lines) at the last training iteration together with the eigenvalue-normalized loss weights (dashed lines). These data suggests that the adaptive weight adjustment successfully equalized the convergence rate of different loss terms, represented by the similar order of magnitude of the dashed lines representing the normalized weights.




Figure 6.5: Eigenvalues of neural kernel associated with each loss term calculated at the last training iteration versus eigenvalue index.

6.3.2 LPM Parameter Inference – Pulmonary Artery Resistance

In a second run, we let the pulmonary resistance parameter R_{pa} be trainable while keeping all other parameters fixed. Pulmonary resistance is particularly relevant in the post-Norwood patient because it can change dynamically due to microthrombi, pulmonary vascular remodeling, or inflammatory processes in the lungs, all of which may critically impact hemodynamics [369]. Since these patients are only a few weeks old and cannot communicate symptoms, non-invasive approaches to infer such changes are essential for timely diagnosis and intervention. The inferred flow and compartment results in Fig. 6.6 show that the PINN is able to produce accurate results, with discrepancies in P_{svb} and P_{slv} comparable to those obtained when all parameters are fixed.




Figure 6.6: Flow rate (Q), pressure (P), and volume (V) inferred by PINN (--) compared to ground-truth (—) with all fixed parameters but trainable R_{pa} .

We evaluate the accuracy of the inferred parameter R_{pa} in 6.7, which displays its value as a function of PINN iteration number. Also, to ensure a stringent benchmark, we used an initial guess of $R_{pa} = 0.1$ for this parameter, almost five times larger than its ground-truth value of $0.0284 \text{ mmHg min/mL}$. The results suggest that R_{pa} is automatically adjusted to satisfy the residual equations of the LPM. During the first $\sim 3 \times 10^5$ iterations, it goes through a phase of rapid change from the initial guess and slowly

converged to the ground-truth value. The final inferred value 0.0336 mmHg min/mL, which is in reasonable agreement with the ground-truth value. When PINN execution stops at 10^6 iterations, the agreement is still slowly improving, suggesting that more accurate inference of R_{pa} could be reached.




Figure 6.7: Convergence of the inferred R_{pa} vs. training iterations compared to the ground-truth value used to generate the synthetic data. Units are mmHg min/mL.

6.3.3 Sensitivity Analysis

Finally, we present results of our sensitivity analysis to facilitate interpreting our results and to further ensure that the PINN can infer pulmonary resistance and other vacular tree parameters in the LPM. Fig. 6.8 shows the normed sensitivities over time of each parameter as it is varied by a small amount ϵ in the synthetic data generation. The capacitance of the systemic venous bed (C_{svb}) is the most sensitive parameter. This result suggests an explanation for why P_{svb} and Q_{svb} are among the least accurately inferred flow signals in the PINN – small residuals to the LPM equation for this component have a stronger effect in the flow.

Other parameters of high sensitivity, in particular R_{shunt} and R_{pa} , are particularly relevant for post-Norwood patients since they can be used as surrogate markers of shunt failure or adverse events. The fact that the LPM-PINN seems to be able to accurately infer these parameters is encouraging for clinical translation of the proposed approach. On the other hand, parameters with low sensitivity, such as the resistance in the systemic arterial bed (R_{sab}), may be harder to infer without additional input data.




Figure 6.8: Norm of the local sensitivities of each parameter over time in descending order.

6.3.4 Conclusion

In this work, we developed a Lumped-Parameter-Model-informed neural network for patient-specific hemodynamic inference in post-Norwood single-ventricle circulation. By combining a physiologically derived synthetic lumped-parameter model with a physics-informed machine learning approach, our

framework can infer flow and compartment variables from clinically measurable data. In addition, LPM-PINN is able to estimate patient-specific parameters while maintaining good agreement for most flow, pressure, and volume waveforms. Our proof-of-concept results demonstrate the feasibility of using LPM-PINN to noninvasively infer flow and patient-specific parameters that are otherwise accessible mainly through catheterization.

Notably, we demonstrate that by combining shunt flow measurable with a device currently under development [375] with systolic and diastolic pressures readily obtainable from standard cuff measurements, it is possible to identify pulmonary and shunt resistance, two critical parameters in the post-Norwood patient. These parameters are highly sensitive to shunt malfunction and adverse events such as microthrombi, pulmonary vascular remodeling, or inflammatory processes in the lungs. In newborns, early detection of such changes, together with the ability to rule out false positives, is crucial for timely intervention while also reducing unnecessary hospital visits and optimizing medical resources.

Chapter 7

CONCLUSION

This thesis focused on developing physics-informed and data-driven methods to infer cardiovascular hemodynamics, thrombosis risk, and patient-specific model parameters from sparse, noninvasive, and clinically accessible measurements. The underlying blood velocity, pressure, residence time, and vascular resistance are closely linked to cardiovascular function and disease progression, but they are often difficult to measure directly and noninvasively in vivo. The work presented in this thesis addresses this challenge by combining medical imaging, computational fluid dynamics, indicator dilution theory, lumped-parameter modeling, and physics-informed neural networks (PINNs). Across the applications studied in this thesis, the central idea is to embed physical constraints into deep learning frameworks, enabling the inference of hidden hemodynamic quantities from sparse measurements that are otherwise inaccessible from standard clinical imaging alone and resource demanding direct numerical simulations.

In Chapter 3, we developed AI-VFM, a Navier–Stokes-informed neural network framework for 2D left ventricular (LV) vector flow mapping (VFM) from color-Doppler echocardiography. Unlike classical VFM methods, which mainly enforce mass conservation to recover cross-beam velocity and require post-processing to estimate pressure, AI-VFM imposes Doppler data consistency, LV wall boundary conditions, mass conservation, and momentum balance within a physics-informed neural network. Validation against CFD ground truth showed accurate recovery of radial velocity, cross-beam velocity, and fluctuating pressure fields in an idealized LV model. AI-VFM also handled practical limitations of clinical Doppler imaging, including phase wrapping, missing data, truncated Doppler sectors, and low frame rates. Clinical applications demonstrated physiologically meaningful flow and pressure reconstructions, showing that AI-VFM can expand standard Doppler echocardiography toward non-invasive inference of intraventricular velocity and pressure.

The AI-VFM framework was then extended in Chapter 4 to time-resolved 3D volumetric reconstruction using triplane color-Doppler imaging. This extension addresses the limitation of 2D VFM, where the planar flow assumption cannot fully resolve inherently three-dimensional LV flow structures.

The 3D AI-VFM framework used sparse radial Doppler measurements from three intersecting planes while enforcing the incompressible Navier–Stokes equations throughout the full LV volume. Validation against CFD simulations showed recovery of dominant volumetric flow features, including filling and ejection patterns, large-scale vortices, kinetic energy trends, vorticity evolution, and pressure fluctuations. These results indicate that sparse triplane Doppler data, combined with physical constraints on a volumetric mesh, can reconstruct flow features in regions without direct measurements. However, the reconstructions were smoother than CFD ground truth, especially for fine-scale vortical structures, and validation was limited to synthetic triplane data. Further sensitivity analysis, training acceleration, and patient-specific and clinical validation are therefore warranted.

Chapter 5 investigated left atrial flow and thrombosis risk from contrast-enhanced 4D CT using a physics-informed neural network and an indicator dilution theory (IDT) framework. Because left atrial appendage (LAA) residence time is difficult to quantify noninvasively but is closely related to blood stasis and thrombosis risk in atrial fibrillation, this study used contrast dynamics to infer the underlying hemodynamics. The PINN framework reconstructed contrast concentration, velocity, pressure, and residence time fields while enforcing transport and Navier–Stokes constraints, whereas the IDT framework rapidly estimated residence time from contrast curves using a gamma-variate transport tank-in-series model. Benchmarking against patient-specific CFD showed that both approaches recovered key LAA stasis patterns and distinguished LAAT/TIA-positive from LAAT/TIA-negative cases based on mean LAA residence time. These results suggest that contrast-enhanced 4D CT contains clinically relevant transport information for patient-specific assessment of LAA stasis and thrombosis risk.

Finally, Chapter 6 extended physics-informed inference to reduced-order modeling of post-Norwood single-ventricle circulation. We developed a lumped-parameter-model-informed neural network that embeds the governing equations of a post-Norwood circulation model into training. Using synthetic data from a physiologically derived LPM, the framework reconstructed pressure, flow, and volume waveforms and inferred patient-specific parameters from clinically measurable quantities. The results showed accurate recovery of hemodynamic states for fixed parameters and successful inference of sensitive parameters such as pulmonary artery resistance when treated as trainable variables. Sensitivity analysis further showed that parameters with stronger influence on measurable waveforms, such as pulmonary and shunt resistances, are more identifiable. This final study broadens the thesis from local intracardiac flow reconstruction to patient-specific inference in system-level cardiovascular physiology.

Several future directions emerge from the frameworks developed in this thesis. For 3D AI-VFM, a comprehensive sensitivity analysis should be conducted using CFD benchmarks to quantify the effects of imaging noise, Doppler aliasing, and spatial and temporal resolution. Clinical validation against independent measurements, such as 4D flow MRI will also be essential to assess its robustness under realistic imaging conditions. For the left atrial CT-based PINN and IDT frameworks, future work should improve the accuracy and training efficiency of the PINN reconstruction, particularly through more efficient architectures and multi-GPU training. The gamma-variate transport assumption used in the IDT framework should also be validated in larger patient cohorts with diverse atrial anatomies, cardiac functions, and contrast injection protocols. In addition, uncertainty quantification should be incorporated into both AI-VFM and LAA-PINN-IDT frameworks to provide confidence measures for inferred velocity, pressure, and residence time. For the LPM-PINN framework, future work should refine simultaneous multi-parameter inference beyond one or two trainable parameters, especially when parameter interactions and identifiability limitations become important. Neural operator methods should also be studied, since they can learn from large synthetic or clinical datasets and will enable rapid patient-specific inference of hemodynamic parameters and flow waveforms without retraining a PINN for each case.

BIBLIOGRAPHY

- [1] R. K. Ritner. The Cardiovascular System in Ancient Egyptian Thought. *Journal of Near Eastern Studies*, 65(2):99–109, 2006.
- [2] Y. C. Fung. *Blood Flow in Arteries*, pages 77–165. Springer New York, New York, NY, 1984.
- [3] K. Patwardhan. The history of the discovery of blood circulation: unrecognized contributions of Ayurveda masters. *Advances in Physiology Education*, 36(2):77–82, 2012.
- [4] N. Lemonnier, G. B. Zhou, B. Prasher, M. Mukerji, Z. Chen, S. K. Brahmachari, D. Noble, C. Auffray, and M. Sagner. Traditional Knowledge-based Medicine: A Review of History, Principles, and Relevance in the Present Context of P4 Systems Medicine. *Progress in Preventive Medicine*, 2(7):e0011, 2017.
- [5] J. Mackay and G. A. Mensah. *The Atlas of Heart Disease and Stroke*. World Health Organization, 2004.
- [6] J. B. West. Ibn al-Nafis, the pulmonary circulation, and the Islamic Golden Age. *Journal of Applied Physiology*, 105(6):1877–1880, 2008.
- [7] Galen and M. T. May. *Galen on the Usefulness of the Parts of the Body: Peri Chreias Morion (De Usu Partium)*. Cornell University Press, Ithaca, NY, 1968.
- [8] C. D. O'Malley, F. N. L. Poynter, and K. F. Russell. *Harvey, William. Lectures on the Whole of Anatomy: An Annotated Translation of Prelectiones Anatomiae Universalis*. University of California Press, 1961.
- [9] I. Cavero, J. M. Guillon, and H. Holzgreffe. Reminiscing about Jan Evangelista Purkinje: a pioneer of modern experimental physiology. *Advances in Physiology Education*, 41:528–538, 2017.
- [10] S. Tawara. *Das Reizleitungssystem des Säugetierherzens: Eine anatomisch–pathologische Studie über das Atrioventrikulärbündel und die Purkinjeschen Fäden*. Verlag von Gustav Fischer, 1906.
- [11] W. His Jr. Die Tätigkeit des embryonalen Herzens und deren Bedeutung für die Lehre von der Herzbewegung beim Menschen. *Journal of the History of Medicine and Allied Sciences*, 4:289–318, 1949.
- [12] J. Roberts and K. D. Keele. *Leonardo da Vinci, Anatomical Drawings from the Royal Library Windsor Castle*. The Metropolitan Museum of Art, New York, 1984.

- [13] OpenStax College. Anatomy & Physiology, 2013.
- [14] K. Menon, Z. Hu, and A. L. Marsden. Cardiovascular fluid dynamics: a journey through our circulation. *Flow*, 4:E7, 2024.
- [15] J. L. Blackshear and J. A. Odell. Appendage obliteration to reduce stroke in cardiac surgical patients with atrial fibrillation. *The Annals of Thoracic Surgery*, 61(2):755–759, 1996.
- [16] S. Yaghi, C. Song, W. A. Gray, K. L. Furie, M. S. V. Elkind, and H. Kamel. Left Atrial Appendage Function and Stroke Risk. *Stroke*, 46(12):3554–3559, 2015.
- [17] M. B. Bastos, D. Burkhoff, J. Maly, J. Daemen, C. A. den Uil, K. Ameloot, M. Lenzen, F. Mahfoud, F. Zijlstra, J. J. Schreuder, and N. M. Van Mieghem. Invasive left ventricle pressure–volume analysis: overview and practical clinical implications. *European Heart Journal*, 41(12):1286–1297, 2019.
- [18] R. M. Lang, M. Cameli, L. E. Sade, F. F. Faletra, F. Fortuni, A. Rossi, and L. Soulat-Dufour. Imaging assessment of the right atrium: anatomy and function. *European Heart Journal – Cardiovascular Imaging*, 23(7):867–884, 2022.
- [19] S. M. Vaziri, M. G. Larson, E. J. Benjamin, and D. Levy. Echocardiographic predictors of nonrheumatic atrial fibrillation: The Framingham Heart Study. *Circulation*, 89(2):724–730, 1994.
- [20] T. S. Tsang, M. E. Barnes, B. J. Gersh, K. R. Bailey, and J. B. Seward. Left atrial volume as a morphophysiologic expression of left ventricular diastolic dysfunction and relation to cardiovascular risk burden. *The American Journal of Cardiology*, 90(12):1284–1289, 2002.
- [21] R. M. Lang, L. P. Badano, V. Mor-Avi, J. Afilalo, A. Armstrong, L. Ernande, F. A. Flachskampf, E. Foster, S. A. Goldstein, T. Kuznetsova, P. Lancellotti, D. Muraru, M. H. Picard, E. R. Rietzschel, L. Rudski, K. T. Spencer, W. Tsang, and J.-U. Voigt. Recommendations for Cardiac Chamber Quantification by Echocardiography in Adults: An Update from the American Society of Echocardiography and the European Association of Cardiovascular Imaging. *Journal of the American Society of Echocardiography*, 28(1):1–39.e14, 2015.
- [22] H. Suga and K. Sagawa. Instantaneous pressure-volume relationships and their ratio in the excised, supported canine left ventricle. *Circulation Research*, 35(1):117–126, 1974.
- [23] M. G. St. John Sutton and N. Sharpe. Left Ventricular Remodeling After Myocardial Infarction. *Circulation*, 101(25):2981–2988, 2000.
- [24] P. G. Newman and G. S. Rozycki. The History of Ultrasound. *Surgical Clinics of North America*, 78(2):179–195, 1998.

- [25] A. Vieillard-Baron, S. J. Millington, F. Sanfilippo, M. Chew, J. Diaz-Gomez, A. McLean, M. R. Pinsky, J. Pulido, P. Mayo, and N. Fletcher. A decade of progress in critical care echocardiography: a narrative review. *Intensive Care Medicine*, 45(6):770–788, 2019.
- [26] Z. T. Al-Sharify, T. A. Al-Sharify, N. T. Al-Sharify, and H. Yahya Naser. A critical review on medical imaging techniques (CT and PET scans) in the medical field. *IOP Conference Series: Materials Science and Engineering*, 870(1):012043, 2020.
- [27] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince. *MRI from Picture to Proton*. Cambridge University Press, 2017.
- [28] B. Maidu, P. Martinez-Legazpi, M. Guerrero-Hurtado, C. M. Nguyen, A. Gonzalo, A. M. Kahn, J. Bermejo, O. Flores, and J. C. del Alamo. Super-resolution left ventricular flow and pressure mapping by Navier–Stokes-informed neural networks. *Computers in Biology and Medicine*, 185:109476, 2025.
- [29] M. Guerrero-Hurtado, M. García-Villalba, A. Gonzalo, E. Durán, P. Martinez-Legazpi, P. Ávila, A. M. Kahn, M. Y. Chen, E. McVeigh, J. Bermejo, J. C. del Álamo, and O. Flores. Hemodynamics affects factor XI/XII anticoagulation efficacy in patient-derived left atrial models. *Computer Methods and Programs in Biomedicine*, 267:108761, 2025.
- [30] P. Dyverfeldt, M. Bissell, A. J. Barker, A. F. Bolger, C. J. Carlhäll, T. Ebbers, C. J. Francios, A. Frydrychowicz, J. Geiger, D. Giese, M. D. Hope, P. J. Kilner, S. Kozerke, S. Myerson, S. Neubauer, O. Wieben, and M. Markl. 4D flow cardiovascular magnetic resonance consensus statement. *Journal of Cardiovascular Magnetic Resonance*, 17(1):72, 2015.
- [31] D. Garcia, J. C. del Álamo, D. Tanné, R. Yotti, C. Cortina, E. Bertrand, J. C. Antoranz, E. Pérez-David, R. Rieu, F. Fernández-Avilés, and J. Bermejo. Two-Dimensional Intraventricular Flow Mapping by Digital Processing Conventional Color-Doppler Echocardiography Images. *IEEE Transactions on Medical Imaging*, 29(10):1701–1713, 2010.
- [32] K. Itatani, T. Okada, T. Uejima, T. Tanaka, M. Ono, K. Miyaji, and K. Takenaka. Intraventricular Flow Velocity Vector Visualization Based on the Continuity Equation and Measurements of Vorticity and Wall Shear Stress. *Japanese Journal of Applied Physics*, 52(7S):07HF16, 2013.
- [33] M. Avesani, B. Degrelle, G. D. Salvo, J. B. Thambo, and X. Iriart. Vector flow mapping: A review from theory to practice. *Echocardiography*, 38(8):1405–1413, 2021.
- [34] K. C. Assi, E. Gay, C. Chnafa, S. Mendez, F. Nicoud, J. F. P. J. Abascal, P. Lantelme, F. Tournoux, and D. Garcia. Intraventricular vector flow mapping—a Doppler-based regularized problem with automatic model selection. *Physics in Medicine and Biology*, 62(17):7131–7147, 2017.
- [35] F. Vixège, A. Berod, Y. Sun, S. Mendez, O. Bernard, N. Ducros, P. Courand, F. Nicoud, and D. Garcia. Physics-constrained intraventricular vector flow mapping by color Doppler. *Physics in Medicine and Biology*, 66(24):10.1088/1361–6560/ac3ffe, 2021.

- [36] G. Galanti, G. Pedrizzetti, S. Pedri, L. Stefani, and L. Toncelli. In vivo assessment of a novel technique for rest and stress intraventricular flow patterns in athletes, 2014.
- [37] B. A. Meyers, C. J. Goergen, P. Segers, and P. P. Vlachos. Colour-Doppler echocardiography flow field velocity reconstruction using a streamfunction–vorticity formulation. *Journal of the Royal Society Interface*, 17(173):20200741, 2020.
- [38] C. Argyriou, G. S. Georgiadis, N. Kontopodis, A. D. Pherwani, J. A. Van Herwaarden, C. E. V. B. Hazenberg, and G. A. Antoniou. Screening for Abdominal Aortic Aneurysm During Transthoracic Echocardiography: A Systematic Review and Meta-analysis. *European Journal of Vascular and Endovascular Surgery*, 55(4):475–491, 2018.
- [39] T. Flohr. CT systems. *Current Radiology Reports*, 1:52–63, 2013.
- [40] T. D. Heseltine, S. W. Murray, B. Ruzsics, and M. Fisher. Latest advances in cardiac CT. *European Cardiology Review*, 15:1–7, 2020.
- [41] M. V. Assen, M. Vonder, G. J. Pelgrim, P. L. Von Knebel Doeberitz, and R. Vliegenthart. Computed tomography for myocardial characterization in ischemic heart disease: a state-of-the-art review. *European Radiology Experimental*, 4:36, 2020.
- [42] D. Ties, Y. M. van der Ende, G. Pundziute, Y. T. van der Schouw, M. L. Bots, C. Xia, et al. Pre-screening to guide coronary artery calcium scoring for early identification of high-risk individuals in the general population. *European Heart Journal – Cardiovascular Imaging*, 24(1):27–35, 2022.
- [43] E. Lim, J.-H. Jang, D. Yoon, Y.-G. Min, and H.-H. Kim. Does exposure to computed tomography contrast media increase risk of end-stage renal disease? *Medical Science Monitor*, 26:e921303–1, 2020.
- [44] Y. Gorelik, H. Yaseen, S. N. Heyman, and M. Khamaisi. Negligible risk of acute renal failure among hospitalized patients after contrast-enhanced imaging with iodinated versus gadolinium-based agents. *Investigative Radiology*, 54:312–318, 2019.
- [45] T. Edvardsen, F. M. Asch, B. Davidson, V. Delgado, A. DeMaria, V. Dilsizian, et al. Non-invasive imaging in coronary syndromes: recommendations of the European Association of Cardiovascular Imaging and the American Society of Echocardiography, in collaboration with the American Society of Nuclear Cardiology, Society of Cardiovascular Computed Tomography, and Society for Cardiovascular Magnetic Resonance. *European Heart Journal – Cardiovascular Imaging*, 23:6–33, 2022.
- [46] A. Azarine, F. Scalbert, and P. Garçon. Cardiac functional imaging. *Presse Médicale*, 51:104119, 2022.

- [47] A. Seraphim, K. D. Knott, J. Augusto, A. N. Bhuva, C. Manisty, and J. C. Moon. Quantitative cardiac MRI. *Journal of Magnetic Resonance Imaging*, 51:693–711, 2020.
- [48] T. K. Foo, E. Laskaris, M. Vermilyea, M. Xu, P. Thompson, G. Conte, et al. Lightweight, compact, and high-performance 3T MR system for imaging the brain and extremities. *Magnetic Resonance in Medicine*, 80:2232–2245, 2018.
- [49] J. Přibíl, A. Přibílová, and I. Frollo. Analysis of energy relations between noise and vibration produced by a low-field MRI device. In *Noise and Vibration Control – From Theory to Practice*, page 97. IntechOpen, Bratislava, Slovak Republic, 2019.
- [50] G. Litjens, F. Ciompi, J. M. Wolterink, B. D. de Vos, T. Leiner, J. Teuwen, et al. State-of-the-art deep learning in cardiovascular image analysis. *JACC: Cardiovascular Imaging*, 12:1549–1565, 2019.
- [51] W. T. Sobol. Recent advances in MRI technology: implications for image quality and patient safety. *Saudi Journal of Ophthalmology*, 26:393–399, 2012.
- [52] P. Lurz, C. Luecke, I. Eitel, F. Föhrenbach, C. Frank, M. Grothoff, et al. Comprehensive cardiac magnetic resonance imaging in patients with suspected myocarditis: the MYO-RACER trial. *Journal of the American College of Cardiology*, 67:1800–1811, 2016.
- [53] C. M. Kramer, J. Barkhausen, C. Bucciarelli-Ducci, S. D. Flamm, R. J. Kim, and E. Nagel. Standardized cardiovascular magnetic resonance imaging (CMR) protocols: 2020 update. *Journal of Cardiovascular Magnetic Resonance*, 22:1–18, 2020.
- [54] N. B. Ameer, N. Khelifa, and T. Kraiem. Parametric images for the assessment of cardiac kinetics by magnetic resonance imaging. In *International Image Processing, Applications and Systems Conference*, pages 1–4. IEEE, 2014.
- [55] V. M. Ferreira, S. K. Piechnik, M. D. Robson, S. Neubauer, and T. D. Karamitsos. Myocardial tissue characterization by magnetic resonance imaging: novel applications of T1 and T2 mapping. *Journal of Thoracic Imaging*, 29:147, 2014.
- [56] P. K. Kim, Y. J. Hong, D. J. Im, Y. J. Suh, C. H. Park, J. Y. Kim, et al. Myocardial T1 and T2 mapping: techniques and clinical applications. *Korean Journal of Radiology*, 18:113–131, 2017.
- [57] H. M. Keijzer, M. M. Verhulst, F. J. Meijer, B. A. Tonino, F. H. Bosch, C. J. Klijn, et al. Prognosis after cardiac arrest: the additional value of DWI and FLAIR to EEG. *Neurocritical Care*, 37:302–313, 2022.
- [58] C. Mekkaoui, T. G. Reese, M. P. Jackowski, H. Bhat, and D. E. Sosnovik. Diffusion MRI in the heart. *NMR in Biomedicine*, 30:e3426, 2017.

- [59] A. R. Patel, M. Salerno, R. Y. Kwong, A. Singh, B. Heydari, and C. M. Kramer. Stress cardiac magnetic resonance myocardial perfusion imaging: JACC review topic of the week. *Journal of the American College of Cardiology*, 78:1655–1668, 2021.
- [60] R. Y. Kwong, Y. Ge, K. Steel, S. Bingham, S. Abdullah, K. Fujikura, et al. Cardiac magnetic resonance stress perfusion imaging for evaluation of patients with chest pain. *Journal of the American College of Cardiology*, 74:1741–1755, 2019.
- [61] C. P. Blanken, E. S. Farag, S. M. Boekholdt, T. Leiner, J. Kluin, A. J. Nederveen, et al. Advanced cardiac MRI techniques for evaluation of left-sided valvular heart disease. *Journal of Magnetic Resonance Imaging*, 48:318–329, 2018.
- [62] A. Kumar, S. Vyas, S. B. Chaluvashetty, and N. Kalra. Magnetic resonance angiography. *Diagnostics in Radiology: Advanced Imaging Technology*, 9:28–36, 2019.
- [63] T. B. Le, F. Sotiropoulos, D. Coffey, and D. Keefe. Vortex formation and instability in the left ventricle. *Physics of Fluids*, 24(9):091110, 2012.
- [64] T. B. Le and F. Sotiropoulos. On the three-dimensional vortical structure of early diastolic flow in a patient-specific left ventricle. *European Journal of Mechanics B: Fluids*, 35:20–24, 2012.
- [65] S. Fortini, G. Querzoli, S. Espa, and A. Cenedese. Three-dimensional structure of the flow inside the left ventricle of the human heart. *Experiments in Fluids*, 54:1–9, 2013.
- [66] J. J. Charonko, R. Kumar, K. Stewart, W. C. Little, and P. P. Vlachos. Vortices formed on the mitral valve tips aid normal left ventricular filling. *Annals of Biomedical Engineering*, 41:1049–1061, 2013.
- [67] G. Pedrizzetti and F. Domenichini. Nature optimizes the swirling flow in the human left ventricle. *Physical Review Letters*, 95(10):108101, 2005.
- [68] J. Zajac, J. Eriksson, P. Dyverfeldt, A. F. Bolger, T. Ebbers, and C.-J. Carlhäll. Turbulent kinetic energy in normal and myopathic left ventricles. *Journal of Magnetic Resonance Imaging*, 41(4):1021–1029, 2015.
- [69] C. Chnafa, S. Mendez, and F. Nicoud. Image-based large-eddy simulation in a realistic left heart. *Computers and Fluids*, 94:173–187, 2014.
- [70] L. Dedè, F. Menghini, and A. Quarteroni. Computational fluid dynamics of blood flow in an idealized left human heart. *International Journal for Numerical Methods in Biomedical Engineering*, 37(11):e3287, 2021.
- [71] C. Chnafa, S. Mendez, and F. Nicoud. Image-based simulations show important flow fluctuations in a normal left ventricle: what could be the implications? *Annals of Biomedical Engineering*, 44(11):3346–3358, 2016.

- [72] A. Azarine, P. Garçon, A. Stansal, N. Canepa, G. Angelopoulos, S. Silvera, D. Sidi, V. Marteau, and M. Zins. Four-dimensional flow MRI: principles and cardiovascular applications. *Radiographics*, 39(3):632–648, 2019.
- [73] J. Garcia, A. J. Barker, and M. Markl. The role of imaging of flow patterns by 4D flow MRI in aortic stenosis. *JACC: Cardiovascular Imaging*, 12(2):252–266, 2019.
- [74] C. M. Augustin, M. A. Gsell, E. Karabelas, E. Willemen, F. W. Prinzen, J. Lumens, E. J. Vigmond, and G. Plank. A computationally efficient physiologically comprehensive 3D–0D closed-loop model of the heart and circulation. *Computer Methods in Applied Mechanics and Engineering*, 386:114092, 2021.
- [75] Y. Cheng, H. Oertel, and T. Schenkel. Fluid-structure coupled CFD simulation of the left ventricular flow during filling phase. *Annals of Biomedical Engineering*, 33(5):567–576, 2005.
- [76] F. Domenichini, G. Pedrizzetti, and B. Baccani. Three-dimensional filling flow into a model left ventricle. *Journal of Fluid Mechanics*, 539:179–198, 2005.
- [77] V. Meschini, M. D. De Tullio, G. Querzoli, and R. Verzicco. Flow structure in healthy and pathological left ventricles with natural and prosthetic mitral valves. *Journal of Fluid Mechanics*, 834:271–307, 2018.
- [78] V. Vedula, J.-H. Seo, A. C. Lardo, and R. Mittal. Effect of trabeculae and papillary muscles on the haemodynamic of the left ventricle. *Theoretical and Computational Fluid Dynamics*, 30(1–2):3–21, 2016.
- [79] J. Lantz, L. Henriksson, A. Persson, M. Karlsson, and T. Ebbers. Patient-specific simulation of cardiac blood flow from high-resolution computed tomography. *Journal of Biomechanical Engineering*, 138(12):121004, 2016.
- [80] F. Sacco, B. Paun, O. Lehmkuhl, T. L. Iles, P. A. Iaizzo, G. Houzeaux, M. Vázquez, C. Butakoff, and J. Aguado-Sierra. Left ventricular trabeculations decrease the wall shear stress and increase the intra-ventricular pressure drop in CFD simulations. *Frontiers in Physiology*, 9:458, 2018.
- [81] V. Meschini, F. Viola, and R. Verzicco. Modeling mitral valve stenosis: a parametric study on the stenosis severity level. *Journal of Biomechanics*, 84:218–226, 2019.
- [82] F. Viola, V. Meschini, and R. Verzicco. Fluid–structure–electrophysiology interaction (FSEI) in the left heart: a multi-way coupled computational model. *European Journal of Mechanics B: Fluids*, 79:212–232, 2020.
- [83] A. Fyrenius, L. Wigström, T. Ebbers, M. Karlsson, J. Engvall, and A. F. Bolger. Three dimensional flow in the human left atrium. *Heart*, 86(4):448–455, 2001.

- [84] K.-H. Park et al. Characterization of the left atrial vortex flow by two-dimensional transesophageal contrast echocardiography using particle image velocimetry. *Ultrasound in Medicine and Biology*, 39(1):62–71, 2013.
- [85] V. Vedula, R. George, L. Younes, and R. Mittal. Haemodynamics in the left atrium and its effect on ventricular flow patterns. *Journal of Biomechanical Engineering*, 137(11):111003, 2015.
- [86] A. Zingaro, L. Dedè, F. Menghini, A. Quarteroni, et al. Haemodynamics of the heart’s left atrium based on a variational multiscale-LES numerical method. *European Journal of Mechanics B: Fluids*, 89:380–400, 2021.
- [87] M. García-Villalba, L. Rossini, A. Gonzalo, D. Vigneault, P. Martinez-Legazpi, E. Durán, O. Flores, J. Bermejo, E. McVeigh, A. M. Kahn, and J. C. del Álamo. Demonstration of patient-specific simulations to assess left atrial appendage thrombogenesis risk. *Frontiers in Physiology*, 12:596596, 2021.
- [88] A. Gonzalo, M. García-Villalba, L. Rossini, E. Durán, D. Vigneault, P. Martínez-Legazpi, O. Flores, J. Bermejo, E. McVeigh, A. M. Kahn, and J. C. del Álamo. Non-Newtonian blood rheology impacts left atrial stasis in patient-specific simulations. *International Journal for Numerical Methods in Biomedical Engineering*, 38(6):e3597, 2022.
- [89] G. M. Bosi, A. Cook, R. Rai, L. J. Menezes, S. Schievano, R. Torii, and G. Burriesci. Computational fluid dynamic analysis of the left atrial appendage to predict thrombosis risk. *Frontiers in Cardiovascular Medicine*, 5:34, 2018.
- [90] A. Masci, L. Barone, L. Dedè, M. Fedele, C. Tomasi, A. Quarteroni, and C. Corsi. The impact of left atrium appendage morphology on stroke risk assessment in atrial fibrillation: a computational fluid dynamics study. *Frontiers in Physiology*, 9:1938, 2019.
- [91] J. Dueñas-Pamplona, J. G. García, J. Sierra-Pallares, C. Ferrera, R. Agujetas, and J. R. Lopez-Minguez. A comprehensive comparison of various patient-specific CFD models of the left atrium for atrial fibrillation patients. *Computers in Biology and Medicine*, 133:104423, 2021.
- [92] T. W. Secomb. Blood flow in the microcirculation. *Annual Review of Fluid Mechanics*, 49:443–461, 2017.
- [93] N. Westerhof, F. Bosman, C. J. De Vries, and A. Noordergraaf. Analog studies of the human systemic arterial tree. *Journal of Biomechanics*, 2(2):121–143, 1969.
- [94] L. Formaggia, F. Nobile, A. Quarteroni, and A. Veneziani. Multiscale modelling of the circulatory system: a preliminary analysis. *Computing and Visualization in Science*, 2(2–3):75–83, 1999.
- [95] L. Formaggia, J. F. Gerbeau, F. Nobile, and A. Quarteroni. On the coupling of 3D and 1D Navier–Stokes equations for flow problems in compliant vessels. *Computer Methods in Applied Mechanics and Engineering*, 191(6–7):561–582, 2001.

- [96] F. Migliavacca, R. Balossino, G. Pennati, G. Dubini, T. Y. Hsia, M. R. De Leval, and E. L. Bove. Multiscale modelling in biofluidynamics: application to reconstructive paediatric cardiac surgery. *Journal of Biomechanics*, 39(6):1010–1020, 2006.
- [97] I. E. Vignon-Clementel, C. Alberto Figueroa, K. E. Jansen, and C. A. Taylor. Outflow boundary conditions for three-dimensional finite element modeling of blood flow and pressure in arteries. *Computer Methods in Applied Mechanics and Engineering*, 195(29–32):3776–3796, 2006.
- [98] H. J. Kim, I. E. Vignon-Clementel, J. S. Coogan, C. A. Figueroa, K. E. Jansen, and C. A. Taylor. Patient-specific modeling of blood flow and pressure in human coronary arteries. *Annals of Biomedical Engineering*, 38(10):3195–3209, 2010.
- [99] S. Sankaran, H. J. Kim, G. Choi, and C. A. Taylor. Uncertainty quantification in coronary blood flow simulations: impact of geometry, boundary conditions and blood viscosity. *Journal of Biomechanics*, 49(12):2540–2547, 2016.
- [100] A. R. Pries, D. Neuhaus, and P. Gaetgens. Blood viscosity in tube flow: dependence on diameter and hematocrit. *American Journal of Physiology - Heart and Circulatory Physiology*, 263(6):H1770–H1778, 1992.
- [101] J.-H. Seo and R. Mittal. Effect of diastolic flow patterns on the function of the left ventricle. *Physics of Fluids*, 25(11):110801, 2013.
- [102] J. Bermejo, Y. Benito, M. Alhama, R. Yotti, P. Martínez-Legazpi, et al. Intraventricular vortex properties in nonischemic dilated cardiomyopathy. *American Journal of Physiology-Heart and Circulatory Physiology*, 306:H718–H729, 2014.
- [103] Y. Benito, J. Bermejo, M. Alhama, R. Yotti, C. Perez del Villar, et al. Heart rate and av delay modify left ventricular filling vortex properties. *Circulation*, 126:A18099, 2012.
- [104] G. Golasch, K. Goscinska-Bis, G. Caracciolo, A. Nakabo, G. Smolka, et al. Crt improves lv filling dynamics: insights from echocardiographic particle imaging velocimetry. *JACC: Cardiovascular Imaging*, 6:704–713, 2013.
- [105] A.F. Bolger, E. Heiberg, M. Karlsson, L. Wigström, J. Engvall, et al. Transit of blood flow through the human left ventricle mapped by cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 9:741–747, 2007.
- [106] J. Eriksson, A.F. Bolger, T. Ebbers, and C.J. Carlhäll. Four-dimensional blood flow-specific markers of LV dysfunction in dilated cardiomyopathy. *European Heart Journal – Cardiovascular Imaging*, 14:417–424, 2013.
- [107] H. Abe, G. Caracciolo, A. Kheradvar, G. Pedrizzetti, B.K. Khandheria, et al. Contrast echocardiography for assessing left ventricular vortex strength in heart failure: a prospective cohort study. *European Heart Journal – Cardiovascular Imaging*, 14:1049–1060, 2013.

- [108] M. Carlsson, E. Heiberg, J. Töger, and H. Arheden. Quantification of left and right ventricular kinetic energy using four-dimensional intracardiac magnetic resonance imaging flow measurements. *American Journal of Physiology – Heart and Circulatory Physiology*, 302:H893–H900, 2012.
- [109] H. Watanabe, S. Sugiura, and T. Hisada. The looped heart does not save energy by maintaining the momentum of blood flowing in the ventricle. *American Journal of Physiology – Heart and Circulatory Physiology*, 294:H2191–H2196, 2008.
- [110] P.C. Adams, M. Cohen, J.H. Chesebro, and V. Fuster. Thrombosis and embolism from cardiac chambers and infected valves. *Journal of the American College of Cardiology*, 8:B76–B87, 1986.
- [111] M. Lengyel, V. Fuster, M. Keltai, R. Roudaut, H.D. Schulte, et al. Guidelines for management of left-sided prosthetic valve thrombosis: a role for thrombolytic therapy. Consensus Conference on Prosthetic Valve Thrombosis. *Journal of the American College of Cardiology*, 30:1521–1526, 1997.
- [112] K. May-Newman, Y.K. Wong, R. Adamson, P. Hoagland, V. Vu, and W. Dembitsky. Thromboembolism is linked to intraventricular flow stasis in a patient supported with a left ventricle assist device. *ASAIO Journal*, 59:452–455, 2013.
- [113] K. Wong, G. Samaroo, I. Ling, W. Dembitsky, R. Adamson, et al. Intraventricular flow patterns and stasis in the LVAD-assisted heart. *Journal of Biomechanics*, 47:1485–1494, 2014.
- [114] N.M. Al-Saady, O.A. Obel, and A.J. Camm. Left atrial appendage: structure, function, and role in thromboembolism. *Heart*, 82(5):547–554, 1999.
- [115] Y. Agmon, B.K. Khandheria, F. Gentile, and J.B. Seward. Echocardiographic assessment of the left atrial appendage. *Journal of the American College of Cardiology*, 34(7):1867–1877, 1999.
- [116] O. Kamp, P. Verhorst, R. Welling, and C. Visser. Importance of left atrial appendage flow as a predictor of thromboembolic events in patients with atrial fibrillation. *European Heart Journal*, 20(13):979–985, 1999.
- [117] E. Turkel. Preconditioning techniques in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 31(1):385–416, 1999.
- [118] W.L. Oberkampf and T.G. Trucano. Verification and validation in computational fluid dynamics. *Progress in Aerospace Sciences*, 38(3):209–272, 2002.
- [119] J.F. Wendt. *Computational fluid dynamics: an introduction*. Springer Science & Business Media, 2008.

- [120] J. Blazek. *Computational fluid dynamics: principles and applications*. Butterworth-Heinemann, 2015.
- [121] J. Bermejo, P. Martínez-Legazpi, and J. C. del Álamo. The clinical assessment of intraventricular flows. *Annual Review of Fluid Mechanics*, 47(1):315–342, 2015.
- [122] E. Balaras, K.S. Cha, B.P. Griffith, and J.S. Gammie. Treatment of aortic stenosis with aortic valve bypass (apicoaortic conduit) surgery: an assessment using computational modeling. *Journal of Thoracic and Cardiovascular Surgery*, 137:680–687, 2009.
- [123] B.E. Griffith, X. Luo, D.M. McQueen, and C.S. Peskin. Simulating the fluid dynamics of natural and prosthetic heart valves using the immersed boundary method. *International Journal of Applied Mechanics*, 1:137–177, 2009.
- [124] J.D. Humphrey. Towards a theory of vascular growth and remodeling. In G.A. Holzapfel and R.W. Ogden, editors, *Mechanics of Biological Tissue*, pages 3–15. Springer, Berlin, 2006.
- [125] E.K. Rodriguez, A. Hoger, and A.D. McCulloch. Stress-dependent finite growth in soft elastic tissues. *Journal of Biomechanics*, 27:455–467, 1994.
- [126] A. Menzel and E. Kuhl. Frontiers in growth and remodeling. *Mechanics Research Communications*, 42:1–14, 2012.
- [127] C.S. Peskin and D.M. McQueen. A three-dimensional computational method for blood flow in the heart. *Journal of Computational Physics*, 81:372–405, 1989.
- [128] B.E. Griffith, R.D. Hornung, D.M. McQueen, and C.S. Peskin. An adaptive, formally second order accurate version of the immersed boundary method. *Journal of Computational Physics*, 223:10–49, 2007.
- [129] R. Faludi, M. Szulik, J. D’Hooge, P. Herijgers, F. Rademakers, et al. Left ventricular flow patterns in healthy subjects and patients with prosthetic mitral valves: an in vivo study using echocardiographic particle image velocimetry. *Journal of Thoracic and Cardiovascular Surgery*, 139:1501–1510, 2010.
- [130] G. Querzoli, S. Fortini, and A. Cenedese. Effect of the prosthetic mitral valve on vortex dynamics and turbulence of the left ventricular flow. *Physics of Fluids*, 22:041901, 2010.
- [131] C. Millington-Sanders, A. Meir, L. Lawrence, and C. Stolinski. Structure of chordae tendineae in the left ventricle of the human heart. *Journal of Anatomy*, 192:573–581, 1998.
- [132] G. Burch. The trabeculae carneae—a thought. *American Heart Journal*, 89:261, 1975.

- [133] A.R. Patel and V. Mor-Avi. Are trabeculae and papillary muscles an integral part of cardiac anatomy: or annoying features to exclude while tracing endocardial boundaries? *JACC: Cardiovascular Imaging*, 5:1124–1126, 2012.
- [134] L. T. Zhang and M. Gay. Characterizing left atrial appendage functions in sinus rhythm and atrial fibrillation using computational models. *Journal of Biomechanics*, 41:2515–2523, 2008.
- [135] R. Koizumi, K. Funamoto, T. Hayase, Y. Kanke, M. Shibata, Y. Shiraishi, et al. Numerical analysis of hemodynamic changes in the left atrium due to atrial fibrillation. *Journal of Biomechanics*, 48:472–478, 2015.
- [136] A. Masci, M. Alessandrini, D. Forti, F. Menghini, L. Dedè, C. Tomasi, et al. A proof of concept for computational fluid dynamic analysis of the left atrium in atrial fibrillation on a patient-specific basis. *Journal of Biomechanical Engineering*, 142:011002, 2020.
- [137] T. Otani, A. Al-Issa, A. Pourmorteza, E.R. McVeigh, S. Wada, and H. Ashikaga. A computational framework for personalized blood flow analysis in the human left atrium. *Annals of Biomedical Engineering*, 44(11):3284–3294, 2016.
- [138] L. Di Biase, P. Santangeli, M. Anselmino, P. Mohanty, I. Salvetti, S. Gili, et al. Does the left atrial appendage morphology correlate with the risk of stroke in patients with atrial fibrillation? Results from a multicenter study. *Journal of the American College of Cardiology*, 60:531–538, 2012.
- [139] F. Bisbal, E. Guieu, N. Calvo, D. Marin, A. Berruezo, E. Arbelo, et al. Left atrial sphericity: a new method to assess atrial remodeling. Impact on the outcome of atrial fibrillation ablation. *Journal of Cardiovascular Electrophysiology*, 24:752–759, 2013.
- [140] M. Yamamoto, Y. Seo, N. Kawamatsu, K. Sato, A. Sugano, T. Machino-Ohtsuka, et al. Complex left atrial appendage morphology and left atrial appendage thrombus formation in patients with atrial fibrillation. *Circulation: Cardiovascular Imaging*, 7:337–343, 2014.
- [141] S. Yaghi, A.D. Chang, R. Akiki, S. Collins, T. Novack, M. Hemendinger, et al. The left atrial appendage morphology is associated with embolic stroke subtypes using a simple classification system: a proof of concept study. *Journal of Cardiovascular Computed Tomography*, 14:27–33, 2020.
- [142] M. Moriche. *A numerical study on the aerodynamic forces and the wake stability of flapping flight at low Reynolds number*. PhD thesis, Universidad Carlos III de Madrid, Getafe, Spain, 2017.
- [143] M. Moriche, O. Flores, and M. Garcia-Villalba. On the aerodynamic forces on heaving and pitching airfoils at low Reynolds number. *Journal of Fluid Mechanics*, 828:395–423, 2017.
- [144] G. Pedrizzetti and G. Tonti. Method of transforming a Doppler velocity dataset into a velocity vector field. U.S. Patent No. 8,858,437, 2014.

- [145] J. Jang, C.Y. Ahn, K. Jeon, J. Heo, D. Lee, C. Joo, J. Choi, and J.K. Seo. A reconstruction method of blood flow velocity in left ventricle using color flow ultrasound. *Computational and Mathematical Methods in Medicine*, 2015:108274, 2014.
- [146] T. Tanaka, T. Okada, T. Nishiyama, and Y. Seki. Relative pressure imaging in left ventricle using ultrasonic vector flow mapping. *Japanese Journal of Applied Physics*, 56:07JF26, 2017.
- [147] F. Vixège, A. Berod, P.-Y. Courand, S. Mendez, F. Nicoud, P. Blanc-Benon, D. Vray, and D. Garcia. Full-volume three-component intraventricular vector flow mapping by triplane color Doppler. *Physics in Medicine and Biology*, 67:095004, 2022.
- [148] Z. Zhuang, G. Liu, W. Ding, A.N.J. Raj, S. Qiu, J. Guo, and Y. Yuan. Cardiac VFM visualization and analysis based on YOLO deep learning model and modified 2D continuity equation. *Computerized Medical Imaging and Graphics*, 82:101732, 2020.
- [149] H.J. Ling, S. Bru, J. Puig, F. Vixège, S. Mendez, F. Nicoud, P.-Y. Courand, O. Bernard, and D. Garcia. Physics-guided neural networks for intraventricular vector flow mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 71:1377, 2024.
- [150] H.S. Wong, W.X. Chan, W. Mao, and C.H. Yap. 3D velocity and pressure field reconstruction in the cardiac left ventricle via physics-informed neural network from echocardiography guided by 3D color Doppler. *Computer Methods and Programs in Biomedicine*, 263:108671, 2025.
- [151] J. H. Henriksen, S. Fuglsang, and H. W. B. Larsson. *Basal Indicator Kinetics: Kinetic Principles in Whole Body and Single Organs*. Forlaget Underskoven, Copenhagen, 2 edition, 2011.
- [152] G. N. Stewart. Researches on the circulation time and on the influences which affect it. *The Journal of Physiology*, 22:159–183, 1897.
- [153] V. Henriques. Über die Verteilung des Blutes vom linken Herzen zwischen dem Herzen und dem übrigen Organismus. *Biochemische Zeitschrift*, 56:230–248, 1913.
- [154] N. A. Lassen and W. Perl. *Tracer Kinetic Methods in Medical Physiology*. Raven Press, New York, 1979.
- [155] J. H. Henriksen. Valdemar Henriques (1864–1936). *Ugeskrift for Læger*, 169:2875, 2007.
- [156] W. F. Hamilton, J. W. Moore, J. M. Kinsman, and R. G. Spurling. Simultaneous determination of the pulmonary and systemic circulation times in man and of a figure related to the cardiac output. *American Journal of Physiology*, 84:338–344, 1928.
- [157] F. Neukirch. *Experimental Circulating Investigations*. Nyt Nordisk Forlag, Copenhagen, 1938.

- [158] F. G. Jr. Ebaugh, E. D. Freis, L. M. Hershenson, R. D. Kovach, L. S. Lilienfield, P. A. Marks, and G. P. Rodnan. The hematocrit of the lesser circulation in man. *Journal of Clinical Investigation*, 35:1385–1392, 1956.
- [159] W. Falholt and J. Fabricius. Application of continuously recording dye colorimeter to determination of cardiac output by the dye injection method. *Danish Medical Bulletin*, 3:55–60, 1956.
- [160] C. Crone and N. A. Lassen. *Capillary Permeability*. Munksgaard, Copenhagen, 1970.
- [161] C. A. Goresky. A linear method for determining liver sinusoidal and extravascular volumes. *American Journal of Physiology*, 204:626–640, 1963.
- [162] S. S. Kety and C. F. Schmidt. The nitrous oxide method for the quantitative determination of cerebral blood flow in man: theory, procedure and normal values. *Journal of Clinical Investigation*, 27:476–483, 1948.
- [163] N. A. Lassen and O. Munck. The cerebral blood flow in man determined by the use of radioactive krypton. *Acta Physiologica Scandinavica*, 33:30–49, 1955.
- [164] L. Sokoloff, M. Reivich, C. Kennedy, M. H. Des Rosiers, C. S. Patlak, K. D. Pettigrew, O. Sakurada, and M. Shinohara. The [14C] deoxyglucose method for the measurement of local cerebral glucose utilization: theory, procedure, and normal values in the conscious and anesthetized albino rat. *Journal of Neurochemistry*, 28:897–916, 1977.
- [165] O. L. Munk, L. Bass, H. Feng, and S. Keiding. Determination of regional flow by use of intravascular PET tracers: microvascular theory and experimental validation for pig livers. *Journal of Nuclear Medicine*, 44:1862–1870, 2003.
- [166] O. L. Munk, S. Keiding, and L. Bass. Capillaries within compartments: microvascular interpretation of dynamic positron emission tomography data. *Journal of Theoretical Biology*, 225:127–141, 2003.
- [167] L. M. Brown, K. D. Liu, and M. A. Matthay. Measurement of extravascular lung water using the single indicator method in patients: research and potential clinical value. *American Journal of Physiology – Lung Cellular and Molecular Physiology*, 297:L547–L558, 2009.
- [168] L. Østergaard, R. M. Weisskoff, D. A. Chesler, C. Gyldensted, and B. R. Rosen. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part I: mathematical approach and statistical analysis. *Magnetic Resonance in Medicine*, 36:715–725, 1996.
- [169] L. Østergaard, A. G. Sørensen, K. K. Kwong, R. M. Weisskoff, C. Gyldensted, and B. R. Rosen. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part II: experimental comparison and preliminary results. *Magnetic Resonance in Medicine*, 36:726–736, 1996.

- [170] H. B. Larsson, F. Courivaud, E. Rostrup, and A. E. Hansen. Measurement of brain perfusion, blood volume, and blood–brain barrier permeability using dynamic contrast-enhanced T1-weighted MRI at 3 tesla. *Magnetic Resonance in Medicine*, 62:1270–1281, 2009.
- [171] D. C. Alsop and J. A. Detre. Multisection cerebral blood flow MR imaging with continuous arterial spin labelling. *Radiology*, 208:410–416, 1998.
- [172] L. M. Severance, A. M. Kahn, J. C. del Alamo, and McVeigh E. R. Estimation of Blood Flow Parameters in the Left Atrial Appendage from 4DCT Dynamic Contrast Enhancement. In *The 8th International Conference on Image Formation in X-Ray Computed Tomography*, pages 388–391, 2024.
- [173] J. T. Ottesen, M. S. Olufsen, and J. K. Larsen. *Applied Mathematical Models in Human Physiology*. Society for Industrial and Applied Mathematics, Philadelphia, 2004.
- [174] Y. Vassilevski, M. Olshanskii, S. Simakov, A. Kolobov, and A. Danilov. *Personalized Computational Hemodynamics*. Academic Press, Cambridge, 2020.
- [175] J. K.-J. Li. *The Arterial Circulation: Physical Principles and Clinical Applications*. Humana Press, New York, 2000.
- [176] N. Westerhof, J.-W. Lankhaar, and B. E. Westerhof. The arterial Windkessel. *Medical & Biological Engineering & Computing*, 47:131–141, 2009.
- [177] S. Zhou et al. A review on low-dimensional physics-based models of systemic arteries: application to estimation of central aortic pressure. *Biomedical Engineering Online*, 18:41, 2019.
- [178] N. Stergiopoulos, B. E. Westerhof, and N. Westerhof. Total arterial inertance as the fourth element of the Windkessel model. *American Journal of Physiology*, 276:H81–H88, 1999.
- [179] B. Deswysen, A. A. Charlier, and M. Gevers. Quantitative evaluation of the systemic arterial bed by parameter estimation of a simple model. *Medical & Biological Engineering & Computing*, 18:153–166, 1980.
- [180] R. M. Goldwyn and T. B. Watt. Arterial pressure pulse contour analysis via a mathematical model for the clinical quantification of human vascular properties. *IEEE Transactions on Biomedical Engineering*, BME-14:11–17, 1967.
- [181] W. C. Rose and A. A. Shoukas. Two-port analysis of systemic venous and arterial impedances. *American Journal of Physiology*, 265:H1577–H1587, 1993.
- [182] H. F. Frasch, J. Y. Kresh, and A. Noordergraaf. Two-port analysis of microcirculation: an extension of the Windkessel. *American Journal of Physiology*, 270:H376–H385, 1996.

- [183] F. Nestler et al. A hybrid mock circulation loop for a total artificial heart: hybrid mock circulation loop. *Artificial Organs*, 38:775–782, 2014.
- [184] E. Kim and M. Capoccia. Synergistic model of cardiac function with a heart assist device. *Bioengineering*, 7:1, 2019.
- [185] A. D. Molfetta et al. Application of a lumped parameter model to study the feasibility of simultaneous implantation of a continuous flow ventricular assist device (VAD) and a pulsatile flow VAD in BIVAD patients. *Artificial Organs*, 41:242–252, 2017.
- [186] E. Benevento, A. Djebbari, Z. Keshavarz-Motamed, R. Cecere, and L. Kadem. Hemodynamic changes following aortic valve bypass: a mathematical approach. *PLoS ONE*, 10:e0123000, 2015.
- [187] J. Y. Kresh, S. K. Brockman, and A. Noordergraaf. Theoretical and experimental analysis of right ventricular bypass and univentricular circulatory support. *IEEE Transactions on Biomedical Engineering*, 37:121–127, 1990.
- [188] F. Migliavacca et al. Modeling of the Norwood circulation: effects of shunt size, vascular resistances, and heart rate. *American Journal of Physiology – Heart and Circulatory Physiology*, 280:H2076–H2086, 2001.
- [189] E. Ben-Assa et al. Ventricular stroke work and vascular impedance refine the characterization of patients with aortic stenosis. *Science Translational Medicine*, 11:eaaw0181, 2019.
- [190] Z. Keshavarz-Motamed et al. Elimination of transcoarctation pressure gradients has no impact on left ventricular function or aortic shear stress after intervention in patients with mild coarctation. *JACC: Cardiovascular Interventions*, 9:1953–1965, 2016.
- [191] Z. Keshavarz-Motamed et al. Mixed valvular disease following transcatheter aortic valve replacement: quantification and systematic differentiation using clinical measurements and image-based patient-specific in silico modeling. *Journal of the American Heart Association*, 9:e015063, 2020.
- [192] R. Sadeghi, S. Khodaei, J. Ganame, and Z. Keshavarz-Motamed. Towards non-invasive computational-mechanics and imaging-based diagnostic framework for personalized cardiology for coarctation. *Scientific Reports*, 10:9048, 2020.
- [193] Z. Keshavarz-Motamed. A diagnostic, monitoring, and predictive tool for patients with complex valvular, vascular and ventricular diseases. *Scientific Reports*, 10:6905, 2020.
- [194] B. Casas et al. Bridging the gap between measurements and modelling: a cardiovascular functional avatar. *Scientific Reports*, 7:6214, 2017.
- [195] Z. Duanmu, M. Yin, X. Fan, X. Yang, and X. Luo. A patient-specific lumped-parameter model of coronary circulation. *Scientific Reports*, 8:874, 2018.

- [196] C. J. Arthurs et al. Reproducing patient-specific hemodynamics in the Blalock–Taussig circulation using a flexible multidomain simulation framework: applications for optimal shunt design. *Frontiers in Pediatrics*, 5:78, 2017.
- [197] M. S. Vieira, C. J. Arthurs, T. Hussain, R. Razavi, and C. A. Figueroa. Patient-specific modeling of right coronary circulation vulnerability post-liver transplant in Alagille’s syndrome. *PLoS ONE*, 13:e0205829, 2018.
- [198] S. M. Shavik, C. Tossas-Betancourt, C. A. Figueroa, S. Baek, and L. C. Lee. Multiscale modeling framework of ventricular–arterial bidirectional interactions in the cardiopulmonary circulation. *Frontiers in Physiology*, 11:1–13, 2020.
- [199] D. Sengupta et al. Image-based modeling of hemodynamics in coronary artery aneurysms caused by Kawasaki disease. *Biomechanics and Modeling in Mechanobiology*, 11:915–932, 2012.
- [200] Y. Shi, P. Lawford, and R. Hose. Review of zero-D and one-D models of blood flow in the cardiovascular system. *Biomedical Engineering Online*, 10:33, 2011.
- [201] E.-O. Jung and W.-H. Lee. Lumped parameter models of cardiovascular circulation in normal and arrhythmia cases. *Journal of the Korean Mathematical Society*, 43:885–897, 2006.
- [202] S. A. Niederer and N. P. Smith. Using physiologically based models for clinical translation: predictive modelling, data interpretation or something in-between? *The Journal of Physiology*, 594:6849–6863, 2016.
- [203] J. W. Holmes and J. Lumens. Clinical applications of patient-specific models: the case for a simple approach. *Journal of Cardiovascular Translational Research*, 11:71–79, 2018.
- [204] M. Mirramezani and S. C. Shadden. A distributed lumped parameter model of blood flow. *Annals of Biomedical Engineering*, 48:2870–2886, 2020.
- [205] K. H. Parker and C. J. H. Jones. Forward and backward running waves in the arteries: analysis using the method of characteristics. *Journal of Biomechanical Engineering*, 112:322–326, 1990.
- [206] S. J. Sherwin, L. Formaggia, J. Peiró, and V. Franke. Computational modelling of 1D blood flow with variable mechanical properties and its application to the simulation of wave propagation in the human arterial system. *International Journal for Numerical Methods in Fluids*, 43:673–700, 2003.
- [207] J. Wan et al. A one-dimensional finite element method for simulation-based medical planning for cardiovascular disease. *Computer Methods in Biomechanics and Biomedical Engineering*, 5:195–206, 2002.

- [208] C. A. Taylor et al. Predictive medicine: computational techniques in therapeutic decision-making. *Computer Aided Surgery*, 4:231–247, 1999.
- [209] X. Xie, M. Zheng, D. Wen, Y. Li, and S. Xie. A new CFD-based non-invasive method for functional diagnosis of coronary stenosis. *Biomedical Engineering Online*, 17:36, 2018.
- [210] M. Mirramezani, S. L. Diamond, H. I. Litt, and S. C. Shadden. Reduced order models for transstenotic pressure drop in the coronary arteries. *Journal of Biomechanical Engineering*, 141:031005, 2019.
- [211] J. F. de Cañete. Object-oriented modeling and simulation of the closed-loop cardiovascular system by using SIMSCAPE. *Computers in Biology and Medicine*, 43:323–333, 2013.
- [212] K. Hassani, M. Navidbakhsh, and M. Rostami. Simulation of the cardiovascular system using equivalent electronic system. *Biomedical Papers*, 150:105–112, 2006.
- [213] S. Vennin et al. Identifying hemodynamic determinants of pulse pressure: a combined numerical and physiological approach. *Hypertension*, 70:1176–1182, 2017.
- [214] R. Pietrabissa, S. Mantero, T. Marotta, and L. Menicanti. A lumped parameter model to evaluate the fluid dynamics of different coronary bypasses. *Medical Engineering & Physics*, 18:477–484, 1996.
- [215] J. P. Mynard, D. J. Penny, and J. J. Smolich. Scalability and in vivo validation of a multiscale numerical model of the left coronary circulation. *American Journal of Physiology – Heart and Circulatory Physiology*, 306:H517–H528, 2014.
- [216] F. Huang and S. Ying. On-line parameter identification of the lumped arterial system model: a simulation study. *PLoS ONE*, 15:e0236012, 2020.
- [217] J. S. Tran, D. E. Schiavazzi, A. B. Ramachandra, A. M. Kahn, and A. L. Marsden. Automated tuning for parameter identification and uncertainty quantification in multi-scale coronary simulations. *Computers & Fluids*, 142:128–138, 2017.
- [218] D. E. Schiavazzi, A. Baretta, G. Pennati, T.-Y. Hsia, and A. L. Marsden. Patient-specific parameter estimation in single-ventricle lumped circulation models under uncertainty. *International Journal for Numerical Methods in Biomedical Engineering*, 33:1–59, 2017.
- [219] C. J. Arthurs et al. A flexible framework for sequential estimation of model parameters in computational hemodynamics. *Advanced Modeling and Simulation in Engineering Sciences*, 7:48, 2020.
- [220] R. B. D’Agostino, R. S. Vasan, M. J. Pencina, et al. General cardiovascular risk profile for use in primary care: the Framingham Heart Study. *Circulation*, 117:743–753, 2008.

- [221] D. C. Goff, D. M. Lloyd-Jones, G. Bennett, et al. 2013 ACC/AHA guideline on the assessment of cardiovascular risk: a report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines. *Circulation*, 129:S49–S73, 2014.
- [222] A. P. DeFilippis, R. Young, C. J. Carrubba, et al. An analysis of calibration and discrimination among multiple cardiovascular risk scores in a modern multiethnic cohort. *Annals of Internal Medicine*, 162:266–275, 2015.
- [223] I. A. Kakadiaris, M. Vrigkas, A. A. Yen, et al. Machine learning outperforms ACC/AHA CVD risk calculator in MESA. *Journal of the American Heart Association*, 7:e009476, 2018.
- [224] M. Chun, R. Clarke, B. J. Cairns, et al. Stroke risk prediction using machine learning: a prospective cohort study of 0.5 million Chinese adults. *Journal of the American Medical Informatics Association*, 28:1719–1727, 2021.
- [225] S. F. Weng, J. Reps, J. Kai, et al. Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLoS ONE*, 12:e0174944, 2017.
- [226] B. Ambale-Venkatesh, X. Yang, C. O. Wu, et al. Cardiovascular event prediction by machine learning: the Multi-Ethnic Study of Atherosclerosis. *Circulation Research*, 121:1092–1101, 2017.
- [227] S. A. Williams, R. Ostroff, M. A. Hinterberg, et al. A proteomic surrogate for cardiovascular outcomes that is sensitive to multiple mechanisms of change in risk. *Science Translational Medicine*, 14:eabj9625, 2022.
- [228] A. Bhuiyan, P. Kanti Roy, T. Bhuiyan, et al. A novel retinal vascular feature and machine learning-based brain white matter lesion prediction model. *medRxiv*, 2021. preprint.
- [229] A. R. Sridhar, A. Z.-H. Chen, J. J. Mayfield, et al. Identifying risk of adverse outcomes in COVID-19 patients via artificial intelligence-powered analysis of 12-lead intake electrocardiogram. *Cardiovascular Digital Health Journal*, 3:62–74, 2022.
- [230] G. Y. H. Lip, R. Nieuwlaat, R. Pisters, et al. Refining clinical risk stratification for predicting stroke and thromboembolism in atrial fibrillation using a novel risk factor-based approach: the Euro Heart Survey on atrial fibrillation. *Chest*, 137:263–272, 2010.
- [231] H. A. Van Den Ham, O. H. Klungel, D. E. Singer, et al. Comparative performance of ATRIA, CHADS2, and CHA2DS2-VASc risk scores predicting stroke in patients with atrial fibrillation: results from a national primary care database. *Journal of the American College of Cardiology*, 66:1851–1859, 2015.
- [232] S. Aspberg, Y. Chang, A. Atterman, et al. Comparison of the ATRIA, CHADS2, and CHA2DS2-VASc stroke risk scores in predicting ischaemic stroke in a large Swedish cohort of patients with atrial fibrillation. *European Heart Journal*, 37:3203–3210, 2016.

- [233] G. Y. H. Lip, G. Tran, A. Genaidy, et al. Improving dynamic stroke risk prediction in non-anticoagulated patients with and without atrial fibrillation: comparing common clinical risk scores and machine learning algorithms. *European Heart Journal – Quality of Care and Clinical Outcomes*, 8:548–556, 2022.
- [234] S. Jung, M. K. Song, E. Lee, et al. Predicting ischemic stroke in patients with atrial fibrillation using machine learning. *Frontiers in Bioscience (Landmark Edition)*, 27:80, 2022.
- [235] K. Kostev, T. Wu, Y. Wang, et al. Predicting the risk of ischemic stroke in patients treated with novel oral anticoagulants: a machine learning approach. *Neuroepidemiology*, 55:387–392, 2021.
- [236] Y. Chahine, M. J. Magoon, B. Maidu, J. C. del Álamo, P. M. Boyle, and N. Akoum. Machine learning and the conundrum of stroke risk prediction. *Arrhythmia & Electrophysiology Review*, 12:e07, 2023.
- [237] 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*, 378:686–707, 2019.
- [238] S. Cuomo, V. S. Di 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*, 92(3):88, 2022.
- [239] Y. Shin, J. Darbon, and G. E. Karniadakis. On the convergence of physics-informed neural networks for linear second-order elliptic and parabolic type PDEs. *Communications in Computational Physics*, 28(5):2042–2074, 2020.
- [240] S. Markidis. The old and the new: can physics-informed deep learning replace traditional linear solvers? *Frontiers in Big Data*, 4:669097, 2021.
- [241] A. Krishnapriyan, A. Gholami, S. Zhe, et al. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 26548–26560. Curran Associates, Inc., 2021.
- [242] Q. He, D. Barajas-Solano, G. Tartakovsky, et al. Physics-informed neural networks for multi-physics data assimilation with application to subsurface transport. *Advances in Water Resources*, 141:103610, 2020.
- [243] G. Kissas, Y. Yang, E. Hwuang, et al. 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*, 358:112623, 2020.
- [244] G. Pang, L. Lu, and G. E. Karniadakis. fPINNs: Fractional physics-informed neural networks. *SIAM Journal on Scientific Computing*, 41(4):A2603–A2626, 2019.

- [245] A. D. Jagtap, E. Kharazmi, and G. E. Karniadakis. Conservative physics-informed neural networks on discrete domains for conservation laws: applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020.
- [246] L. D. McClenny, M. A. Haile, and U. M. Braga-Neto. Tensordiffeq: scalable multi-GPU forward and inverse solvers for physics-informed neural networks. *arXiv preprint*, 2021.
- [247] S. Wang, Y. Teng, and P. Perdikaris. Understanding and Mitigating Gradient Flow Pathologies in Physics-Informed Neural Networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.
- [248] E. Kharazmi, Z. Zhang, and G. E. Karniadakis. Variational physics-informed neural networks for solving partial differential equations. *arXiv preprint*, 2019.
- [249] K. Xu and E. Darve. ADCME: Learning spatially-varying physical fields using deep neural networks. *arXiv preprint*, 2020.
- [250] S. Wang, S. Sankaran, and P. Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421:116813, 2024.
- [251] G. Kutyniok. The mathematics of artificial intelligence. *arXiv preprint*, 2022.
- [252] T. De Ryck and S. Mishra. Error analysis for physics-informed neural networks (PINNs) approximating Kolmogorov PDEs. *Advances in Computational Mathematics*, 48(6):79, 2022.
- [253] S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mechanica Sinica*, 37(12):1727–1738, 2021.
- [254] M. Raissi, A. Yazdani, and G. E. Karniadakis. Hidden fluid mechanics: learning velocity and pressure fields from flow visualizations. *Science*, 367(6481):1026–1030, 2020.
- [255] M. Sarabian, H. Babaei, and K. Laksari. Physics-informed neural networks for brain hemodynamic predictions using medical imaging. *IEEE Transactions on Medical Imaging*, 41(9):2285–2303, 2022.
- [256] A. Arzani, J. X. Wang, and R. M. D’Souza. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. *Physics of Fluids*, 33(7):071905, 2021.
- [257] M. Yin, X. Zheng, J. D. Humphrey, et al. Non-invasive inference of thrombus material properties with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 375:113603, 2021.
- [258] J. Eriksson, C. J. Carlhäll, P. Dyverfeldt, J. Engvall, A. F. Bolger, and T. Ebbers. Semi-automatic quantification of 4D left ventricular blood flow. *Journal of Cardiovascular Magnetic Resonance*, 12(1):9, 2010.

- [259] J. Eriksson, P. Dyverfeldt, J. Engvall, A. F. Bolger, T. Ebbers, and C. J. Carlhäll. Quantification of presystolic blood flow organization and energetics in the human left ventricle. *American Journal of Physiology – Heart and Circulatory Physiology*, 300(6):H2135–H2141, 2011.
- [260] P. D. Gatehouse, J. Keegan, L. A. Crowe, S. Masood, R. H. Mohiaddin, K. F. Kreitner, and D. N. Firmin. Applications of phase-contrast flow and velocity imaging in cardiovascular MRI. *European Radiology*, 15(10):2172–2184, 2005.
- [261] A. E. Arai, C. C. III Gaither, F. H. Epstein, R. S. Balaban, and S. D. Wolff. Myocardial velocity gradient imaging by phase-contrast MRI with application to regional function in myocardial ischemia. *Magnetic Resonance in Medicine*, 42(1):98–109, 1999.
- [262] D. T. Wymer, K. P. Patel, W. F. Burke, and V. K. Bhatia. Phase-contrast MRI: physics, techniques, and clinical applications. *RadioGraphics*, 40(1):122–140, 2020.
- [263] A. Raisinghani, P. Rafter, P. Phillips, M. A. Vannan, and A. N. DeMaria. Microbubble contrast agents for echocardiography: rationale, composition, ultrasound interactions, and safety. *Cardiology Clinics*, 22(2):171–180, 2004.
- [264] P. P. Sengupta, B. K. Khandheria, J. Korinek, A. Jahangir, S. Yoshifuku, I. Milosevic, and M. Belohlavek. Left ventricular isovolumic flow sequence during sinus and paced rhythms: new insights from use of high-resolution Doppler and ultrasonic digital particle imaging velocimetry. *Journal of the American College of Cardiology*, 49(8):899–908, 2007.
- [265] A. Falahatpisheh, G. Pedrizzetti, and A. Kheradvar. Three-dimensional reconstruction of cardiac flows based on multi-planar velocity fields. *Experiments in Fluids*, 55(11):1848, 2014.
- [266] C. Prinz, R. Faludi, A. Walker, M. Amzulescu, H. Gao, T. Uejima, A. G. Fraser, and J.-U. Voigt. Can echocardiographic particle image velocimetry correctly detect motion patterns as they occur in blood inside heart chambers? A validation study using moving phantoms. *Cardiovascular Ultrasound*, 10(1):24, 2012.
- [267] A. S. Daae, M. S. Wigen, S. Fadnes, L. Løvstakken, and A. Støylen. Intraventricular vector flow imaging with blood speckle tracking in adults: feasibility, normal physiology and mechanisms in healthy volunteers. *Ultrasound in Medicine and Biology*, 47(12):3501–3513, 2021.
- [268] S. A. Nytnes, S. Fadnes, M. S. Wigen, L. Mertens, and L. Løvstakken. Blood speckle-tracking based on high-frame rate ultrasound imaging in pediatric cardiology. *Journal of the American Society of Echocardiography*, 33(4):493–503.e5, 2020.
- [269] J. Zhang, S. Gajjala, P. Agrawal, G. H. Tison, L. A. Hallock, L. Beussink-Nelson, M. H. Lassen, E. Fan, M. A. Aras, C. Jordan, K. E. Fleischmann, M. Melisko, A. Qasim, S. J. Shah, R. Bajcsy, and R. C. Deo. Fully automated echocardiogram interpretation in clinical practice. *Circulation*, 138(16):1623–1635, 2018.

- [270] Z. I. Attia, P. A. Noseworthy, F. Lopez-Jimenez, S. J. Asirvatham, A. J. Deshmukh, B. J. Gersh, R. E. Carter, X. Yao, A. A. Rabinstein, B. J. Erickson, S. Kapa, and P. A. Friedman. An artificial intelligence-enabled ECG algorithm for the identification of patients with atrial fibrillation during sinus rhythm: a retrospective analysis of outcome prediction. *The Lancet*, 394(10201):861–867, 2019.
- [271] T. Grandits, S. Pezzuto, F. S. Costabal, P. Perdikaris, T. Pock, G. Plank, and R. Krause. Learning atrial fiber orientations and conductivity tensors from intracardiac maps using physics-informed neural networks. In *Functional Imaging and Modeling of the Heart (FIMH)*, pages 650–658. Springer, 2021.
- [272] H. Nahas, J.S. Au, T. Ishii, B.YS. Yiu, A.JY. Chee, and CH. Alfred. A deep learning approach to resolve aliasing artifacts in ultrasound color flow imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(12):2615–2628, 2020.
- [273] E. Ferdian, D.J. Dubowitz, C.A. Mauger, A. Wang, and A.A. Young. WSSNet: aortic wall shear stress estimation using deep learning on 4D flow MRI. *Frontiers in cardiovascular medicine*, 8:1969, 2022.
- [274] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [275] M. F. Fathi, I. Perez-Raya, A. Baghaie, P. Berg, G. Janiga, A. Arzani, and R. M. D’Souza. Super-resolution and denoising of 4D-flow MRI using physics-informed deep neural networks. *Computer Methods and Programs in Biomedicine*, 197:105729, 2020.
- [276] P. Ramachandran, B. Zoph, and Q.V. Le. Searching for Activation Functions. *ArXiv*, abs/1710.05941, 2018.
- [277] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256. PMLR, 2010.
- [278] T. Salimans and D.P. Kingma. Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks. *Advances in neural information processing systems*, 2016.
- [279] A.G. Baydin, B.A. Pearlmutter, A.A. Radul, and J.M. Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18:1–43, 2018.
- [280] S. Wang, X. Yu, and P. Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.
- [281] D.P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *ArXiv*, 2017.

- [282] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, et al. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283. USENIX Association, 2016.
- [283] D. Muraru, L.P. Badano, D. Peluso, L.D. Bianco, S. Casabianca, G. Kocabay, G. Zoppellaro, and S. Iliceto. Comprehensive analysis of left ventricular geometry and function by three-dimensional echocardiography in healthy adults. *Journal of the American Society of Echocardiography: Official Publication of the American Society of Echocardiography*, 26(6):618–28, 2013.
- [284] H.F.J. Mannaerts, J.A. van der Heide, O. Kamp, M.G. Stoel, J. Twisk, and C.A. Visser. Early identification of left ventricular remodelling after myocardial infarction, assessed by transthoracic 3D echocardiography. *European Heart Journal*, 25(8):680–7, 2004.
- [285] M. Guerrero-Hurtado, J. M. Catalán, M. Moriche, A. Gonzalo, and O. Flores. A Python-based flow solver for numerical simulations using an immersed boundary method on single GPUs. *Computers & Fluids*, 288:106511, 2025.
- [286] M. Uhlmann. An immersed boundary method with direct forcing for the simulation of particulate flows. *Journal of Computational Physics*, 209(2):448–476, 2005.
- [287] E. Khalili, C. Daversin-Catty, A.L. Olivares, J. Mill, O. Camara, and K. Valen-Sendstad. On the importance of fundamental computational fluid dynamics toward a robust and reliable model of left atrial flows. *International Journal for Numerical Methods in Biomedical Engineering*, 40(4):e3804, 2024.
- [288] M. Loecher, E. Schrauben, K.M. Johnson, and O. Wieben. Phase unwrapping in 4D MR flow with a 4D single-step laplacian algorithm. *Journal of Magnetic Resonance Imaging*, 43(4):833–842, 2016.
- [289] J.H. Seo, V. Vedula, T. Abraham, A.C. Lardo, F. Dawoud, H. Luo, and R. Mittal. Effect of the mitral valve on diastolic flow patterns. *Physics of Fluids*, 26(12):121901, 2014.
- [290] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. In *International conference on machine learning*, pages 5301–5310. PMLR, 2019.
- [291] M. Henglin, G. Stein, P.V. Hushcha, J. Snoek, A.B. Wiltchko, and S. Cheng. Machine Learning Approaches in Cardiovascular Imaging. *Circulation: Cardiovascular Imaging*, 10(10):e005614, 2017.
- [292] G. Quer, R. Arnaout, M. Henne, and R. Arnaout. Machine Learning and the Future of Cardiovascular Care. *Journal of the American College of Cardiology*, 77(3):300–313, 2021.

- [293] S.L. Brunton, B.R. Noack, and P. Koumoutsakos. Machine Learning for Fluid Mechanics. *Annual Review of Fluid Mechanics*, 52(1):477–508, 2020.
- [294] X. Morales Ferez, J. Mill, K.A. Juhl, C. Acebes, X. Iriart, B. Legghe, H. Cochet, O. De Backer, R.R. Paulsen, and O. Camara. Deep Learning Framework for Real-Time Estimation of in-silico Thrombotic Risk Indices in the Left Atrial Appendage. *Frontiers in Physiology*, 12:694945, 2021.
- [295] G. Hasanuzzaman, H. Eivazi, S. Merbold, C. Egbers, and R. Vinuesa. Enhancement of PIV measurements via physics-informed neural networks. *Measurement Science and Technology*, 34(4):044002, 2023.
- [296] C. Herrero Martin, A. Oved, R.A. Chowdhury, E. Ullmann, N.S. Peters, A.A. Bharath, and M. Varela. EP-PINNs: Cardiac Electrophysiology Characterisation Using Physics-Informed Neural Networks. *Frontiers in Cardiovascular Medicine*, 8, 2022.
- [297] S. Buoso, T. Joyce, and S. Kozerke. Personalising left-ventricular biophysical models of the heart using parametric physics-informed neural networks. *Medical Image Analysis*, 71:102066, 2021.
- [298] B. Maidu, M. Guerrero-Hurtado, C. Nguyen, P. Martinez-Legazpi, A. Kahn, J. Bermejo, O. Flores, and J.C. del Alamo. All-In-One Left Ventricular Vector Flow, Pressure, & Clotting Risk Mapping by Multi-Physics-Informed Neural Network. *Bulletin of the American Physical Society*, 2023.
- [299] F. Isensee, P.F. Jaeger, S.A.A. Kohl, J. Petersen, and K. H. Maier-Hein. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, 2021.
- [300] M. Löcke, J.E. Garay Labra, P. Franco, S. Uribe, and C. Bertoglio. A comparison of phase unwrapping methods in velocity-encoded MRI for aortic flows. *Magnetic Resonance in Medicine*, 90(5):2102–2115, 2023.
- [301] S. Muth, S. Dort, I.A. Sebag, M. Blais, and D. Garcia. Unsupervised dealiasing and denoising of color-Doppler data. *Medical Image Analysis*, 15(4):577–588, 2011.
- [302] N.R. Shah, C. Cevik, A. Hernandez, I.D. Gregoric, O.H. Frazier, and Ra.F. Stainback. Transthoracic Echocardiography of the HeartWare Left Ventricular Assist Device. *Journal of Cardiac Failure*, 18(9):745–748, 2012.
- [303] R.A. Bafghi and M. Raissi. PINNs-TF2: Fast and User-Friendly Physics-Informed Neural Networks in TensorFlow V2. *ArXiv*, 2023.
- [304] L. Rossini, P. Martinez-Legazpi, V. Vu, L. Fernández-Friera, C. Pérez del Villar, S. Rodríguez-López, Y. Benito, M.G. Borja, D. Pastor-Escuredo, R. Yotti, M.J. Ledesma-Carbayo, A.M. Kahn, B. Ibáñez, F. Fernández-Avilés, K. May-Newman, J. Bermejo, and J.C. del Álamo. A clinical method for mapping and quantifying blood stasis in the left ventricle. *Journal of Biomechanics*, 49(11):2152–2161, 2016.

- [305] M. Guerrero-Hurtado, M. Garcia-Villalba, A. Gonzalo, P. Martinez-Legazpi, A.M. Kahn, E. McVeigh, J. Bermejo, J.C. del Alamo, and O. Flores. Efficient multi-fidelity computation of blood coagulation under flow. *PLOS Computational Biology*, 19:1–25, 2023.
- [306] P. P. Sengupta, G. Pedrizetti, and J. Narula. Multiplanar Visualization of Blood Flow Using Echocardiographic Particle Imaging Velocimetry. *JACC: Cardiovascular Imaging*, 5(5):566–569, 2012.
- [307] J. Voorneveld, H. Saaïd, C. Schinkel, N. Radeljic, B. Lippe, F. J. H. Gijssen, A. F. W. van der Steen, N. de Jong, T. Claessens, H. J. Vos, S. Kenjeres, and J. G. Bosch. 4-D Echo-Particle Image Velocimetry in a Left Ventricular Phantom. *Ultrasound in Medicine & Biology*, 46(3):805–817, 2020.
- [308] P. Chen, R. J. G. van Sloun, S. Turco, H. Wijkstra, D. Filomena, L. Agati, P. Houthuizen, and M. Mischi. Blood Flow Patterns Estimation in the Left Ventricle with Low-Rate 2D and 3D Dynamic Contrast-Enhanced Ultrasound. *Computer Methods and Programs in Biomedicine*, 198:105810, 2021.
- [309] A. Gomez, K. Pushparajah, J. M. Simpson, D. Giese, T. Schaeffter, and G. Penney. A Sensitivity Analysis on 3D Velocity Reconstruction from Multiple Registered Echo Doppler Views. *Medical Image Analysis*, 17(6):616–631, 2013.
- [310] A. Gomez, A. de Vecchi, M. Jantsch, W. Shi, K. Pushparajah, J. M. Simpson, N. P. Smith, D. Rueckert, T. Schaeffter, and G. P. Penney. 4D Blood Flow Reconstruction Over the Entire Ventricle From Wall Motion and Blood Velocity Derived From Ultrasound Data. *IEEE Transactions on Medical Imaging*, 34(11):2298–2308, 2015.
- [311] T. Grønli, M. Wigen, P. Segers, and L. Lovstakken. A Fast 4D B-Spline Framework for Model-Based Reconstruction and Regularization in Vector Flow Imaging. In *2018 IEEE International Ultrasonics Symposium (IUS)*, pages 1–9, 2018.
- [312] B. Maidu, A. Gonzalo, M. Guerrero-Hurtado, C. Bargellini, P. Martinez-Legazpi, J. Bermejo, F. Contijoch, O. Flores, M. García-Villalba, E. McVeigh, A. M. Kahn, and J. C. del Alamo. Left Atrial Flow and Thrombosis Risk from 4D CT Contrast Dynamics by Physics-Informed Neural Network and Indicator Dilution Theory. *bioRxiv*, 2026.
- [313] S. Wang, B. Li, Y. Chen, and P. Perdikaris. PirateNets: Physics-informed Deep Learning with Residual Adaptive Networks. *Journal of Machine Learning Research*, 25(402):1–51, 2024.
- [314] S. Wang, B. Li, P. Perdikaris, et al. Gradient alignment in physics-informed neural networks: A second-order optimization perspective. *Advances in Neural Information Processing Systems*, 38:168482–168532, 2026.
- [315] E. Benjamin, P. Muntner, A. Alonso, and et al. Heart Disease and Stroke Statistics-2019 Update: A Report From the American Heart Association. *Circulation*, 139(10):e56–e528, 2019.

- [316] M. Lamassa, A. Di Carlo, G. Pracucci, A. M. Basile, G. Trefoloni, P. Vanni, S. Spolveri, M. C. Baruffi, G. Landini, A. Ghetti, C. D. Wolfe, and D. Inzitari. Characteristics, outcome, and care of stroke associated with atrial fibrillation in Europe: data from a multicenter multinational hospital-based registry (The European Community Stroke Project). *Stroke*, 32(2):392–398, 2001.
- [317] G.Y.H. Lip, R. Nieuwlaat, R. Pisters, D. A. Lane, H.J.G.M. Crijns, et al. Refining Clinical Risk Stratification for Predicting Stroke and Thromboembolism in Atrial Fibrillation Using a Novel Risk Factor-Based Approach: The Euro Heart Survey on Atrial Fibrillation. *Chest*, 137(2):263–272, 2010.
- [318] T. J. Siddiqi, M. S. Usman, I. Shahid, J. Ahmed, S. U. Khan, L. Ya’qoub, C. S. Rihal, and M. Alkhouli. Utility of the CHA₂DS₂-VASc score for predicting ischaemic stroke in patients with or without atrial fibrillation: a systematic review and meta-analysis. *European Journal of Preventive Cardiology*, 29(4):625–631, 2022.
- [319] M. Markl, D. C. Lee, J. Ng, M. Carr, J. Carr, and J. J. Goldberger. Left Atrial 4-Dimensional Flow Magnetic Resonance Imaging: Stasis and Velocity Mapping in Patients With Atrial Fibrillation. *Investigative Radiology*, 51(3):147–154, 2016.
- [320] B. T. Costello, A. Voskoboinik, A. M. Qadri, M. Rudman, M. C. Thompson, F. Touma, A. La Gerche, J. L. Hare, S. Papapostolou, J. M. Kalman, P. M. Kistler, and A. J. Taylor. Measuring atrial stasis during sinus rhythm in patients with paroxysmal atrial fibrillation using 4 Dimensional flow imaging: 4D flow imaging of atrial stasis. *International Journal of Cardiology*, 315:45–50, 2020.
- [321] M. Spartera, G. Pessoa-Amorim, A. Stracquadanio, A. Von Ende, A. Fletcher, P. Manley, S. Neubauer, V. M. Ferreira, B. Casadei, A. T. Hess, and R. S. Wijesurendra. Left atrial 4D flow cardiovascular magnetic resonance: a reproducibility study in sinus rhythm and atrial fibrillation. *Journal of Cardiovascular Magnetic Resonance*, 23(1):29, 2021.
- [322] P. Martinez-Legazpi, L. Rossini, C. P. del Villar, Y. Benito, C. Devesa-Cordero, R. Yotti, A. Delgado-Montero, A. Gonzalez-Mansilla, A. M. Kahn, F. Fernandez-Avilés, J. C. del Álamo, and J. Bermejo. Stasis Mapping Using Ultrasound: A Prospective Study in Acute Myocardial Infarction. *JACC: Cardiovascular Imaging*, 11(3):514–515, 2018.
- [323] E. Rodríguez-González, P. Martínez-Legazpi, T. Mombiela, A. González-Mansilla, A. Delgado-Montero, J. A. Guzmán-De-Villoria, F. Díaz-Otero, R. Prieto-Arévalo, M. Juárez, M. C. G. Del Rey, et al. Stasis imaging predicts the risk of cardioembolic events related to acute myocardial infarction: the ISBITAMI study. *Revista Española de Cardiología (English Edition)*, 2024.
- [324] E. Rodríguez-González, P. Martinez-Legazpi, A. González-Mansilla, M. Espinosa, T. Mombiela, J. A. Guzmán De-Villoria, M. G. Borja, F. Díaz-Otero, R. Gómez de Antonio, P. Fernández-García, et al. Cardiac stasis imaging, stroke, and silent brain infarcts in patients with nonischemic dilated cardiomyopathy. *American Journal of Physiology-Heart and Circulatory Physiology*, 327(2):H446–H453, 2024.

- [325] T. Inoue and Y. Suematsu. Poor enhancement pattern of left atrial appendage in cardiac computed tomography is associated with stroke in persistent atrial fibrillation patients. *Journal of Thoracic Disease*, 11(8):3315, 2019.
- [326] X.-W. Liu, L.-q. Yan, G.-s. Wang, Y.-h. Sun, W.-j. Bao, H.-w. Zhang, L. Yang, C. Chen, M.-Y. Li, L. Yang, et al. Left atrial appendage filling defects restricted to the early phase of cardiac computed tomography is significantly associated with ischemic stroke. *Clinical Imaging*, 98:16–21, 2023.
- [327] S. S. Nio, L. A. Rinkel, O. N. Cramer, Z. B. Özata, C. FP. Beemsterboer, V. Guglielmi, B. J. Bouma, S. M. Boekholdt, N. HJ. Lobé, L. FM. Beenen, et al. Left atrial appendage opacification on cardiac computed tomography in acute ischemic stroke: the clinical implications of slow-flow. *Journal of the American Heart Association*, 13(17):e034106, 2024.
- [328] G. Garcia-Isla, A. L. Olivares, E. Silva, M. Nuñez-Garcia, C. Butakoff, D. Sanchez-Quintana, H. G. Morales, X. Freixa, J. Noailly, T. De Potter, and O. Camara. Sensitivity analysis of geometrical parameters to study haemodynamics and thrombus formation in the left atrial appendage. *Int J Numer Method Biomed Eng*, 34(5):e3100, 2018.
- [329] S. Sanatkhan, S. Nedios, P. G. Menon, A. Bollmann, G. Hindricks, and S. G. Shroff. Subject-Specific Calculation of Left Atrial Appendage Blood-Borne Particle Residence Time Distribution in Atrial Fibrillation. *Frontiers in Physiology*, 12, 2021.
- [330] R. Fang, Z. Wang, X. Zhao, J. Wang, Y. Li, Y. Zhang, Q. Chen, J. Wang, Q. Liu, M. Chen, et al. Stroke risk evaluation for patients with atrial fibrillation: Insights from left atrial appendage with fluid-structure interaction analysis. *Comput. Biol. Med.*, 148:105897, 2022.
- [331] Z. Zhang, J. Zhu, M. Wu, M. Neidlin, W. Wu, and P. Wu. Computational modeling of hemodynamics and risk of thrombosis in the left atrial appendage using patient-specific blood viscosity and boundary conditions at the mitral valve. *Biomech. Model. Mechanobiol.*, 22(4):1447–1457, 2023.
- [332] S. Bäck, I. Skoda, J. Lantz, L. Henriksson, L. O. Karlsson, A. Persson, C. J. Carlhäll, and T. Ebbens. Elevated atrial blood stasis in paroxysmal atrial fibrillation during sinus rhythm: a patient-specific computational fluid dynamics study. *Front. cardiovasc. med.*, 10:1219021, 2023.
- [333] N. Paliwal, H. C. Park, Y. Mao, S. J. Hong, Y. Lee, D. D. Spragg, H. Calkins, and N. A. Trayanova. Slow blood-flow in the left atrial appendage is associated with stroke in atrial fibrillation patients. *Heliyon*, 10(5), 2024.
- [334] A. Qureshi, P. Melidoro, M. Balmus, G.Y.H. Lip, D.A. Nordsletten, S.E. Williams, O. Aslanidi, and A. de Vecchi. MRI-based modelling of left atrial flow and coagulation to predict risk of thrombogenesis in atrial fibrillation. *Medical Image Analysis*, 101:103475, 2025.

- [335] H. A. Kjeldsberg, C. Albors, J. Mill, D. V. Medel, O. Camara, J. Sundnes, and K. Valen-Sendstad. Impact of left atrial wall motion assumptions in fluid simulations on proposed predictors of thrombus formation. *International Journal for Numerical Methods in Biomedical Engineering*, page e3825, 2024.
- [336] E. Durán, M. García-Villalba, P. Martínez-Legazpi, A. Gonzalo, E. McVeigh, A. M. Kahn, J. Bermejo, O. Flores, and J. C. Del Álamo. Pulmonary vein flow split effects in patient-specific simulations of left atrial flow. *Comput. Biol. Med.*, 163:107128, 2023.
- [337] E. Khalili, C. Daversin-Catty, AL Olivares, J. Mill, O. Camara, and K. Valen-Sendstad. On the importance of fundamental computational fluid dynamics toward a robust and reliable model of left atrial flows. *International Journal for Numerical Methods in Biomedical Engineering*, pages e3804–e3804, 2024.
- [338] X. Liu, H. Lin, X. Liu, J. Qian, S. Cai, H. Fan, and Q. Gao. LAFlowNet: A dynamic graph method for the prediction of velocity and pressure fields in left atrium and left atrial appendage. *Engineering Applications of Artificial Intelligence*, 136:108896, 2024.
- [339] Jens H. Henriksen, Gorm B. Jensen, and Henrik B. W. Larsson. A century of indicator dilution technique. *Clinical Physiology and Functional Imaging*, 34(1):1–9, 2014.
- [340] V. Mor-Avi, D. David, S. Akselrod, Y. Bitton, and I. Choshniak. Myocardial regional blood flow: Quantitative measurement by computer analysis of contrast enhanced echocardiographic images. *Ultrasound in Medicine & Biology*, 19(8):619–633, 1993.
- [341] P. S. Tofts. Modeling tracer kinetics in dynamic Gd-DTPA MR imaging. *Journal of Magnetic Resonance Imaging*, 7(1):91–101, 1997.
- [342] K. A. Miles, M. Hayball, and A. K. Dixon. Colour perfusion imaging: a new application of computed tomography. *The Lancet*, 337(8742):643–645, 1991.
- [343] M. T. Madsen. A simplified formulation of the gamma variate function. *Physics in Medicine & Biology*, 37(7):1597, 1992.
- [344] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, 31(3):1116–1128, 2006.
- [345] C.-W. Shu. Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*, pages 325–432, 1998.
- [346] L. Lu, R. Pestourie, W. Yao, Z. Wang, F. Verdugo, and S. G. Johnson. Physics-Informed Neural Networks with Hard Constraints for Inverse Design. *SIAM Journal on Scientific Computing*, 43(6):B1105–B1132, 2021.

- [347] J. L. Stephenson. Theory of the measurement of blood flow by the dilution of an indicator. *Bulletin of Mathematical Biophysics*, 10:117–121, 1948.
- [348] José M. González-Fernández. Theory of the Measurement of the Dispersion of an Indicator in Indicator-Dilution Studies. *Circulation Research*, 10(3):409–428, 1962.
- [349] H. K. Thompson Jr, C. F. Starmer, R. E. Whalen, and H. D. McIntosh. Indicator transit time considered as a gamma variate. *Circulation research*, 14(6):502–515, 1964.
- [350] James B. Bassingthwaight and Francis H. Ackerman. Mathematical linearity of circulatory transport. *Journal of Applied Physiology*, 22(5):879–888, May 1967.
- [351] K. Zierler. Indicator dilution methods for measuring blood flow, volume, and other properties of biological systems: a brief history and memoir. *Annals of Biomedical Engineering*, 28(8):836–848, 2000.
- [352] K. A. Miles. Measurement of tissue perfusion by dynamic computed tomography. *British Journal of Radiology*, 64(761):409–412, 1991.
- [353] F. Calamante, D. G. Gadian, and A. Connelly. Quantification of Perfusion Using Bolus Tracking Magnetic Resonance Imaging in Stroke. *Stroke*, 33(4):1146–1151, 2002.
- [354] L. Østergaard. Principles of cerebral perfusion imaging by bolus tracking. *Journal of Magnetic Resonance Imaging*, 22(6):710–717, 2005.
- [355] A. Arzani, JX. Wang, and R.M. D’Souza. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. *Physics of Fluids*, 33(7), 2021.
- [356] 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*, 358:112623, 2020.
- [357] M. Yin, X. Zheng, J.D. Humphrey, and G.E. Karniadakis. Non-invasive inference of thrombus material properties with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 375:113603, 2021.
- [358] C. Garay, J. Dunstan, F. Uribe, and F. Costabal. Model-uncertainty-aware physics-informed neural networks for parameter estimation in cardiovascular models. *Comput. Biol. Med.*, page 108706, 2024.
- [359] R. L.M. van Herten, A. Chiribiri, M. Breeuwer, M. Veta, and C. M. Scannell. Physics-informed neural networks for myocardial perfusion mri quantification. *Medical Image Analysis*, 78:102399, 2022.

- [360] L. Rotkopf, C. Ziener, N. von Knebel-Doeberitz, S. Wolf, A. Hohmann, W. Wick, M. Bendszus, H.-P. Schlemmer, D. Paech, and F. Kurz. A physics-informed deep learning framework for dynamic susceptibility contrast perfusion MRI. *Medical Physics*, 51(12):9031–9040, 2024.
- [361] B. Zapf, J. Haubner, M. Kuchta, G. Ringstad, P. K. Eide, and K.-A. Mardal. Investigating molecular transport in the human brain from MRI with physics-informed neural networks. *Scientific Reports*, 12(1):15475, 2022.
- [362] H. J. Ling, S. Bru, J. Puig, F. Vixège, S. Mendez, F. Nicoud, P.-Y. Courand, O. Bernard, and D. Garcia. Physics-Guided Neural Networks for Intraventricular Vector Flow Mapping. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 71(11):1377–1388, 2024.
- [363] U.S. Food and Drug Administration. Drug label: Nda 020608, supplement s036 (prescribing information). FDA Drugs@FDA, 2017.
- [364] C. Xu, B. T. Cao, Y. Yuan, and G. Meschke. Transfer learning based physics-informed neural networks for solving inverse problems in engineering structures under different loading scenarios. *Computer Methods in Applied Mechanics and Engineering*, 405:115852, 2023.
- [365] Y. Stöcker, M. Guerrero-Hurtado, E. Durán, A. Gonzalo, Z. Ristić, Å. Telle, A. Kassar, R. Haykal, N. Akoum, P. M. Boyle, et al. Effect of wall motion sampling on cfd-derived left atrial flow metrics. *bioRxiv*, pages 2025–10, 2025.
- [366] J. Guo, G. S. Khurana, A. Gonzalo Grande, J. C. del Alamo, and F. Contijoch. Computed Tomography (CT)-derived Cardiovascular Flow Estimation Using Physics-Informed Neural Networks Improves with Sinogram-based Training: A Simulation Study. *arXiv*, 2025.
- [367] D. J. Barron, M. D. Kilby, B. Davies, J. G. Wright, T. J. Jones, and W. J. Brawn. Hypoplastic left heart syndrome. *Lancet*, 374(9689):551–564, 2009.
- [368] W. I. Norwood, J. K. Kirklin, and S. P. Sanders. Hypoplastic left heart syndrome: experience with palliative surgery. *The American Journal of Cardiology*, 45(1):87–91, 1980.
- [369] J. A. Feinstein, D. W. Benson, A. M. Dubin, and et al. Hypoplastic left heart syndrome: current considerations and expectations. *Journal of the American College of Cardiology*, 59(1 Suppl):S1–S42, 2012.
- [370] U. Bartram, J. Grunenfelder, and R. Van Praagh. Causes of death after the modified norwood procedure: a study of 122 postmortem cases. *The Annals of Thoracic Surgery*, 64(6):1795–1802, 1997.
- [371] J. P. Mynard, M. R. Davidson, D. J. Penny, and J. J. Smolich. A simple, versatile valve model for use in lumped parameter and one-dimensional cardiovascular models. *International Journal for Numerical Methods in Biomedical Engineering*, 28(6-7):626–641, 2012.

- [372] E. L. Bove, F. Migliavacca, M. R. de Leval, R. Balossino, G. Pennati, T. R. Lloyd, S. Khambadkone, T.-Y. Hsia, and G. Dubini. Use of mathematic modeling to compare and predict hemodynamic effects of the modified blalock–taussig and right ventricle–pulmonary artery shunts for hypoplastic left heart syndrome. *The Journal of Thoracic and Cardiovascular Surgery*, 136(2):312–320.e2, 2008.
- [373] G. Biglino, A. Giardini, T.-Y. Hsia, R. Figliola, A. M. Taylor, and S. Schievano. Modelling single ventricle physiology: Review of engineering tools to study first stage palliation of hypoplastic left heart syndrome. *Frontiers in Pediatrics*, 1:31, 2013.
- [374] A. M. Taylor-LaPole, L. M. Paun, D. Lior, J. D. Weigand, C. Puelz, and M. S. Olufsen. Parameter selection and optimization of a computational network model of blood flow in single-ventricle patients. *Journal of the Royal Society Interface*, 22:20240663, 2025.
- [375] R. Chen, J. Liang, M. Li, and E. Meng. Recent progress in blood flow sensing. *Sensors and Actuators A: Physical*, 387:116457, 2025.