

ADVANCING TIME-RESOLVED PHASE-CONTRAST MRI  
ANALYSIS: DEVELOPMENT AND APPLICATION OF THE  
INPUT-PARAMETERIZED PHYSICS-INFORMED NEURAL  
NETWORK (IP-PINN)

by

Amin Pashaei Kalajahi

A Dissertation Submitted in  
Partial Fulfillment of the  
Requirements for the Degree of

Doctor of Philosophy  
in Engineering

at

The University of Wisconsin–Milwaukee

August 2025

# ABSTRACT

## ADVANCING TIME-RESOLVED PHASE-CONTRAST MRI ANALYSIS: DEVELOPMENT AND APPLICATION OF THE INPUT-PARAMETERIZED PHYSICS-INFORMED NEURAL NETWORK (IP-PINN)

by

Amin Pashaei Kalajahi

The University of Wisconsin–Milwaukee, 2025  
Under the Supervision of Professor Roshan M. D’Souza

Time-Resolved Three-Dimensional Phase-Contrast MRI (4D Flow MRI) is a powerful non-invasive technique for quantitatively assessing cardiovascular hemodynamics. Despite its potential, the clinical application of 4D Flow MRI is constrained by coarse spatio-temporal resolution, acquisition noise, and artifacts including velocity aliasing and eddy current induced phase offsets. These limitations compromise the accuracy and reliability of hemodynamic assessments, particularly in complex vascular structures. This thesis proposes a novel deep learning-based framework called the Input-Parameterized Physics-Informed Neural Network (IP-PINN) to address these challenges. By integrating advanced machine learning with the underlying physics of blood flow, the IP-PINN framework enhances low-resolution 4D Flow MRI data, attenuates acquisition noise, and mitigates velocity aliasing and phase offset artifacts. The framework leverages a ResNet-based convolutional neural network to encode input data into a latent vector, which is then utilized by a feedforward neural network to produce a continuous spatio-temporal representation of the variables of interest. The IP-PINN’s unique ability to generalize across different datasets by parameterizing the solution with respect to the input velocity-encoded images significantly reduces the need for time-intensive retraining. The IP-PINN does not require ground-truth labels and pre-training with either real low-resolution image data or synthetic data from computational fluid dynamics (CFD) simulations enhances the framework’s applicability. The IP-PINN

preserves the continuous spatio-temporal representation and the ability to generate truncation error-free derivatives characteristic of PINNs, while significantly expediting processing time. Benchmark tests against simulated datasets demonstrate that IP-PINN is over an order of magnitude faster than traditional PINNs, achieving superior accuracy. Additionally, the method generates high-resolution magnitude images for lumen boundary segmentation, relying solely on velocity-encoded scans and negating the need for reference scans. Building on an initial implementation that operated in complex image space and relied on the three velocity-encoded scans, the thesis extends IP-PINNs in two directions. First, reconstruction of three-component, three-dimensional (3D-3C) velocity fields and high-resolution spin-density maps in the vicinity of the imaging plane using data from minimally altered 2D PC-MRI sequence. Second, reconstruction of 3D-3C velocity maps and spin-density maps from a *pseudo one-point* 4D-Flow MRI sequence, that acquires only one velocity-encoded dataset per slice, reducing the raw data burden by 75%. For this extremely sparse regime, the data-fidelity term is reformulated directly in  $k$ -space, which preserves the exact acquisition physics, avoids the convolution blurring inherent in image-space. With a dramatically reduced execution time of approximately two minutes, operating on undersampled acquired data, and simplified operational requirements (no need for specifying geometry and boundary conditions), the IP-PINN promises to advance the state-of-the-art in hemodynamic assessment, offering a robust and efficient solution for enhancing time-resolved PCMRI data, with significant implications for both clinical practice and cardiovascular research.

© Copyright by Amin Pashaei Kalajahi, 2025  
All Rights Reserved

To my beloved parents, for their unwavering love, sacrifices, and lifelong support that have shaped every step of my journey.

And to my wife, for her endless patience, encouragement, and companionship through every challenge and success.

This work would not have been possible without your strength and belief in me.

# TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	The Role of Hemodynamics in Cardiovascular Diseases . . . . .	2
1.3	Assessment Techniques for Hemodynamics . . . . .	3
1.4	In-vivo Blood Flow Measurements . . . . .	4
1.4.1	Invasive Methods . . . . .	5
1.4.2	Doppler Echocardiography . . . . .	5
1.4.3	Ultrasound Imaging Velocimetry (Echo-PIV) . . . . .	9
1.4.4	Time-Resolved Phase-Contrast MRI . . . . .	10
1.4.5	4D Flow MRI . . . . .	12
1.5	Computational Methods for Hemodynamics Assessment . . . . .	14
1.5.1	Patient-Specific Computational Fluid Dynamics (CFD) . . . . .	15
1.5.2	Data Assimilation . . . . .	16
1.6	Thesis Outline and Objectives . . . . .	29
1.6.1	Objectives . . . . .	29
1.6.2	Thesis Outline . . . . .	31
<b>2</b>	<b>Fundamentals of Time-Resolved Phase-Contrast MRI</b>	<b>34</b>
2.1	Velocity Encoding and Signal Acquisition . . . . .	35
2.2	Image Reconstruction . . . . .	37
2.3	Limitations . . . . .	38
2.3.1	Phase Offset Errors . . . . .	39
2.3.2	Velocity aliasing . . . . .	41
2.3.3	Random noise in k-space . . . . .	42
<b>3</b>	<b>Development of the Input-Parametrized Physics-Informed Neural Network</b>	<b>44</b>
3.1	Introduction . . . . .	44
3.2	Background . . . . .	45
3.2.1	Convolutional Neural Networks (CNNs) . . . . .	46
3.2.2	Multilayer Perceptrons (MLPs) . . . . .	47
3.2.3	Activation Functions . . . . .	48
3.2.4	Physics-Informed Neural Networks (PINNs) . . . . .	49
3.2.5	Loss Function . . . . .	50
3.2.6	Optimization Process . . . . .	51
3.2.7	Back-Propagation . . . . .	52

3.3	Input Parameterized Physics Informed Neural Net (IP-PINN)	53
3.3.1	Flow Image Encoder	54
3.3.2	Augmented PINN	55
3.3.3	Loss Function	57
3.4	Conclusion	59
<b>4</b>	<b>Application of IP-PINN for Advanced 4D Flow MRI</b>	<b>61</b>
4.1	Introduction	61
4.2	Generation of Synthetic Training and Test Data	62
4.2.1	Generating Synthetic 4D-Flow MRI from CFD Simulations	64
4.3	Error Metrics	65
4.4	Pre-Training IP-PINN	66
4.5	Applying Pre-Trained IP-PINN to Unseen Data	67
4.5.1	Time Interpolation	68
4.5.2	Luminal Boundary Prediction	71
4.5.3	Performance Evaluation on Noisy Synthetic 4D-Flow MRI with Velocity Aliasing	71
4.5.4	Comparison with Physics-Informed Neural Network (PINN) and Divergence Free Radial Basis Functions (DF-RBFs)	74
4.6	Discussion	76
4.7	Conclusion	77
<b>5</b>	<b>IP-PINN for Time-Resolved 3D Blood Flow Velocity Reconstruction and Wall Shear Stress Calculation Using Data from Modified 2D PCMRI</b>	<b>79</b>
5.1	Introduction	79
5.2	Proposed Modified 2D PC-MRI Sequence	82
5.3	IP-PINN Adoption for Modified 2DPCMR Sequence	83
5.3.1	Overview of the 2D Multi-Directional Data	83
5.3.2	ROI Selection and Input Parameterization	84
5.4	Wall Shear Stress Calculation	85
5.5	Generation of Synthetic 2D PC-MRI Data	87
5.5.1	High-Resolution Reference Fields	87
5.5.2	Synthetic 2D PC-MRI Acquisition	88
5.6	Error Metrics	89
5.7	Results	89
5.7.1	Pre-Training the IP-PINN	89
5.7.2	Fine-Tuning on Unseen MCA Data	90
5.7.3	Temporal Super-Resolution Capability of IP-PINN	94
5.7.4	Luminal Boundary Prediction	95
5.7.5	Wall Shear Stress Calculation	96
5.8	Discussion	97
5.9	Conclusion	100
<b>6</b>	<b>Accelerating 4D Flow MRI with IP-PINN: A pseudo 1-point acquisition and k-space reconstruction approach</b>	<b>101</b>
6.1	Introduction	101

6.2	Proposed Pseudo 1-Point 4D Flow MRI Sequence . . . . .	103
6.3	IP-PINN Adoption for Pseudo 1-point 4D Flow MRI sequence . . . . .	106
6.3.1	ROI Selection, Input Tensor, and Input Parameterization . . . . .	107
6.3.2	Implementing $k$ -space Data Fidelity Loss Function . . . . .	108
6.4	Generation of Synthetic Training and Test Data . . . . .	109
6.4.1	Generating Synthetic Pseudo 1-point 4D Flow MRI Data . . . . .	110
6.5	Error Metrics . . . . .	111
6.6	Results . . . . .	112
6.6.1	Pre-training IP-PINN on Pseudo 1-Point k-Space Data . . . . .	112
6.6.2	Applying Pre-trained IP-PINN to Unseen Pseudo 1-Point Data . . . . .	112
6.6.3	Temporal Super-Resolution . . . . .	113
6.6.4	Luminal Boundary Prediction . . . . .	114
6.6.5	Performance Evaluation on Pseudo 1-Point Data with Velocity Aliasing . . . . .	114
6.7	Discussion . . . . .	115
6.8	Conclusion . . . . .	116
<b>7</b>	<b>Concluding Remarks</b>	<b>122</b>
7.1	Summary of Contributions . . . . .	122
7.2	Future Work . . . . .	125
	<b>Bibliography</b>	<b>127</b>

# LIST OF FIGURES

1.1	An overview of the wide-ranging techniques used to assess hemodynamics, spanning invasive procedures, advanced imaging, and sophisticated computational methods. . . . .	5
1.2	Flow velocity vector fields obtained by vector flow mapping. The flow velocity fields are superimposed on the color Doppler images [156]. . . . .	7
1.3	Flow streamlines of the left ventricle derived by Echo-PIV [156] . . . . .	9
1.4	4D-Flow MRI data acquisition for 3D velocity and magnitude images. [106] .	13
2.1	4D-Flow MRI acquisition. At cardiac phase $t = t_l$ , and for each coronal plane slice, a minimum of 4 acquisitions corresponding to a flow compensated scan ( $j = 0$ ), followed by flow encoded scans in $x, y, z$ directions ( $j = 1, 2, 3$ ) are acquired. The acquisition is spread over several cardiac cycles, with each cardiac cycle accounting for a small portion of the $\mathbf{k}$ -space. Here $\Delta t$ is the time span of the cardiac phase $t_l$ , i.e., the time for acquiring the data in $\mathbf{k}$ -space for that particular cardiac phase. A $\text{FFT}^{-1}$ of the $\mathbf{k}$ -space acquisition generates the complex $\mathbf{x}$ -space image. The magnitude (computed on a per voxel basis) of the $\mathbf{x}$ -space image is typically used in examining anatomical structures. The phase contains tissue velocity information. . . . .	36
2.2	Correction of Phase Offsets in 4D Flow MRI: (a, e) Acquired magnitude images; (b, c, f, g) Phase offset errors due to tissue motion; (d, h) Phase difference images obtained after subtracting consecutive phase images with toggled bipolar gradients, effectively canceling time-dependent artifacts while leaving magnetic field inhomogeneities and non-linear gradient fields uncorrected. [135]	41
2.3	Illustration of velocity aliasing in PC-MRI: The left panel shows flow images acquired with low $v_{enc}$ , resulting in aliasing artifacts. The right panel shows the corresponding high $v_{enc}$ acquisition, where aliasing is mitigated. The diagram on the far right visualizes the phase wrapping effect that leads to velocity aliasing when the actual flow exceeds the $v_{enc}$ threshold. [107] . . . .	42
3.1	IP-PINN usage: The user selects a spatio-temporal region of interest (ROI) in the flow region to process. The input to the trained IP-PINN is the image data consisting of the real and imaginary parts of the three flow-encoded scans in $\mathbf{x}$ -space and the spatio-temporal query point $(\mathbf{x}, z, t)$ . The IP-PINN predicts the underlying point-wise 3D velocity $(u_{pred}, v_{pred}, w_{pred})$ , relative pressure $p$ , and image magnitude $M_{pred}$ which can be segmented to obtain the luminal boundary. . . . .	54

3.2	Schematic representation of the proposed IP-PINN architecture. A section $\mathbf{U}$ of the three velocity encoded complex Caterisian images is processed by a convolutional neural network consisting of ResNet encoder blocks to generate a latent vector $\mathbf{L}$ . The secondary components (Multilayer Perceptron) take in the latent vector and local spatio-temporal coordinates to approximate a continuous function mapping, resulting in normalized flow velocity components, pressure, magnitude, and directional phase offsets. . . . .	56
4.1	Schematic representation of the proposed algorithm’s workflow. (a) Pre-training phase. (b) Fine-tuning phase. . . . .	62
4.2	Visualization of aneurysm geometries for model pre-training and testing. (a) Posterior cerebral artery (PCA) aneurysm geometry is used for generating pre-training dataset. (b) MCA aneurysm geometry is used for generating test dataset. The regions highlighted in red represent the volume from which flow velocities were used for training and testing. . . . .	63
4.3	Process of generating synthetic 4D-Flow MRI with noise and phase offset artifacts from CFD simulations. . . . .	65
4.4	Results showcase the impact of varying noise levels in the region of interest: (a) noiseless, median absolute percentage errors stand at $MDAPE_u = 3.747$ , $MDAPE_v = 1.222$ , and $MDAPE_w = 6.522$ ; (b) with an SNR of 20 dB, the corresponding median absolute percentage errors are $MDAPE_u = 6.178$ , $MDAPE_v = 2.267$ , and $MDAPE_w = 8.360$ . . . . .	69
4.5	Comparison of performance with noise-free input and noisy input with 20dB SNR. It is clear that there is good agreement with the ground truth in both scatter plots for noise-free and noisy input. The performance degrades slightly as indicated in the table below. The COSE plot show directional agreement with the ground truth (COSE $\approx 0$ ) . . . . .	70
4.6	Time interpolation of velocity. Here the synthetic 4D-Flow MRI low resolution velocities are available at $t + 2\Delta t$ , $t + 6\Delta t$ , $t + 10\Delta t$ , $t + 14\Delta t$ , $t + 18\Delta t$ , and $t + 22\Delta t$ where $\Delta t = 10ms$ . $u_{pred}, v_{pred}, w_{pred}$ are the predictions of the IP-PINN algorithm. $u_{GT}, v_{GT}, w_{GT}$ are the ground truth snapshots. . . . .	70
4.7	Luminal boundary prediction. The left most image is the low resolution magnitude image which is generated from the voxelized lumen boundary in the spatial resolution of the synthetic 4D-Flow MRI data set. This is typically the shapes seen while segmenting 4D-Flow MRI magnitude images. The center is the voxelization of the luminal boundary into the grid into which the ground truth CFD flow data was interpolated. The right most image is segmentation of the magnitude image prediction by the neural net. . . . .	71
4.8	Velocity fields predicted by IP-PINN models for input exhibiting velocity aliasing at SNR of 10 dB. White arrow indicate the velocity aliasing artifacts resulting from low $v_{enc}$ setting. . . . .	72

4.9	Comparison between IP-PINN, PINN, and DF-RBF. ‘LR’ is the low resolution velocity obtained by traditional processing of the synthetic 4D-Flow MRI scans (3 velocity encoded and 1 reference scan). ‘GT’ represents the ground truth which was used to generate the synthetic 4D-Flow MRI data. Subsequent columns represent the output of the IP-PINN, PINN, and DF-RBF methods with their absolute errors. The white arrows point to fine feature in the ground truth that are replicated faithfully by IP-PINN but not by PINN as well as DF-RBF. All velocities are in cm/s. The LR images are 8x8x8, where the voxel dimension is 1.05mm. All the super-resolved results are 24x24x24, where the voxel dimension is 0.35mm. . . . .	73
5.1	Comparison of conventional (a) and proposed modified (b) 2D phase-contrast MRI acquisition schemes. In (a), The red and blue lines represent reference and velocity-encoded scans, respectively, collected at different phases. Only a single velocity-encoding direction is captured in each cardiac phase. In (b), the red, blue, and green lines illustrate how velocity encodes in three perpendicular directions (e.g., $x$ , $y$ , and $z$ ) are alternated across cardiac phases. Here, multiple directional encodings are interleaved across phases. This interleaving enables the proposed IP-PINN framework to reconstruct the 3D-3C velocity fields in the vicinity of the 2D imaging plane, albeit at a reduced temporal sampling rate for each velocity direction. . . . .	81
5.2	Illustration of the limitations associated with assuming orthogonality of the vessel walls’ normal vectors to the 2D interrogation plane, neglecting for the geometric complexity and variability of vessel anatomy. . . . .	86
5.3	Results of processing velocity-encoded data at $20dB$ with IP-PINN. Each row of the input block corresponds to one velocity-encoding direction at a given cardiac phase (velocity encoded $u$ at $t_{l-1}$ , velocity encoded $v$ at $t_l$ , and velocity encoded $w$ at $t_{l+1}$ ). The IP-PINN framework generates all three components of the velocity fields at every cardiac phase. . . . .	92
5.4	IP-PINN performance comparison of for processing noise-free and noisy input with SNR of $20dB$ , illustrated through scatter and cosine distance plots. Both cases show strong agreement with the ground truth in terms of velocity magnitude (scatter plots) and flow direction ( $COSE \approx 0$ ). . . . .	93
5.5	Super-resolution of 3D velocity fields in time by the IP-PINN framework. The figure demonstrates the predicted ( $u_{pred}$ , $v_{pred}$ , and $w_{pred}$ ) and corresponding ground truth ( $u_{GT}$ , $v_{GT}$ , $w_{GT}$ ) velocity components across cardiac phases $t_{l-1}$ , $t_{l-1/2}$ , $t_l$ , $t_{l+1/2}$ , and $t_{l+1}$ . The input data includes velocity-encoded scans in the $x$ -direction at $t_{l-1}$ , $y$ -direction at $t_l$ , and $z$ -direction at $t_{l+1}$ . Each row represent corresponding velocity components for different cardiac phases. . .	95

5.6	Comparison of predicted low-resolution (left), predicted (center), and ground truth (right) lumen boundaries. The highlighted red area in the low-resolution lumen boundary indicates the spatial location (2D interrogation plane). In the prediction and ground truth figures highlighted red region shows the small region around the interrogation plane, while the gray regions illustrate out-of-plane volumetric boundaries. The IP-PINN framework recovers the high-resolution lumen structure near the plane but exhibits larger deviations farther from the slice, where no direct measurements are available. . . . .	96
5.7	Comparison of wall shear stress (WSS) magnitudes along the vessel boundary computed using the IP-PINN framework (black solid line), ground truth from CFD simulations (red dashed line), and the cubic B-spline interpolation method (blue dash-dot line) from [164]. The boundary length for the IP-PINN method is measured from the 3D faces of the marching cubes algorithm applied to the reconstructed lumen, whereas for the B-spline interpolation method, it is calculated from the contour points in the 2D imaging plane. . . . .	98
6.1	Comparison of conventional (a) and proposed pseudo 1-point (b) 4D Flow MRI acquisition schemes. In (a), at each cardiac phase, four separate 3D datasets are acquired—one reference scan plus three velocity-encoded acquisitions—to determine blood flow velocities along the $x$ , $y$ , and $z$ axes. In (b), The bipolar velocity-encoding gradients alternate across successive slices in the $z$ -direction, such that only one velocity-encoded dataset is acquired per slice. In this way, the different encoding directions are interleaved slice-by-slice, reducing the spatial sampling rate for each velocity component. . . . .	104
6.2	Super-resolution results on unseen pseudo 1-point MCA data with noise-free input. . . . .	118
6.3	Super-resolution results on unseen pseudo 1-point MCA data with 20 dB SNR noise. . . . .	119
6.4	Comparison of scatter and cosine-error distributions for noise-free and 20 dB SNR inputs. Close alignment to the $45^\circ$ line and low cosine-error values indicate accurate magnitude and directional recovery. . . . .	120
6.5	Temporal super-resolution on pseudo 1-point MCA data. Low-resolution velocity-encoded data are available at $t_0 + 2\Delta t$ , $t_0 + 6\Delta t$ , and $t_0 + 10\Delta t$ (with $\Delta t = 10$ ms). Columns $u_{\text{pred}}, v_{\text{pred}}, w_{\text{pred}}$ show the IP-PINN predictions at 10 ms intervals, while $u_{\text{GT}}, v_{\text{GT}}, w_{\text{GT}}$ are the CFD ground-truth snapshots. . . . .	120
6.6	Luminal boundary prediction from pseudo 1-point acquisition data. <b>Left:</b> low-resolution magnitude image generated from the voxelized lumen geometry at the synthetic 4D-Flow MRI resolution. <b>Center:</b> IP-PINN-predicted segmentation of the magnitude image. <b>Right:</b> ground-truth lumen boundary voxelized on the CFD grid used for interpolation. . . . .	121
6.7	Velocity field reconstruction using IP-PINN on pseudo 1-point data with aliasing artifacts and 20dB SNR noise. Red arrow highlights the region affected by velocity aliasing under low $v_{\text{enc}}$ . . . . .	121

## LIST OF TABLES

4.1	Parameters used for simulating blood flow in PCA and MCA arteries to generate synthetic velocity encoded images. Here the blood dynamic viscosity $\mu = 0.004$ Kg/ms and density $\rho = 1000$ Kg/m <sup>3</sup> . . . . .	63
4.2	The loss terms of pre-trained IP-PINN for the unsteady posterior cerebral artery aneurysm (PCA) case and their empirically chosen weights. . . . .	67
4.3	A summary of total relative percent error ( <i>TRPE</i> ), coefficient of determination ( $R^2$ ), and median absolute percentage error of $u^*$ , $v^*$ , and $w^*$ for processing the noiseless input data and input data with SNR of 20 dB. IP-PINN required 1.5 minutes for processing the input data . . . . .	68
4.4	Summary of performance metrics for the proposed IP-PINN model when processing noisy input data with velocity aliasing artifacts at signal-to-noise ratios (SNR) of 10 dB. Metrics include total relative percent error ( <i>TRPE</i> ), coefficient of determination ( $R^2$ ), and median absolute percentage error of $u_{pred}$ ( <i>MDAPE<sub>u</sub></i> ), $v_{pred}$ ( <i>MDAPE<sub>v</sub></i> ), $w_{pred}$ ( <i>MDAPE<sub>w</sub></i> ) . . . . .	72
4.5	A comparison of total relative percent error ( <i>TRPE</i> ), coefficient of determination ( $R^2$ ), and median absolute percentage error of $u^*$ , $v^*$ , and $w^*$ achieved by the PINN models and the proposed IP-PINN model for processing noise-free input data. IP-PINN required 1.5 minutes for processing the same input data, while PINN required 30 minutes . . . . .	75
5.1	Sample loss values and weighting factors for the pre-trained IP-PINN, focused on unsteady blood flow in a PCA aneurysm. . . . .	90
5.2	A summary of total relative percent error ( <i>TRPE</i> ), coefficient of determination ( $R^2$ ), and median absolute percentage error ( <i>MDAPE</i> ) for processing the noiseless input data and input data with SNR of 20 dB. IP-PINN required 1.5 minutes for processing the input data . . . . .	93
6.1	Pre-training loss terms and weights for the pseudo 1-point k-space PCA dataset. 112	
6.2	Quantitative performance on unseen pseudo 1-point MCA data: <i>TRPE</i> , $R^2$ , <i>MDAPE</i> for velocity magnitude, and fine-tuning time. . . . .	113
6.3	Performance metrics of IP-PINN on pseudo 1-point input data with 20 dB SNR and velocity aliasing. Metrics include <i>TRPE</i> , $R^2$ , and <i>MDAPE</i> for velocity magnitude. . . . .	115

# Chapter 1

## Introduction

### 1.1 Motivation

Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, accounting for approximately 17.9 million deaths annually, according to the World Health Organization (WHO) [124]. The complexity and progressive nature of these diseases necessitate early and accurate diagnosis to improve patient outcomes. A critical factor in diagnosing, treating, and preventing CVDs is understanding the underlying hemodynamics, which refers to the dynamics of blood flow [155, 65]. Hemodynamics significantly influences the development, progression, and management of conditions such as atherosclerosis, aneurysms, and heart valve disorders [16, 45, 111, 157, 29, 40]. However, accurately assessing hemodynamics in clinical practice remains challenging due to the complex and patient-specific nature of blood flow patterns.

Recent advancements in imaging technologies and computational methods have enabled more detailed investigations into cardiovascular hemodynamics. Techniques such as 4D Flow MRI [104] provide non-invasive ways to visualize and quantify blood flow in three dimensions over time, offering critical insights into complex flow patterns that are difficult to capture with traditional methods. Despite these advancements, several limitations persist, including low spatio-temporal resolution, long acquisition times, acquisition noise, and various image artifacts, all of which hinder the clinical applicability of these techniques. This thesis aimed at overcoming these limitations, with the ultimate goal of enhancing the accuracy and reliability of hemodynamic assessments, thereby improving the diagnosis and treatment of cardiovascular diseases.

## 1.2 The Role of Hemodynamics in Cardiovascular Diseases

Hemodynamics plays a pivotal role in the pathophysiology of cardiovascular diseases. The interaction between blood flow, vascular geometry, and endothelial function significantly impacts the development, progression, and outcomes of various cardiovascular conditions. Aberrations in hemodynamic forces, such as wall shear stress (WSS), pressure gradients, and vascular resistance [164, 39], are closely linked to the pathogenesis of diseases like atherosclerosis, aneurysms, and heart valve disorders.

Wall shear stress, the tangential force exerted by blood flow on the endothelial surface, is particularly crucial in vascular health. Regions of low WSS are often associated with the development of atherosclerotic plaques, as they promote endothelial dysfunction, inflammation, and lipid accumulation [16, 111]. Conversely, high WSS is implicated in the formation and rupture of aneurysms, where elevated mechanical forces lead to vascular remodeling and wall weakening [157, 29]. Understanding these hemodynamic factors is essential for identifying high-risk areas within the vasculature and improving the prediction and prevention of adverse cardiovascular events.

Moreover, structural variations such as heart valve defects, chamber geometry, and wall motion abnormalities can profoundly influence the hemodynamic response. These anatomical changes not only impact blood flow dynamics but also contribute to the vascular remodeling process [138]. The intricate relationship between blood flow alterations and vascular changes underscores the relevance of hemodynamics as a critical biomarker in clinical practice. Hemodynamic parameters can guide patient diagnosis, monitor disease progression, and inform surgical planning. Additionally, hemodynamic analysis is invaluable in evaluating and optimizing the design of medical devices that interact with blood, ensuring their efficacy and safety [185].

Abnormalities in cerebral hemodynamics can lead to severe conditions such as cerebral

aneurysms, underscoring the importance of accurate hemodynamic assessment [119]. In the context of cerebral aneurysms, understanding hemodynamics is crucial, as the stresses exerted by blood flow, including wall shear stress (WSS), significantly influence aneurysm initiation, evolution, and stability. These hemodynamic factors not only impact the structural integrity of the vessel wall but also play a key role in determining the success of surgical interventions such as endovascular flow diverter placement, which aim to redirect flow and promote thrombosis within the aneurysm sac. Accurate pre- and post-operative hemodynamic assessments are therefore essential for guiding the proper sizing and positioning of flow diverters, ensuring optimal treatment outcomes. While the precise role of WSS in aneurysm formation remains a subject of ongoing research, it is evident that WSS, along with other markers like Oscillatory Shear Index and Spatial WSS gradient, provides valuable insights into aneurysm progression and the identification of high-risk areas [94, 111, 21]. The direct involvement of hemodynamics in vascular remodeling and its impact on blood flow further underscores its significance as a biomarker in both clinical and research settings.

### 1.3 Assessment Techniques for Hemodynamics

Advancements in cardiovascular imaging and computational technologies have significantly expanded our ability to assess hemodynamics non-invasively. Among the key non-invasive imaging techniques are Doppler echocardiography [188, 36, 110, 176], Echo Particle Image Velocimetry (Echo-PIV) [84, 38, 153, 83], and time-resolved phase-contrast magnetic resonance imaging (PC-MRI) [113, 165, 163, 5, 116]. Additionally, Computational Fluid Dynamics (CFD) offers a powerful computational tool for simulating and analyzing blood flow within the cardiovascular system, providing detailed insights that are often difficult to obtain through direct measurements [122, 64, 109]. When combined with imaging modalities like PC-MRI, CFD allows for the creation of patient-specific models, enhancing the accuracy and applicability of hemodynamic assessments [10, 148, 114]. Furthermore, deep learning and machine learning methods are increasingly being applied to hemodynamic data,

enabling the automatic extraction of complex patterns and relationships within the data [86, 55, 54, 50, 52]. These advanced computational techniques can enhance the interpretation of hemodynamic assessments, provide predictive analytics, and support personalized treatment planning. Below are detailed explanations of hemodynamics assessment techniques along with their limitations.

## 1.4 In-vivo Blood Flow Measurements

In-vivo blood flow measurement techniques are essential for understanding the dynamics of blood flow within living organisms, offering critical insights into physiological and pathological processes in real time. Unlike in-vitro methods [96, 72, 152, 77], which assess blood flow in controlled or artificial settings, in-vivo methods capture the complexities of blood flow within the intact cardiovascular system, accounting for variables such as vessel geometry, blood viscosity, and interactions with surrounding tissues. Over time, various techniques have been developed to measure in-vivo blood flow, ranging from traditional invasive methods to advanced non-invasive imaging technologies. While invasive methods provide direct and precise measurements, they come with significant risks and are less suitable for routine use. In contrast, non-invasive techniques have gained prominence for their ability to assess hemodynamics without requiring surgical intervention, making them more practical in clinical settings.

This section explores a range of in-vivo blood flow measurement techniques, beginning with traditional invasive methods and advancing to non-invasive imaging technologies. Techniques covered include Doppler echocardiography [6, 188, 115], Echo Particle Image Velocimetry (Echo-PIV) [83, 38], Time-Resolved Phase-Contrast MRI (PC-MRI) [104, 138], and 4D Flow MRI [105]. Fig. 1.1 presents a categorization of hemodynamic assessment techniques, illustrating the spectrum of approaches available, from in-vivo methods to computational modeling. A comprehensive understanding of these techniques, along with their strengths and limitations, is crucial for addressing current challenges in cardiovascular diag-

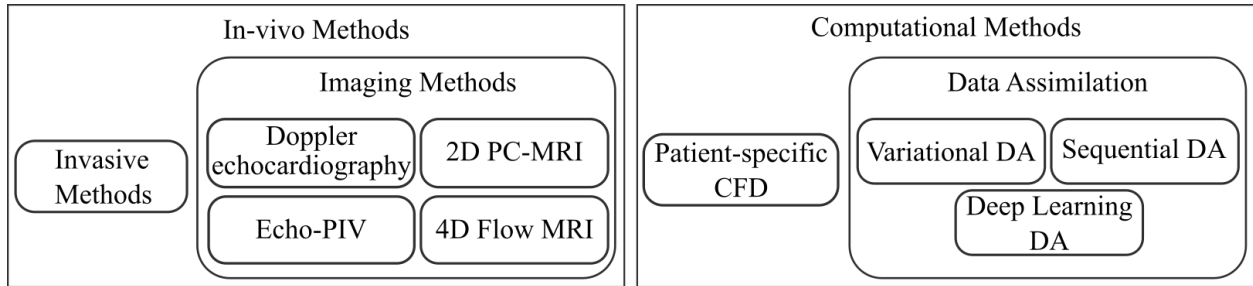


Figure 1.1: An overview of the wide-ranging techniques used to assess hemodynamics, spanning invasive procedures, advanced imaging, and sophisticated computational methods.

nostics and recognizing opportunities for innovation, particularly through the integration of deep learning algorithms to enhance imaging resolution and diagnostic accuracy.

### 1.4.1 Invasive Methods

Invasive methods, such as catheter-based techniques, have traditionally been used to measure blood flow directly within vessels [102]. These approaches typically involve the insertion of a catheter into the cardiovascular system to obtain measurements of flow velocity and pressure. However, these methods are associated with significant risks, including infection, vessel damage, hemiplegia, and transient ischemic attack [3]. Although they offer good temporal resolution, their spatial resolution and measurement localization are often inadequate. Additionally, these techniques generally provide only sparse spatial measurements of velocity magnitude and pressure, making it challenging to compute higher-order metrics such as pressure and velocity gradients. These limitations have spurred the development of non-invasive techniques that aim to deliver similar insights with reduced risk.

### 1.4.2 Doppler Echocardiography

Doppler echocardiography is a widely used non-invasive technique that combines ultrasound imaging with Doppler effect principles to assess blood flow [117]. The Doppler effect, which underpins this technique, involves the change in frequency of sound waves as they reflect off moving red blood cells. By calculating the frequency shift ( $\Delta f$ ) between the transmitted fre-

quency ( $f_t$ ) and the received frequency ( $f_r$ ), the velocity of blood flow ( $v$ ) can be determined using the Doppler equation [6]:

$$v = \frac{c \cdot \Delta f}{2 \cdot f_t \cdot \cos(\theta)} \quad (1.1)$$

, where  $c$  is the speed of sound in blood and  $\theta$  is the angle between the ultrasound beam and the direction of blood flow.

Doppler echocardiography encompasses two distinct modalities: continuous wave Doppler and pulsed wave Doppler. In continuous wave Doppler, two piezoelectric crystals are employed: one continuously transmits ultrasound waves towards the moving red blood cells, while the other continuously receives a wide spectrum of the reflected wave frequencies from blood motion. This continuous measurement of Doppler shifts resulting from blood flow generates sound signals from the returning echoes, enabling observation of flow along the ultrasound beam. However, the absence of echo receive time measurement precludes this technique from determining velocity range resolution, flow direction, and the depth from which the echo signals originate [117, 6].

In pulsed wave Doppler, a single piezoelectric crystal functions as both transmitter and receiver of sound beams. Ultrasound signals are emitted in short bursts or pulses, and the echoes reflected from moving red blood cells are received during the intervals between transmitted pulses. This enables the calculation of distances within the tissue based on the echo receive time. Through range gating, pulsed wave Doppler can selectively acquire Doppler information from a specific location within the cardiovascular region of interest using a sample volume. pulsed wave Doppler examinations are frequently combined with simultaneous two-dimensional echocardiography, allowing for visual display of the precise sample volume location [117].

Building upon the principles of pulsed wave Doppler, Doppler color flow imaging visually represent blood flow using a color map based on velocity and direction. This is achieved by analyzing Doppler shifts at numerous points along multiple ultrasound beams. The




Figure 1.2: Flow velocity vector fields obtained by vector flow mapping. The flow velocity fields are superimposed on the color Doppler images [156].

measured shifts are then converted into a digital format and assigned colors according to a predetermined scheme, creating a color-coded overlay on the 2D echocardiogram. Blood moving towards the transducer is depicted in shades of red, signifying a positive Doppler shift, while blood flowing away appears in shades of blue, indicating a negative shift. The intensity of the color corresponds to the velocity of the blood flow, with lighter shades representing higher speeds [6]. Fig. 1.2 shows a example of flow velocity fields, achieved by combining measured axial velocities with the estimated radial velocities based on physical principles, superimposed on the color Doppler images [156].

Spectral Doppler offers a detailed spectrum or waveform that displays the range of blood flow velocities at a specific location. This is achieved by placing a Doppler gate, which defines a specific region within the sample volume where the Doppler shift is measured. The Doppler shift signals reflected from blood cells within the Doppler gate are processed using Fast Fourier Transform (FFT) to create a spectral display—a graphical representation of blood flow velocity over time. In this display, the vertical axis represents velocity, and the

horizontal axis represents time, allowing to assess both the velocity and direction of blood flow at the sampled site. The accuracy of the spectral Doppler measurement depends on the alignment of the Doppler angle, which is the angle between the ultrasound beam and the direction of blood flow [117, 6].

While Doppler echocardiography is primarily used for cardiac and vascular assessments, similar Doppler principles are applied in other clinical settings. One such specialized application is transcranial Doppler (TCD), which uses ultrasound to measure blood flow velocity in the major intracranial arteries. Using low-frequency (usually 2 MHz or less) ultrasound waves, TCD assesses the velocity of blood flow through basal cerebral arteries with a high temporal resolution [115].

Despite its widespread utility, Doppler echocardiography has inherent limitations. The accuracy of Doppler measurements is limited by various aspects such as the technical skill of the operator, probe orientation, and the angle of insonation; optimal accuracy is achieved when the ultrasound beam is aligned parallel to the direction of blood flow. Angles greater than 20 degrees can result in significant underestimation of velocities [70, 130]. Turbulent or non-laminar flow, which is common in conditions like severe valve regurgitation or stenosis, can complicate the spectral Doppler signal and hinder precise measurements. While, 2D velocity maps can be reconstructed using techniques like vector Doppler [121], traditional Doppler echocardiography techniques are limited to measuring only the 1D velocity component along the ultrasonic beam. Moreover, while Bernoulli's principle allows for the estimation of the pressure field, computing other hemodynamic quantities using traditional Doppler echocardiography is challenging [138]. Additionally, the maximum Doppler shift measured by Pulse Wave Doppler is constrained by the Nyquist limit, which can lead to aliasing for frequency shifts higher than this limit [6]. Thus, while Doppler ultrasound is an invaluable tool for hemodynamic assessment, these limitations necessitate careful interpretation and, in some cases, complementary invasive measurements for accurate diagnosis and management.

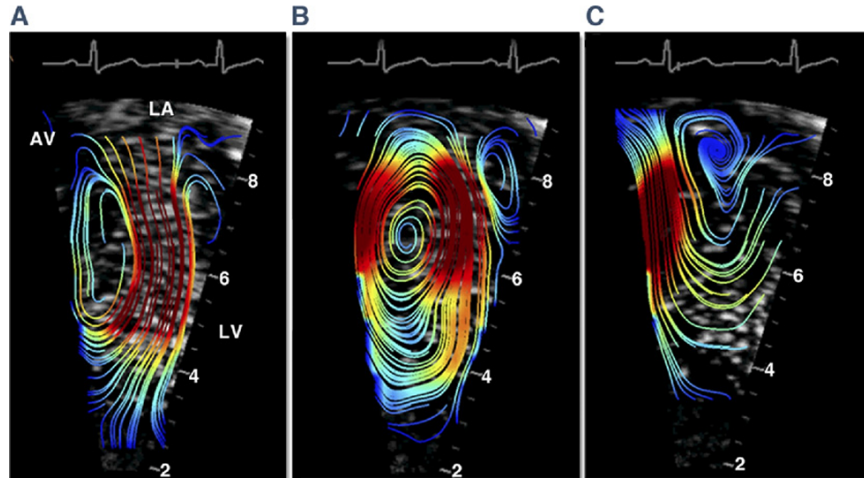


Figure 1.3: Flow streamlines of the left ventricle derived by Echo-PIV [156]

### 1.4.3 Ultrasound Imaging Velocimetry (Echo-PIV)

Ultrasound Imaging Velocimetry (UIV), commonly known as echo-PIV, is a whole-field velocity measurement technique based on ultrasound imaging, specifically designed to address the limitations of traditional optical flow measurement methods, such as Particle Image Velocimetry (PIV) in opaque fluids. In ultrasound imaging images are created by sending out sound waves and analyzing the echoes that return. These echoes are generated when the sound waves encounter variations in the acoustic impedance of the materials they pass through. Since ultrasound can pass through materials that are not transparent to light, it allows for measurements in flows where optical methods are not feasible due to limited visibility. Echo-PIV is particularly valuable in applications like *in vivo* cardiovascular flow measurements [84, 133, 134, 166, 1]. Fig. 1.3 show the flow streamlines along the long-axis view of the left ventricle [156].

The basic working principle involves generating ultrasound pulses using piezoelectric transducers, which then propagate through the fluid medium. These pulses interact with scatterers within the fluid, such as microbubbles or particles, causing the ultrasound waves to reflect back to the transducers. The received signals are processed to form two-dimensional gray-scale brightness-mode (B-mode) images that displays the intensity of echoes received from tissue, where the brightness of each pixel corresponds to the strength of the echo.

In the B-mode images, the speckle patterns generated by the scatterers are tracked across consecutive frames. These speckle patterns are processed using correlation algorithms similar to those in optical PIV to extract velocity data. The correlation algorithms identify the most likely displacement between patterns in successive frames, enabling the calculation of velocity vectors [112, 133, 167, 18, 84].

Despite its advantages, echo-PIV has several limitations that must be considered. Echo-PIV is capable of measuring two-dimensional, two-component velocity fields (in-plane components) within the ultrasound imaging plane. However, it does not directly measure the out-of-plane velocity component. To achieve full three-dimensional velocity measurements, techniques such as multiple view angles or advanced three-dimensional ultrasound imaging modalities would be necessary, but these add complexity and may affect the accuracy and resolution. The spatial resolution of echo-PIV is generally lower than that of optical PIV due to the constraints of probe specifications such as ultrasound frequency, wavelength of the incident echo beam, the number of cycles in the excitation pulse, beam width, number of elements in the transducer, imaging frequency, and focal depth. In high scatterer concentration conditions, speckle patterns can decorrelate quickly due to flow gradients, which increases noise and limits the measurable velocity range. Additionally, the technique is constrained by the maximum detectable velocity, which is influenced by the frame rate of the ultrasound system and the rate of speckle decorrelation [133, 118, 84, 93].

#### **1.4.4 Time-Resolved Phase-Contrast MRI**

Phase-Contrast Magnetic Resonance Imaging (PC-MRI) is a powerful non-invasive imaging modality used to measure and visualize hemodynamics by capturing the velocity of moving fluids, such as blood and cerebrospinal fluid (CSF). Unlike conventional MRI, which primarily relies on signal amplitude, PC-MRI distinguishes itself by leveraging the sensitivity of the MR signal phase to motion [105].

In PC-MRI, velocity encoding is achieved by applying bipolar gradients along the di-

rection of flow. These gradients induce a phase shift in the MRI signal that is directly proportional to the velocity of moving spins within the imaging slice [105, 165]. The phase shift ( $\Delta\phi$ ) is mathematically described by the equation [186]:

$$\Delta\phi = \gamma \cdot g \cdot u \quad (1.2)$$

where  $\gamma$  is the gyromagnetic ratio,  $g$  represents the first moments of the gradient, and  $u$  is the proton/spin velocity.

The technique utilizes bipolar velocity-encoding gradients to acquire two images with different velocity-dependent signal phases while keeping other sequence parameters identical. By subtracting these two phase images, a phase difference image is obtained, from which velocity maps can be derived [116, 186].

2D PC-MRI is the simplest phase-contrast method and is widely used in monitoring volumetric blood flow in large blood vessels like the aorta. In a standard 2D PC-MRI clinical procedure, the imaging slice is aligned perpendicular to the center line of the vessel lumen. The data acquisition generally involves measuring a single velocity component in the through-plane direction and is carried out during a breath-hold lasting 10 to 20 seconds [165, 19].

One significant limitation of 2D PC-MRI is its sensitivity to misalignment between the imaging plane and the direction of blood flow, which can lead to inaccuracies in flow measurement. Additionally, 2D PC-MRI provides flow information in only a single plane, potentially overlooking complex three-dimensional flow patterns and interactions within the vessel lumen. Furthermore, for patients with certain cardiovascular conditions, such as congestive heart failure or shortness of breath, the required breath-hold duration may not be feasible [19].

### 1.4.5 4D Flow MRI

4D Flow MRI is an advanced extension of phase-contrast MRI (PC-MRI) that significantly enhances the ability to visualize and quantify cardiovascular hemodynamics. Unlike traditional 2D PC-MRI, which provides velocity information in only a single plane and direction, 4D Flow MRI captures three-directional velocity-encoded images across a three-dimensional volume over time. This capability allows for detailed, retrospective analysis of blood flow dynamics within entire vascular territories, such as the heart, aorta, and cerebral arteries. It provides comprehensive insights into cardiovascular biomarkers, such as relative pressure fields [45] and wall shear stress [136], thereby enhancing the understanding of cardiovascular health and diseases.

One of the major advantages of 4D Flow MRI is its ability to visualize and quantify hemodynamic parameters throughout the entire acquired volume, without the need for precise pre-acquisition plane placement. This feature contrasts with 2D PC-MRI, where the accuracy of flow measurements is heavily dependent on the correct alignment of the imaging plane with the flow direction [113, 57]. Studies have shown that 4D Flow MRI can provide more accurate and reproducible measurements of peak velocity, peak flow, and net flow volume compared to 2D PC-MRI, especially in regions with complex flow, such as the aorta [19, 23, 27, 46, 59, 164].

The data acquisition process involves the application of bipolar velocity-encoding gradients in three orthogonal directions, resulting in phase shifts proportional to the velocity components of blood flow. These phase shifts are combined to reconstruct a comprehensive map of blood flow velocities at each voxel within the imaging volume. Fig. 1.4 illustrates the data acquisition process utilized in the 4D-Flow MRI technique. Due to the sheer amount of volumetric data that has to be acquired for a given cardiac phase, the acquisition is spread over multiple cardiac cycles using ECG gating, providing time-resolved imaging that captures the dynamic patterns of blood flow. This approach allows 4D Flow MRI to deliver a four-dimensional view of cardiovascular hemodynamics, offering a significant advantage over

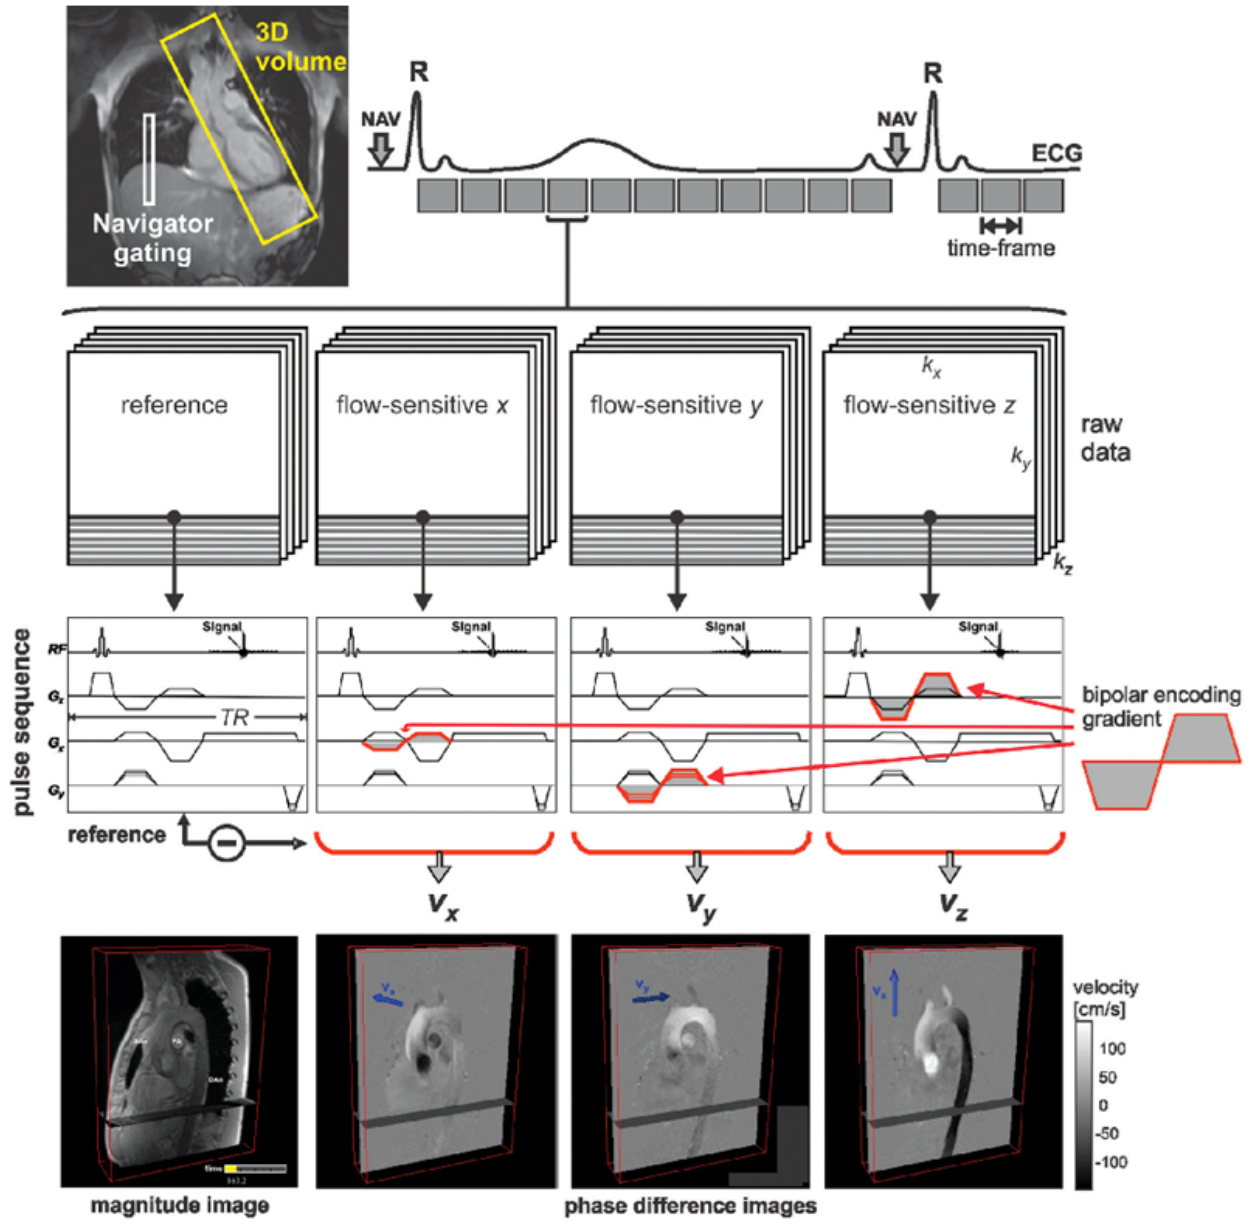


Figure 1.4: 4D-Flow MRI data acquisition for 3D velocity and magnitude images. [106]

traditional imaging methods [106, 30].

While 4D Flow MRI offers numerous benefits, it is not without limitations, including low spatio-temporal resolution, acquisition artifacts such as phase offset errors, and the need for prolonged scan times. These factors can impact the accuracy of flow measurements, particularly near arterial walls and in small vessels. Moreover, the method’s dependency on the velocity encoding parameter ( $v_{enc}$ ) requires careful calibration to avoid velocity aliasing, a common artifact that occurs when actual flow velocities exceed the set  $v_{enc}$  threshold [138, 113, 107, 135, 88].

## 1.5 Computational Methods for Hemodynamics Assessment

While in-vivo measurements such as PC-MRI and 4D Flow MRI provide valuable insights into cardiovascular hemodynamics, they are often limited by issues such as low spatial and temporal resolution, noise, and artifacts, particularly in complex vascular structures [6, 105, 131, 165]. To overcome these challenges and provide a more detailed analysis of blood flow, numerical methods have become essential tools in hemodynamic research. These methods, including Computational Fluid Dynamics (CFD) and various data assimilation techniques, complement in-vivo measurements by simulating and refining blood flow data.

CFD provides high-resolution simulations of blood flow within patient-specific anatomical models, offering detailed insights into flow patterns and hemodynamic parameters that are often difficult to capture with imaging alone [28, 122, 177]. However, the accuracy of CFD is contingent upon accurate boundary conditions and geometric models, which can introduce uncertainties [187]. To enhance the predictive accuracy of these simulations, data assimilation techniques are employed [44, 89]. These methods integrate CFD simulations with experimental or clinical data, refining the models and reducing discrepancies between simulated and observed data. This section explores the key numerical methods used in hemo-

dynamics assessment, including CFD and data assimilation techniques such as variational data assimilation methods, sequential data assimilation methods, and deep learning-based approaches, which collectively enhance our understanding and assessment of cardiovascular function.

### 1.5.1 Patient-Specific Computational Fluid Dynamics (CFD)

Patient-specific CFD is an alternative method to estimate high Spatio-temporal resolution 3D time-resolved velocity flow patterns and ultimately other hemodynamics parameters, in various vascular geometries, by solving the Navier-Stokes equations under particular assumptions on flow models [109, 146, 20, 177, 28]. Typically, vascular geometry is derived from segmented scans using angiography techniques like Computed Tomography Angiography (CTA) or Contrast-Enhanced Magnetic Resonance Angiography (CE-MRA), which requires contrast agent injection. Alternatively, Time of Flight (TOF-MRA) can be used, though it introduces potential errors like in-plane saturation and shine-through artifacts. [132, 147, 64, 9] These geometries are then discretized (meshed) into finite elements or volumes for CFD analysis.

Patient-specific inlet and outlet flow conditions for CFD are derived from 2D PC-MRI or 4D-Flow MRI data by positioning planes perpendicular to the blood vessel centerline at inlets and outlets, and integrating velocity data across these areas. However, this method inherits the limitations of 4D-Flow MRI, such as noise and suboptimal temporal resolution. Furthermore, the model typically assumes rigid vessel walls and a no-slip condition at these surfaces for the simulations [180, 80].

CFD simulations, by iteratively solving the Navier-Stokes equations for each mesh element under specific flow assumptions, generate high-resolution 3D time-resolved velocity flow patterns and other hemodynamic parameters. [109, 146, 20, 177, 28]. The general Navier-Stokes equations [15], are adapted for incompressible Newtonian fluids with constant viscosity and density, and within rigid vessel walls. The modified equation, represented in

Eq. 1.3, is as follows:

$$\rho \left( \frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V} \right) = -\nabla \mathbf{p} + \mu \nabla^2 \mathbf{V} \quad (1.3)$$

here  $\rho$  is the fluid density,  $\mathbf{V} = (u, v, w)$  represents the velocity vector,  $t$  is time,  $\nabla$  is the gradient operator,  $\nabla \mathbf{p}$  denotes the pressure gradient term,  $\mu$  is the dynamics viscosity, and  $\mu \nabla^2 \mathbf{V}$  is the diffusion term.

Additionally, for incompressible fluid flow, the continuity equation is a fundamental requirement, ensuring that the flow of the fluid is consistent with the principle of mass conservation. This equation, succinctly stated as [15]:

$$\nabla \cdot \mathbf{V} = 0 \quad (1.4)$$

CFD simulations allow for generating noiseless and high resolution fluid flow within patient-specific anatomical models. In such a method, the accuracy of the produced velocity fields is heavily affected by various factors, such as blood rheological properties assumptions, the assumptions on vessel properties and interactions between blood and vessel walls, and the segmentation accuracy of vascular geometries driven from imaging modalities. For in vivo vascular flows, such assumptions are not entirely accurate. As a result, despite the potential of the utility of CFD in generating aesthetic hemodynamics data, application of the pure CFD in vascular disease managements and endovascular treatments could be problematic due to its inherent limitations [78].

## 1.5.2 Data Assimilation

Data assimilation is a mathematical approach used to enhance the accuracy and reliability of numerical simulations by integrating experimental or observational data into computational models. The primary objective of data assimilation is to reduce the uncertainty inherent in computational models by continuously updating the model with real-world data [81, 120].

In the context of hemodynamics assessment, data assimilation techniques combine in-vivo measurements, such as 4D Flow MRI, with computational models like CFD to create a more accurate representation of blood flow dynamics within the cardiovascular system.

Two main components of a data assimilation problem include a dynamic system and its observations. The dynamic system is represented by a state vector,  $X_t$ , which evolves over time based on a mathematical model (e.g. the Navier-Stokes equations with velocity state variables in the case of cardiovascular modeling). The second component involves noisy observations of the system's states, such as velocity data from 4D Flow MRI, which can be described as [62]:

$$Y_t^{\text{data}} = \mathcal{H}(X_t) + \epsilon_t \quad (1.5)$$

Here  $\epsilon_t$  represents the observations errors, and  $\mathcal{H}$  is the observation operator that captures how the true system states  $X_t$  relates to the the observed data, accounting for the influence of measurement dynamics. The goal of data assimilation is to optimize a cost function that minimizes the discrepancy between the model predictions and the observed data [120].

Data assimilation methods have shown significant potential in improving the spatio-temporal resolution of hemodynamic assessments, enhancing the accuracy of pressure and velocity fields, and facilitating patient-specific modeling. Following subsections delve into the main data assimilation techniques used in hemodynamics, including variational methods, sequential methods, and deep learning-based approaches and references to recent applications on hemodynamics assessment.

## Variational Data Assimilation Methods

Variational data assimilation techniques optimize the mathematical models of the cardiovascular system with observation data, such as 4D Flow MRI blood flow measurements, using adjoint flow equations [120, 172]. These techniques rely on solving a minimization problem aimed at reducing a cost function that quantifies the discrepancy between the model's

predictions and the observational data, along with contributions from regularization term. [22]. In hemodynamics, variational data assimilation is particularly useful for estimating uncertain parameters, reconstructing unmeasured states, and correcting initial conditions, thereby enhancing the predictive power of simulations [60, 44]. Like CFD, they rely on specified geometry and mesh alignment. Recent advancements reconstruct both velocity field and geometry [90] using immersed boundary adjoint formulation on a rectangular grid. This method requires an initial guess of the geometry with delineated inflow/outflow and no-slip regions but is currently limited to steady-state flows.

Recent studies have demonstrated the effectiveness of variational data assimilation methods in enhancing hemodynamic simulations through the integration of clinical imaging data. In their study, Rispoli et al. [148] presented a variational data assimilation approach that integrates patient-specific Computational Fluid Dynamics (CFD) simulations with 4D Flow MRI data to enhance the accuracy of hemodynamic predictions. The authors utilized a modified Semi-Implicit Method for Pressure Linked Equations (SIMPLER) algorithm [126] combined with Tikhonov regularization [171] to align CFD-generated velocity fields with 4D Flow MRI observations. This approach was applied to a model of an aneurysm, by employing boundary conditions on the CFD simulation that derived from 4D-Flow MRI data and assuming Newtonian rheological properties and rigid vessel walls. They also used a weighting strategy to adjust the relative contribution of the 4D-Flow MRI observations and the CFD simulation for the generated data to satisfy the Navier-Stokes equations. The study highlighted the potential of variational data assimilation to refine hemodynamic assessments by incorporating real-world measurements. However, the approach's reliance on uniform grid meshes identical to the MRI grids which, in contrast with refined mesh in the proximity of the vessel walls, is not a suitable approach for the cardiovascular problems [7].

Bakhshinejad et al. [10] developed a method to merge patient specific CFD and 4D-Flow MRI, using proper orthogonal decomposition (POD) [82] and ridge regression [49]. Snap shots from an approximate patient specific CFD simulation using vascular geometry

obtained from CTA and boundary conditions from 4D-Flow MRI and used to generate high resolution POD basis. Low resolution POD basis is obtained by volume averaging. The POD coefficients in low resolution one obtained by projecting the 4D-Flow MRI data onto the low resolution POD basis. Next, high resolution coefficients are obtained from low resolution coefficients using ridge regression. High resolution coefficients along with the high resolution POD basis are then used to generate noise and artifact-free super resolved 4D-Flow MRI.

Fathi et al. [51] applied Least Absolute Shrinkage and Selection Operator (LASSO) [169] regression to augment patient specific CFD with 4D-Flow MRI data. An 'approximate' patient specific CFD simulation is done using boundary flow conditions from 4D-Flow MRI and geometry from CTA. A proper orthogonal basis is generated from the CFD snapshots. The true solution is sparse in this POD basis. Consequently, a LASSO regression approach is used to find the sparse coefficients. Here the data fidelity term minimizes the distance between predicted low-resolution velocities and the 4D-Flow MRI data. Both these approaches are limited by the need for a CFD simulation, requiring segmented vascular geometry and estimating boundary condition and also the need for localization of 4D-Flow MRI data in the CFD mesh. While they can super resolve spatially, the temporal resolution is limited to that of 4D-Flow MRI.

Funke et al. [61] merged CFD simulation and 4D-Flow MRI data by formulating the data assimilation problem as an optimization problem of discretized Navier-Stokes equations and observations for optimizing unknown model parameters like initial and boundary conditions. Authors proposed using an adjoint-based gradient computation for CFD solver and linearly interpolating the 4D-Flow MRI observations onto the CFD mesh nodes. The method leverages the transient Navier-Stokes equations and regularization techniques to ensure robustness against noise, demonstrating accurate reconstruction of blood flow even with sparse and noisy measurements. However, due to high computation expense of the full-fledged adjoint simulation, an in-house low-order CFD solver is employed, which limits the efficiency and results in reduced spatio-temporal resolution of the numerical model.

When compared to a single CFD simulation run, the proposed linear and adjoint approach is 50-100 times more expensive.

The research conducted by Toger et al. [173] proposed a variational data assimilation framework called 4D-flow HIKING, which optimizes high-resolution computational fluid dynamics (CFD) simulations to match 4D-flow MRI data for improved hemodynamic assessments. The approach uses an adjoint-based optimization framework that incorporates the Navier-Stokes equations, aligning CFD simulations with measured 4D-flow data by adjusting boundary conditions to minimize discrepancies. The framework recovers a forward model that projects the CFD solution onto the MRI data geometry, incorporating spatio-temporal smoothing to account for noise and resolution limitations. The authors evaluated 4D-flow HIKING on synthetic 2D 4D-flow MRI data of cerebral arteries with varying noise levels, 4D-flow MRI data of a physical flow phantom with pulsatile inflow, and in-vivo data from intracranial arteries. The proposed framework showed visually comparable flow fields to reference CFD data in synthetic cases and accurately predicted peak inflow velocity. However, the approach demands iterative forward simulations of computationally expensive CFD models and accurate image registration.

Klemens et al. [88] combined MRI velocity observations with CFD to de-noise the MRI measurements by modeling the measured velocity data within the Navier-Stokes filter through a full-order topology optimization using the standard lattice Boltzmann method (LBM) [92]. The optimization problem is defined to minimize the distance between the observed velocity field and the simulated data using the BGK-Boltzmann equation as the side condition. Consequently, the CFD simulation is repeated until the simulated data matches the observations. The authors investigated the proposed method on synthetic noisy MRI data with Gaussian noise in a pipe with a cylindrical object inside, considering a Poiseuille inflow, free outflow, and no-slip condition at walls. They tested the proposed algorithm on data sets with different noise levels and indicated that the approach is not overly reliant on noise intensity. However, it requires multiple CFD simulations with a high computational

cost of optimal control problems. Also, the algorithm is not tested on complex geometries, and the limited resolution of the MRI data is yet to be addressed.

Zhang et al. [190] proposed a method for de-noising and super-resolution of 4D-Flow MRI data of cerebral aneurysms by leveraging data sparsity in a high-resolution flow library constructed from patient-specific CFD simulations and in-vitro particle tracking velocimetry (PTV) measurements. The approach involves building a flow library using velocity fields from CFD simulations and PTV data, and then reconstructing the individual velocity components from noisy, low-resolution synthetic MRI data by calculating a coefficient vector. This vector, along with an observation matrix, establishes a linear relationship between the flow library and the reconstructed velocity field. The optimization problem, involving the L2-norm of fitting residuals with LASSO regularization [169], is solved to determine the coefficient vector. The observation matrix maps high-resolution velocity fields to the 4D-Flow MRI grid, with its coefficients defined by normalized kernel functions (truncated sinc-functions) based on the distance between MRI voxel centers and fine grid points. The flow library includes velocity fields derived from PTV data and CFD simulations, and 4D-Flow MRI flow-rate waveforms as boundary conditions. The results indicate that underestimated velocities and wall shear stress (WSS) in synthetic 4D-Flow MRI data are enhanced in the reconstructed data using this approach. However, the method requires an extensive flow library to capture possible flow fields, making it computationally intensive due to the necessary CFD simulations and the need for expert operators for accurate PTV data acquisition. Additionally, uncertainties and limitations in CFD simulations and PTV significantly affect the outcomes.

## **Sequential Data Assimilation Methods**

Sequential data assimilation methods progressively adjust the state of a dynamic system by incorporating new observations as they are received. Unlike variational approaches, which match the entire state trajectory to observations over a time window, sequential methods use a recursive process where observations are integrated into the model as they occur during

a forward time step. This makes them particularly well-suited for managing processes that evolve over time, such as cardiovascular flows. Unlike the adjoint-based variational methods, the sequential approach only considers observations gathered up to the current time step, without incorporating any future observations. [120, 47, 48].

The Kalman filter [79], which assumes linear dynamics and Gaussian noise, is the foundational algorithm upon which all subsequent sequential data assimilation methods have been built and evolved [183]. It functions through two primary steps: prediction and update. In the prediction step, the filter projects the current state estimate forward in time using the system’s dynamic model, while also predicting the associated uncertainty. During the update step, the filter refines this predicted state by incorporating new, often noisy, observations. This refinement is achieved by calculating a weighted average between the model’s prediction and the observed data, with the weights determined by the Kalman gain. The Kalman gain balances the influence of the prediction and the measurements, taking into account the predicted error covariance (representing the uncertainty of the model’s prediction), the measurement noise covariance (reflecting the uncertainty in the observations), and the observation model (which links the measurements to the system state) [47, 48].

The standard Kalman Filter is optimal only when the assumptions of linearity and Gaussian noise hold true. For systems with significant nonlinearity, the Kalman filter’s estimates can become inaccurate. The Extended Kalman Filter (EKF) extends the Kalman filter to nonlinear systems by linearizing the system around the current estimate using Jacobian matrices, which approximate the dynamics locally as linear. While this allows the EKF to handle some nonlinearity, its performance degrades rapidly as nonlinearity increases. In contrast, the Ensemble Kalman Filter (EnKF) is specifically designed for nonlinear systems and avoids linearization by using an ensemble of state vectors to represent the system’s state distribution. Each ensemble member is propagated through the full nonlinear model, and updates are made based on the ensemble statistics, allowing the EnKF to better capture the complexities of nonlinear dynamics [47, 48].

In the context of cardiovascular modeling, recent studies explored the application of Kalman Filtering in enhancing the 4D Flow MRI measurements. Gaidzik et al. [62] employed a variant of EnKF called the local ensemble transform Kalman filter (LETKF) as a data assimilation algorithm to merge 4D PC-MRI flow data with CFD simulations in order to increase the spatial resolutions of the hemodynamic flow estimations in the subject-specific Circle of Willis geometry and to reduce the uncertainty of the data. The proposed data assimilation approach used a spatial binning operator to map the variables to the lower resolution space. Also, the velocity variance of the 4D PC-MRI measurements was considered as an indication of the uncertainty in the data. Gaidzik et al. [62] compared the peak systolic flow obtained by data assimilation, CFD simulation, and measured by PC-MRI. The reported vascular velocity fields, estimated after data assimilation, were in the range of the CFD simulation and PC-MRI measurements for most of the test cases. Furthermore, results claim a considerable reduction in the estimated velocity variances.

Habibi et al. [68] developed a reduced-order Kalman filter (ROM-KF), combining the dynamic mode decomposition [175] and a linear Kalman filter, incorporating the CFD simulation data and 4D-Flow MRI measurements to de-noise and super-resolution the measurements and enhance the near-wall hemodynamics evaluation in aneurysm geometry. The proposed method employs forward-backward dynamic mode decomposition (fbDMD) to discover a discrete and linear relation between any arbitrary consequent snapshots of the CFD data. The authors used the proper orthogonal decomposition (POD) to compute the low-dimensional state vector of the CFD data. Then the full-order measurements and the low-dimensional CFD data are related by a reduced-order observation matrix and normally distributed measurement noise vector at a particular snapshot. Consequently, the Kalman filter is applied on the reduced-order model, followed by a backward smoother. Eventually, data in full-order space is reconstructed using the predicted low-dimensional state vector and the POD modes matrix of the CFD data. The authors have tested the proposed algorithm on three synthetic cases, including 1D Womersley flow, 2D idealized cerebral aneurysm, and 3D

patient-specific cerebral aneurysm. In the 3D test case, the reported results show that there was approximately a 70% improvement relative to the ground truth when compared to the non-filtered scenario. Also, the method’s capability in improving the near-wall hemodynamics, even in cases with missing the near-wall experimental data, is demonstrated. However, the Kalman filter approach needs an accurate uncertainty assessment for the CFD simulation and an accurate estimation of the noise covariance matrix for the measurement data, which is an ongoing challenge for the real MRI acquired data.

Sequential methods (the Kalman filter and its variants) provide real-time updates and are better suited for handling dynamic processes, but they assume Gaussian errors and may perform poorly if the true noise characteristics are non-Gaussian. Additionally, sequential methods can become computationally expensive when applied to high-dimensional systems typical of hemodynamic studies, such as large-scale 3D blood flow simulations. Both variational and sequential data assimilation approaches can also suffer from inadequate representation of the true error covariance, especially when the underlying models are simplified or when observational data is sparse or highly noisy, limiting their ability to accurately capture the full complexity of hemodynamic phenomena [120]. Moreover, data assimilation algorithms may inherit uncertainties from both CFD simulations and imaging modalities; CFD-related uncertainties often arise from assumptions regarding blood rheological properties, while imaging modalities contribute uncertainties in the form of geometric inaccuracies and boundary condition errors due to limitations in resolution, noise, and artifacts.

## **Deep Learning Based Methods**

Deep learning-based data assimilation for hemodynamics assessment is an promising approach that leverages the strengths of deep learning models to integrate observational data with computational simulations, thereby improving the accuracy and resolution of hemodynamic predictions. These advanced computational techniques are increasingly being utilized to enhance image reconstruction [53, 56, 54], noise reduction [123], antialiasing [14] in

4D-Flow MRI data. Deep learning algorithms, especially those based on neural networks, are capable of efficiently processing complex 4D-Flow MRI datasets, extracting meaningful patterns and correlations that might not be immediately obvious. These methods offer complementary solutions to traditional CFD methods in the field of medical imaging.

Deep learning-based data assimilation can be broadly categorized into two approaches: implicit physics methods and explicit physics methods. Implicit physics methods rely on deep learning models trained on large datasets generated from CFD simulations, where the physical laws have already been applied. In this approach, the neural network is not directly constrained by physical equations during training. Instead, it learns patterns and relationships within the data that reflect the physics indirectly. These networks often employ convolutional neural networks (CNNs) and structured with encoder-decoder architectures [56, 54], learn to enhance imaging data by performing tasks like super-resolution and denoising.

On the other hand, explicit physics methods (physics-informed methods), such as Physics-Informed Neural Networks (PINNs), directly integrate the governing equations of fluid dynamics, like the Navier-Stokes equations, into the training process [142]. These models typically use fully connected neural network architectures to produce continuous solutions within the computational domain, which allows for the use of automatic differentiation to embed physical laws as constraints in the loss function. This approach makes training data-efficient by directly integrating physics into the learning process. By minimizing the residuals of the physics equations, these models ensure that the output flow fields adhere to the fundamental laws of physics. One notable advantage of the PINNs is their capability in solving inverse problems agnostic to the boundary conditions, unlike patient-specific CFD simulations that depend on accurate domain boundaries data [144, 52].

Ferdian et al. [56] developed a convolutional neural network model, called 4DFlowNet, based on the SR generative adversarial network architecture [95] to enhance the spatial resolution of the 4D-Flow MRI images. This model, combining CFD simulations with 4D-Flow MRI, trains on synthetic 4D Flow MRI data to learn the mapping from noisy, low-

resolution data to de-noised, high-resolution data. The 4DFlowNet uses bilinear interpolation for upsampling, connecting low-resolution residual blocks that denoise the input to high-resolution blocks that refine upsampling estimations. This approach requires both low-resolution magnitude and velocity phase images. While the model effectively denoises and super-resolves 4D flow phase images, it is computationally intensive and doesn't ensure fidelity to flow physics.

In their subsequent study, Ferdian et al. [54] enhanced the 4DFlowNet model to analyze cerebrovascular flow, combining it with the vWERP algorithm [108] for improved pressure estimates. This updated model uses smaller input patches and a linear output activation function, omitting gradient terms in the loss function. Similar to its predecessor, the revised 4DFlowNet is a CNN-based super-resolution method transforming input data - velocity component images and magnitude images - into high-resolution velocity images. Like other CNN approaches, it outputs images at a fixed resolution. This model requires a combination of high-resolution patient-specific CFD simulation data and low-resolution synthetic 4D-Flow MRI data for training. The validation on in-silico and in-vivo data showed effective denoising and super-resolution, albeit with limitations in near-wall flow accuracy and reliance on extensive training datasets.

Rutkowski et al. [150] developed a convolutional neural network (CNN) based, CFD and 4D-Flow MRI augmentation method to rebuild the de-noised and super-resolution cerebral 4D-Flow MRI velocity fields. First, the authors trained the proposed block-wise residual architecture based CNN using pair of high-resolution uncorrupted and corrupted (added by white Gaussian noise in k-space) patient-specific CFD simulation data. Subsequently, the trained convolutional network is used to enhance the 4D-Flow MRI observations. The authors tested the proposed method on CFD simulated data, patient-specific in-vitro models, and in-vivo cerebral time-averaged 4D flow MRI data. The results show the proposed algorithm effectively de-noised the MRI data and reduced velocity inaccuracy. The CNN based methods of Ferdian et al. [54] and Rutkowski et al. [150] simply map a given low-resolution

velocity image obtained by processing 4D-Flow MRI to a high-resolution images. The spatial super-resolution is fixed. (2x, 4x, ...) depending on the training. They cannot super-resolve 4D-Flow MRI in time. They also do not explicitly enforce flow physics. They cannot handle velocity aliasing and eddy current induced phase offsets.

Shit et al. [159] developed a deep learning-based super-resolution method for 4D-flow MRI data, addressing the limitations of traditional techniques like cubic B-spline and sinc interpolation. In their study, Shit et al. introduced SRflow, a convolutional neural network incorporating a direction-sensitive mutually projected  $\mathcal{L}1$  (*mp* $\mathcal{L}1$ ) loss function, enhancing the peak velocity-to-noise ratio and flow consistency compared to conventional methods. Authors validated the proposed method on both synthetic and in vivo datasets across various anatomical regions, including cerebrovascular and cardiovascular data. However, the study also highlighted some limitations, including the dependency on finite difference schemes for gradient computation, which restricts the accuracy to the maximum super-resolution factor learned during training. Additionally, the current implementation only explored spatial super-resolution, without temporal super-resolution. Also, other artifacts such as velocity aliasing and phase offset errors are not considered.

Fathi et al. [53] developed a patient-specific PINN to denoise the input 4D-Flow MRI data and reconstruct the local blood pressure and velocity fields. In their approach, the 3D velocity components, the relative pressure, and the magnitude image are all represented using Multilayer Preceptron (MLP), allowing for continuous spatio-temporal field approximations as functions of space and time. These predicted quantities are used to compute the corresponding predicted complex Cartesian 4D-Flow MRI images. Training is conducted directly in the complex Cartesian image space, where the neural network minimizes a loss function comprising data fidelity terms in this space. The loss function is further regularized using the incompressible Navier-Stokes equations and mass conservation principles. By operating in the complex image space and incorporating physical laws, the method naturally handles noise and velocity aliasing artifacts. An advantage of this approach is that it

does not necessitate the specification of geometric boundaries or boundary conditions. The authors validated the proposed method on CFD generated synthetic steady-state 2D flow within solid boundaries of an intracranial aneurysm with one inlet and two outlets, synthetic unbounded unsteady 2D flow for the wake behind the cylinder, as well as an in-vitro test data obtained by 4D-Flow MRI scans. The results indicate a significant improvement of spatio-temporal resolutions and noise levels in the synthetic tests. However, for the in-vitro test, while generated flow fields noticeably improved visually, the quantitative error metrics remained relatively constant.

Kissas et al. [87] proposed combining 4D-flow MRI and CFD using a domain-decomposition and PINN based method, constructing a single network to solve the momentum and mass conservation laws and recover absolute pressure from the 4D-Flow MRI data of the 1D arterial blood flow. The proposed method involves the decomposition of the arterial tree into non-overlapping segments connected by interface boundary conditions (at the arterial tree bifurcations) and modeled using 1D fluid-structure interaction system of PDEs, and consequently, discretized by a PINN. The resultant PINNs are trained to solve the optimization problem with a single loss function, which includes data fidelity term, physical constraints fidelity term, and an interfaces continuity term that couples the individual PINNs at the interface boundaries. The authors tested the proposed algorithm in three different cases: synthetic carotid bifurcation with three vessels and one bifurcation, idealized pelvic geometry with seven vessels and three bifurcations, and a real clinical case with 4D-Flow MRI data of carotid bifurcation of a healthy subject. This method is applicable to unidirectional flows and is not applicable in regions of complex 3D flow, such as cerebral arteries, where the low resolution and low signal to noise ratio of the input 4D-Flow MRI become a more significant challenge. Also, the noise statistics of the real data may not be modeled only by Gaussian noise, which is used in this study.

Shone et al. [160] employed a physics-informed neural network (PINN) with sine activation functions, incorporating the Navier-Stokes equations and the no-slip boundary condition

on the endocardium as constraints within the loss function, to enhance cardiac 4D-flow MRI data and address challenges such as low spatio-temporal resolution and significant noise artifacts in moving domains. The study validated the model using synthetic data generated from computational fluid dynamics (CFD) simulations of both 2D idealized and 3D patient-specific left ventricles. The results demonstrated that the PINN-based approach effectively upsampled the data at extreme rates, achieving normalized velocity RMSE values below 16% even under significant data degradation. The model also exhibited robustness to various levels of noise and demonstrated superiority over conventional interpolation methods and previous PINN implementations by preserving fine-scale flow features.

Aghaee and Khan [2] explored the performance of Fourier-based activation functions in PINNs for modeling patient-specific cardiovascular flows, addressing the spectral bias limitation that PINNs often face. Spectral bias, where neural networks tend to learn simpler, low-frequency functions first [141], poses challenges in accurately modeling complex cardiovascular flows that feature high spatial frequency variations, such as in aneurysms and stenoses. The study compared the performance of a sine activation function, known for its ability to reduce spectral bias, against conventional activation functions like tanh and swish in various cardiovascular flow scenarios, including a 1D advection-diffusion problem, a 2D eccentric stenosis model, and a 3D patient-specific aortic model. The results demonstrated that the sine activation function offered superior performance in terms of error reduction, faster convergence, and better representation of high-frequency data compared to the other activation functions.

## 1.6 Thesis Outline and Objectives

### 1.6.1 Objectives

Existing hemodynamic assessment methods, including time-resolved PC-MRI, patient-specific computational fluid dynamics (CFD) simulations, and traditional deep learning techniques,

suffer from significant drawbacks. Time-resolved PC-MRI, despite its capabilities, faces challenges such as inadequate spatio-temporal resolution, background phase offset errors, acquisition noise, and velocity aliasing, all of which compromise the precision of flow measurements.

Patient-specific CFD simulations, while providing detailed insights, require accurate flow assumptions, segmented geometry, and boundary conditions. Errors in these inputs can lead to substantial inaccuracies, and the computational intensity of CFD limits its practicality for routine clinical use.

Physics-Informed Neural Networks (PINNs) present a promising alternative by integrating physical laws into the learning process. However, their application in medical imaging is limited by a lack of generalizability, especially given the wide variability in anatomical structures, flow patterns, and physiological conditions across datasets. The training time for a given dataset is also inordinately high.

To overcome these challenges, this thesis proposes the input-parametrized physics-informed neural network (IP-PINN) framework, which integrates the generalizability of convolutional neural network (CNN)-based super-resolution techniques with the continuous spatio-temporal field representation inherent in PINNs. By incorporating a novel data fidelity term based on PC-MRI acquisition physics, this unified framework seeks to overcome the limitations of 4D Flow MRI, including the elimination of background phase offsets, reduction of random noise, and eliminating errors due to velocity aliasing. Moreover, the algorithm operates directly on MRI images, generating noise- and artifact-free, super-resolved flow images as well as high-quality magnitude images without requiring reference scans.

The results presented in this thesis demonstrate the IP-PINN’s capability to generalize across different patient-specific datasets, significantly reducing the need for time-intensive training for processing new data. The framework also achieves considerable improvements in velocity field prediction, boundary segmentation, and processing time compared to traditional methods like PINN and DF-RBF, highlighting its potential for clinical application.

Ultimately, this thesis showcases the successful application of IP-PINN in super-resolving and post-processing patient-specific 3D time-resolved PC-MRI (4D Flow MRI) data, and post-processing 2D PCMRI data and reconstructing three-dimensional three-component velocity maps in the vicinity of the interrogation plane of the 2D PCMRI. The advancements made through this research will be discussed in detail in subsequent chapters, underscoring the IP-PINN framework’s potential to revolutionize hemodynamic assessment in medical imaging and contribute to better understanding and management of cardiovascular diseases.

## 1.6.2 Thesis Outline

This thesis is structured into seven main chapters, each building on the previous ones to advance the understanding and application of the Input-Parameterized Physics-Informed Neural Network (IP-PINN) for enhancing time-resolved phase-contrast MRI data.

Chapter 2 delves into the core principles and techniques of phase-contrast MRI, detailing the acquisition process, velocity encoding, and signal acquisition. This chapter highlights the unique advantages of 4D Flow MRI in capturing comprehensive, time-resolved, three-directional velocity data over entire vascular volumes. It also explores the limitations, including low spatio-temporal resolution and common acquisition artifacts, setting the stage for the development of the IP-PINN framework.

Chapter 3 presents the innovative IP-PINN framework, which combines convolutional neural networks and multilayer perceptrons with physics-informed neural network principles to address the limitations of traditional hemodynamic assessment methods. This chapter provides a comprehensive explanation of the network architecture, data fidelity terms, and physics-informed constraints, showcasing how IP-PINN enhances 4D Flow MRI by super-resolving images and correcting artifacts, without the need for ab-initio retraining on new datasets.

Chapter 4 constitutes the first large-scale validation of the IP-PINN concept on synthetic patient-specific 4D Flow MRI datasets. Starting from synthetic aneurysm data pro-

duced with high-fidelity CFD, the network is pre-trained once on a posterior-cerebral-artery (PCA) model and then rapidly fine-tuned on an unseen middle-cerebral-artery (MCA) case, illustrating the framework’s generalizability. By formulating its data-consistency term in complex image space, IP-PINN simultaneously super-resolves the three-component velocity field, denoises the data, removes phase offsets, and unwraps severe velocity aliasing, all without access to reference scans or ground-truth labels. The continuous, analytical representation delivered by the network further supports temporal interpolation and automatic reconstruction of high-resolution magnitude images, from which accurate lumen boundaries are extracted. Quantitative comparisons show that IP-PINN outperforms classical fully connected PINNs by an order of magnitude in both accuracy and run-time. This chapter therefore establishes IP-PINN as a fast, generalizable alternative to traditional physics-informed solvers for standard 4D-Flow MRI acquisitions.

Chapter 5 adapts the IP-PINN concept to two-dimensional phase-contrast MRI (2D PC-MRI). Because standard 2D PC-MRI encodes only through-plane velocity, it cannot resolve in-plane flow components, 3D vascular geometry, or provide accurate wall-shear-stress (WSS) information. To overcome this limitation, the chapter first introduces a minimally modified acquisition in which the bipolar velocity-encoding gradients are cycled through the spatial directions over successive cardiac phases, eliminating the separate reference scan. The resulting single-direction, single-slice data are then fed to an IP-PINN to reconstruct super-resolved, noise- and alias-free 3D-3C velocity fields and spin-density maps in the neighborhood of the imaging slice. The results show that the method recovers all three velocity components, produces temporally super-resolved flow fields, delineates high-resolution lumen boundaries, and computes WSS with near-CFD accuracy—all within minutes and without any ground-truth labels. This chapter demonstrates that IP-PINN can transform a conventional 2D PC-MRI scan into a volumetric hemodynamic dataset, providing a fast and clinically practical alternative to full 4D Flow MRI.

Chapter 6 pushes the IP-PINN concept to the extreme undersampled pseudo 1-point

4D-Flow MRI sequence and a k-space-based reconstruction strategy. The proposed sequence acquires just one velocity-encoded dataset per slice, rotating the encode direction slice-by-slice and omitting the reference scan altogether, thereby slashing the raw data burden by roughly 75% relative to the conventional four-point protocol. To recover fully sampled, time-resolved 3D-3C velocity and magnitude images from these highly sparse measurements, the chapter reformulates the IP-PINN data-consistency term directly in Fourier space, preserving exact acquisition physics and avoiding the convolution blurring that complicates image-domain fidelity. The revised IP-PINN produces super-resolved, denoised, alias-free flow fields throughout the volume. The results show that, even at 20 dB noise levels or with deliberate velocity aliasing, the k-space IP-PINN reconstructs accurate velocities and lumen geometry, converging in about two minutes. Thus, Chapter 6 demonstrates that IP-PINN can enable radically accelerated 4D-Flow MRI acquisitions without sacrificing quantitative fidelity.

Chapter 7 concludes the thesis by summarizing the major contributions and findings across all chapters. It discusses the significance of the IP-PINN framework in advancing time-resolved phase-contrast MRI and evaluates its performance across various acquisition settings. The chapter also reflects on the limitations of the current implementation, including challenges related to turbulent flow regimes and moving boundaries. Building on these reflections, several directions for future research are proposed, including in-vitro validation studies, extensions to turbulent and pulsatile flows, applications to dynamic anatomical geometries, and improved boundary extraction techniques using level-set methods. These future directions aim to enhance the clinical translation and versatility of IP-PINN in real-world cardiovascular imaging applications.

## Chapter 2

### Time-Resolved Phase-Contrast MRI

The time-resolved PC-MRI acquisition process leverages the phase sensitivity of MRI to the motion of magnetic moments hydrogen protons (spins) [43, 104]. When subjected to an external magnetic field, these spins align and precess, generating the MRI signal that is sensitive to their movement, allowing the measurement of flow and motion within the body. The development of the first 3D cine phase-contrast MRI sequence was pioneered by Firmin et al. [58] This innovation marked a significant advancement in the field of cardiovascular imaging by allowing for comprehensive velocity measurements across three spatial dimensions, capturing the complex flow patterns within the cardiovascular system.

Considering the pulsatile nature of the cardiovascular flows, the cardiac cycle could be divided into a number of cardiac phases and due to the time taken to complete the volumetric scans and fill the k-space for a given cardiac phase, the process of acquisition is spread over several cardiac cycles through ECG gating [165]. At each cardiac phase a minimum of four scans are required: three velocity-encoded scans in the  $x$ ,  $y$ , and  $z$  directions, and one reference scan without velocity encoding. The bipolar velocity-encoding gradients induce phase shifts in the MRI signal, which are proportional to the velocity components of the moving spins in the tissue. A bipolar gradient consists of a sequence of two gradient pulses—one negative and one positive—with the same amplitude. These gradients are particularly advantageous in PC-MRI because they result in zero phase shift for static spins and non-zero phase shift for moving spins between the end of the radio frequency (RF) excitation and echo time (TE), where the signal is collected [138]. These gradients are applied sequentially in each spatial direction ( $x$ ,  $y$ , and  $z$ ) [106, 30]. By combining phase shifts from each acquisition and subtracting the reference phase (from the non-encoded scan) from the

velocity-encoded phases to isolate the velocity-induced phase shifts, the technique reconstructs the time-resolved full 3D velocity vector at each voxel. There are the following three main assumptions that are used in this process:

- The velocity flow patterns are periodic with a time period equal to that of the cardiac cycle.
- The cardiac cycle-to-cycle variations is negligible.
- The velocity is unchanged over the duration of the cardiac phase.

The description in the next subsections is valid for data acquired as part of a single cardiac phase

## 2.1 Velocity Encoding and Signal Acquisition

The velocity-encoded scan MRI signal in the  $j^{th}$  direction is given by:

$$\hat{C}^j(\mathbf{k}, z, t) = \mathcal{F}(\sigma^j(\mathbf{x}, z, t)) + \hat{\eta}^j(\mathbf{k}, z, t) \quad (2.1)$$

where  $\mathbf{x} = (x, y)$  are the coordinates in the coronal plane,  $z$  is the distance of the coronal plane from the MRI coordinate system (Fig. 2.1),  $t$  is the time/cardiac phase,  $\sigma(\cdot)$  is the apparent spin density,  $\hat{\eta}(\cdot)$  is the acquisition noise, and  $\mathcal{F}(\cdot)$  is the Fourier transform operator. Note that the selected slice is parallel to the coronal plane and therefore we separate out the coordinates as  $(x, y, z, t) = (\mathbf{x}, z, t)$ . The transformed complex  $\mathbf{x}$ -space image (a complex number at each spatio-temporal location) is therefore given by:

$$C^j(\mathbf{x}, z, t) = \mathcal{F}^{-1}\hat{C}^j(\mathbf{k}, z, t) = \sigma^j(\mathbf{x}, z, t) + \eta^j(\mathbf{x}, z, t) \quad (2.2)$$

Since  $\mathcal{F}^{-1}$  is a linear operation, the noise in  $\mathbf{x}$ -space,  $\eta^j(\cdot)$ , has the same characteristics (0 mean Gaussian) as  $\mathbf{k}$ -space noise,  $\hat{\eta}^j(\cdot)$ . The apparent spin density is given by [138, 116]:

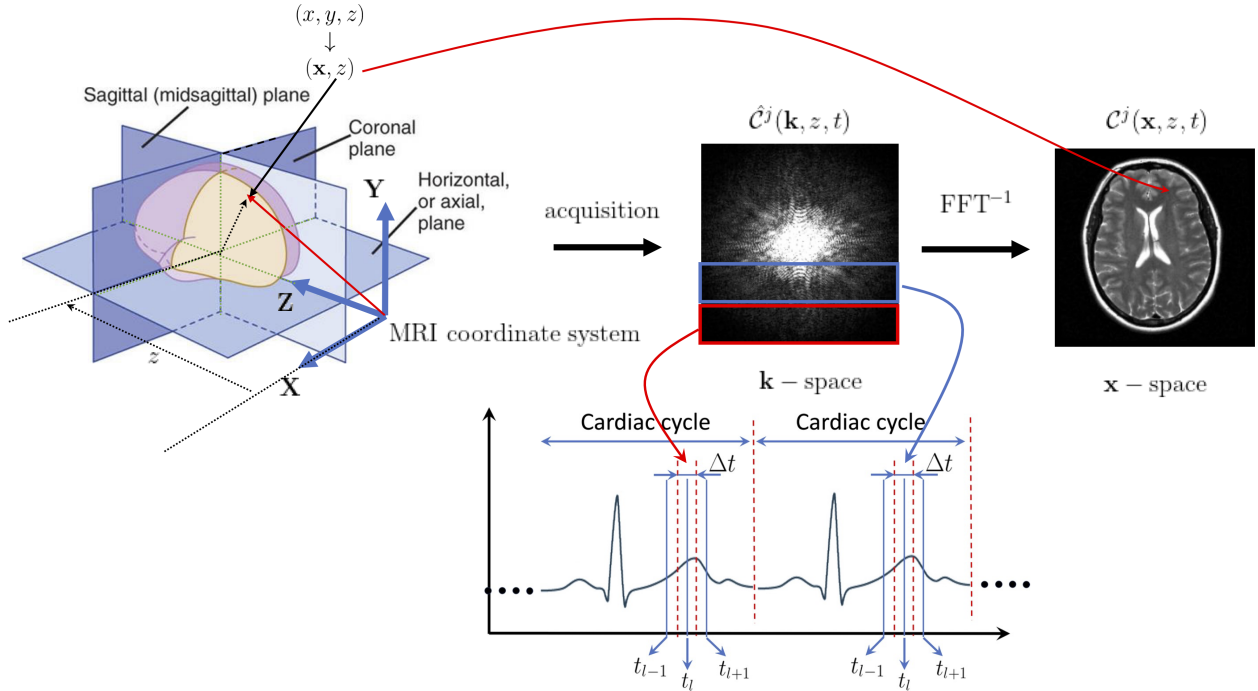


Figure 2.1: 4D-Flow MRI acquisition. At cardiac phase  $t = t_l$ , and for each coronal plane slice, a minimum of 4 acquisitions corresponding to a flow compensated scan ( $j = 0$ ), followed by flow encoded scans in  $x, y, z$  directions ( $j = 1, 2, 3$ ) are acquired. The acquisition is spread over several cardiac cycles, with each cardiac cycle accounting for a small portion of the  $\mathbf{k}$ -space. Here  $\Delta t$  is the time span of the cardiac phase  $t_l$ , i.e., the time for acquiring the data in  $\mathbf{k}$ -space for that particular cardiac phase. A  $\text{FFT}^{-1}$  of the  $\mathbf{k}$ -space acquisition generates the complex  $\mathbf{x}$ -space image. The magnitude (computed on a per voxel basis) of the  $\mathbf{x}$ -space image is typically used in examining anatomical structures. The phase contains tissue velocity information.

$$\sigma^j(\mathbf{x}, z, t) = \omega_0 \Lambda |M_{xy}(\mathbf{x}, z, 0)| |C_{xy}^r(\mathbf{x}, z)| e^{i(\phi_0 - \phi_r(\mathbf{x}, z) - \phi_c(\mathbf{x}, z))} e^{i(-\phi_u^j(\mathbf{x}, z, t) - \phi_e^j(\mathbf{x}, z))} \quad (2.3)$$

where  $\omega_0$  is the Larmor frequency,  $\Lambda$  are the amplification factors of the signal processing electronics,  $|M_{xy}(\mathbf{x}, z, 0)|$  is the magnitude of the transverse magnetization at the end of the RF pulse,  $|C_{xy}^r(\mathbf{x}, z)|$  is the magnitude of the receiver coil sensitivity,  $\phi_0$  is the phase of the transverse magnetization at the end of the RF pulse,  $\phi_r(\mathbf{x}, z)$  is the phase due to the receiver coils,  $\phi_c(\mathbf{x}, z)$  is the phase error due to changes in the main static field  $B_0$ , chemical shift [34], susceptibility [184], and field inhomogeneity [179],  $\phi_u^j(\mathbf{x}, z, t)$  is the phase induced by

tissue velocity [103],  $\phi_e^j(\mathbf{x}, z)$  is bi-polar gradient dependent phase errors due to eddy currents [178], concomitant fields [168], and gradient nonlinearities [24]. The phase induced by tissue velocity is given by:

$$\phi_{\mathbf{u}}^j(\mathbf{x}, z, t) = \gamma \mathbf{u}(\mathbf{x}, z, t) \cdot \mathbf{g}^j \quad (2.4)$$

where  $\gamma$  is the gyromagnetic ratio,  $\mathbf{u}(\mathbf{x}, z, t)$  is the tissue velocity at  $t$ , and  $\mathbf{g}^j$  is the first moment of the bipolar gradient given by:

$$\mathbf{g}^j = \|\mathbf{g}^j\| \hat{\mathbf{g}}^j = \int_0^{t_v} \tau \mathbf{G}^j(\tau) d\tau \quad (2.5)$$

Here  $\mathbf{G}^j(\tau)$  is the  $j^{\text{th}}$  bipolar gradient function,  $t_v$  is the time for which the bipolar gradient is switched on. By adjusting the direction of the bipolar gradient  $\mathbf{G}^j$ , various kinds of velocity encoded scans can be accomplished (e.g., 4-point, 4-point balanced, or 5-point balanced scan) [75]. Defining  $v_{enc}$  as the magnitude of the velocity projection along  $\mathbf{g}^j$  that causes  $\pi$  phase, we can write:

$$v_{enc} = \frac{\pi}{\gamma \|\mathbf{g}^j\|} \quad (2.6)$$

## 2.2 Image Reconstruction

A 4-point scan involves an interleaved reference scan followed by three velocity encoded scans in the  $x, y, z$  directions. The unit gradient vectors are therefore given by:

$$\hat{\mathbf{g}}^x = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \hat{\mathbf{g}}^y = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad \hat{\mathbf{g}}^z = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \quad \hat{\mathbf{g}}^z = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad (2.7)$$

Using these 4 scans, the phase due to velocity can be isolated and the component-wise velocities can be calculated. We rewrite equation 2.3 as:

$$\sigma^j(\mathbf{x}, z, t) = M(\mathbf{x}, z) e^{i(\phi_B^j(\mathbf{x}, z) - \phi_{\mathbf{u}}^j(\mathbf{x}, z, t))} \quad (2.8)$$

where

$$M(\mathbf{x}, z) = \omega_0 \Lambda |M_{xy}(\mathbf{x}, z, 0)| |C_{xy}^r(\mathbf{x}, z)| \quad (2.9)$$

is the image magnitude and

$$\phi_B^j(\mathbf{x}, z) = \phi_0 - \phi_r(\mathbf{x}, z) - \phi_C(\mathbf{x}, z) - \phi_e^j(\mathbf{x}, z) \quad (2.10)$$

is the composite phase. In practice, the MRI image is discrete, and the value of the apparent spin density at a voxel centered at  $(\mathbf{x}_k, z_k)$  at cardiac phase  $t_l$  is given by:

$$\sigma_D^j(\mathbf{x}_k, z_k, t_l) = \frac{1}{\Delta t \Delta z} \int_{t_l - \frac{\Delta t}{2}}^{t_l + \frac{\Delta t}{2}} \int_{z_k - \frac{\Delta z}{2}}^{z_k + \frac{\Delta z}{2}} [\sigma^j(\cdot, z, t) * h(\cdot)](\mathbf{x}_k) dz dt \quad (2.11)$$

where  $\Delta z$  is the slice thickness and  $\Delta t$  is the time slice thickness of the cardiac phase. The action of the convolution of the apparent spin density  $\rho^j(\cdot)$  with the acquisition point spread function (PSF)  $h(\cdot)$  could be conducted with a pointwise multiplication in k-space given by:

$$[\sigma^j(\cdot, z, t) * h(\cdot)](\mathbf{x}_k, z, t) = \mathcal{F}_D^{-1} [\hat{\sigma}^j(\mathbf{k}, z, t) \cdot u(\mathbf{k}) \cdot W(\mathbf{k})] \quad (2.12)$$

where  $u(\mathbf{k})$  is a sampling pulse train with pulse width  $\Delta k$  and  $W(\mathbf{k})$  is the windowing function in the frequency domain that cuts out high frequency components beyond  $\pm k^{max}$  in the  $x$  and  $y$  directions. Consequently, the discrete Cartesian image after processing the  $k$  space acquisition is given by:

$$C_D^j(\cdot, t_l) = \sigma_D^j(\cdot, t_l) + \mathcal{F}^{-1} \hat{\eta}(\mathbf{k}, t_l) \quad (2.13)$$

where  $\mathcal{F}_D$  is the discrete Fourier transform.

## 2.3 Limitations

While time-resolved PC-MRI offers advantages over traditional methods, its limitations include low spatio-temporal resolution and acquisition artifacts such as phase offset errors,

which can impact accuracy, especially near arterial walls and in small vessels or cardiovascular conditions like brain aneurysms. In the literature, the spatial and temporal resolutions of MRI are reported to be in the order of 0.8-3.5 mm and 20-50 ms, respectively [51, 69, 137, 11]. Therefore, the applicability of this method in complex cardiovascular regions with numerous smaller vessels is highly dependent on spatial resolution and the user-defined parameter, velocity encoding ( $v_{enc}$ ). Additionally, The need for multiple velocity-encoded images can result in long scan times, which may be challenging for certain patient populations. Following are the common artifacts observed in the time-resolved PC-MRI acquired data:

### 2.3.1 Phase Offset Errors

Phase offset errors in time-resolved PC-MRI data arise from various non-idealities within the MRI system, leading to inaccuracies in the measurement of velocity and thereby affecting the flow data's fidelity. Such errors can stem from several factors:

- **Chemical Shifts:** The resonance frequencies of protons vary based on their molecular environment, creating a local magnetic field that conflicts with the external field. This phenomenon, known as chemical shift, results in minor variations in resonance frequency due to the shielding effect, impacting the accuracy of phase data in MRI [34, 138].
- **Magnetic Susceptibility:** Different materials either disperse or concentrate the surrounding magnetic field, causing distortions at the interface between two tissues. These variations in magnetic susceptibility can influence the phase data, especially at tissue interfaces [184].
- **Hardware Inconsistencies:** Factors like receiver coil sensitivities, inhomogeneities in the static magnetic field, and non-linear gradients along the acquisition axes can affect phase signals, introducing local field phase errors [103, 179].

- **Eddy Currents:** In time-resolved PC-MRI, eddy currents can cause phase offset errors. These currents are generated by the rapid switching of magnetic gradients in the MRI machine. The changing magnetic fields induce secondary currents in the conductive parts of the scanner, which in turn produce their own magnetic fields. These induced fields can distort the primary magnetic field, leading to inaccuracies in the phase data acquired during the imaging process. This issue is particularly significant when high-speed gradients are used, as they can exacerbate the effects of these eddy currents [178, 181].

The most of these phase offset errors can be mitigated through a combination of pre-processing steps and calibration techniques, such as reference scan subtraction or phase-encoding schemes. The reference scan subtraction method involves acquiring a static reference image without any flow encoding gradients, which captures only the phase due to non-velocity-related factors. By subtracting this reference phase from the flow-encoded images, the non-flow-related phase offsets are removed, isolating the true velocity-induced phase shifts. However, this approach is impractical for routine clinical use due to the additional time required. An alternative is the static-tissue interpolation method, which uses data points from static tissue to fit a polynomial model to the background phase in non-flow regions; this model can be used to subtract the background phase from the phase of the dynamic tissue [182, 26]. On the other hand, phase-encoding schemes correct phase offset errors by strategically varying the phase encoding direction and amplitude across different acquisitions. One common method is the use of flow compensation techniques, where phase-encoding gradients are adjusted or duplicated in opposite directions, allowing phase shifts due to flow to be differentiated from those due to other sources like gradient imperfections. By averaging or comparing these acquisitions, the true flow-induced phase can be isolated while canceling out the systematic phase errors. [135, 165, 24, 41]. Although these approaches effectively corrects phase offset artifacts due to most of the error sources, it does not address phase offset errors due to the eddy currents [165]. Fig. 2.2 shows the phase




Figure 2.2: Correction of Phase Offsets in 4D Flow MRI: (a, e) Acquired magnitude images; (b, c, f, g) Phase offset errors due to tissue motion; (d, h) Phase difference images obtained after subtracting consecutive phase images with toggled bipolar gradients, effectively canceling time-dependent artifacts while leaving magnetic field inhomogeneities and non-linear gradient fields uncorrected. [135]

offset errors (b, c, f, and g) observed in the 4D Flow MRI data. Aforementioned subtraction techniques ultimately results in a phase difference image, as shown in Fig. 2.2 (d and h). Additionally, Fig. 2.2 (a and e) displays the corresponding anatomical structure without velocity information acquired from magnitude images [135].

### 2.3.2 Velocity aliasing

In PC-MRI, the phase of moving spins is encoded within a range from  $-\pi$  and  $+\pi$ . The velocity encoding setting  $v_{enc}$  must be carefully chosen by the operator to match the maximum expected flow velocity in the specific direction, with the highest anticipated velocity corresponding to a phase shift of  $-\pi$ . When the actual velocity in a flow region exceeds the predefined  $v_{enc}$  threshold in PC-MRI, velocity aliasing occurs. This leads to the phase difference wrapping around  $\pi$  introducing a negative velocity reading. Such a scenario is often characterized by sudden and abrupt changes in velocity within the flow region [165]




Figure 2.3: Illustration of velocity aliasing in PC-MRI: The left panel shows flow images acquired with low  $v_{enc}$ , resulting in aliasing artifacts. The right panel shows the corresponding high  $v_{enc}$  acquisition, where aliasing is mitigated. The diagram on the far right visualizes the phase wrapping effect that leads to velocity aliasing when the actual flow exceeds the  $v_{enc}$  threshold. [107]

(Fig. 2.3). Adjusting  $v_{enc}$  involves changing the strength and duration of velocity encoding gradients. Lower velocities require larger gradient areas, leading to longer echo and repetition times. [113] Also, The velocity noise in PC MRI is proportional to  $v_{enc}$  and inversely related to the signal-to-noise ratio (SNR) of the magnitude images.[127] Increasing  $v_{enc}$  can prevent velocity aliasing but may also heighten velocity noise in the velocity images.

### 2.3.3 Random noise in k-space

Random noise in k-space is a significant concern in PC-MRI as it affects the quality and accuracy of the reconstructed images. This noise originates from various sources like electronic interference in the MRI system and thermal electron agitation in the receiver coils. The quality of PC-MRI measurements is largely influenced by settings such as the velocity encoding  $v_{enc}$  and the duration of the acquisition. A higher  $v_{enc}$  tends to lower the signal-to-noise ratio (SNR). Likewise, a higher spatial resolution may lead to longer acquisition times or a reduced SNR, underscoring the trade-offs in PC-MRI settings. [88] In k-space, noise appears as random variations with a zero-mean Gaussian distribution. Post Fourier transformation to complex Cartesian image space, the noise characteristics remain constant

due to the linearity of the inverse Fourier transform. However, the conversion of complex image space data into phase and magnitude images is a non-linear process, which alters the noise characteristics. [53, 88]

## Chapter 3

# Development of the Input-Parametrized Physics-Informed Neural Network

### 3.1 Introduction

Input-Parameterized Physics-Informed Neural Network (IP-PINN) represents an innovative approach to overcoming the limitations associated with current hemodynamic assessment methods, including time-resolved PC-MRI, patient-specific CFD simulations, and conventional deep learning models. By integrating convolutional neural networks (CNNs) and multilayer perceptrons (MLPs) with the principles of physics-informed neural networks (PINNs), IP-PINN creates a robust framework capable of generalizing across diverse datasets while preserving high fidelity to the underlying physical principles of hemodynamics.

Designed to enhance time-resolved PC-MRI data, IP-PINN directly processes MRI images to produce super-resolved, noise-free, and artifact-free flow fields. In contrast to traditional PINNs [142], which are typically constrained to single datasets and require costly ab-initio training for new, unseen data, IP-PINN parametrizes the solution based on input data, accommodating variations in patient-specific anatomy and flow conditions. This adaptability enables the model to address a broader range of clinical scenarios without extensive retraining, making it a more practical tool for clinical applications.

The key features of IP-PINN are as listed below:

- **Integration of CNNs and MLPs:** IP-PINN employs CNNs for their ability to extract features from complex patterns and MLPs for their capability to approximate continuous spatio-temporal field representation. Through a CNN-based encoder architecture, IP-PINN generalizes conventional PINNs by parametrizing the output with

input data, enhancing the model’s versatility.

- **Physics-Informed Constraints:** By embedding the governing physical laws directly into the neural network’s loss function, IP-PINN ensures that the predicted flow fields adhere to the fundamental principles of fluid dynamics. This explicit physics-based approach allows the network to correct common artifacts in 4D Flow MRI, such as velocity aliasing.
- **Data Fidelity Term Based on PC-MRI Acquisition Physics:** A novel aspect of IP-PINN is the inclusion of a data fidelity term that is specifically tailored to the physics of PC-MRI acquisition. This term helps align the network’s predictions with the observed data, enabling the network to be trained using both real and synthetic data and enhancing the accuracy and reliability of the resulting flow fields.

This chapter provides a comprehensive exploration of the IP-PINN framework, detailing the principles behind its components, the architecture of the network, and the key innovations that differentiate it from existing hemodynamic assessment methods.

## 3.2 Background

The development of the Input-Parameterized Physics-Informed Neural Network (IP-PINN) framework relies on a strong foundation in deep learning architectures, particularly Convolutional Neural Networks (CNNs) and Multilayer Perceptrons (MLPs). These neural network models form the backbone of modern data-driven approaches for complex tasks like image analysis and time-series prediction, which are critical for advancing hemodynamic assessments in medical imaging. This section provides an overview of the principles behind these neural networks, exploring their key components, such as convolutional and pooling layers in CNNs, the structure and function of MLPs, and the role of activation functions. Additionally, it delves into the integration of physics-based constraints, as seen in PINNs, which allow these networks to incorporate domain-specific knowledge into the learning process.

Understanding these foundational concepts is crucial for appreciating the innovations and capabilities of the IP-PINN framework in addressing the limitations of traditional hemodynamic assessment methods.

### 3.2.1 Convolutional Neural Networks (CNNs)

Convolutional Neural Networks (CNNs) are deep learning architectures used extensively in image and volumetric data analysis due to their ability to capture spatial hierarchies and patterns. The various structures of the CNNs often involve very similar component, such as convolutional layers, and pooling layers. The convolutional layers often consist of several filter or kernels  $K$ , aim to learn different feature maps through convolution operations [66]. In a typical CNN, the convolution operation involves applying a kernel to an input  $I$  to produce a feature map  $C$ . For a 2D convolution layer, this operation is mathematically represented as:

$$C(x, y) = \sum_{i=-ks}^{ks} \sum_{j=-ks}^{ks} I(x+i, y+j) \cdot K(i, j) \quad (3.1)$$

where  $C(x, y)$  is the output,  $I$  is the input image,  $K$  is the kernel, and  $ks$  is the kernel size.

Similarly, a 3D convolution extends this concept to three dimensions, which is particularly useful for volumetric data such as MRI scans. The 3D convolution operation is given by:

$$C(x, y, z) = \sum_{i=-ks}^{ks} \sum_{j=-ks}^{ks} \sum_{l=-ks}^{ks} I(x+i, y+j, z+l) \cdot K(i, j, l) \quad (3.2)$$

where  $C(x, y, z)$  is the output feature map,  $I$  is the input volume,  $K$  is the 3D kernel, and  $ks$  defines the kernel's dimensions. In both 2D and 3D cases, the convolution operation is critical in capturing spatial relationships within the data.

Pooling layers, another essential component of CNNs, are used to reduce the spatial dimensions of feature maps, thereby lowering the computational load and mitigating overfitting

[66]. The most common types of pooling are max pooling and average pooling:

Max pooling takes the maximum value from each patch of the feature map, thereby preserving the most prominent features. Average pooling calculates the average of each patch, smoothing the feature maps. Mathematically, pooling can be expressed as:

$$P(x, y, z) = \text{pool}(C(x + i, y + j, z + l) \mid i, j, l \in R) \quad (3.3)$$

where  $P(x, y, z)$  represents the output after pooling,  $C$  is the input feature map, and  $R$  defines the pooling region. Pooling functions, such as max or average, are applied over local neighborhoods defined by  $R$  to generate pooled outputs. Pooling layers contribute to shift-invariance and robustness of the model by aggregating features from patches, allowing the network to effectively summarize spatial information from its inputs [66].

By integrating convolutional layers with pooling layers, CNNs efficiently extract hierarchical features, enabling the network to capture both low-level details like edges and corners, as well as high-level features such as textures and shapes, which are crucial for tasks like image classification, segmentation, and object detection.

### 3.2.2 Multilayer Perceptrons (MLPs)

Multilayer Perceptrons (MLPs) are fundamental building blocks of deep learning architectures, consisting of interconnected layers of neurons. An MLP typically includes an input layer, one or more hidden layers, and an output layer, where each layer is composed of nodes (neurons) that apply linear transformations followed by non-linear activation functions. The architecture enables MLPs to learn complex mappings from input to output by adjusting the weights and biases through backpropagation during training [32].

Each neuron in a hidden layer takes inputs from the previous layer, processes them through a weighted sum, adds a bias, and applies an activation function to introduce non-linearity. Mathematically, this can be expressed as:

$$o = f\left(\sum_i (w_i \cdot x_i) + b\right) \quad (3.4)$$

here,  $o$  is the output of the perceptron,  $f$  is the activation function,  $w_i$  are the weights,  $x_i$  are the inputs, and  $b$  is the bias term. The MLPs adjust these weights and biases to minimize the difference between its output and the target values through the training process. MLPs are particularly powerful due to their universal approximation capability, which allows them to approximate any continuous function given sufficient neurons and layers [32].

### 3.2.3 Activation Functions

Activation functions are critical components in neural network architectures, introducing non-linearity into the network to enable it to learn complex patterns and relationships in data. Without activation functions, the network would simply be a linear transformation of the inputs, limiting its capacity to model real-world phenomena.

Among the various activation function, below is the list of commonly used activation functions in physics-informed neural networks:

- **Rectified Linear Unit (ReLU):** Outputs the input directly if it is positive; otherwise, it outputs zero. This function is widely used due to its simplicity and effectiveness in accelerating convergence during training.
- **Sigmoid:** Maps the input into a range between 0 and 1, often used in the output layer for binary classification tasks.
- **Swish:** The Swish activation function, defined as  $Swish(x) = x \cdot sigmoid(x)$ , is a smooth and non-monotonic function. It allows small negative values, which can help the network explore a broader range of features and improve learning dynamics.
- **Tanh:** Maps the input to a range between -1 and 1, providing a zero-centered output which can help in the convergence of models.

- **Sine:** The sinusoidal (sine) activation function, prominently used in SIRENs [161], leverages its periodic nature to capture fine details in the data and maintain information about signal derivatives. This function excels in applications that require continuous and differentiable signal representation.

These functions enable the neurons to activate based on their inputs, facilitating the deep networks' ability to approximate complex, non-linear functions. Traditional activation functions, like ReLU and tanh, have fixed forms that may not optimally fit the specific patterns or complexities within the data, which can limit their performance, especially in more complex tasks. Also, the choice of activation function isn't straightforward and depends on the problem. To address this, adaptive approaches have been proposed [74, 139, 42, 97, 158, 189]. Adaptive activation functions dynamically adjust their parameters during training, allowing the model to optimize the activation's slope or shape based on the data. This adaptability enhances the model's learning capacity, particularly for complex tasks involving differential equations, like those found in physics-informed neural networks (PINNs). By modifying the function's characteristics during training, adaptive activation functions help improve convergence rates and overall performance, especially in networks dealing with highly variable data distributions [74].

### 3.2.4 Physics-Informed Neural Networks (PINNs)

Physics-Informed Neural Networks (PINNs) are a class of neural networks that incorporate physical laws into the learning process by embedding differential equations directly into the loss function [142]. Unlike traditional data-driven neural networks, PINNs leverage both observational data and prior knowledge of physical systems to improve model accuracy.

In PINNs, the neural network approximates the solution to a physical problem by minimizing a loss function that includes two main components:

1. **Data Fidelity Term ( $\mathcal{L}_{\text{data}}$ ):** Measures the discrepancy between the network's predictions and the observed data.

2. **Physics-Based Regularization Term ( $\mathcal{L}_{\text{physics}}$ ):** Enforces compliance with the governing physical laws, typically expressed as differential equations.

By integrating the physical laws directly into the loss function, PINNs ensure that the network’s predictions are not only consistent with the data but also adhere to the fundamental principles governing the system. This approach is particularly advantageous in scenarios with limited or noisy data, as it allows the network to leverage known physics to guide the learning process.

### 3.2.5 Loss Function

The loss function in neural networks quantifies the difference between the predicted outputs and the actual target values, guiding the optimization of network parameters. In the context of PINNs, the loss function is designed to balance data fidelity with adherence to physical laws.

- **Data Fidelity Term** The data fidelity term ( $\mathcal{L}_{\text{data}}$ ) measures how well the network’s predictions fit the observed data. For example, if using MSE (Mean Squared Error), it’s formulated as:

$$\mathcal{L}_{\text{fidelity}} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (3.5)$$

where,  $y_i$  is the observed value,  $\hat{y}_i$  is the predicted value, and  $n$  is the number of data points.

- **Physics Regularization Term** The physics-based regularization term ( $\mathcal{L}_{\text{physics}}$ ) ensures that the network’s predictions satisfy the governing physical laws. For a physical system described by a differential equation  $\mathcal{F}(\hat{y}) = 0$ , this term can be expressed as:

$$\mathcal{L}_{\text{physics}} = \frac{1}{m} \sum_{j=1}^m [\mathcal{F}(\hat{y}(x_j))]^2 \quad (3.6)$$

where,  $x_j$  are spatio-temporal co-location points at which the differential equation must be satisfied, and  $m$  is the number of such points.

- **Total Loss Function** The total loss function in a PINN combines both terms, typically with weighting coefficients to balance their contributions:

$$\mathcal{L}_{\text{total}} = \alpha_1 \cdot \mathcal{L}_{\text{data}} + \alpha_2 \cdot \mathcal{L}_{\text{physics}} \quad (3.7)$$

where,  $\alpha_1$  and  $\alpha_2$  are the loss terms weighting coefficients that control the relative importance of each term. By minimizing  $\mathcal{L}_{\text{total}}$ , the network learns to produce outputs that not only fit the observed data but also satisfy the underlying physical laws.

### 3.2.6 Optimization Process

Optimization algorithms play a crucial role in minimizing the loss function. Among the most renowned optimization algorithms is ADAM, as described in Kingma and Ba’s work [85]. The ADAM optimizer represents a hybrid of the *Momentum* optimization algorithm [37] and the Root Mean Square Propagation (*RMSprop*) method [170]. *Momentum* accelerates the gradient descent process by leveraging past gradients to refine updates. *RMSprop*, on the other hand, dynamically adjusts the learning rate for each parameter by dividing the learning rate by an exponentially decaying average of squared gradients, facilitating adaptive learning rate modifications. ADAM synthesizes these methodologies effectively. It maintains a moving average of both the gradients, as in *Momentum*, and the squared gradients, as in *RMSprop*. This dual approach enables ADAM to fine-tune the learning rate per parameter, incorporating the concept of momentum [85].

The moving average of the gradients  $m_t$  and the moving average of the squared gradients  $v_t$  can be formulated as:

$$m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot gr_t \quad (3.8)$$

$$v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot gr_t^2 \quad (3.9)$$

where,  $gr_t$  is the gradient at optimization time step  $t$ , and  $\beta_1$  and  $\beta_2$  are hyperparameters controlling the exponential decay rates of these moving averages. Since  $m_t$  and  $v_t$  are initialized as vectors of zeros, they are biased towards zero, especially during the initial time steps and when the decay rates are small. ADAM performs bias correction to counteract this:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} \quad (3.10)$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t} \quad (3.11)$$

The final step of this process is the update of the model parameters, including weights and biases in a neural network, guided by these refined and balanced gradients, expressed as:

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t \quad (3.12)$$

where  $\theta_t$  are the network parameters at optimization time step  $t$ ,  $\eta$  is the learning rate, and  $\epsilon$  is a small number added to prevent division by zero.

### 3.2.7 Back-Propagation

Back-propagation is a method used for calculating the gradient of the loss function with respect to each weight in the network. It effectively involves performing a forward pass to compute outputs and errors, and then a backward pass to compute gradients. The gradient with respect to a weight in a simple network can be computed as:

$$\frac{\partial \mathcal{L}_{total}}{\partial \theta} = \frac{\partial \mathcal{L}_{total}}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial \theta} \quad (3.13)$$

here,  $\mathcal{L}_{total}$  is the total loss function, and  $\hat{y}$  is the output of the neural network.

### 3.3 Input Parameterized Physics Informed Neural Net (IP-PINN)

The IP-PINN framework seeks to super-resolve time-resolved PC-MRI images and simultaneously address issues of acquisition noise, and artifacts due to velocity aliasing and phase offset errors. It is a generalization of PINNs to avoid costly ab-initio training for unseen data sets. As the name implies, IP-PINN, enables generalization of PINNs by parameterizing the output with the input that is representative of the data being processed which in this case is a spatio-temporal region of interest (ROI) in the velocity encoded scans in complex-Cartesian image space ( $\mathbf{x}$ -space). Fig. 3.1 shows the operation of the proposed method. The user selects a ROI in the flow region. The ROI may have vessel boundaries and encompasses a volume of  $n_x \times n_y \times n_z$  voxels. The IP-PINN takes as input the data within the ROI comprising of the real and imaginary parts of the 3 velocity encoded scans in  $\mathbf{x}$ -space over a temporal region given by  $n_t$  time steps centered around the cardiac phase  $t = t_l$  along with the spatio-temporal coordinate  $(\mathbf{x}, z, t)$  and generates point-wise velocities  $(u, v, w)$ , relative pressure  $p$ , and magnitude images  $M$ .

Fig. 3.2 illustrates the architecture of the IP-PINN. It has two distinct parts:

- A flow image encoder that takes in as input the complex data from a spatio-temporal ROI in the three flow encoded scans and produces a latent vector  $\mathbf{L}$  that captures the important features in the flow images.
- An augmented PINN which takes in as input the latent vector generated by the flow encoder and the spatio-temporal coordinate to generate flow velocities, pressure, image magnitude, and direction-wise composite phases.

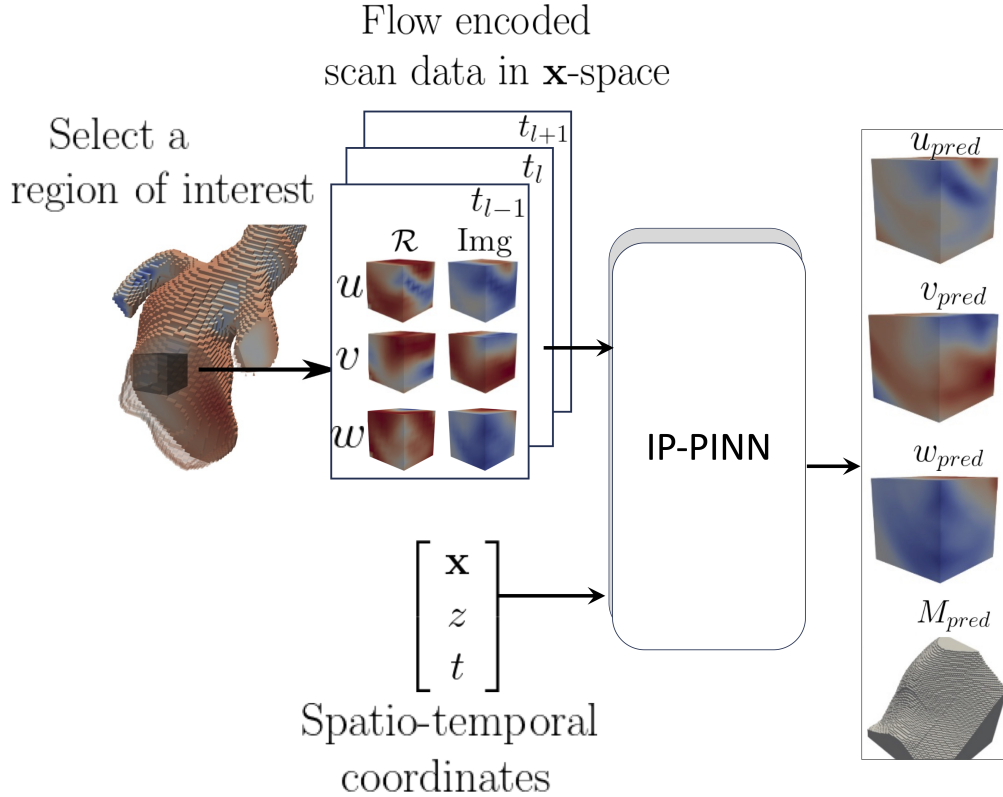


Figure 3.1: IP-PINN usage: The user selects a spatio-temporal region of interest (ROI) in the flow region to process. The input to the trained IP-PINN is the image data consisting of the real and imaginary parts of the three flow-encoded scans in  $\mathbf{x}$ -space and the spatio-temporal query point  $(\mathbf{x}, z, t)$ . The IP-PINN predicts the underlying point-wise 3D velocity  $(u_{pred}, v_{pred}, w_{pred})$ , relative pressure  $p$ , and image magnitude  $M_{pred}$  which can be segmented to obtain the luminal boundary.

### 3.3.1 Flow Image Encoder

The flow encoder enables the parametrization of the PINN on the input flow encoded scan data. Consequently, unlike typical PINNs, it allows the evaluation of flow velocities for different input flow encoded images. The encoder is given by

$$\mathbf{L} = \mathcal{G}_{\Theta_1}(\mathcal{U}) \quad (3.14)$$

where  $\mathcal{U} \in \mathbb{C}^{n_x \times n_y \times n_z \times n_t \times 3}$  is the input data from the velocity encoded scans in  $\mathbf{x}$ -space within the selected ROI, and  $\mathbf{L} \in \mathbb{R}^m$  is the output latent vector. As shown in Fig. 3.2, the structure of  $\mathcal{G}_{\Theta_1}(\cdot)$  is inspired by the contracting path of the U-Net architecture [149]. The

input data to the encoder is arranged as a 4-dimensional tensor with dimensions  $n_x \times n_y \times n_z \times n_t$ , where the channels  $c$  represent the real and imaginary parts of the flow encoded scans in the  $x, y, z$  directions. The encoder architecture has five different layers:

- The first layer is a pure 3D ResNet [71] block applied to the spatial dimensions alone with 16 filters.
- The second and third layers are 3D ResNet blocks followed by Max Pool, once again applied to the spatial dimensions alone with 64 and 128 channels, respectively.
- The fourth is a pure 3D ResNet Block with 128 channels. Layers 1-4 have padding equal to 1.
- The fifth layer includes a flattening of the spatio-temporal dimensions followed by a 1D convolution followed by Max Pool with 256 filters.
- Finally, the latent vector is a simple flattening of the output of the fifth layer.

By employing the robust feature extraction capabilities of 3D convolutional ResNet blocks, this component efficiently captures hierarchical spatio-temporal information within the input data block  $\mathcal{U}$ .

### 3.3.2 Augmented PINN

The latent vector  $\mathbf{L}$  along with the normalized spatio-temporal coordinates is the input to two fully connected nets (FCNs). These two FCNs are given by

$$\begin{bmatrix} u^* \\ v^* \\ w^* \\ p^* \\ M \end{bmatrix} = \mathcal{M}_{\Theta_2}(\mathbf{L}, \mathbf{x}^*, z^*, t^*) \quad \begin{bmatrix} \phi_B^x \\ \phi_B^y \\ \phi_B^z \end{bmatrix} = \mathcal{M}_{\Theta_3}(\mathbf{L}, \mathbf{x}^*, z^*, t^*) \quad (3.15)$$

where

$$\begin{aligned} \mathbf{x}^* &= \frac{\mathbf{x} - \mathbf{x}_{min}}{L} & z^* &= \frac{z - z_{min}}{L} & t^* &= \frac{t}{T} = \frac{tU}{L} \\ u^* &= \frac{u}{U} & v^* &= \frac{v}{U} & w^* &= \frac{w}{U} & p^* &= \frac{p}{\rho U^2} \end{aligned} \quad (3.16)$$

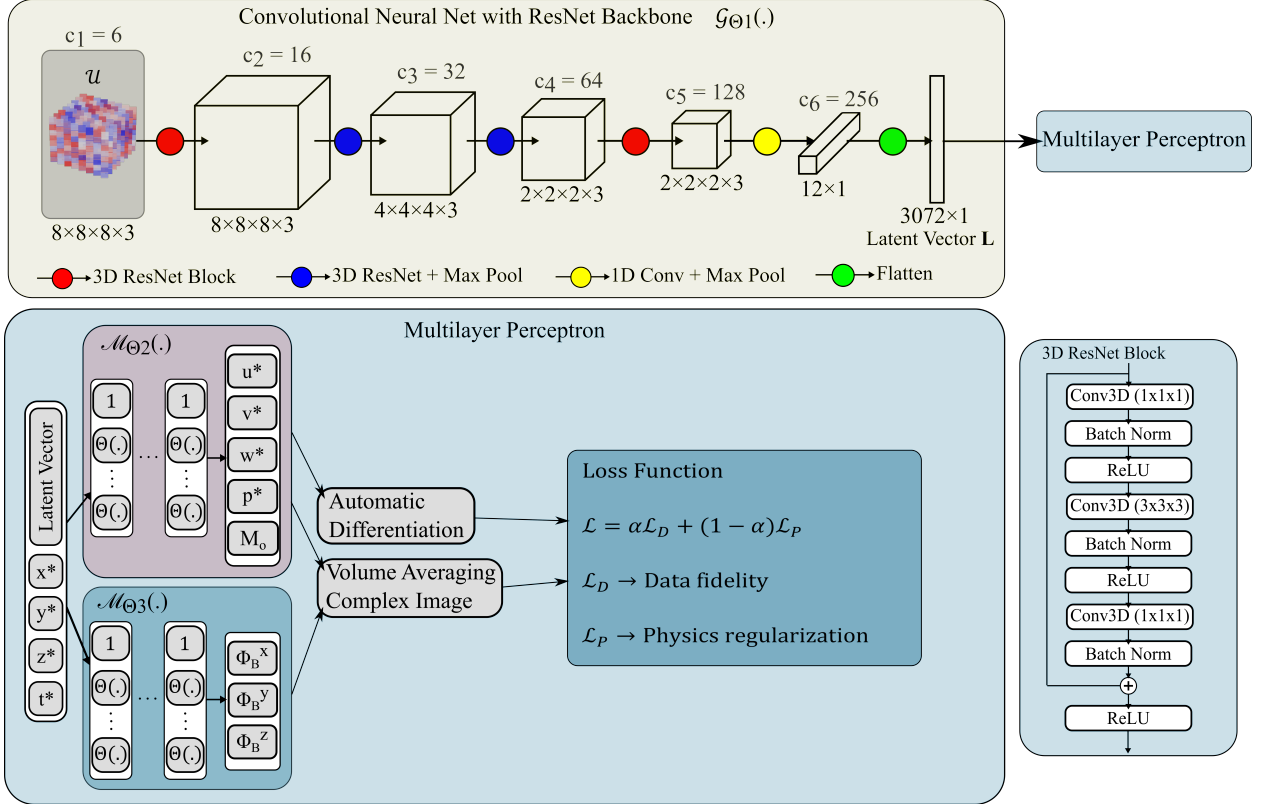


Figure 3.2: Schematic representation of the proposed IP-PINN architecture. A section  $\mathcal{U}$  of the three velocity encoded complex Caterisian images is processed by a convolutional neural network consisting of ResNet encoder blocks to generate a latent vector  $\mathbf{L}$ . The secondary components (Multilayer Perceptron) take in the latent vector and local spatio-temporal coordinates to approximate a continuous function mapping, resulting in normalized flow velocity components, pressure, magnitude, and directional phase offsets.

Here  $(\mathbf{x}^*, z^*, t^*)$  are the normalized spatio-temporal coordinates,  $u^*, v^*, w^*, p^*$  are the non-dimensional velocity components and pressure,  $U, L, T$  represent the characteristic velocity, length, and time respectively, and  $(\mathbf{x}_{min}, z_{min})$  represents the coordinates of the lower left corner of the ROI. The FCN  $\mathcal{M}_{\Theta_2}(\cdot)$  is equipped with 60 neurons in each layer with six layers. The FCN  $\mathcal{M}_{\Theta_3}$  has 48 neurons per layer for its first three layers followed by 32 neurons

for its remaining three layers. We incorporate neuron-wise Sigmoid Linear Unit adaptive activation function (AdaptiveSiLU), which enhances the flexibility of the traditional Sigmoid Linear Unit activation by enabling each neuron to independently optimize the slope of its activation function during training. Partial derivatives with respect to the spatiotemporal coordinates required to impose physics constraints can be computed through automatic differentiation. The Adam Optimizer was used with an initial learning rate of  $10^{-3}$  which decayed exponentially to the final value of  $10^{-8}$ , employing a ‘ReduceLROnPlateau’ learning rate scheduler.

Unlike traditional PINNs [144, 143, 52], the proposed architecture enables pre-training through parametrization with respect to the input image data from the ROI. Simultaneously, it retains the ability to adhere to governing physics constraints, just like PINNs. Since  $\mathcal{M}_{\Theta_2}(\cdot)$  is a continuous function with respect to the normalized spatiotemporal coordinate, the super-resolution ratio is not fixed, unlike other works that use pure CNNs [55, 151]. The neural network parameters  $\Theta_1$ ,  $\Theta_2$ , and  $\Theta_3$  are trained by minimizing a loss function that will be described below. To ensure that the IP-PINN framework performs optimally across different scenarios, a carefully designed loss function is employed. This function balances data fidelity with adherence to the governing physical laws, enabling the network to achieve high accuracy while respecting the underlying physics of blood flow.

### 3.3.3 Loss Function

The loss function used in this work is given by

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_D + \alpha\mathcal{L}_P \quad (3.17)$$

here  $\mathcal{L}_D$  is the data fidelity term, which minimize difference between the neural net output and data, and  $\mathcal{L}_P$  is the physics regularization term.  $\alpha$  is a hyperparameter that controls the relative strength of the two loss terms and is currently tuned manually for convergence.

Data Fidelity: Typical deep learning-based methods for supresolving 4D-Flow MRI ac-

compish data fidelity by minimizing the error between predicted velocities and velocities obtained by processing the phase images of 4D-Flow MRI. This has several drawbacks: first, velocity aliasing due to low  $v_{enc}$  setting can cause discontinuities in the velocity data which cannot be represented by the neural net. Second, the noise in the velocity is non Gaussian because of the nonlinear operation that must be used to compute the phase from  $\mathbf{x}$ -space images. Finally, the composite phase offset errors must also be handled.

To avoid the problem of velocity aliasing, data fidelity is conducted by computing the predicted flow encoded image in the  $\mathbf{x}$ -space. The neural net  $\mathcal{M}_{\Theta_2}(\cdot)$  predicts non-dimensionalized velocities  $u_{pred}^*, v_{pred}^*, w_{pred}^*$ , pressure  $p_{pred}^*$ , and image magnitude  $M_{pred}$ . Using Eq. 3.16, the predicted velocities are computed as  $u_{pred} = Uu_{pred}^*, v_{pred} = Uv_{pred}^*, w_{pred} = Uw_{pred}^*$ . These velocities are used to calculate the predicted velocity induced phases, given the  $v_{enc}$  parameter using Eqns 2.4 and 2.6 as

$$\phi_{\mathbf{u}_{pred}}^j(\mathbf{x}, z, t) = \frac{\pi}{v_{enc}} \begin{bmatrix} u_{pred} \\ v_{pred} \\ w_{pred} \end{bmatrix} \cdot \hat{\mathbf{g}}^j \quad (3.18)$$

The predicted apparent spin density  $\sigma^j(\mathbf{x}, z, t)_{pred}$  in each encoding direction is given by Eqn. 2.8. Note that the predicted composite phases  $\phi_B^j(\cdot)_{pred}$ , which account for eddy current induced phase offsets, are outputs of neural net  $\mathcal{M}_{\Theta_3}(\cdot)$  at the same spatio-temporal location. We then approximate the predicted discretized spin density at the spatio temporal location  $(\mathbf{x}_k, z_k, t_l)$  by spatio-temporally averaging the prediction over the voxel as

$$\sigma_D^j(\mathbf{x}_k, z_k, t_l)_{pred} = \frac{1}{\Delta t \Delta V} \iiint \sigma^j(\mathbf{x}, z, t)_{pred} dV dt \quad (3.19)$$

where  $\Delta t$  is temporal length of the cardiac phase, and  $\Delta V$  is the volume of the voxel. We employ 8 point tensor product Gaussian quadrature to numerically evaluate Eq. 3.19. Finally,

we have the data loss term given by

$$\mathcal{L}_D = \text{mean} \left[ \sum_{l=1}^{n_t} \sum_{k=1}^{n_x \times n_y \times n_z} \sum_{j=x,y,z} \|\sigma_D^i(\mathbf{x}_k, z_k, t_l)_{pred} - C_D^j(\mathbf{x}_k, z_k, t_l)\|^2 \right] \quad (3.20)$$

Physics Regularization: Assuming blood to be an incompressible fluid, imposes the governing equations of incompressible fluid flow on the neural net  $\mathcal{M}_{\theta_2}(\cdot)$  which generates as outputs the non-dimensionalized velocities. The residual terms that constitute the regularization loss term are therefore

$$\mathcal{E}_1(\mathbf{x}_k, z_k, t_k) = \left\| \frac{\partial \mathbf{u}^*}{\partial t^*} + (\mathbf{u}^* \cdot \nabla^*) \mathbf{u}^* + \nabla^* p^* - \frac{1}{\text{Re}} \nabla^* \cdot (\nabla^* \mathbf{u}^*) \right\|^2 \quad (3.21)$$

$$\mathcal{E}_2(\mathbf{x}_k, z_k, t_k) = \|\nabla^* \cdot \mathbf{u}^*\|^2$$

Here  $\text{Re} = \frac{\rho UL}{\mu}$  is the Reynolds number, where  $\rho$  and  $\mu$  are the blood density and viscosity, respectively,  $\mathbf{u}^*(\mathbf{x}^*, z^*, t^*)$  is the non-dimensional 3D velocity,  $p(\mathbf{x}^*, z^*, t^*)$  is the pressure, and  $\nabla^*$  is the non-dimensional 3D spatial derivative operator. The residual terms  $\mathcal{E}_1, \mathcal{E}_2$  are evaluated at  $n_P$  collocation/physics regularization points placed randomly in the spatio-temporal domain. The physics loss is therefore given by

$$\mathcal{L}_P = \text{mean} \sum_{k=1}^{n_P} \mathcal{E}_1(\mathbf{x}_k, z_k, t_k) + \mathcal{E}_2(\mathbf{x}_k, z_k, t_k) \quad (3.22)$$

## 3.4 Conclusion

In this chapter, we introduced the Input-Parameterized Physics-Informed Neural Network (IP-PINN), a novel framework designed to overcome the limitations of existing hemodynamic assessment methods, including 4D Flow MRI and patient-specific CFD simulations. By integrating CNNs and MLPs with physics-informed constraints, IP-PINN effectively generates super-resolved, noise-free flow fields directly from MRI images, while maintaining adherence

to fundamental fluid dynamics principles. The IP-PINN architecture leverages a CNN-based encoder to extract features from input data, and an augmented PINN component to provide continuous spatio-temporal field representations. This design allows IP-PINN to generalize across diverse datasets, accommodating variations in anatomy and flow conditions without extensive retraining, making it a robust tool for clinical applications.

The subsequent chapters translate the algorithmic blueprint laid down here into three progressively more data-efficient MRI scenarios, with the network architecture itself slightly altered to match the dimensionality of each application, while preserving the rapid ( $< 2$  minutes) fine-tuning time on unseen data. Chapter 4 deploys the original, image-space IP-PINN on conventional 4D-Flow datasets that retain only the three velocity-encoded volumes; after a single synthetic pre-training run, a 90 seconds fine-tune delivers super-resolved, noise-free velocity and magnitude fields that correct aliasing and phase offsets. Chapter 5 adapts the framework to a lightly modified 2D PC-MRI sequence in which the encoding direction is cycled across cardiac phases: because each input now consists of three complex images on a single 2D plane, the ROI dimensions is changed to  $n_x \times n_y \times n_z \times n_t = 16 \times 16 \times 1 \times 3$  voxels, enabling the network to lift a volumetric, three-component flow field, produce high-resolution spin-density maps, and support reliable wall-shear-stress estimation, all without a reference acquisition. Chapter 6 pushes sparsity to its limit with a *pseudo one-point* 4D-Flow protocol that acquires only one velocity-encoded dataset per slice. In this chapter, the ROI size is changed to  $n_x \times n_y \times n_z \times n_t = 14 \times 14 \times 6 \times 3$ , the number of feature maps in each ResNet block is altered, and a new k-space data-fidelity term is introduced. Even with a data load 75% smaller than the standard four-point scheme, IP-PINN reconstructs accurate three-component velocities and spin-density images. Collectively, these implementations confirm that the framework developed in this chapter is a robust, architecture-adaptable engine for recovering hemodynamically meaningful information from a wide spectrum of accelerated phase-contrast MRI acquisitions.

## Chapter 4

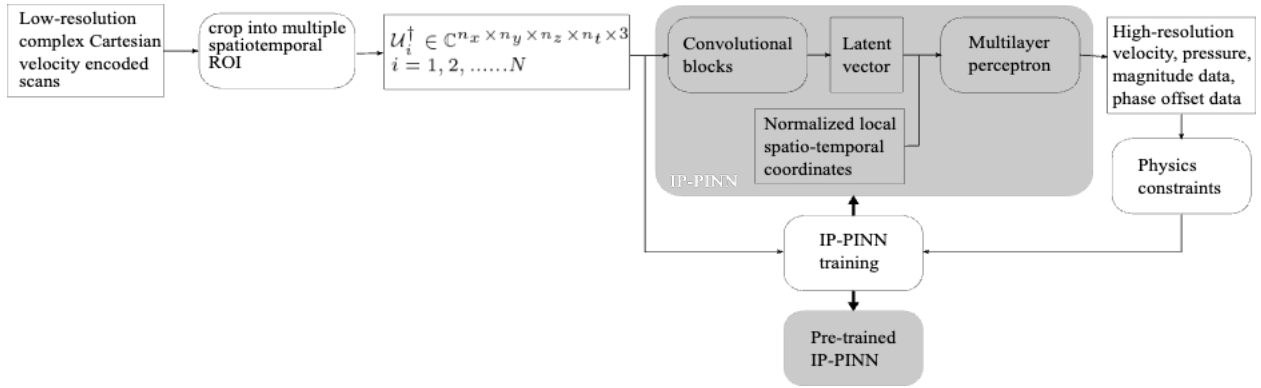
# Application of IP-PINN for Advanced 4D Flow MRI

### 4.1 Introduction

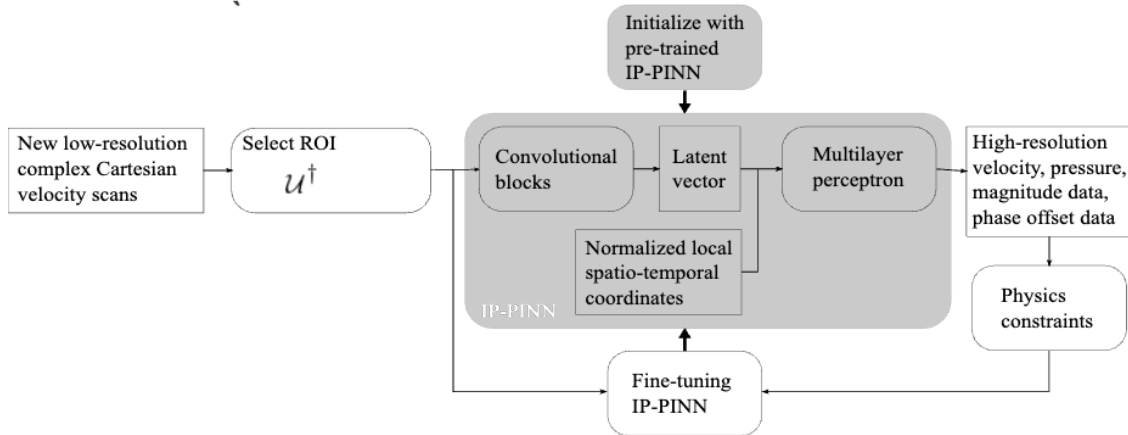
This chapter presents the outcomes derived from the implementation of the proposed Input-Parameterized Physics-Informed Neural Network (IP-PINN) framework on patient-specific 4D Flow MRI datasets. Building upon the methodological foundation established in Chapters 2 and 3, where the acquisition and processing of 4D Flow MRI data and the development of the IP-PINN were detailed, this chapter focuses on the application and evaluation of these techniques. The results demonstrate the effectiveness of the IP-PINN in super-resolving MRI images, mitigating noise, and correcting artifacts such as velocity aliasing and phase offsets.

Figure 4.1 illustrates the workflow of the proposed IP-PINN algorithm, divided into the pre-training phase and the fine-tuning phase. In the pre-training phase, a large set of time-resolved phase-contrast MRI images, either simulated or actual, is used to train the network to satisfy both data fidelity and physics constraints. The pre-training and fine-tuning phases do not require ground truth data, as the data fidelity term is formulated in the complex Cartesian image space. The pre-training is conducted using data from a single aneurysm, setting the neural network parameters  $\Theta_1$ ,  $\Theta_2$ , and  $\Theta_3$ .

In the fine-tuning/use phase, the pre-trained network is applied to unseen data, where a forward evaluation is conducted using the latent vector and desired spatiotemporal coordinates. Purely relying on the forward evaluation of the neural net will not generate accurate results. Consequently, the results are refined through a brief mini-training session, leveraging transfer learning [4, 31] to achieve high accuracy (this is illustrated in our results) (Fig. 4.1 (b)). This process ensures that the algorithm generalizes well to new vascular conditions,



(a) Pre-training phase



(b) Fine-tuning/use phase, for new unseen data

Figure 4.1: Schematic representation of the proposed algorithm’s workflow. (a) Pre-training phase. (b) Fine-tuning phase.

enabling efficient and precise analysis of 4D Flow MRI data. While the pre-training phase typically takes roughly 20-30 hrs, the mini-training session typically lasts about 1.5 minutes. Moreover, several ROIs can be simultaneously fed into the same mini-training session to reduce the time for analysis.

## 4.2 Generation of Synthetic Training and Test Data

Two distinct 3D unsteady CFD simulations were conducted for Posterior Cerebral Artery (PCA) and Middle Cerebral Artery (MCA) aneurysms, as depicted in Fig. 4.2. Synthetic flow image scan data derived from the PCA simulation was used for training, while MCA data was used for testing. SimVascular was the numerical solver for both simulations. For PCA, a

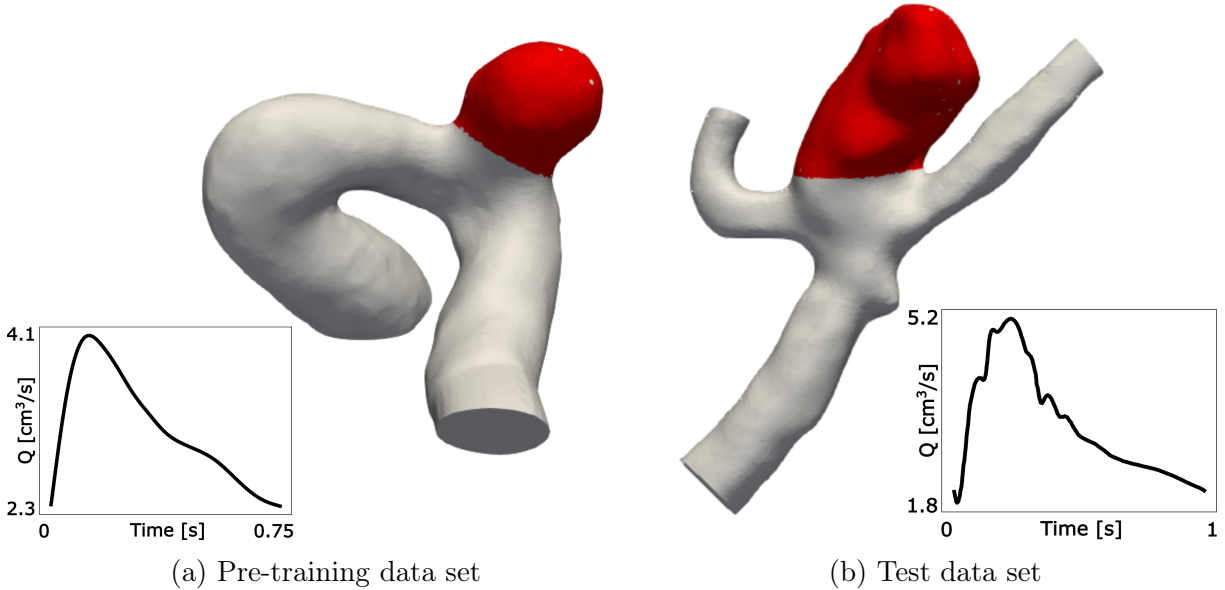


Figure 4.2: Visualization of aneurysm geometries for model pre-training and testing. (a) Posterior cerebral artery (PCA) aneurysm geometry is used for generating pre-training dataset. (b) MCA aneurysm geometry is used for generating test dataset. The regions highlighted in red represent the volume from which flow velocities were used for training and testing.

12.5M-element mesh with rigid vessel walls and Newtonian blood rheology was utilized. Inlet conditions were derived from in-vivo MRI data, with a parabolic profile. Outlet conditions featured a resistance with a reference pressure of 100 mmHg. MCA geometry came from the Anuerisk dataset [154], with a 7.3M-element mesh. A pulsatile inlet flow rate waveform was used, and Murray’s law was applied at outlets [145]. Both simulations were run for three cardiac cycles to settle out the transients due to the default initial conditions used and data was collected from the fourth cardiac cycle simulation. Table 4.1 listed the parameters used in the two simulations.

	Cardiac Cycle Period (s)	Simulation $\Delta t$ (s)	Peak Re	# Snapshots Collected
PCA	0.75	$7.5 \cdot 10^{-5}$	300	50
MCA	1	$10^{-4}$	640	1000

Table 4.1: Parameters used for simulating blood flow in PCA and MCA arteries to generate synthetic velocity encoded images. Here the blood dynamic viscosity  $\mu = 0.004 \text{ Kg/ms}$  and density  $\rho = 1000 \text{ Kg/m}^3$

### 4.2.1 Generating Synthetic 4D-Flow MRI from CFD Simulations

The CFD simulation data from both the PCA and MCA aneurysms was interpolated onto a fine uniform grid with isotropic grid size  $\Delta x_f = 0.35mm$ . For the MCA data set, every 15<sup>th</sup> snapshot was selected to set the time resolution to 15ms which is the same as the PCA dataset. This represented the reference data set/ground truth against which we compared the results of our algorithm.

Velocity in the voxels outside the geometry was set to zero. The velocities were converted to  $\mathbf{x}$ -space representation according to Eqn. 2.8. The  $v_{enc}$  parameter was set to be 20% lower than the maximum directional velocity to induce phase wrapping (velocity-aliasing) artifacts. The composite phase  $\phi_B(\mathbf{x}, z)$  was set to be a linear function for the PCA data set, and quadratic function for the MCA dataset. This was done to simulate eddy current induced phase offset in the data [24, 35, 43]. The image magnitude  $M(\mathbf{x}, z)$  was set to for the flow region inside the lumen and outside to represent different proton densities in moving tissue and static tissue.

The data set was spatio-temporally down sampled by first spatial averaging over 3 voxels in the  $z$  direction (resulting in a 1.05 mm voxel in  $z$  direction) followed by temporal averaging over 4 time steps to reduce the temporal resolution to 60ms. Each averaged slice in the  $\mathbf{x}$ -space was transformed to  $\mathbf{k}$ -space using FFT. Down sampling in the coronal plane (see Fig. 1) was done by cutting off the high spatial frequencies by windowing at  $\pm k_{max}$  where  $k_{max} = (2(3)\Delta x_f)^{-1} = 0.667m^{-1}$ . Spatial downsampling resulted in a coarse resolution of 1.05mm. Acquisition noise was simulated by adding zero mean Gaussian noise of 20dB SNR to the real and imaginary channels in  $\mathbf{k}$ -space. This  $\mathbf{k}$ -space image was transformed back to  $\mathbf{x}$ -space using an  $FFT^{-1}$  operation. As illustrated in Fig. 4.3, this operation generated the synthetic 4D-Flow MRI data set.

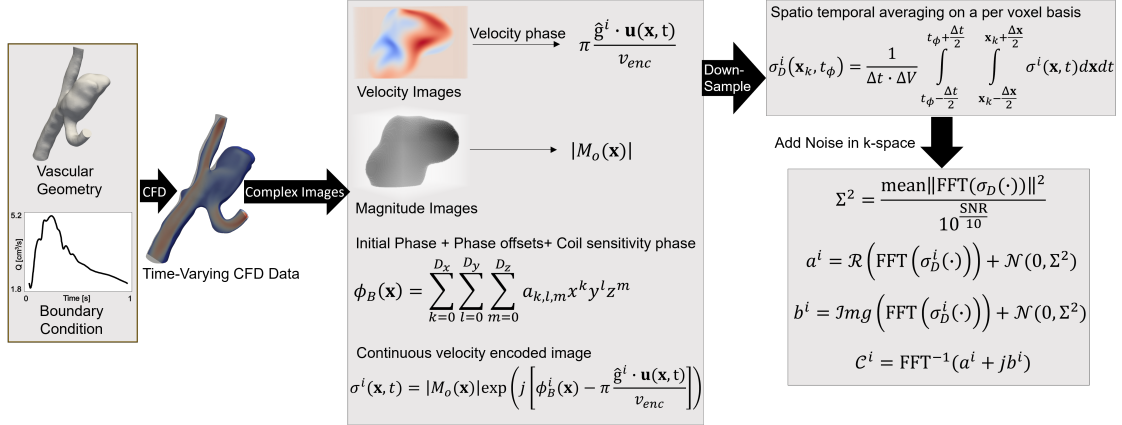


Figure 4.3: Process of generating synthetic 4D-Flow MRI with noise and phase offset artifacts from CFD simulations.

### 4.3 Error Metrics

We employed a comprehensive set of error metrics between the high resolution ground truth velocity  $\mathbf{u}_{GT}$  obtained from CFD simulation results interpolated to the fine grid and the predicted velocity  $\mathbf{u}_{pred}$  from the IP-PINN. These metrics are defined as below:

1. Total relative percent error (TRPE)

$$\text{TRPE} = \frac{\|\mathbf{X}_{pred} - \mathbf{X}_{GT}\|}{\|\mathbf{X}_{GT}\|} \quad (4.1)$$

where  $\mathbf{X} = \begin{bmatrix} \mathbf{U} & \mathbf{V} & \mathbf{W} \end{bmatrix}^T$  are vectorized component-wise velocities in the region of interest

2. Plot of point-wise direction/cosine error given by

$$\text{COSE}_i = 1 - \frac{\mathbf{u}_{pred}^i \cdot \mathbf{u}_{GT}^i}{\|\mathbf{u}_{pred}^i\| \|\mathbf{u}_{GT}^i\|} \quad (4.2)$$

3. Point-wise absolute error in x, y, z velocity directions

$$u_e = \|\mathbf{u}_{pred} - \mathbf{u}_{GT}\| \quad (4.3)$$

$$v_e = \|v_{pred} - v_{GT}\| \quad (4.4)$$

$$w_e = \|w_{pred} - w_{GT}\| \quad (4.5)$$

4. Scatter plot between  $\|\mathbf{u}_{GT}^i\|$  and  $\|\mathbf{u}_{pred}^i\|$
5. Coefficient of determination  $R^2$  of velocity magnitude given by

$$R^2 = 1 - \frac{\text{RSS}}{\text{TSS}} \quad (4.6)$$

where

$$\text{RSS} = \sum_{i=1}^N \|\mathbf{u}_{pred}^i - \mathbf{u}_{GT}^i\|^2$$

$$\text{TSS} = \sum_{i=1}^N \|\mathbf{u}_{GT}^i - \bar{\mathbf{u}}\|^2$$

here,  $\bar{\mathbf{u}}$  is the mean ground truth velocity magnitude

6. Median absolute percentage error (MDAPE) given by

$$\text{MDAPE}_x = \text{median} \left( \frac{|x_{pred}^i - x_{GT}^i|}{\|x_{GT}^i\|} \right) \times 100 \quad (4.7)$$

where  $x = u, v, w$

## 4.4 Pre-Training IP-PINN

The IP-PINN model is pre-trained with input data blocks of size  $n_x \times n_y \times n_z \times n_t$ , from the low-resolution image data generated from the flow simulation of the training aneurysm shown in Fig. 4.2 (a). We used roughly 11000 data blocks and 3000 collocation points per block for physics regularization. The normalization factors were  $L = 0.8$  cm,  $U = 93.6494$  cm/s, and  $T = 0.0085$  s. In the pre-training phase, we employed a batch size consisting of two input data blocks and completed a total of 250 epochs. At the end of training, a sample of

the final values of the loss terms, along with the empirically determined loss term weights for the pre-training phase when the ROI was of size  $n_x = 8, n_y = 8, n_z = 8, n_t = 3$ , corresponding to an image resolution of 1.0 mm, is shown in Table 4.2.

Loss term	Weight	Loss value
Physics Loss ( $\mathcal{L}_P$ )	0.0009	$3.972 \times 10^{-6}$
Data Fidelity Loss ( $\mathcal{L}_D$ )	0.9991	$8.573 \times 10^{-5}$
Total Loss (weighted sum)		$8.970 \times 10^{-5}$

Table 4.2: The loss terms of pre-trained IP-PINN for the unsteady posterior cerebral artery aneurysm (PCA) case and their empirically chosen weights.

The successful pre-training of the IP-PINN using the synthetic PCA dataset established a strong foundation for further evaluation. The network’s ability to generalize was then tested by applying it to unseen data, specifically from the MCA dataset, to assess its performance under new and varied conditions.

## 4.5 Applying Pre-Trained IP-PINN to Unseen Data

Following pre-training with simulated 4D-Flow MRI data derived from CFD simulation on the PCA, we assessed the performance using simulated 4D-Flow MRI data derived from CFD simulation on the MCA (see Fig. 4.2). We also assessed the robustness of IP-PINN to noise in the input by adding noise with 20dB SNR in  $\mathbf{k}$ -space. The unseen MCA data set required ‘fine tuning’ training, wherein we repeat the training process with the unseen input data set. The execution time therefore is the time required for the fine tuning. Fig. 4.4 presents the results of super-resolving noiseless (a) and noisy (b) inputs. The first two columns show the real and imaginary channels of the velocity encoded scans in  $\mathbf{x}$ -space which form the input to the IP-PINN. The third column shows the low resolution velocities obtained using conventional processing. The fourth column represents prediction of the IP-PINN. The

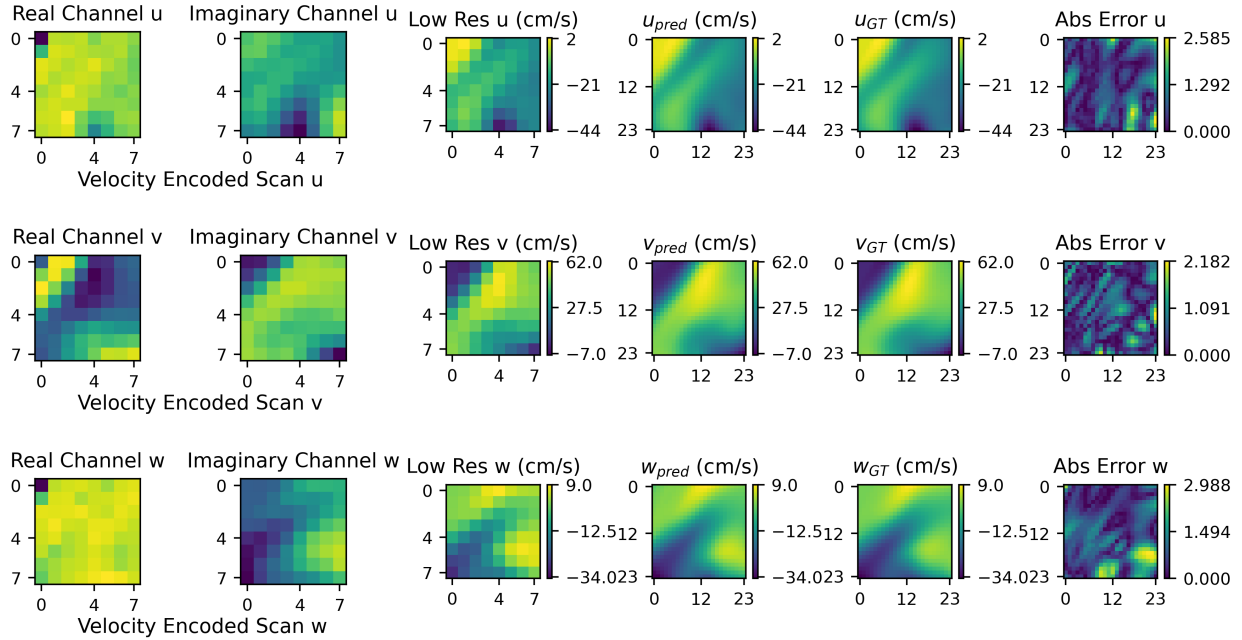
fifth column represents the ground truth. The final column represents the absolute error as defined in the error metrics. The MDAPE ranges from 3.7% to 6.5% for noise free input and 6.2% to 8.4% for input with 20dB SNR. Fig. 4.5 shows the scatter plot and point-wise cosine error plot for both noise-free input and input with 20dB SNR noise. Low magnitude error will show low deviation from the  $45^\circ$  line in the scatter plot. Low direction error will show small cosine error. It can clearly be seen that IP-PINN predictions are accurate in magnitude and direction. Table 4.3 lists the total relative percent error ( $TRPE$ ), the coefficient of determination ( $R^2$ ), and the MDAPE in  $u$ ,  $v$ , and  $w$ , corresponding to the noiseless and noisy conditions. The performance of the algorithm degrades with noise.

$SNR$	$TRPE$	$R^2$	$MDAPE_u$	$MDAPE_v$	$MDAPE_w$	Processing Time
Noiseless	3.65%	0.997	4.068	1.343	6.764	1.5 minutes
20 <i>dB</i>	5.40%	0.992	6.178	2.267	8.360	1.5 minutes

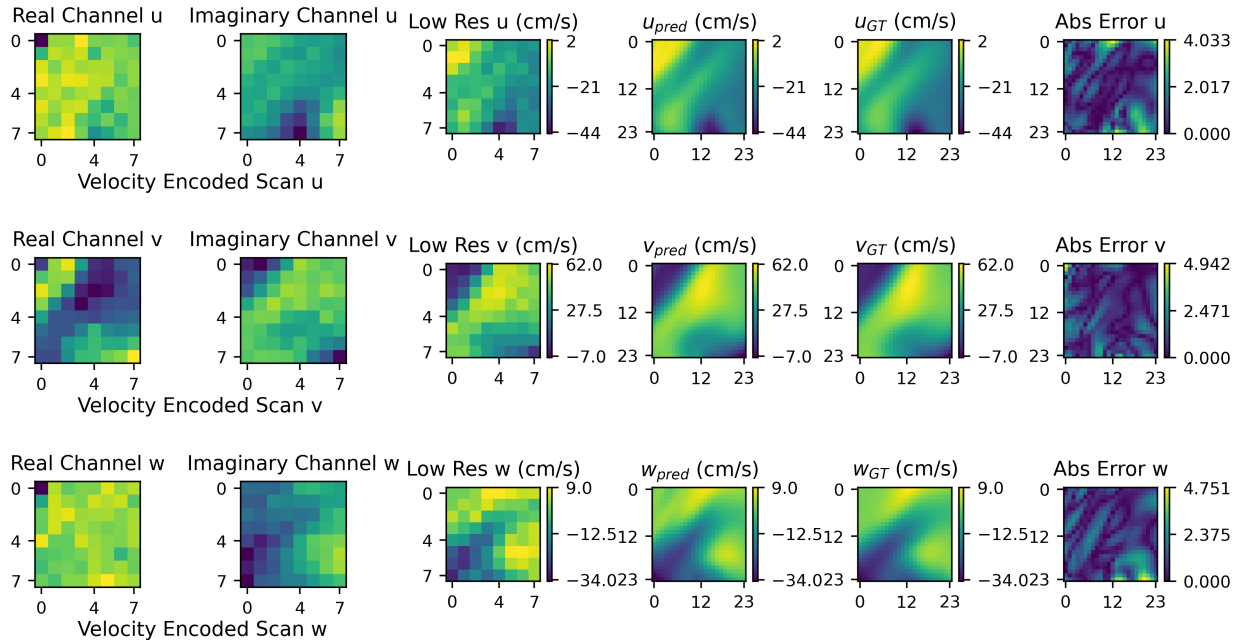
Table 4.3: A summary of total relative percent error ( $TRPE$ ), coefficient of determination ( $R^2$ ), and median absolute percentage error of  $u^*$ ,  $v^*$ , and  $w^*$  for processing the noiseless input data and input data with SNR of 20 dB. IP-PINN required 1.5 minutes for processing the input data

### 4.5.1 Time Interpolation

In this section, we assess the ability of the IP-PINN algorithm to super resolve the flow in time. Once the IP-PINN undergoes fine tuning, we evaluated the neural net at 4x the time frequency of the simulated 4D-Flow MRI. Fig. 4.6 shows the results. Here, low-resolution synthetic 4D-Flow MRI velocity encoded scans that are processed to obtain the  $u$ ,  $v$ ,  $w$  velocities are available only at  $t_o+2\Delta t$ ,  $t_o+6\Delta t$ ,  $t_o+10\Delta t$ ,  $t_o+14\Delta t$ ,  $t_o+18\Delta t$ , and  $t_o+22\Delta t$  where  $t_o$  is the starting time of the block of data being processed, and  $\Delta = 10ms$  is sampling period of the predicted solution. Columns  $u_{pred}, v_{pred}, w_{pred}$  represent the velocity predictions of the IP-PINN in high resolution. Columns  $u_{GT}, v_{GT}, w_{GT}$  represent the ground truth velocity from which the synthetic 4D-Flow MRI data were derived. It can be seen clearly, that the predictions closely match the ground truth. In this particular case, the time resolution of



(a)



(b)

Figure 4.4: Results showcase the impact of varying noise levels in the region of interest: (a) noiseless, median absolute percentage errors stand at  $MDAPE_u = 3.747$ ,  $MDAPE_v = 1.222$ , and  $MDAPE_w = 6.522$ ; (b) with an SNR of 20 dB, the corresponding median absolute percentage errors are  $MDAPE_u = 6.178$ ,  $MDAPE_v = 2.267$ , and  $MDAPE_w = 8.360$ .




Figure 4.5: Comparison of performance with noise-free input and noisy input with 20dB SNR. It is clear that there is good agreement with the ground truth in both scatter plots for noise-free and noisy input. The performance degrades slightly as indicated in the table below. The COSE plot show directional agreement with the ground truth (COSE  $\approx 0$ )

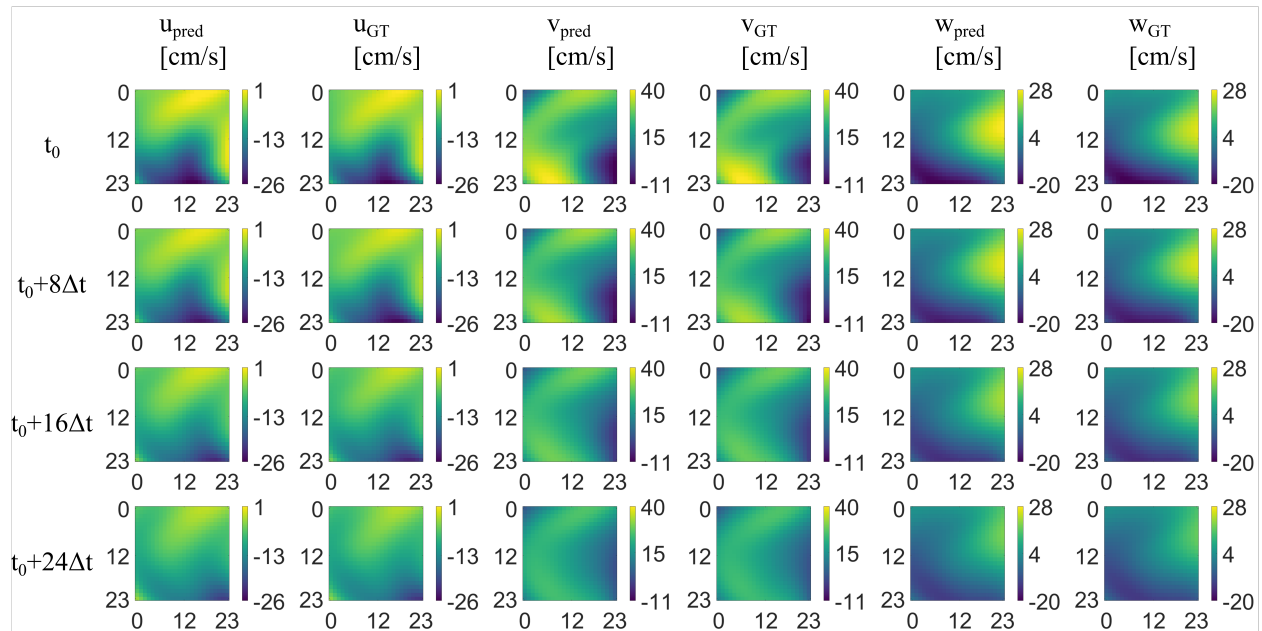


Figure 4.6: Time interpolation of velocity. Here the synthetic 4D-Flow MRI low resolution velocities are available at  $t + 2\Delta t$ ,  $t + 6\Delta t$ ,  $t + 10\Delta t$ ,  $t + 14\Delta t$ ,  $t + 18\Delta t$ , and  $t + 22\Delta t$  where  $\Delta t = 10ms$ .  $u_{pred}, v_{pred}, w_{pred}$  are the predictions of the IP-PINN algorithm.  $u_{GT}, v_{GT}, w_{GT}$  are the ground truth snapshots.

the 4D-Flow MRI was 40 ms while as the time resolution of the prediction was 10 ms.

## 4.5.2 Luminal Boundary Prediction




Figure 4.7: Luminal boundary prediction. The left most image is the low resolution magnitude image which is generated from the voxelized lumen boundary in the spatial resolution of the synthetic 4D-Flow MRI data set. This is typically the shapes seen while segmenting 4D-Flow MRI magnitude images. The center is the voxelization of the luminal boundary into the grid into which the ground truth CFD flow data was interpolated. The right most image is segmentation of the magnitude image prediction by the neural net.

One pivotal innovation in our approach is the network’s ability to predict not only high-resolution velocity fields but also high-resolution magnitude images using exclusively time-resolved phase  $\mathbf{x}$ -space velocity-encoded images as inputs, thereby eliminating the need for reference scan. Fig. 4.7 demonstrates the luminal boundary for a selected ROI. For comparison, the ground truth magnitude image is also presented. The figure additionally displays low-resolution MRI magnitude images of the ROI, which are synthetically generated using the corresponding reference CFD mesh data. These images result from processing six overlapping input blocks of flow-encoded scan data. As can be discerned from the figure, the IP-PINN model adeptly identifies the subtle demarcation of the luminal boundary.

## 4.5.3 Performance Evaluation on Noisy Synthetic 4D-Flow MRI with Velocity Aliasing

In this section, we demonstrate the ability of the IP-PINN algorithm to overcome velocity aliasing artifacts that are routinely present due to low  $v_{enc}$  parameter setting. In the test data, the  $v_{enc}$  parameter was set to 0.7 times the maximum of component wise velocity in

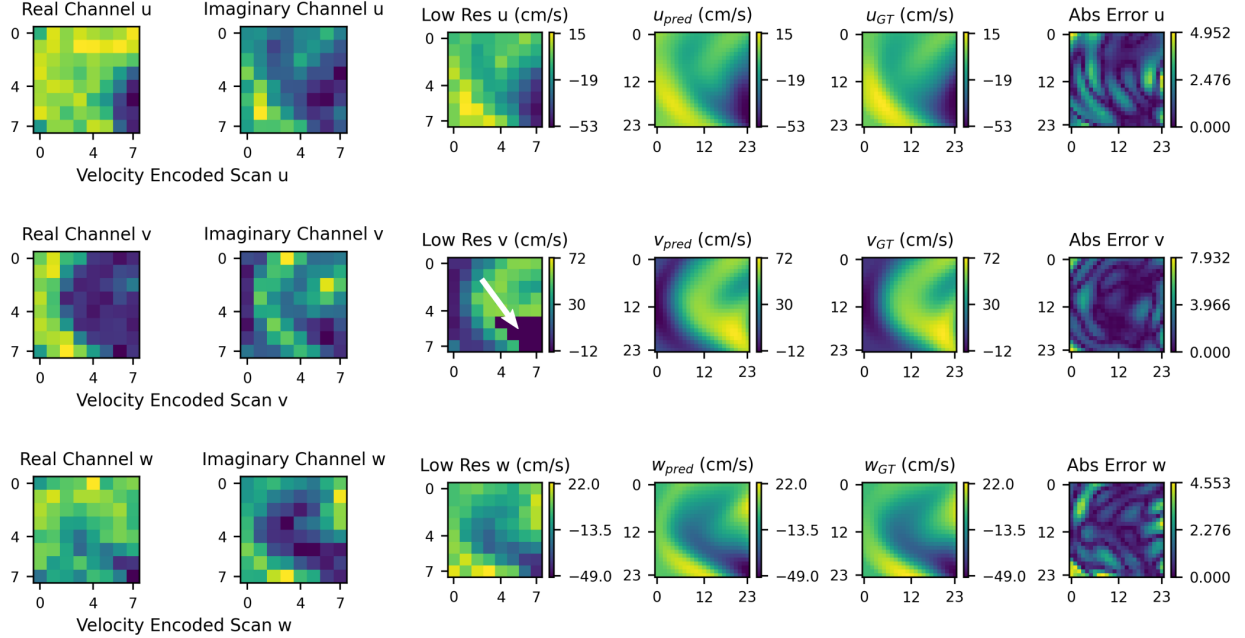


Figure 4.8: Velocity fields predicted by IP-PINN models for input exhibiting velocity aliasing at SNR of 10 dB. White arrow indicate the velocity aliasing artifacts resulting from low  $v_{enc}$  setting.

the ground truth data set. Furthermore, acquisition noise with 10dB SNR in  $\mathbf{k}$ -space was also added. Consequently, conventional processing of 4D-Flow MRI data will cause velocity aliasing (Fig. 4.8, Low res  $v$ ). The IP-PINN algorithm, because of its unique data fidelity loss defined in  $\mathbf{x}$ -space matching velocity encoded scans, is able to overcome the aliasing artefact (see Fig. 4.8,  $v_{pred}$ ). Table 4.4 summarizes the final total relative percent error ( $TRPE$ ), the coefficient of determination ( $R^2$ ), and the median absolute percentage errors of predictions for u- velocity  $MDAPE_u$ , v-velocity  $MDAPE_v$ , and w-velocity  $MDAPE_w$ . As demonstrated, the IP-PINN algorithm is able to accurately recover the velocity field even when there is significant noise and velocity aliasing artefacts.

SNR	TRPE	$R^2$	$MDAPE_u$	$MDAPE_v$	$MDAPE_w$
10 dB	8.03%	0.983	14.890	5.084	11.393

Table 4.4: Summary of performance metrics for the proposed IP-PINN model when processing noisy input data with velocity aliasing artifacts at signal-to-noise ratios (SNR) of 10 dB. Metrics include total relative percent error ( $TRPE$ ), coefficient of determination ( $R^2$ ), and median absolute percentage error of  $u_{pred}$  ( $MDAPE_u$ ),  $v_{pred}$  ( $MDAPE_v$ ),  $w_{pred}$  ( $MDAPE_w$ )

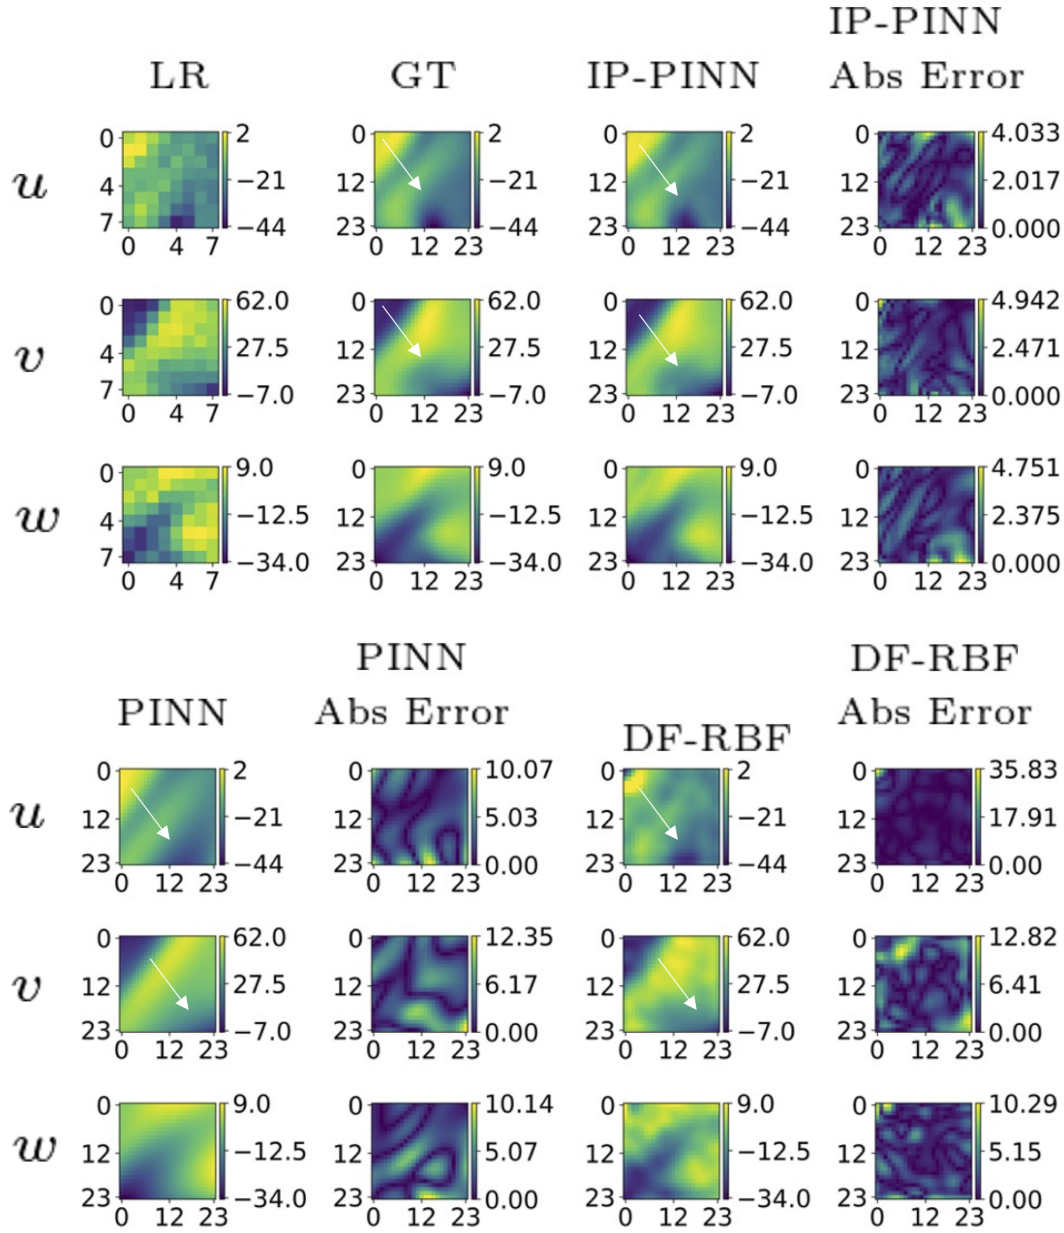


Figure 4.9: Comparison between IP-PINN, PINN, and DF-RBF. ‘LR’ is the low resolution velocity obtained by traditional processing of the synthetic 4D-Flow MRI scans (3 velocity encoded and 1 reference scan). ‘GT’ represents the ground truth which was used to generate the synthetic 4D-Flow MRI data. Subsequent columns represent the output of the IP-PINN, PINN, and DF-RBF methods with their absolute errors. The white arrows point to fine feature in the ground truth that are replicated faithfully by IP-PINN but not by PINN as well as DF-RBF. All velocities are in cm/s. The LR images are 8x8x8, where the voxel dimension is 1.05mm. All the super-resolved results are 24x24x24, where the voxel dimension is 0.35mm.

#### 4.5.4 Comparison with Physics-Informed Neural Network (PINN) and Divergence Free Radial Basis Functions (DF-RBFs)

In this section, we evaluate the performance of the IP-PINN framework in comparison to the conventional fully-connected physics-informed neural network (PINN), a concept previously developed by the authors [52], and divergence-free radial basis function (DF-RBF) approximations [25]. DF-RBFs incorporate the incompressibility constraint in approximating the flow field. This approach is superior to ordinary polynomial interpolation as the solutions are guaranteed to obey the divergence-free condition applicable to incompressible fluids.

We exploit the generalizability capabilities of the IP-PINN algorithm by employing a pre-trained model to process a novel dataset unseen by the pre-trained model, resulting in a substantial reduction in processing time. Conversely, a PINN model must be trained from scratch for each new dataset due to the absence of generalizability capabilities in traditional PINNs.

To evaluate the IP-PINN’s performance in comparison to PINN and DF-RBF models, we trained PINN and DF-RBF models for the same region of interest examined in Section 4.5, with inputs at a noise level corresponding to a signal-to-noise ratio (SNR) of 20 dB. We employed identical neural network hyperparameters for the PINN model as those used for the multilayer perceptron component of the IP-PINN framework, which includes network architecture, neuron activation functions, loss term weighting, optimizer, and learning rates. For the DF-RBF method, we used a Gaussian kernel with a standard deviation  $\sigma = \Delta x$ , where  $\Delta x = 1.0$  mm is the 4D-Flow MRI voxel dimension.

As seen in Fig. 4.9, visually, IP-PINN very closely approximates the ground truth. It appears that PINN captures the general features accurately but smooths out the details. DF-RBF picks up a lot of the noise from the input data, resulting in high-frequency oscillations in the results. Note that for the DF-RBF method, the input is the low-resolution velocities, which have to be computed from the reference and velocity-encoded scans. The IP-PINN and PINN methods, due to their formulation, only use the velocity-encoded scans. Table

4.5 juxtaposes the final total relative percent error ( $TRPE$ ) of predictions, the coefficient of determination ( $R^2$ ), and the median absolute percentage error of the  $u$ ,  $v$  and  $w$  velocities achieved by the PINN models with those of the proposed IP-PINN model. As outlined in Table 4.5, the training duration for each PINN model is 30 minutes, which is nearly twenty times longer than the processing time for each IP-PINN model, which is only 1.5 minutes. The processing time for DF-RBF is not included because it was run on the CPU. These findings emphasize the enhanced performance of the IP-PINN method over the PINN and DF-RBF methods in terms of accuracy and processing time when handling noiseless and noisy input data.

Method	SNR	$TRPE$	$R^2$	$MDAPE_u$	$MDAPE_v$	$MDAPE_w$	Processing Time
IP-PINN	Noiseless	3.65%	0.997	4.068	1.343	6.764	1.5 minutes
IP-PINN	20 dB	5.40%	0.992	6.178	2.267	8.360	1.5 minutes
PINN	Noiseless	10.14%	0.964	10.231	6.889	16.845	30 minutes
PINN	20 dB	13.33%	0.913	13.066	8.626	21.144	30 minutes
DF-RBF	Noiseless	10.98%	0.924	8.675	8.563	10.565	
DF-RBF	20 dB	17.60%	0.749	17.316	9.514	17.647	

Table 4.5: A comparison of total relative percent error ( $TRPE$ ), coefficient of determination ( $R^2$ ), and median absolute percentage error of  $u^*$ ,  $v^*$ , and  $w^*$  achieved by the PINN models and the proposed IP-PINN model for processing noise-free input data. IP-PINN required 1.5 minutes for processing the same input data, while PINN required 30 minutes

The results presented in this chapter demonstrate the effectiveness of the IP-PINN framework in enhancing 4D Flow MRI data through super-resolution, noise reduction, and artifact correction. The algorithm’s ability to generalize across different vascular conditions, as evidenced by its performance on unseen datasets, underscores its potential for broad clinical applicability. The IP-PINN not only achieves significant improvements in velocity field prediction and boundary segmentation but also outperforms traditional PINN and DF-RBF methods in terms of accuracy and processing time. These findings lay a solid foundation for further exploration and validation of the IP-PINN framework in real-world clinical scenarios.

## 4.6 Discussion

Results from this chapter demonstrate that the IP-PINN method achieves super-resolution of velocities within 15% of the reference high-resolution CFD simulation in the worst-case scenario, where the simulated phase contrast MRI has a spatial resolution of 1.0mm with an SNR of 10dB. A typical 4D-Flow MRI scan has an SNR of 20-40 dB [103, 43], highlighting the robustness and clinical relevance of the proposed method.

A key factor contributing to the success of the method is the formulation of the data fidelity term in the loss function of the IP-PINN in the complex Cartesian image space ( $\mathbf{x}$ -space). This formulation naturally handles velocity aliasing artifacts, enabling fine-tuning of results on unseen data with minimal training. Furthermore, the framework's ability to generalize across different datasets without requiring labeled data underscores its potential for ongoing, adaptive training as new data becomes available. The significant reduction in processing time compared to traditional PINN methods also positions IP-PINN as a promising tool for advancing hemodynamic analysis in a clinical setting.

An exciting possibility is the potential reduction in scan time. Currently, 4D-Flow MRI requires a minimum of 4 scans (1 reference and 3 velocity-encoded scans) to deduce time-resolved 3D velocities [105]. The results provided in this chapter has shown that IP-PINN can infer not only velocities (phase) but also magnitude (transverse magnetization) using only the velocity-encoded scan data. This has the potential to reduce scan time by 25%. Furthermore, the resolution of 4D-Flow MRI is often too coarse, particularly in cerebrovascular applications, to be used for image segmentation to delineate the lumen boundary [13]. This chapter has demonstrated that IP-PINN can generate high-resolution magnitude images from velocity-encoded scans of 4D-Flow MRI, which can be utilized for accurate lumen boundary segmentation. This is crucial for the precise computation of near-wall blood velocity-dependent hemodynamic parameters [33], such as wall shear stress, which has shown a high correlation with disease prognosis.

This discussion aligns with the broader objectives of this thesis by demonstrating how

the IP-PINN framework contributes to the advancement of time-resolved phase-contrast MRI analysis. The work presented here lays a foundation for further exploration, particularly in integrating IP-PINN with other emerging methodologies.

## 4.7 Conclusion

Chapter 4 demonstrates, for the first time, the feasibility of enhancing low-resolution, artifact-laden 4D-Flow MRI solely from the three velocity-encoded scans by means of the proposed Input-Parameterized Physics-Informed Neural Network (IP-PINN). Using synthetic data derived from high-fidelity CFD simulations of a posterior cerebral aneurysm for pre-training and a dissimilar middle cerebral aneurysm for validation, the network (i) super-resolves the velocity field by a factor of three in each spatial direction, (ii) suppresses acquisition noise, (iii) corrects eddy-current-induced phase offsets, (iv) corrects velocity aliasing, and (v) simultaneously recovers a high-resolution magnitude image suitable for lumen segmentation, all without requiring a reference scan or ground-truth labels.

The quantitative analysis shows that, after a brief fine-tuning of  $\sim 1.5$  min per unseen dataset, the IP-PINN attains a total relative percentage error below 6% and  $R^2$  values exceeding 0.99 in noise-free conditions; even at an SNR of 20 dB and under deliberate aliasing, accuracy remains within clinically acceptable limits ( $R^2 \approx 0.98$ ). Compared with a conventional fully connected PINN, the proposed method is at least an order of magnitude faster, delivers markedly lower errors, and preserves fine hemodynamic features that the competing approaches either smooth away or corrupt with noise. The continuous spatio-temporal representation intrinsic to the IP-PINN further enables four-fold temporal super-resolution and the computation of noise-free derivatives.

Beyond these technical achievements, the chapter underscores two clinically significant implications. First, because the framework obviates the reference scan while faithfully reproducing both phase and magnitude information, it offers a potential 25% reduction in scan time for standard 4D-Flow protocols. Second, by boosting spatial resolution to a level

suitable for boundary delineation, the method addresses a long-standing limitation of cerebrovascular 4D-Flow MRI, namely the under-resolved vessel wall.

Taken together, the results establish IP-PINN as a robust, generalisable, and computationally efficient alternative to existing physics-constrained learning strategies for 4D-Flow MRI. The work provides a solid foundation for the subsequent chapters, where the framework is adapted to minimally modified 2D PC-MRI and to an aggressively undersampled *pseudo one-point* 4D-Flow acquisition, further pushing the limits of scan-time reduction and data efficiency.

## Chapter 5

# IP-PINN for Time-Resolved 3D Blood Flow Velocity Reconstruction and Wall Shear Stress Calculation Using Data from Modified 2D PCMRI

### 5.1 Introduction

Building on the IP-PINN framework’s application to 4D Flow MRI in the previous chapter, this chapter adapts IP-PINN for two-dimensional phase-contrast MRI (2D PC-MRI). Although 4D Flow MRI provides volumetric, time-resolved velocity data, its prolonged acquisition times and lower spatial resolution hinder routine clinical use [43]. In contrast, 2D PC-MRI—with its shorter scan durations and superior in-plane resolution—remains the clinical standard for non-invasive flow quantification [19]. However, conventional 2D PC-MRI encodes only the through-plane velocity component in a single imaging slice, limiting its ability to resolve complex three-dimensional flow patterns and precise lumen boundary. In addition, the imaging plane must be oriented perpendicular to the vessel centerline, and misalignment can lead to underestimation or overestimation of flow rates. [100, 174].

2D PC-MRI constraints can lead to inaccuracies in vessel geometries with complex flow dynamics, such as bifurcations or aneurysms, where in-plane velocities are significant. In these cardiovascular conditions, accurately calculating WSS is challenging because WSS estimation requires detailed knowledge of the velocity gradients near the vessel wall in all three spatial dimensions, as well as an accurate 3D lumen boundary around the interrogation plane, which are not possible through conventional 2D PC-MRI acquisition [135, 164]. As discussed in chapter1, other PC-MRI limitations such as velocity aliasing, noise, low spatio-temporal

resolution, and phase offset errors due to background phase shifts caused by magnetic field inhomogeneities, eddy currents, and hardware imperfections [52, 50, 98, 75, 24, 184, 179] further degrade the accuracy of hemodynamic parameter estimation. Although computational fluid dynamics (CFD) can estimate WSS, its clinical translation is impeded by uncertainties in vessel geometry, boundary conditions, model assumptions, and the need for specialized expertise and long simulation times [78].

To address the limitations of standard 2D PC-MRI and to enhance hemodynamic parameter estimation based on 2D PC-MRI data, we propose an adaptation of the IP-PINN framework tailored for 2D PC-MRI. The IP-PINN framework aims to reconstructs high-resolution, artifact-free three-directional-three-component (3D–3C) velocity maps and 3D spin-density maps in the vicinity of the 2D PC-MRI interrogation plane. However, it requires a slightly modified acquisition sequence compared with conventional 2D PC-MRI. In standard 2D PC-MRI, only the through-plane velocity is encoded at each cardiac phase [100, 67], whereas our modified sequence retains the same plane-selection scheme but alternates the directional bipolar gradients in the  $x$ –,  $y$ –, and  $z$ –directions across cardiac phases. Although cycling through three separate scans (one for each velocity direction) reduces temporal resolution by a factor of three, the total amount of acquired data remains unchanged while providing the minimal directional diversity needed for volumetric reconstruction. The IP-PINN combines a convolutional encoder that extracts a latent representation from this undersampled, single-scan data with a physics-informed multilayer perceptron that enforces blood flow and MRI signal formation physics. During training, the network matches predicted phase-encoded images to acquired measurements and penalizes deviations from the Navier–Stokes and magnetization-encoding equations, all without requiring separate reference scans or high-resolution ground truth.

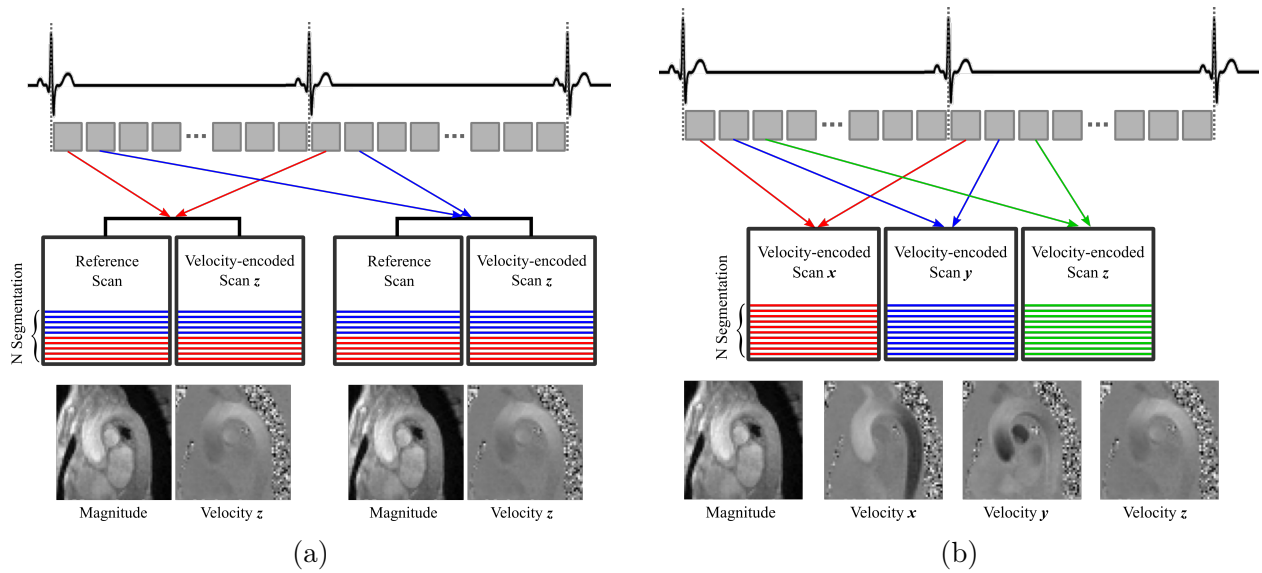


Figure 5.1: Comparison of conventional (a) and proposed modified (b) 2D phase-contrast MRI acquisition schemes. In (a), The red and blue lines represent reference and velocity-encoded scans, respectively, collected at different phases. Only a single velocity-encoding direction is captured in each cardiac phase. In (b), the red, blue, and green lines illustrate how velocity encodes in three perpendicular directions (e.g.,  $x$ ,  $y$ , and  $z$ ) are alternated across cardiac phases. Here, multiple directional encodings are interleaved across phases. This interleaving enables the proposed IP-PINN framework to reconstruct the 3D-3C velocity fields in the vicinity of the 2D imaging plane, albeit at a reduced temporal sampling rate for each velocity direction.

## 5.2 Proposed Modified 2D PC-MRI Sequence

To enable three-component, volumetric flow reconstruction from a single imaging plane using IP-PINN framework, we introduce a minimally altered 2D PC-MRI protocol that alternates the bipolar velocity-encoding gradients along each of the three orthogonal directions  $x$ ,  $y$ , and  $z$  axes across successive cardiac phases (Fig. 5.1b). This approach preserves the standard slice-placement and overall number of scans per phase, yet provides phase data in all three directions within each triplet of phases. Although each individual component is sampled at one third the native temporal rate and the modified sequence does not directly yield three-dimensional (3D) velocity fields or spin-density maps, the interleaved acquisition provides the multi-directional data diversity within the interrogation plane required by our IP-PINN framework to reconstruct high-resolution 3D-3C velocity fields and spin-density maps in the vicinity of the 2D plane.

Because reference scans are omitted in favor of a single velocity-encoded acquisition per phase, the total data volume—and thus scan time—can be reduced by up to 50 % relative to the conventional protocol. The multi-directional phase measurements are then processed by the IP-PINN, enforcing the governing Navier–Stokes equations and MRI signal–formation physics to generate a continuous volumetric representation of flow and spin density in the region immediately surrounding the imaging plane.

### Advantages and Trade-Offs

- *Volumetric Flow Information Reconstruction:* In combination with the IP-PINN framework, this method acquires velocity data in all three spatial directions near the 2D plane, enabling more precise calculation of hemodynamic parameters such as wall shear stress.
- *Comparable Velocity-Encoding Burden:* Although multiple directional velocity-encoded scans are acquired, the total number of velocity-encoded scans per cardiac phase remains the same as in conventional 2D PC-MRI.

- *Shorter Overall Scan Time:* Because the IP-PINN framework does not require a reference scan, the modified acquisition sequence potentially reduces the total scan time by up to 50%.
- *Reduced Temporal Sampling for Each Velocity Direction:* Distributing velocity encodes across consecutive cardiac phases ( $x$ ,  $y$ , and  $z$ ) lowers the effective temporal sampling for each velocity component, which can limit temporal fidelity in rapidly changing flow conditions. However, the IP-PINN framework mitigates this issue by applying super-resolution to all velocity components.

The following sections, detail the integration of modified 2DPCMR sequence into the IP-PINN framework to generate super-resolved 3D-3C velocity fields and spin-density maps with reduced noise and imaging artifacts, followed by validation results on synthetic 2D CINE-3dir PCMR data.

### 5.3 IP-PINN Adoption for Modified 2DPCMR Sequence

This section describes how the Input-Parameterized Physics-Informed Neural Network (IP-PINN) developed in Chapter 3 is adapted to process modified 2D phase-contrast MRI (2D PC-MRI) data and reconstruct the 3-dimensional 3-component super-resolve velocity maps and 3D spin-density maps around the interrogation plane. Although the core architecture and loss functions remain unchanged, the ROI selection and network’s input tensor for the 2D PC-MRI acquisition described as follows.

#### 5.3.1 Overview of the 2D Multi-Directional Data

Unlike a conventional 2D PC-MRI acquisition, which encodes only the through-plane velocity for every cardiac phase, the proposed modified 2D PC-MRI acquires separate velocity-encoded scans in the same 2D plane but at three different time steps (or cardiac phases). At the first time step, a velocity-encoding gradient is applied in the  $x$ -direction; at the second

time step, in the  $y$ -direction; and at the third time step, in the  $z$ -direction. This procedure then repeats for subsequent triplets of cardiac phases. Although each individual acquisition remains two-dimensional, their combination provides enough directional information for the IP-PINN to infer the 3D flow field near the plane (see Section 5.2). The IP-PINN uses this multi-directional 2D dataset as input to reconstruct super-resolved three-component velocity fields and spin-density maps in the region surrounding the 2D plane.

We store these velocity-encoded scans in complex  $x$ -space, yielding six channels:

1. Real part of the  $x$ -direction velocity-encoded data (first time step).
2. Imaginary part of the  $x$ -direction velocity-encoded data (first time step).
3. Real part of the  $y$ -direction velocity-encoded data (second time step).
4. Imaginary part of the  $y$ -direction velocity-encoded data (second time step).
5. Real part of the  $z$ -direction velocity-encoded data (third time step).
6. Imaginary part of the  $z$ -direction velocity-encoded data (third time step).

Thus, each pixel in the 2D plane has its own six-channel complex signal corresponding to velocity encoded in  $x$ ,  $y$ , and  $z$  at different cardiac phases, respectively.

### 5.3.2 ROI Selection and Input Parameterization

A region of interest (ROI) is selected by the user within the 2D interrogation plane, which may include vessel boundaries. This ROI spans spatial dimensions  $(n_x \times n_y)$  and  $n_t$  time steps for each of the three velocity-encoded scan directions, resulting in a data block  $\mathcal{U} \in \mathbb{C}^{n_x \times n_y \times n_t(\text{time}) \times 6(\text{channels})}$ , which represents the complex input data from three temporally direction-altering velocity encoded scans in  $\mathbf{x}$ -space for a selected region of interest (ROI) in the 2D interrogation plane. Specifically, these 3 scans correspond to three consecutive cardiac phases where velocity is encoded along the  $x$ -,  $y$ -, and  $z$ -axes, and each scan has both real and imaginary parts. Consequently, the input to the encoder is a 4-dimensional tensor

with dimensions  $n_x \times n_y \times n_z \times n_t = 16 \times 16 \times 1 \times 3$ , where each element of the input tensor has 6 channels (real and imaginary parts of the each flow-encoded scans). The IP-PINN leverages these data to produce pointwise velocities  $(u, v, w)$  and spin-density distribution in a localized neighborhood around the extracted region, guided by physical constraints from both fluid dynamics and MR signal equations.

Following the idea of input parameterization, the network’s flow image encoder  $\mathcal{G}_{\Theta_1}$  encodes this six-channel spatio-temporal ROI into a latent representation that will serve as a learnable “parameter” for the physics-informed neural network  $\mathcal{M}_{\Theta_2}$  and  $\mathcal{M}_{\Theta_3}$  (see chapter 3 section 3.3). By enforcing relevant physical constraints (such as MR phase equations, continuity, or momentum conservation) across the reconstructed volume, the IP-PINN can fill in the third dimension around the 2D plane and produce super-resolved 3D–3C velocity fields.

## 5.4 Wall Shear Stress Calculation

Wall shear stress (WSS) is a critical hemodynamic parameter that quantifies the tangential force exerted by blood flow on the vessel walls [191]. Using the reconstructed super-resolved 3D-3C velocity fields and 3D lumen boundary derived from the IP-PINN framework, the WSS was computed in the region of interest (ROI) following the methodology outlined below.

### Definition of Wall Shear Stress

In an incompressible Newtonian fluid, the viscous part of the wall-traction vector,  $\mathbf{t}_{\text{visc}}$ , is obtained by contracting the rate-of-strain tensor with the wall normal [8, 164]:

$$\mathbf{t}_{\text{visc}} = \boldsymbol{\sigma}_{\text{visc}} \cdot \mathbf{n} = 2\mu \dot{\boldsymbol{\epsilon}} \cdot \mathbf{n}, \quad (5.1)$$

where  $\mu$  is blood dynamic viscosity and  $\dot{\boldsymbol{\epsilon}} = \frac{1}{2}(\nabla \mathbf{v} + \nabla \mathbf{v}^T)$  is the rate-of-strain tensor,  $\mathbf{v}$  is the velocity vector, and  $\mathbf{n}$  is the unit normal vector to the vessel wall.




Figure 5.2: Illustration of the limitations associated with assuming orthogonality of the vessel walls' normal vectors to the 2D interrogation plane, neglecting for the geometric complexity and variability of vessel anatomy.

Only the tangential component of this traction acts as the wall-shear-stress (WSS) vector.

Projecting  $\mathbf{t}_{\text{visc}}$  onto the vessel wall gives

$$\boldsymbol{\tau}_w = \mathbf{t}_{\text{visc}} - (\mathbf{t}_{\text{visc}} \cdot \mathbf{n}) \mathbf{n}, \quad (5.2)$$

## Computational Steps

**1. Normal Vectors Calculation** The predicted 3D lumen boundary was segmented to identify the vessel wall within the ROI. The marching cubes algorithm [99] was applied to extract a triangular surface mesh, enabling calculation of the unit normal vectors  $\mathbf{n}$  at each boundary point. Incorporating these normal vectors, calculated from the 3D lumen boundary, mitigates errors associated with the common assumption that vessel walls are orthogonal to the 2D interrogation plane, which neglects geometric variations in vessel anatomy (see Fig.5.2).

**2. Rate-of-Strain Tensor Calculation** Since IP-PINN outputs a continuous function to predict the velocity fields, we reconstructed 3D-3C velocity components  $(u, v, w)$  and employed the auto-differentiation [12] to compute the velocity gradient tensor and subsequently

the rate-of-strain tensor at the boundary points.

**3. WSS Vector Calculation** First compute the viscous traction via Eq. (5.1); then apply the projection in Eq. (5.2) to recover the WSS vector  $\boldsymbol{\tau}_w$  at every boundary point.

This approach enables accurate WSS computation directly from 2D PC-MRI velocity-encoded scans, eliminating the need for full 3D acquisitions or reference scans. The robust reconstruction capabilities of the IP-PINN framework, which mitigate noise and imaging artifacts, facilitated the derivation of 3D-3C velocity fields and the corresponding 3D lumen boundary. These reconstructions enabled accurate computation of vessel wall normals and subsequent WSS.

The predicted WSS values were compared against the ground truth derived from CFD and also against the WSS calculation method implemented by Stalder et al. [164], where the normal vector to the vessel wall and the velocity gradients were estimated from the measured three-component velocities within the 2D imaging plane using B-spline interpolation. The ability to incorporate full 3D boundary curvature is expected to provide more reliable WSS estimates.

## 5.5 Generation of Synthetic 2D PC-MRI Data

To evaluate the adapted IP-PINN on 2D PC-MRI, we generated paired training and test datasets from the same CFD simulations used in Chapter 4 (PCA for training, MCA for testing; see Fig. 4.2 and Table 4.1).

### 5.5.1 High-Resolution Reference Fields

The unsteady 3D velocity fields from each CFD simulation were interpolated onto a uniform Cartesian grid with isotropic voxel size  $\Delta x_f = 0.35$  mm. For the MCA case, every 15<sup>th</sup> time step was selected to achieve a 15 ms temporal resolution matching the PCA data. These high-resolution fields serve as the “ground truth” for both training and quantitative evaluation.

### 5.5.2 Synthetic 2D PC-MRI Acquisition

To mimic our modified 2D PC-MRI sequence, each cardiac phase encodes velocity along exactly one direction ( $x$ ,  $y$ , or  $z$ ), cycling through them in order.

**Complex  $\mathbf{x}$ -space conversion and phase wrap** At each time step, the velocity components ( $u, v, w$ ) were combined into a complex-valued image via the standard PC-MRI signal model. The velocity-encoding parameter  $V_{\text{enc}}$  was set to 80% of the maximum directional speed to enforce aliasing. A composite background phase  $\phi_B(\mathbf{x}, z)$ —linear for MCA, quadratic for PCA—was added to simulate eddy-current offsets [24, 43]. The spin-density magnitude  $M(\mathbf{x}, z)$  was assigned as 1.0 inside the lumen and 0.1 outside to reflect tissue contrast.

**Spatio-temporal downsampling and noise** Each complex image block was first spatially averaged over three voxels in the out-of-plane ( $z$ ) direction, yielding a 1.05 mm effective slice thickness. Four consecutive time frames were then temporally averaged to a 60 ms cadence. The downsampled  $\mathbf{x}$ -space slices were Fourier-transformed to  $\mathbf{k}$ -space via FFT, low-pass filtered by truncating all frequencies outside  $\pm k_{\text{max}}$ , where

$$k_{\text{max}} = (2 \times 3 \Delta x_f)^{-1} = 0.667 \text{ m}^{-1},$$

and corrupted with zero-mean Gaussian noise at 20dB SNR in both real and imaginary channels. An inverse FFT returned the final synthetic 2D PC-MRI images.

**Interrogation-plane extraction** Each voxel slab along  $z$  was treated as an independent imaging plane. For each plane, three separate velocity-encoded scans were arranged so that at cardiac phase  $t$ , only the  $j$ th directional encoding ( $j \in \{x, y, z\}$ ) was present, cycling  $x \rightarrow y \rightarrow z$ . The resulting dataset thus replicates the modified acquisition: one velocity direction per phase, with realistic aliasing, noise, and phase offsets.

## 5.6 Error Metrics

The reconstruction accuracy is quantified using the same set of metrics defined in Section 4.3 of Chapter 4 (TRPE, COSE, component-wise absolute errors, velocity-magnitude scatter plots,  $R^2$ , and MDAPE).

## 5.7 Results

### 5.7.1 Pre-Training the IP-PINN

To initialize our model, we first conducted a pre-training phase using low-resolution synthetic data derived from a flow simulation of the PCA aneurysm shown in Fig. 4.2(a). Specifically, we used roughly 40,000 distinct data blocks (ROIs) out of the total 119680 data blocks of size  $n_x \times n_y \times n_z \times n_t = 16 \times 16 \times 1 \times 3$  with a spatial resolution of 1.0 mm, each accompanied by 7,500 collocation points for imposing physics-based regularization. The physical normalization factors were set as follows:

$$L = 0.8 \text{ cm}, \quad U = 93.6494 \text{ cm/s}, \quad T = 0.0085 \text{ s}.$$

For each data block, the non-dimensional  $z$ -coordinates were sampled according to the desired out-of-plane coverage. In particular, we evaluated five out-of-plane slices centered around the primary interrogation plane, treating the plane of interest as the midpoint. This configuration ensures that our model learns flow variations across multiple slices, even though it remains anchored to a single 2D PC-MRI interrogation plane at acquisition.

A batch size of two data blocks was used for each gradient update, and the model was trained for 500 epochs. Table 5.1 provides a representative snapshot of the loss terms, along with their empirically assigned weights, at the end of the pre-training phase.

Loss Term	Weight	Final Value
Physics Loss ( $\mathcal{L}_P$ )	0.1	$3.327 \times 10^{-5}$
Data Fidelity Loss ( $\mathcal{L}_D$ )	0.9	$4.124 \times 10^{-5}$
Total Weighted Loss	—	$7.451 \times 10^{-5}$

Table 5.1: Sample loss values and weighting factors for the pre-trained IP-PINN, focused on unsteady blood flow in a PCA aneurysm.

Overall, this pre-training process enables the IP-PINN to learn a broad representation of hemodynamic features from the PCA aneurysm data prior to fine-tuning on new cases. As we show in the subsequent sections, this strategy accelerates convergence on unseen datasets while maintaining robust reconstruction performance.

### 5.7.2 Fine-Tuning on Unseen MCA Data

After pre-training on synthetic 2D PC-MRI scans derived from the PCA aneurysm simulation, we evaluated the IP-PINN’s performance using a new dataset generated from a middle cerebral artery (MCA) aneurysm (Fig. 4.2). To further challenge the model, we introduced additional noise at 20 dB SNR within the  $\mathbf{k}$ -space domain.

Figure 5.3 illustrates representative reconstructions for a noisy dataset at an SNR of 20 dB, demonstrating how the IP-PINN framework extends a single-direction, 2D PC-MRI velocity-encoded scan into full three-component velocity fields at each cardiac phase. In this figure:

- *Columns 1 and 2:* Show the real and imaginary components of the velocity-encoded scans in  $\mathbf{x}$ -space, which serve as the raw input to the IP-PINN. Each cardiac phase of the input data includes only one directional component of the velocity.
- *Column 3:* Depicts velocities at a lower resolution, recovered via conventional 2D PC-MRI processing.

- *Column 4:* Displays the IP-PINN's predicted velocity fields. IP-PINN predict three directional velocity components at every cardiac cycle.
- *Column 5:* Shows the ground truth from the CFD simulation.
- *Column 6:* Presents the point-wise absolute error distribution.

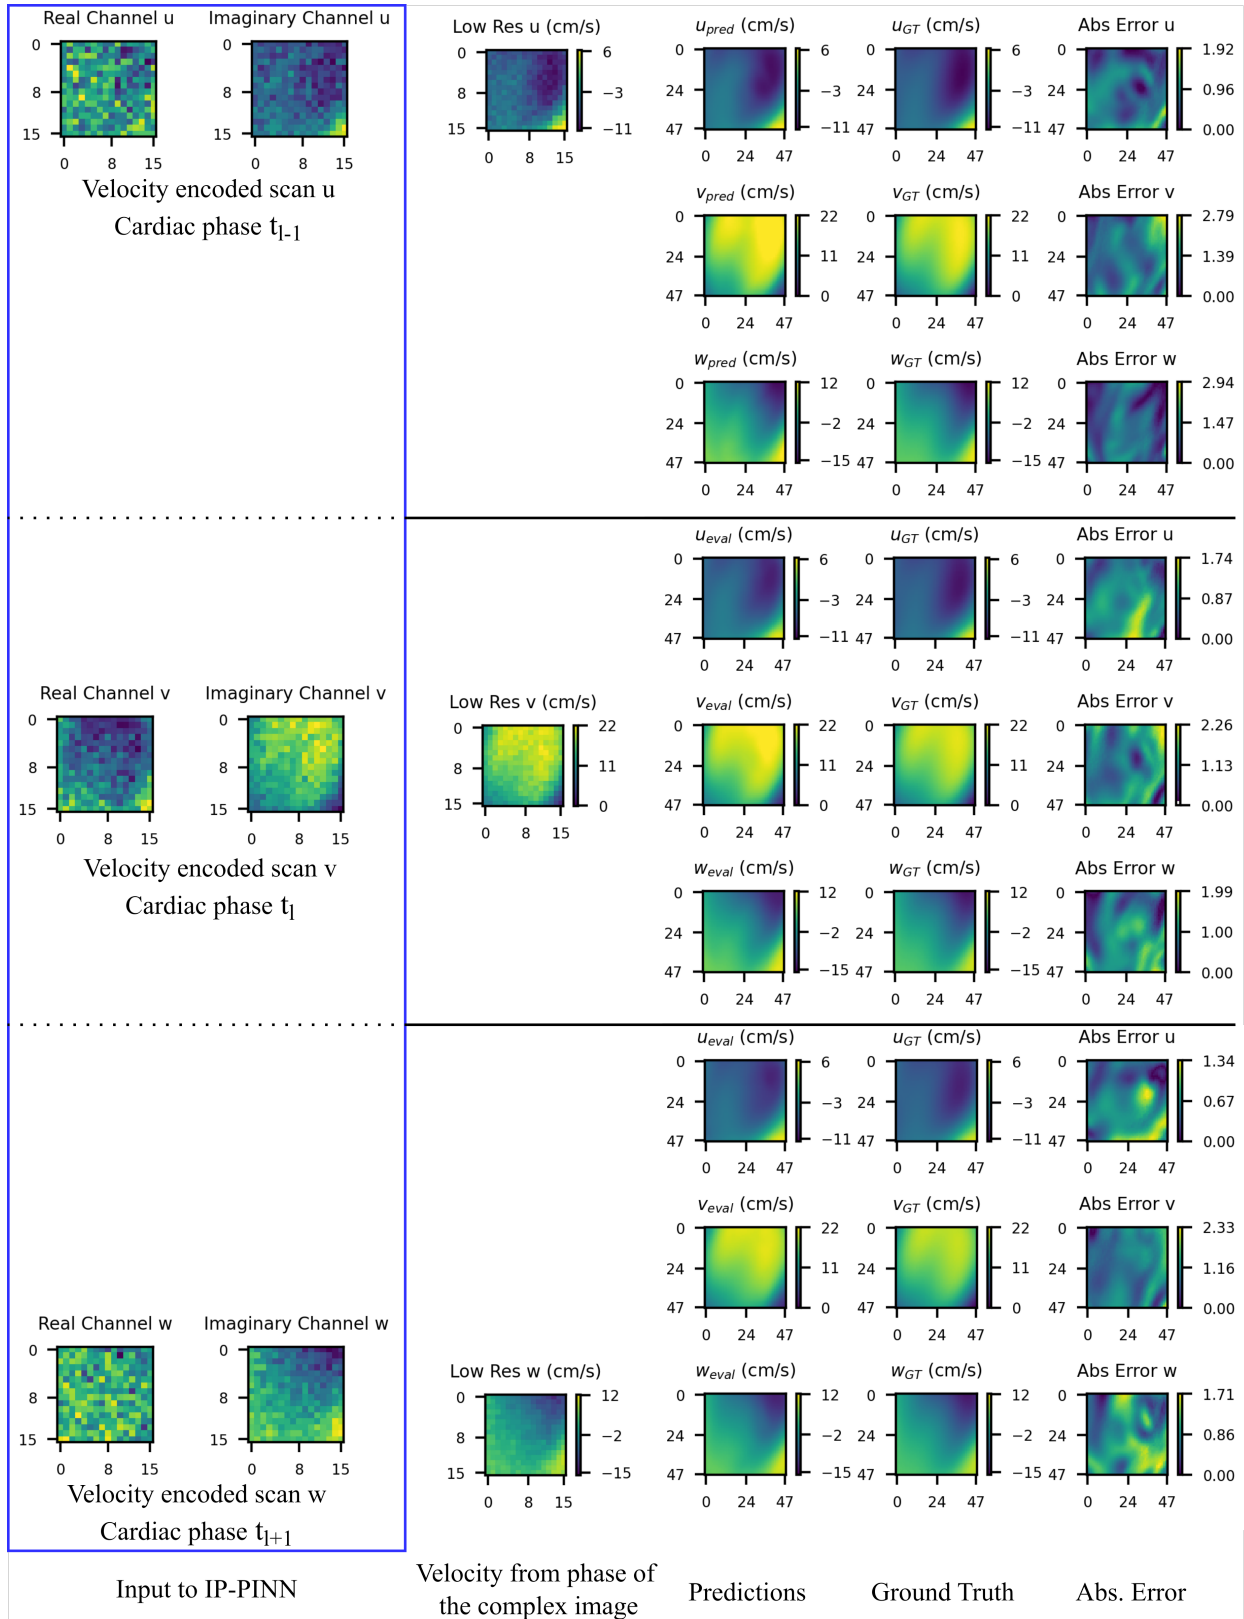


Figure 5.3: Results of processing velocity-encoded data at 20dB with IP-PINN. Each row of the input block corresponds to one velocity-encoding direction at a given cardiac phase (velocity encoded  $u$  at  $t_{l-1}$ , velocity encoded  $v$  at  $t_l$ , and velocity encoded  $w$  at  $t_{l+1}$ ). The IP-PINN framework generates all three components of the velocity fields at every cardiac phase.

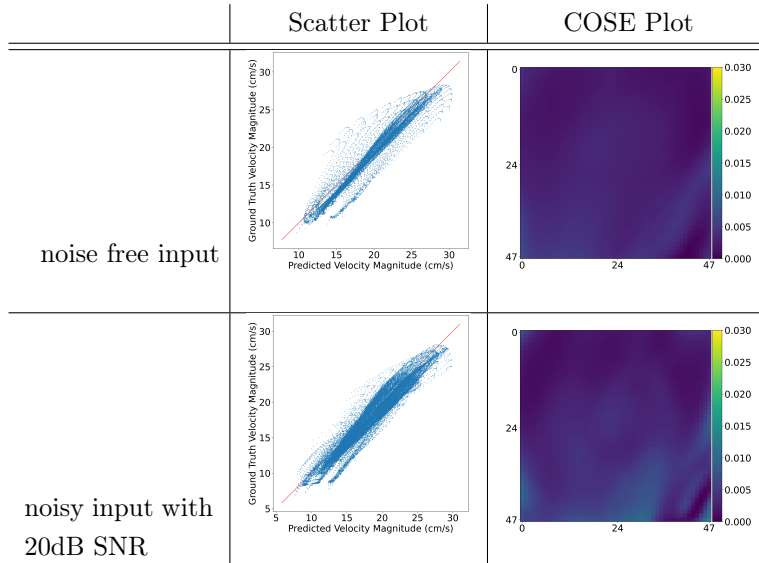


Figure 5.4: IP-PINN performance comparison of for processing noise-free and noisy input with SNR of  $20dB$ , illustrated through scatter and cosine distance plots. Both cases show strong agreement with the ground truth in terms of velocity magnitude (scatter plots) and flow direction (COSE  $\approx 0$ ).

$SNR$	$R^2$	$TRPE$	$MDAPE$	Processing Time
Noiseless	0.942	9.25%	4.34%	1.5 minutes
20 $dB$	0.927	11.51%	5.12%	1.5 minutes

Table 5.2: A summary of total relative percentage error ( $TRPE$ ), coefficient of determination ( $R^2$ ), and median absolute percentage error ( $MDAPE$ ) for processing the noiseless input data and input data with SNR of 20 dB. IP-PINN required 1.5 minutes for processing the input data

Table 5.2 summarizes error metrics achieved by IP-PINN processing noiseless input data as well as input data with SNR of  $20dB$ . These metrics include the coefficient of determination ( $R^2$ ), total relative percentage error ( $TRPE$ ), and the median absolute percentage error ( $MDAPE$ ) of the velocity magnitude. As anticipated, noise slightly degrades the accuracy, but overall, the fine-tuned IP-PINN successfully preserves key flow features even under challenging conditions. Since the MCA dataset was not part of the pre-training process, we performed a “fine-tuning” step by retraining the IP-PINN specifically on these unseen data. Note that the total execution time reported in table 5.2 corresponds to this fine-tuning process.

Figure 5.4 contrasts these results via scatter plots (for velocity magnitude) and cosine error plots (for flow direction). Points clustering near the  $45^\circ$  line in the scatter plots indicate strong agreement in magnitude, whereas low cosine errors reflect accurate directional alignment. In both cases—noisy and noiseless—the IP-PINN demonstrates robust performance, albeit with a modest decrease in accuracy when noise is present.

### 5.7.3 Temporal Super-Resolution Capability of IP-PINN

An important feature of the proposed IP-PINN framework is its ability to generate super-resolved, artifact-free 3D velocity fields for cardiac phases between those present in the input data. This time-interpolation capability is critical for achieving temporally resolved hemodynamic analysis.

Figure 5.5 demonstrates the results of super-resolving 3D velocity fields between the cardiac phases provided as input. Specifically, the input includes velocity-encoded scans in a 2D interrogation plane for three cardiac phases:

- $t_{l-1}$ : Velocity encoded scan in the  $x$ -direction.
- $t_l$ : Velocity encoded scan in the  $y$ -direction.
- $t_{l+1}$ : Velocity encoded scan in the  $z$ -direction.

Using the provided inputs, the IP-PINN framework predicts super-resolved 3D velocity fields at cardiac phases  $t_{l-1}$ ,  $t_{l-1/2}$ ,  $t_l$ ,  $t_{l+1/2}$ , and  $t_{l+1}$ . The results are visualized in Figure 5.5, where each row corresponds to a cardiac phase and displays the 3 component predicted and ground truth velocity components.

The IP-PINN successfully reconstructs temporally resolved 3D velocity fields, capturing smooth transitions between input cardiac phases and eliminating noise and artifacts. This capability underscores the robustness of the framework in super-resolving both spatial and temporal dynamics from 2D PC-MRI data.




Figure 5.5: Super-resolution of 3D velocity fields in time by the IP-PINN framework. The figure demonstrates the predicted ( $u_{pred}$ ,  $v_{pred}$ , and  $w_{pred}$ ) and corresponding ground truth ( $u_{GT}$ ,  $v_{GT}$ ,  $w_{GT}$ ) velocity components across cardiac phases  $t_{l-1}$ ,  $t_{l-1/2}$ ,  $t_l$ ,  $t_{l+1/2}$ , and  $t_{l+1}$ . The input data includes velocity-encoded scans in the  $x$ -direction at  $t_{l-1}$ ,  $y$ -direction at  $t_l$ , and  $z$ -direction at  $t_{l+1}$ . Each row represent corresponding velocity components for different cardiac phases.

### 5.7.4 Luminal Boundary Prediction

Alongside high-resolution, three-component velocity fields, IP-PINN reconstruct three-dimensional magnitude images around the interrogation plane, using only the 2D PC-MRI phase-encoded input. By eliminating the need for a separate reference scan, this approach significantly streamlines the acquisition process. Figure 5.6 illustrates the luminal boundary reconstruction in a selected region of interest (ROI). In the left panel, a synthetically generated low-resolution magnitude image is shown, with the 2D interrogation slice highlighted in red. The middle panel depicts the IP-PINN’s predicted volumetric lumen boundary, while the right panel shows the ground truth for comparison. The red slice in both the predicted and




Figure 5.6: Comparison of predicted low-resolution (left), predicted (center), and ground truth (right) lumen boundaries. The highlighted red area in the low-resolution lumen boundary indicates the spatial location (2D interrogation plane). In the prediction and ground truth figures highlighted red region shows the small region around the interrogation plane, while the gray regions illustrate out-of-plane volumetric boundaries. The IP-PINN framework recovers the high-resolution lumen structure near the plane but exhibits larger deviations farther from the slice, where no direct measurements are available.

ground-truth boundaries marks the same small region near the interrogation plane, whereas the grey volume indicates out-of-plane tissue. Although the model accurately delineates the lumen boundary near the imaging plane, prediction errors increase at greater out-of-plane distances, reflecting the lack of observational data beyond the central 2D slice.

### 5.7.5 Wall Shear Stress Calculation

To evaluate the accuracy of the wall shear stress (WSS) predictions obtained using the proposed IP-PINN framework, we compared the computed WSS magnitudes along the vessel boundary with both the ground truth from CFD simulations and the cubic B-spline interpolation method used in [164]. Figure 5.7 presents the WSS magnitude as a function of arc length along the boundary for the three methods.

For the IP-PINN framework, the boundary length was determined based on the distance between 3D faces obtained from applying the marching cubes algorithm on the reconstructed 3D lumen boundary near the interrogation plane. In contrast, for the cubic B-spline interpolation method, the boundary length was computed using contour points extracted from


the 2D imaging plane.

As highlighted in Fig. 5.7, the WSS calculated using IP-PINN predictions (solid black line) follows the ground truth (red dashed line), capturing most of the local variations in WSS with high fidelity. However, The cubic B-spline interpolation method (blue dash-dot line) exhibits smoother variations and underestimates peak WSS values compared to both the IP-PINN and ground truth. This approach is constrained by in-plane 2D normal vectors and velocity gradients, fails to capture the full complexity of the hemodynamic forces acting on the vessel wall. The IP-PINN framework, by leveraging the reconstructed 3D lumen boundary, accurately incorporates out-of-plane velocity gradients, leading to a more precise representation of WSS, particularly in regions of sharp WSS variations. These findings demonstrate the advantage of using a 3D lumen boundary reconstruction for WSS estimation, as done in the IP-PINN framework, over conventional 2D-based approaches.

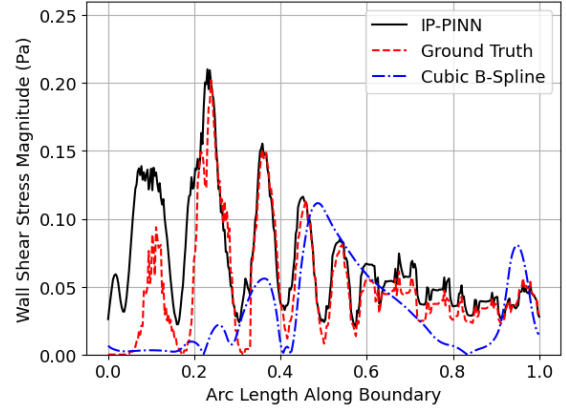
## 5.8 Discussion

The proposed input-parametrized physics-informed neural network (IP-PINN) framework, combined with a modified 2D phase-contrast MRI (2D PC-MRI) acquisition protocol, enables the reconstruction of high-resolution, artifact-free volumetric (3D) three-component (3C) velocity fields and corresponding 3D spin-density maps from 2D PC-MRI data. This approach effectively extends the applicability of 2D PC-MRI to volumetric hemodynamic analysis by addressing the limitations of conventional methods that rely solely on through-plane velocity measurements. Our results highlight several key advantages of the IP-PINN framework over standard 2D PC-MRI analysis techniques (including the typical cubic B-spline interpolation method for WSS estimation):


- Volumetric velocity field reconstruction: It recovers all three velocity components at every cardiac phase, effectively producing a 3D flow field in the vicinity of the imaging plane (rather than only a single through-plane velocity component).



(a)  $t = t_{l-1}$



(b)  $t = t_l$



(c)  $t = t_{l+1}$

Figure 5.7: Comparison of wall shear stress (WSS) magnitudes along the vessel boundary computed using the IP-PINN framework (black solid line), ground truth from CFD simulations (red dashed line), and the cubic B-spline interpolation method (blue dash-dot line) from [164]. The boundary length for the IP-PINN method is measured from the 3D faces of the marching cubes algorithm applied to the reconstructed lumen, whereas for the B-spline interpolation method, it is calculated from the contour points in the 2D imaging plane.

- Accurate wall geometry and WSS estimation: It reconstructs the three-dimensional lumen boundary (instead of assuming the vessel wall lies entirely within the imaging plane), yielding more accurate wall-normal vectors and thus improving wall shear stress (WSS) estimation.
- Temporal super-resolution: It interpolates the velocity fields between discrete cardiac phases to generate temporally super-resolved data, enhancing the fidelity of time-dependent hemodynamic assessments.
- Reduced scan time: It eliminates the need for separate reference scans, reducing the total scan time by up to 50% compared to traditional 2D PC-MRI protocols.
- Robustness to noise and artifacts: Its physics-informed formulation inherently mitigates noise, velocity aliasing, and phase-offset errors, leading to more reliable velocity reconstructions under realistic imaging conditions.

One of the most notable improvements with the IP-PINN framework is in the estimation of wall shear stress. Conventional WSS estimation from 2D PC-MRI data typically involves segmenting the lumen and fitting a cubic B-spline to the measured velocities along the vessel wall. A key limitation of this approach is the assumption that the imaging plane is exactly orthogonal to the vessel and that the vessel wall normal vectors lie entirely within that plane [164]. This assumption effectively neglects any out-of-plane velocity gradients (arising from vessel curvature or complex three-dimensional flow), leading to inaccuracies in WSS calculations—especially in regions with significant curvature or non-planar flow patterns. In contrast, the IP-PINN method imposes no such in-plane assumption. It reconstructs the full 3D lumen geometry and explicitly incorporates out-of-plane velocity components, resulting in significantly more accurate WSS estimates. In our comparisons, the WSS distributions obtained with IP-PINN closely matched the computational fluid dynamics (CFD) ground truth, capturing fine-scale spatial variations and correctly identifying peak WSS regions. Meanwhile, the conventional planar B-spline interpolation method produced a much smoother

WSS profile and substantially underestimated the peak shear values. By capturing the true magnitude and distribution of WSS, the IP-PINN framework provides critical insights that would be missed by planar methods. This improvement in WSS estimation is particularly relevant for studies of vascular diseases (e.g., aneurysms), where localized changes in WSS are believed to contribute to disease progression.

## 5.9 Conclusion

This chapter adopts the developed IP-PINN in chapter 3 for reconstructing high-resolution, artifact-free 3D-3C velocity fields and 3D lumen boundaries from 2D PC-MRI acquisitions. By integrating a modified acquisition sequence that alters bi-polar velocity encoding gradients across cardiac phases, the adopted IP-PINN effectively reconstructs volumetric hemodynamic information while maintaining the efficiency and accessibility of 2D PC-MRI. The results demonstrate that the IP-PINN framework significantly improves the accuracy of hemodynamic parameter estimation, particularly in wall shear stress (WSS) calculations, by leveraging full 3D lumen boundary reconstruction and out-of-plane velocity gradients.

The IP-PINN framework successfully reconstructed 3D velocity fields and lumen boundaries from limited 2D PC-MRI data, overcoming the inherent limitations of conventional approaches that rely solely on through-plane velocity measurements. Compared to existing methods such as cubic B-spline interpolation, the proposed approach more accurately captures localized WSS variations and peak values. Additionally, the ability to generate temporally super-resolved velocity fields enhances the fidelity of time-dependent hemodynamic assessments. The framework demonstrated robust performance in the presence of noise and phase artifacts, ensuring reliable velocity reconstruction under realistic imaging conditions. Furthermore, by eliminating the need for reference scans, the proposed method significantly reduces acquisition time, enhancing the practicality of high-fidelity hemodynamic analysis PC-MRI protocols.

## Chapter 6

# Accelerating 4D Flow MRI with IP-PINN: A pseudo 1-point acquisition and k-space reconstruction approach

### 6.1 Introduction

As discussed in Chapter 2, standard 4D Flow MRI acquisitions employ an interleaved four-point phase-contrast encoding: for each cardiac timeframe, one non-velocity-encoded reference scan and three velocity-encoded scans (with velocity encoding in the orthogonal x, y, and z directions) are acquired [105]. This velocity encoding scheme inherently multiplies scan duration (often requiring at least four times the scan time of a single conventional cine acquisition, on the order of 5–20 minutes for a typical examination) [165, 129]. Such prolonged acquisitions introduce challenges including patient fatigue, elevated risk of motion artifacts (necessitating respiratory gating or breath-holding), and increased cost. A variety of scan acceleration strategies have been explored to mitigate these issues. Hardware advances (e.g. multi-channel coils and high-performance gradients) and parallel imaging have been combined with advanced undersampling techniques – such as k-t acceleration (kt-BLAST, kt-SENSE, kt-GRAPPA) and compressed sensing – to shorten 4D Flow MRI scan times [76, 165, 91, 63, 101, 128, 73]. Despite these improvements, 4D Flow MRI acquisitions can still range from several minutes up to 20 minutes in practice, motivating the development of even more data-efficient acquisition methods.

This chapter of the thesis introduces a novel pseudo 1-point 4D Flow MRI acquisition scheme designed to drastically reduce the data requirements (and thus scan time) for IP-PINN framework to reconstruct time-resolved three-dimensional three-component velocity maps. In a conventional 4D Flow sequence, every slice or volume element is scanned with

all three velocity-encoding directions (plus a non-encoded reference), resulting in four acquisitions per slice per timeframe [105]. By contrast, the proposed pseudo 1-point sequence acquires only one velocity-encoded scan per slice and omits the separate reference scan entirely. The velocity encoding direction is varied across slices in a repeating pattern (for instance, the first slice is acquired with  $x$ -direction flow encoding, the next slice with  $y$ -encoding, the next with  $z$ -encoding, and then cycling back to  $x$  for subsequent slices). In essence, each slice yields a single velocity-encoded image without a corresponding non-encoded image. This scheme reduces the total acquired datasets by a factor of four compared to the fully sampled 4-point acquisition approach – a transformative decrease in data volume and scan time. The results provided in 4 demonstrated that eliminating the reference scan alone is feasible with IP-PINN reconstruction, yielding 3D-3C velocity maps and 3D magnitude image with minimal error. The pseudo 1-point method expands on this concept, pushing data reduction further by not only skipping the reference but also sparsely sampling different velocity directions across slices.

Reconstructing high-fidelity flow fields from such highly sparse data demands a reconstruction approach that rigorously enforces the underlying flow physics, the MRI physics, and measurement consistency. This chapter modifies the IP-PINN framework to impose the data fidelity directly in the  $k$ -space domain (the raw measurement domain) rather than in complex image space. This means that during reconstruction, the estimated velocity-encoded images are Fourier-transformed and compared to the acquired  $k$ -space samples, ensuring that the solution is exactly consistent with the recorded MRI signals. Enforcing fidelity in  $k$ -space has important advantages: The Fourier transform of a function with finite support results in a function with infinite support [17]. As discussed in Chapter 2, MRI is limited by how far out in  $k$ -space we can measure. In practice, MRI collect data only within a finite  $k$ -space region around the center (usually a maximal  $k$ -value in each direction determined by gradient limits, scan time, and desired resolution). This is equivalent to multiplying the object’s continuous spectrum by a rectangular windowing function (rect-function) that truncates  $k$ -

space beyond some  $\pm k^{max}$  (see Chapter 2, Eq. 2.12). Multiplication of the rect-function function with apparent spin-density of the object results in convolution of the apparent spin density maps with a sinc-function with infinite support. Therefore, any local discrepancy in the image could originate from errors in many  $k$ -space frequencies. Moreover, effects such as  $T_2$  and  $T_2^*$  decay during readout are baked into the image [140]. Enforcing fidelity in image space requires the network to implicitly learn and invert all these combined effects, with no direct access to the underlying physics. In contrast, a  $k$ -space fidelity term can explicitly incorporate known decay envelopes or sensitivity maps, preserving mathematical clarity and interpretability.

To reconstruct the fully 3D-3C velocity field and the accompanying magnitude image of the selected ROI from the limited pseudo 1-point data, we employ an IP-PINN. Since the Fourier transform is linear, one can isolate the effective  $k$ -space of the ROI by zero padding the ROI to retain the FOV of the imaging domain. Therefore, the employed IP-PINN operates on the zero padded ROI with a rephased frequency data to recenter the ROI in the imaging FOV according to the Fourier shift theorem [125]. The ROI consists of directional altering velocity-encoded scans per slice basis. The IP-PINN framework, embedding physical knowledge into the reconstruction process by incorporating a priori known physics equations and acquisition models as loss functions, reconstructs super-resolved, denoised, and artifact free, 3D-3C velocity maps and magnitude image from under-sampled measurements. The following sections will detail the proposed the pseudo 1-point sequence and the IP-PINN reconstruction algorithm, and evaluate their performance in reconstructing high-quality flow fields from exceptionally sparse data.

## 6.2 Proposed Pseudo 1-Point 4D Flow MRI Sequence

Conventional 4D Flow MRI acquires four phase-contrast datasets per slice per cardiac phase (three velocity-encoded along orthogonal axes plus one non-encoded reference) to compute volumetric three-component velocity vectors via phase differencing [105]. The proposed

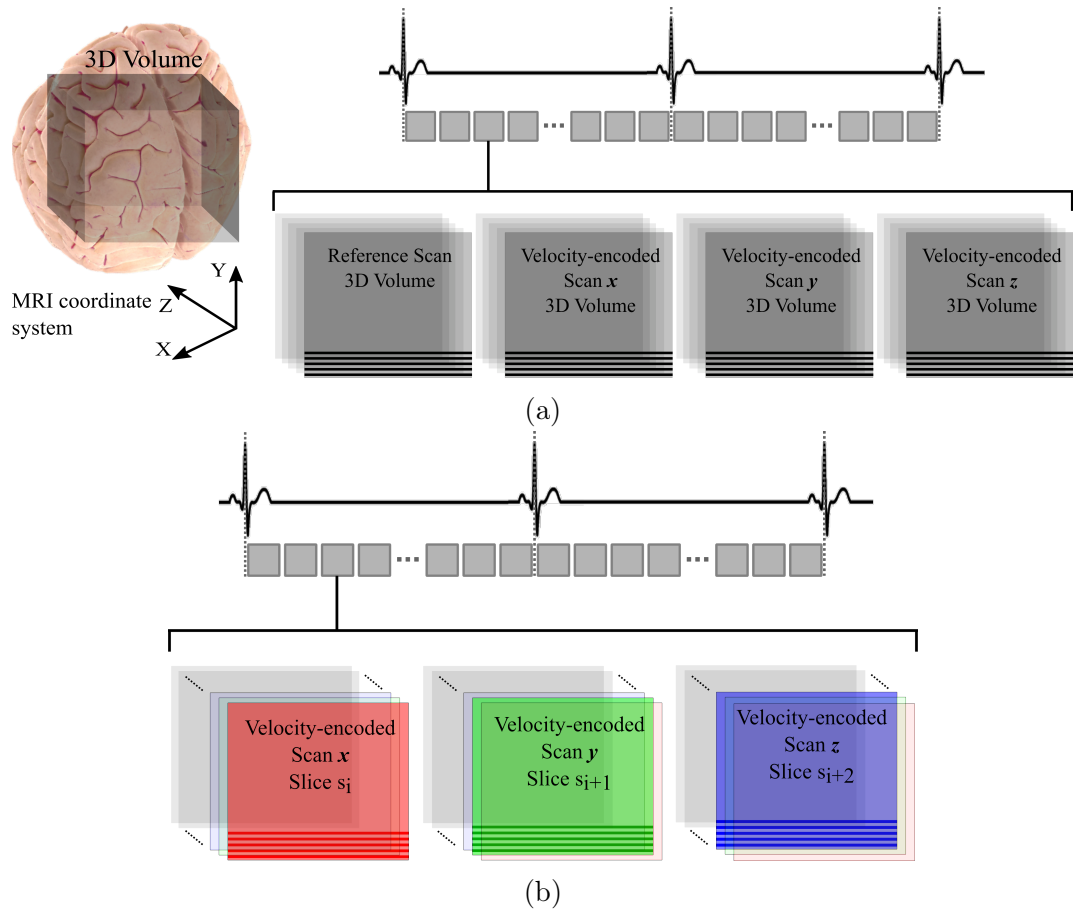


Figure 6.1: Comparison of conventional (a) and proposed pseudo 1-point (b) 4D Flow MRI acquisition schemes. In (a), at each cardiac phase, four separate 3D datasets are acquired—one reference scan plus three velocity-encoded acquisitions—to determine blood flow velocities along the  $x$ ,  $y$ , and  $z$  axes. In (b), The bipolar velocity-encoding gradients alternate across successive slices in the  $z$ -direction, such that only one velocity-encoded dataset is acquired per slice. In this way, the different encoding directions are interleaved slice-by-slice, reducing the spatial sampling rate for each velocity component.

pseudo 1-point sequence instead alternates the bipolar velocity-encoding gradients along each of the three orthogonal directions  $x$ ,  $y$ , and  $z$  axes across successive spatial slices and collects only one velocity-encoded dataset per slice and no explicit reference. Figure 6.1 By encoding only a single velocity direction at each slice location and rotating the encoding direction on a slice-by-slice basis the total number of scanned volumes is reduced from four to one for each slice. This design cuts the raw data acquisition load (and thus scan time) by a factor of four.

Let  $s$  index physical slices along the head–feet ( $z$ ) direction. We define the velocity-encoding direction for slice  $s$  as:

$$\text{VENC\_dir}(s) = \begin{cases} x, & s \bmod 3 = 0, \\ y, & s \bmod 3 = 1, \\ z, & s \bmod 3 = 2. \end{cases}$$

This acquisition scheme reduces the spatial sampling rate of each velocity component by a factor of three, and therefore does not directly produce fully sampled volumetric velocity fields. However, when integrated with the IP-PINN framework, these sparse measurements enable the reconstruction of high-resolution, denoised, artifact-free three-dimensional, three-component velocity maps alongside corresponding 3D spin-density images.

### Advantages and Trade-Offs

- *Volumetric Flow Information Reconstruction:* In combination with the IP-PINN framework, this method acquires high-resolution, denoised, artifact-free 3D-3C velocity maps and 3D spin-density maps within the ROI.
- *Hardware and Pulse program modifications:* For each excited slice, the sequence applies a bipolar velocity-encoding gradient only along the designated axis. The gradients along the other two axes are set to zero. Standard slice-selective excitation and Cartesian readout are used, preserving typical echo times (TE) and repetition times (TR)

except that only one phase encode per slice is required. This sequence requires minor pulse program modification but leverages standard hardware and readout modules, facilitating integration into clinical scanners.

- *Shorter Overall Scan Time:* By eliminating the separate reference acquisition, the pseudo 1-point sequence reduces the total scan duration by approximately 75%.
- *Reduced Spatial Sampling for Each Velocity Direction:* Interleaving velocity encodes across successive slices lowers the effective spatial sampling density for each flow component, which may compromise fidelity of the acquired velocity fields. However, the IP-PINN framework compensates for this undersampling by applying super-resolution reconstruction to all three velocity components.

The following sections describe how the pseudo 1-point 4D Flow MRI sequence is incorporated into the IP-PINN framework to produce super-resolved 3D-3C velocity fields and spin-density maps with minimized noise and imaging artifacts, and then present validation results obtained from synthetic data.

### **6.3 IP-PINN Adoption for Pseudo 1-point 4D Flow MRI sequence**

This section describes how the developed IP-PINN in Chapter 3 is adapted to process proposed pseudo 1-point 4D Flow MRI data and reconstruct the super-resolved 3D-3C velocity maps and 3D spin-density maps within the ROI volume. Although the core architecture remains unchanged, the ROI selection, network’s input tensor, and k-space data loss function are described as follows.

### 6.3.1 ROI Selection, Input Tensor, and Input Parameterization

As described in Section 6.2, the pseudo 1-point 4D Flow MRI sequence encodes velocity by alternating the bipolar velocity-encoding gradient for each slice along the  $z$ -axis, acquiring only one velocity-encoded dataset per slice. A region of interest (ROI) is then selected in  $x$ -space with spatial dimensions  $n_x \times n_y \times n_z$  and  $n_t$  time frames. Under the pseudo 1-point scheme, each transverse slice of the ROI contains just one velocity-encoded scan. Consequently, for each velocity-encoded acquisition, the ROI spans spatial dimensions

$$n_x \times n_y \times \frac{n_z}{3}$$

over  $n_t$  time steps. This produces a complex input tensor

$$\mathcal{U} \in \mathbb{C}^{n_x \times n_y \times (n_z/3) \times n_t \times 6},$$

where the six channels correspond to the real and imaginary parts of the three velocity-encoded scans. In this chapter, we employ an ROI of size

$$n_x \times n_y \times n_z \times n_t = 14 \times 14 \times 6 \times 3.$$

The IP-PINN framework then uses  $\mathcal{U}$  to reconstruct pointwise velocity components  $(u, v, w)$  and the spin-density distribution within the ROI, enforcing constraints from both fluid-dynamics equations and the MRI signal model.

Following the input parameterization paradigm, the network’s flow image encoder  $\mathcal{G}_{\Theta_1}$  maps this ROI data block into a latent vector that will serve as a learnable “parameter” for the physics-informed neural network  $\mathcal{M}_{\Theta_2}$  and  $\mathcal{M}_{\Theta_3}$  (see chapter 3 section 3.3).

### 6.3.2 Implementing $k$ -space Data Fidelity Loss Function

In the following section, we turn our attention to the precise formulation and implementation of the  $k$ -space data-fidelity loss within the IP-PINN reconstruction. Rather than comparing reconstructed and reference images in the complex domain, we enforce consistency directly on the Fourier coefficients via an  $\ell_1$  penalty.

Since the Fourier transform is a linear transform, in order to isolate the correct  $k$ -space data of the selected ROI, we zero-pad the selected ROI to achieve the same spatial dimension size (the same field-of-view) as the imaging slices ( $N_x \times N_y$ ). Moreover, according to IP-PINN architecture, the physics-informed neural network branches  $\mathcal{M}_{\Theta_2}$  and  $\mathcal{M}_{\Theta_3}$  (see chapter 3 section 3.3) input the local spatio-temporal coordinates. Therefore, to maintain the consistency of the local coordinates of the various ROI inputs, the ROI is spatially translated to the center of the zero-padded full FOV. Due to the linearity and shift-invariance properties of the Fourier transform, the centered ROI and off-center ROI's Fourier transforms ( $k$ -space data) differ by a frequency-dependent linear phase term [162]. Specifically, according to the Fourier shift theorem, the relationship between their Fourier transforms is given by:

$$F_{\text{off}}(u, v) = \mathcal{F}\{I_{\text{off}}(x, y)\} = \exp\left[-i2\pi\left(\frac{u\Delta_x}{N_x} + \frac{v\Delta_y}{N_y}\right)\right] F_{\text{ROIcentered}}(u, v) \quad (6.1)$$

where  $(\Delta_x, \Delta_y)$  is the spatial offset between the two ROI placements, and  $(u, v)$  are the spatial-frequency indices in the Fourier domain.

To achieve equivalence between the  $k$ -space representations, a phase correction was applied in the frequency domain after computing the Fourier transform of the centered ROI. The phase correction is implemented as follows:

$$\hat{\mathcal{C}}_{\text{ROI,centered,corrected}}(u, v, z, t) = \hat{\mathcal{C}}_{\text{ROI,centered}}(u, v, z, t) \cdot \exp\left[-i2\pi\left(\frac{u\Delta_x}{N_x} + \frac{v\Delta_y}{N_y}\right)\right] \quad (6.2)$$

This frequency-domain phase correction method effectively aligns the two off-center and centered ROI's Fourier data, thus isolating the ROI's intrinsic spectral characteristics in-

dependent of its spatial location. Subsequently, the resultant  $k$ -space acquisition data  $\hat{C}_{ROI,centered,corrected}^j(k, z, t)$  are used to match the IP-PINN predicted frequency domain predicted apparent spin-density data  $\hat{\sigma}_D^j(\mathbf{k}, z, t)_{pred}$ . To this end, the IP-PINN framework predicts the point-wise velocity data and spin-density maps for the entire FOV. This follows by computing the predicted discretized spin density given the  $v_{enc}$  parameter and predicted composite phases  $\phi_B^j(\cdot)_{pred}$  according to Eqn. 2.8 and followed by converting to the Fourier domain and applying a windowing function to cut off the high-frequency components beyond  $\pm k^{max}$ , and applying the frequency-domain phase correction using the frequency-dependent linear phase term (Eqn. 6.1). Finally, we have the data fidelity loss function term as:

$$\mathcal{L}_D = \text{mean} \left[ \sum_{l=1}^{n_t} \sum_{(\mathbf{k}_m, z_m) \in \Omega} \left| \hat{\sigma}_{pred,centered,corrected}^{VENC\_dir(z_m)}(\mathbf{k}_m, z_m, t_l) - \hat{C}_{ROI,centered,corrected}(\mathbf{k}_m, z_m, t_l) \right| \right] \quad (6.3)$$

where,  $\Omega$  denotes the set of acquired  $(k_m, z_m)$  locations under the pseudo 1-point scheme,  $\hat{\sigma}_{pred,centered,corrected}^{VENC\_dir(z_m)}(\mathbf{k}_m, z_m, t_l)$  is the IP-PINN's predicted k-space value at slice  $z_m$  encoded along the direction  $VENC\_dir(z_m)$ ,  $\hat{C}_{ROI,centered,corrected}(\mathbf{k}_m, z_m, t_l)$  is the corresponding acquired k-space sample, and  $n_t$  is the number of cardiac time frames within the ROI.

The physics loss term is the same as Eqn. 3.22, which includes non-dimensional momentum balance and mass balance equations, imposed at an arbitrary number of points within the ROI domain.

## 6.4 Generation of Synthetic Training and Test Data

Two distinct 3D unsteady CFD simulations of Posterior Cerebral Artery (PCA) and Middle Cerebral Artery (MCA) aneurysms were conducted using SimVascular (see Chapter 3, Fig. 4.2). The PCA dataset was used for training and the MCA dataset for testing. Synthetic

4D-Flow MRI data were generated following the pipeline described in Chapter 4, section 4.2, with two key modifications:

1. The high-resolution magnitude image  $M(\mathbf{x}, z)$  was set to 1 within the vessel lumen and 0.5 in the surrounding static tissue.
2. A pseudo 1-point acquisition scheme was adopted, wherein only a single velocity-encoded scan per  $z$ -slice is acquired, cycling through the  $x$ ,  $y$ , and  $z$  directions on a per-slice basis.

### 6.4.1 Generating Synthetic Pseudo 1-point 4D Flow MRI Data

The CFD velocity fields were interpolated onto a uniform grid with isotropic voxel size  $\Delta x_f = 0.35$  mm. For the MCA simulation, every 15th time step was selected to achieve a 15 ms temporal resolution. The magnitude image was then assigned by

$$M(\mathbf{x}, z) = \begin{cases} 1, & (\mathbf{x}, z) \in \text{lumen}, \\ 0.5, & (\mathbf{x}, z) \notin \text{lumen}. \end{cases}$$

For each encoding direction  $j \in \{x, y, z\}$ , the complex signal was modeled as in Eqn. 2.8:

$$\sigma^j(\mathbf{x}, z, t) = M(\mathbf{x}, z) e^{i(\phi_B^j(\mathbf{x}, z) - \phi_u^j(\mathbf{x}, z, t))} \quad (6.4)$$

where  $\phi_B^j(\mathbf{x}, z)$  denotes the spatially varying non-velocity phase bias (see Chapter 2, Eqn. 2.10), and the velocity-induced phase shift is

$$\phi_u^j(\mathbf{x}, z, t) = \pi \frac{\mathbf{u}_j(\mathbf{x}, z, t)}{v_{\text{enc}}}. \quad (6.5)$$

To promote phase-wrapping (velocity-aliasing) artifacts, the encoding velocity parameter was set to

$$v_{\text{enc}} = 0.8 \max_{\mathbf{x}, z, t} |\mathbf{u}_j(\mathbf{x}, z, t)|.$$

The bias  $\phi_B$  was prescribed as a linear spatial function in the PCA dataset and a quadratic function in the MCA dataset.

Spatio-temporal downsampling followed the procedure of Chapter 4.2.1: first, spatial averaging over three adjacent voxels in the  $z$ -direction (resulting in an effective slice thickness of 1.05 mm), and then temporal averaging over four consecutive frames (yielding a 60 ms interval). Each downsampled slice in  $\mathbf{x}$ -space was Fourier-transformed to  $k$ -space, truncated beyond

$$k_{\max} = [2(3\Delta x_f)]^{-1} = 0.667 \text{ m}^{-1},$$

to achieve a coarse in-plane resolution of 1.05 mm. Zero-mean Gaussian noise (20 dB SNR) was added independently to both real and imaginary  $k$ -space channels.

After generating the low-resolution, artifact-laden  $k$ -space datasets as described above, the pseudo 1-point acquisition protocol of Section 6.2 was applied. For each slice  $s$ , only the  $k$ -space lines corresponding to the encoding direction  $\text{VENC\_dir}(s)$  were preserved, while all other channels were discarded. The resulting synthetic pseudo 1-point 4D-Flow MRI datasets thus emulate the proposed sequence and serve as the inputs and the training datasets to the IP-PINN reconstruction and evaluation presented in the following sections.

## 6.5 Error Metrics

To assess the fidelity of the reconstructed fields, we employ the suite of quantitative measures introduced in Section 4.3 of Chapter 4. Specifically, we report the total relative percentage error (TRPE), cosine similarity error (COSE), component-wise absolute error, velocity-magnitude scatter plots, the coefficient of determination ( $R^2$ ), and the median absolute percentage error (MDAPE).

## 6.6 Results

### 6.6.1 Pre-training IP-PINN on Pseudo 1-Point k-Space Data

The IP-PINN was first pre-trained using low-resolution pseudo 1-point k-space datasets generated from the PCA aneurysm simulation (Fig. 4.2(a)) following the procedure in Section 6.2.  $k$ -space data of input blocks of size  $n_x \times n_y \times n_z \times n_t = 14 \times 14 \times 6 \times 3$  were extracted. In total, approximately 4800 ROI blocks were used, with 5000 collocation points per block to enforce the Navier–Stokes physics constraint in the loss. The normalization scales remained

$$L = 0.8 \text{ cm}, \quad U = 93.65 \text{ cm/s}, \quad T = 0.0085 \text{ s},$$

and a batch size of two blocks was employed over 250 training epochs.

Loss term	Weight	Final value
Physics loss $\mathcal{L}_P$	0.09	$2.2526 \times 10^{-6}$
Data-fidelity loss $\mathcal{L}_D$	0.91	$3.0056 \times 10^{-4}$
Total (weighted sum)		$2.7371 \times 10^{-4}$

Table 6.1: Pre-training loss terms and weights for the pseudo 1-point k-space PCA dataset.

This pre-training on PCA aneurysm dataset establishes the baseline IP-PINN parameters before subsequent fine-tuning on the MCA dataset.

### 6.6.2 Applying Pre-trained IP-PINN to Unseen Pseudo 1-Point Data

After pre-training on the synthetic pseudo 1-point PCA dataset, the IP-PINN model was evaluated on the unseen pseudo 1-point dataset of the MCA aneurysm (Fig. 4.2(b)). To assess robustness to measurement noise, we tested the framework on both noise-free inputs and inputs corrupted with 20 dB SNR in k-space. Figures 6.2 and 6.3 present the reconstruction results under these two conditions. In each figure, the first column displays the k-space magnitude of the selected slices; the second column shows the corresponding low-resolution

velocity maps obtained from the phase of the complex images; and the third through fifth columns illustrate (i) the IP-PINN–predicted velocity fields, (ii) the CFD ground truth, and (iii) the point-wise absolute error between prediction and ground truth. Notably, although the pseudo 1-point acquisition provides only a single velocity-encoded direction per slice, the IP-PINN successfully recovers all three orthogonal velocity components for every slice.

Figure 6.4 shows the velocity-magnitude scatter plots and cosine-distance distributions for both noise-free and 20 dB SNR inputs. In the scatter plots, minimal magnitude error corresponds to close clustering about the  $45^\circ$  identity line, while low cosine-distance error indicates accurate directional reconstruction. Table 6.2 summarizes the total relative percentage error (TRPE), coefficient of determination ( $R^2$ ), and median absolute percentage error (MDAPE) for the velocity magnitude under each noise condition. As expected, performance degrades modestly in the presence of k-space noise. The reported processing time corresponds to the duration of the fine-tuning stage.

SNR	TRPE	$R^2$	MDAPE	Time
Noiseless	5.56%	0.992	2.15%	2 min
20 dB	8.88%	0.978	4.35%	2 min

Table 6.2: Quantitative performance on unseen pseudo 1-point MCA data: TRPE,  $R^2$ , MDAPE for velocity magnitude, and fine-tuning time.

### 6.6.3 Temporal Super-Resolution

To evaluate the IP-PINN’s ability to refine temporal resolution, we applied the fine-tuned model to predict flow fields at four times the sampling frequency of the input pseudo 1-point 4D-Flow MRI. The low-resolution under-sampled input at  $t_0 + 2\Delta t$ ,  $t_0 + 6\Delta t$ , and  $t_0 + 10\Delta t$  ( $\Delta t = 10$  ms) serve as the only observed time points. Here, where  $t_0$  is the starting time of the block of data being processed. Figure 6.5 compares the high-resolution predictions ( $u_{\text{pred}}, v_{\text{pred}}, w_{\text{pred}}$ ) against the CFD ground truth ( $u_{\text{GT}}, v_{\text{GT}}, w_{\text{GT}}$ ). The close correspondence across all velocity components demonstrates that the IP-PINN accurately interpolates the temporal dynamics, improving effective temporal resolution from 40 ms to 10 ms.

### 6.6.4 Luminal Boundary Prediction

This section demonstrates the IP-PINN’s ability to reconstruct high-resolution magnitude images that enable precise lumen boundary segmentation. Fig. 6.6 demonstrates the luminal boundary for a selected ROI. For comparison, the ground truth magnitude image is also presented. The figure additionally displays low-resolution MRI magnitude images of the ROI, which are synthetically generated using the corresponding reference CFD mesh data. These results are predicted by processing only velocity-encoded scan of the pseudo 1-point sequence. Notably, these results are obtained using only the single-direction, undersampled, velocity-encoded scans of the pseudo 1-point sequence, without any explicit volumetric magnitude or even volumetric velocity-encoded data. Despite this limited data, the IP-PINN framework successfully recovers the full volumetric magnitude field and accurately delineates the luminal boundary.

### 6.6.5 Performance Evaluation on Pseudo 1-Point Data with Velocity Aliasing

This section demonstrates the capability of the IP-PINN framework to resolve velocity aliasing artifacts in pseudo 1-point 4D-Flow MRI data. To simulate aliasing, the velocity encoding parameter  $v_{enc}$  was deliberately set to 70% of the maximum velocity magnitude in the MCA ground truth dataset. Additionally, zero-mean Gaussian noise was added to the k-space data at an SNR of  $20dB$ . Under these conditions, conventional MRI processing of the velocity-encoded scans yields phase-wrapped velocity fields with clear aliasing artifacts, particularly in regions of high flow (Fig. 6.7, Low Res  $v$ ).

Despite these challenges, the IP-PINN—trained with a k-space fidelity loss and physics-based regularization—successfully recovers the aliased components. As shown in Fig. 6.7, the predicted velocity map ( $v_{pred}$ ) correctly resolves the phase wraps and matches the unaliased ground truth. Table 6.3 reports the performance metrics, including total relative percentage error (TRPE, coefficient of determination ( $R^2$ ), and MDAPE values for the velocity mag-

nitude. These results confirm the model’s robustness against combined aliasing and noise, even in the undersampled pseudo 1-point acquisition setting.

SNR	TRPE	$R^2$	MDAPE
20 dB	7.83%	0.925	4.85%

Table 6.3: Performance metrics of IP-PINN on pseudo 1-point input data with 20 dB SNR and velocity aliasing. Metrics include TRPE,  $R^2$ , and MDAPE for velocity magnitude.

## 6.7 Discussion

This chapter presented a significant advancement in the development and application of the Input-Parameterized Physics-Informed Neural Network (IP-PINN) by adapting it to a novel 4D Flow MRI acquisition protocol: the *pseudo 1-point sequence*. Unlike conventional acquisition strategies that require three orthogonal velocity-encoded scans and one reference scan per slice, the proposed approach acquires only a single directional velocity-encoded scan per slice, with encoding directions interleaved along the z-dimension. This design reduces acquisition time by approximately 75%, directly addressing a key barrier to the clinical adoption of 4D Flow MRI.

A central innovation in this chapter is the formulation of the data fidelity loss in k-space, in contrast to the complex image-space fidelity used in Chapter 4. This k-space fidelity approach allows the model to process raw acquisition data and makes it compatible with the pseudo 1-point scheme, which yields highly undersampled and potentially aliased information. When coupled with physics-based constraints, this loss formulation enables the IP-PINN to recover physically plausible, high-fidelity volumetric flow fields from minimal directional input data.

We validated the proposed framework using synthetically generated pseudo 1-point datasets. Pre-training was conducted on the PCA aneurysm simulation, while fine-tuning was applied to the unseen MCA aneurysm data. Despite the sparse and directionally limited inputs, the IP-PINN successfully reconstructed full 3D-3C velocity fields and spin-density magnitude

maps with high accuracy. The framework was tested under several challenging conditions, including:

- Noise-free and noisy inputs (20 dB SNR),
- Temporal super-resolution (prediction at 10 ms intervals, 4× the input rate),
- Accurate segmentation of the vessel lumen based on reconstructed magnitude images,
- Correction of velocity aliasing due to deliberately low  $v_{enc}$  settings ( $0.7 v_{max}$ ) and added noise.

Quantitative performance metrics confirmed that IP-PINN accurately reconstructed flow fields, even in the presence of significant undersampling, noise, and aliasing. The model also demonstrated strong temporal interpolation capabilities and accurate recovery of anatomical boundaries from sparse data.

While the results are promising, several limitations merit consideration. First, the interleaving scheme applied on a per-slice basis ( $VENC\_dir(z_m)$ ) may introduce practical constraints in real-world acquisition. Future work could explore alternative strategies, such as cycling the velocity-encoding directions on a per-timeframe basis across the entire volume, to improve flexibility and reduce potential biases. Lastly, although synthetic testing provides a strong proof of concept, further validation on in-vitro data is necessary to assess the framework’s generalizability under more realistic and experimentally controlled conditions.

## 6.8 Conclusion

In this chapter, we proposed a novel 4D Flow MRI acquisition strategy, the pseudo 1-point sequence, and extended our IP-PINN framework to reconstruct high-resolution 3D-3C velocity and magnitude fields from undersampled velocity-encoded scans. The key contributions of this work include:

- Introduction of a pseudo 1-point acquisition protocol that reduces scan time by 75% compared to conventional 4D Flow MRI,
- Reformulation of the IP-PINN’s data fidelity loss in k-space, enabling direct use of raw acquisition data,
- Robust reconstruction of volumetric velocity fields and spin-density maps from only one directional velocity encoding per slice,
- Accurate performance in the presence of noise, aliasing, and sparse temporal sampling.

The results demonstrate that the IP-PINN framework not only improves flow image quality but also supports novel acquisition paradigms aimed at reducing clinical acquisition time and scanner burden. These findings highlight the potential of physics-informed learning to serve as a bridge between efficient data acquisition and high-fidelity flow imaging, thereby advancing the clinical utility of 4D Flow MRI.

The successful implementation of this framework in both conventional and pseudo 1-point contexts positions IP-PINN as a flexible and generalizable tool. In the next chapter, we synthesize the findings of this dissertation and outline future directions for clinical translation, real-data validation, and sequence-integrated model deployment.

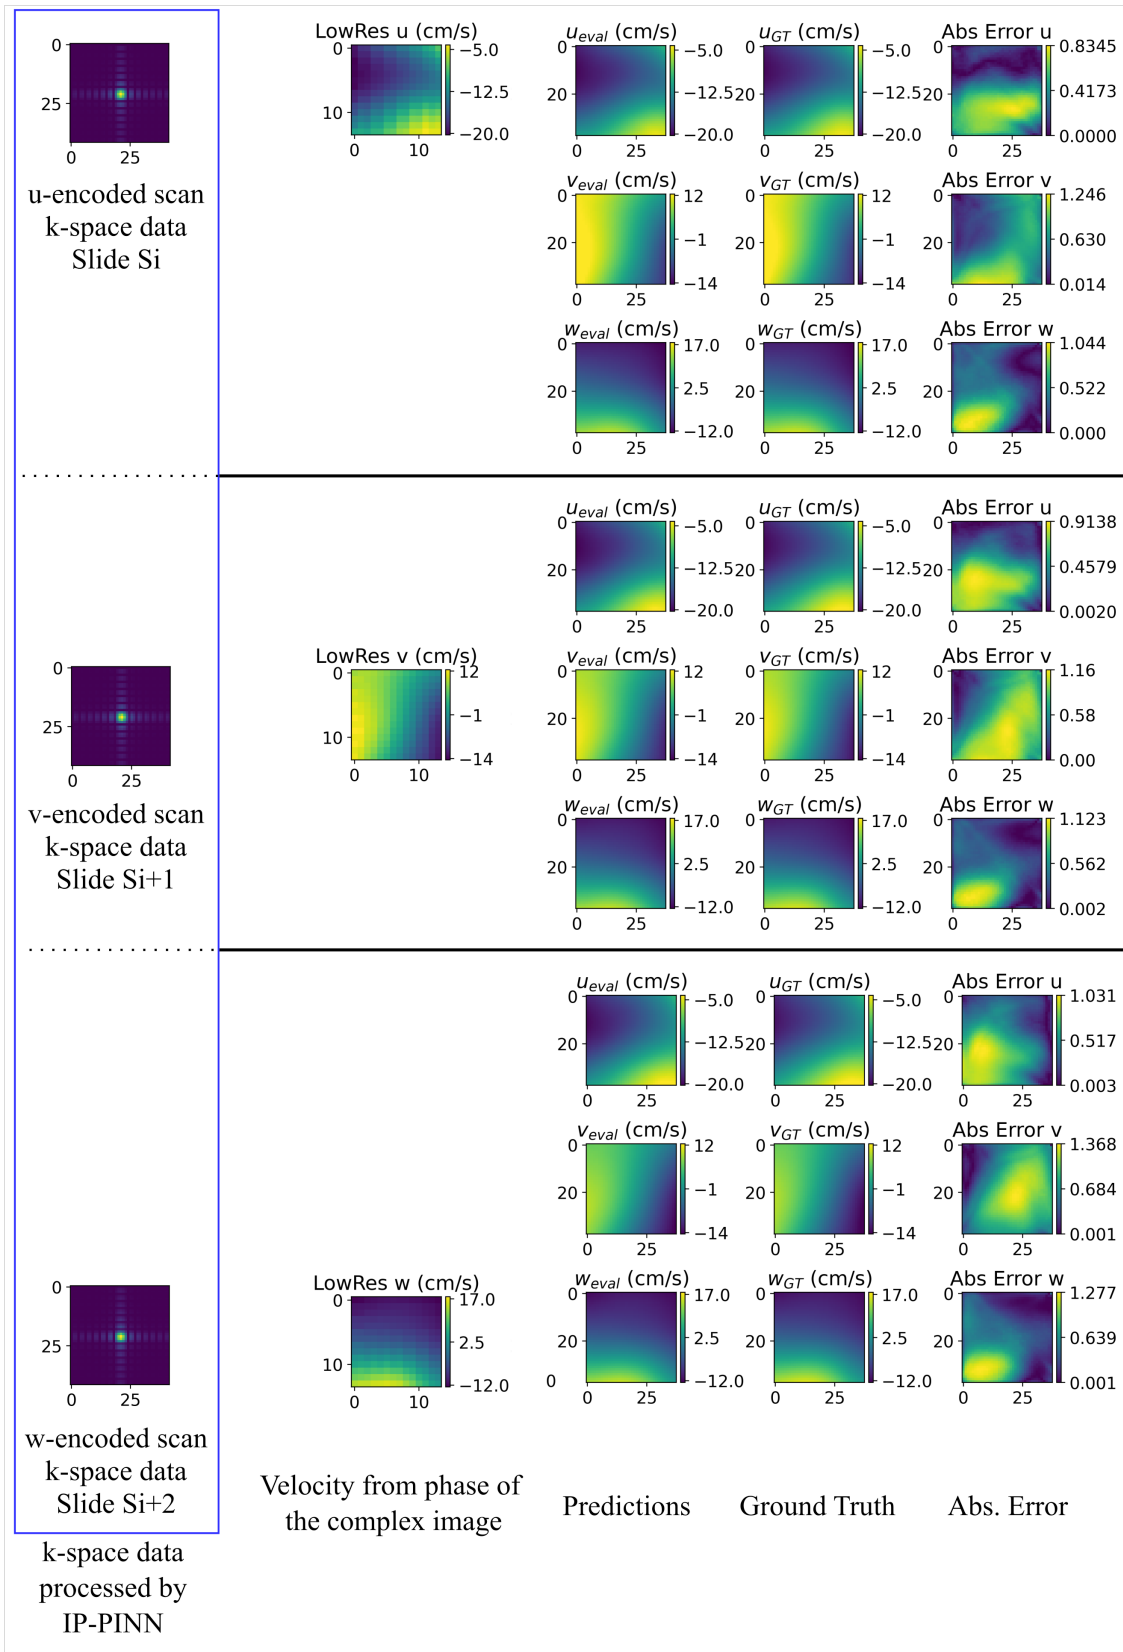


Figure 6.2: Super-resolution results on unseen pseudo 1-point MCA data with noise-free input.

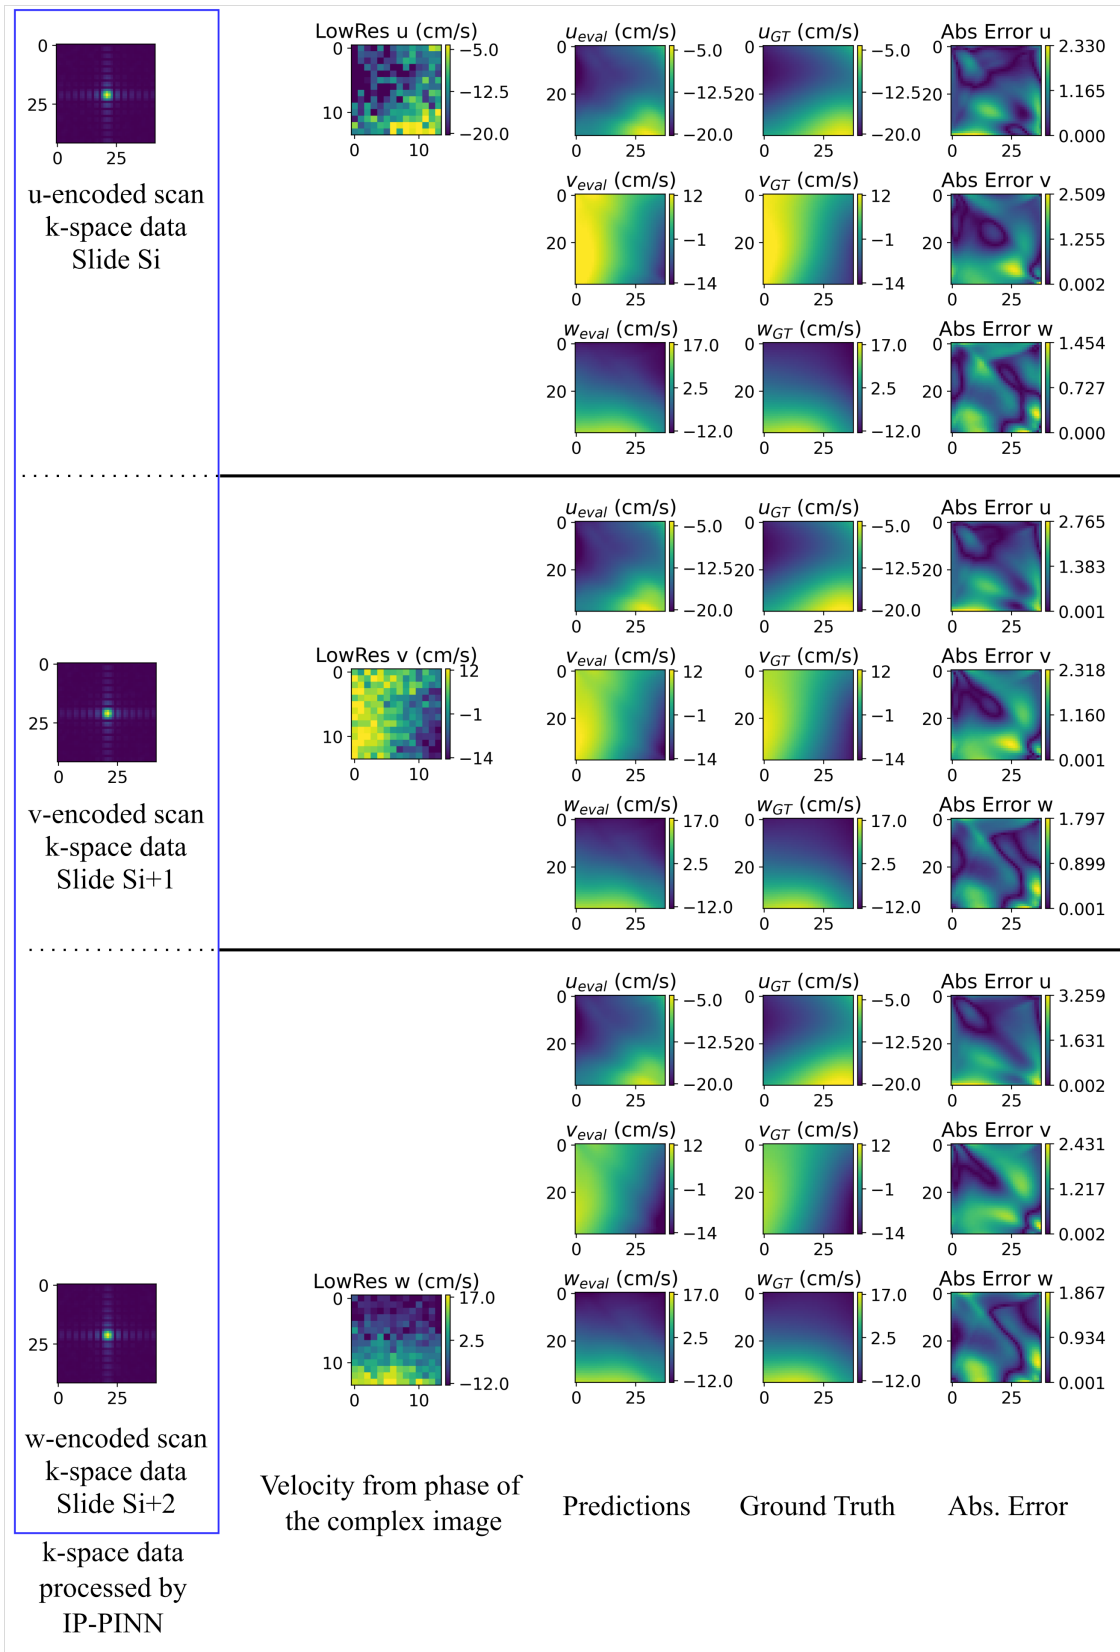


Figure 6.3: Super-resolution results on unseen pseudo 1-point MCA data with 20 dB SNR noise.

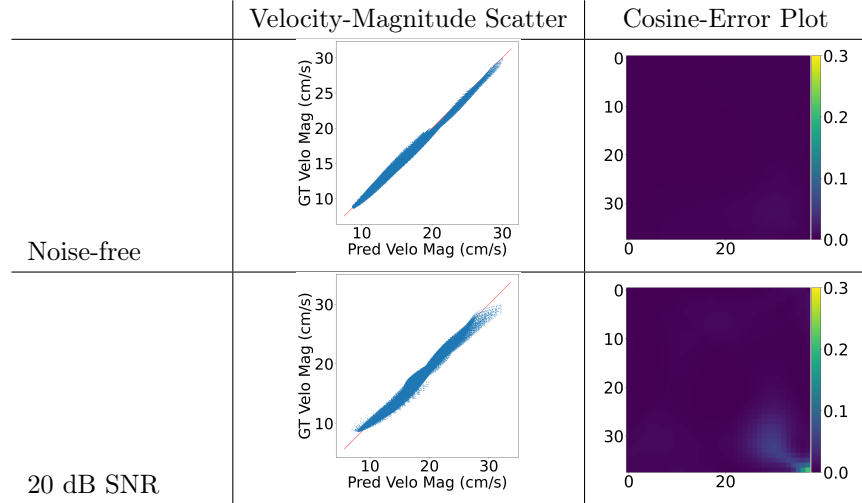


Figure 6.4: Comparison of scatter and cosine-error distributions for noise-free and 20 dB SNR inputs. Close alignment to the 45° line and low cosine-error values indicate accurate magnitude and directional recovery.

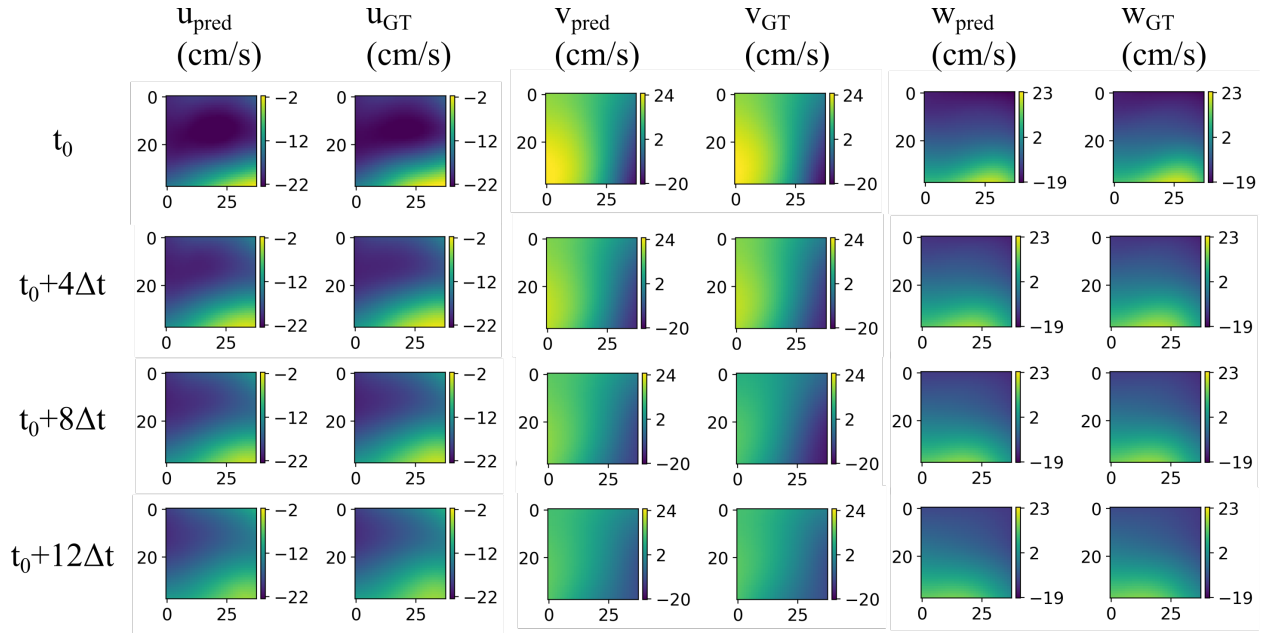


Figure 6.5: Temporal super-resolution on pseudo 1-point MCA data. Low-resolution velocity-encoded data are available at  $t_0 + 2\Delta t$ ,  $t_0 + 6\Delta t$ , and  $t_0 + 10\Delta t$  (with  $\Delta t = 10$  ms). Columns  $u_{\text{pred}}, v_{\text{pred}}, w_{\text{pred}}$  show the IP-PINN predictions at 10 ms intervals, while  $u_{\text{GT}}, v_{\text{GT}}, w_{\text{GT}}$  are the CFD ground-truth snapshots.




Figure 6.6: Luminal boundary prediction from pseudo 1-point acquisition data. **Left:** low-resolution magnitude image generated from the voxelized lumen geometry at the synthetic 4D-Flow MRI resolution. **Center:** IP-PINN-predicted segmentation of the magnitude image. **Right:** ground-truth lumen boundary voxelized on the CFD grid used for interpolation.




Figure 6.7: Velocity field reconstruction using IP-PINN on pseudo 1-point data with aliasing artifacts and  $20dB$  SNR noise. Red arrow highlights the region affected by velocity aliasing under low  $v_{enc}$ .

# Chapter 7

## Concluding Remarks

### 7.1 Summary of Contributions

In this thesis, we addressed several critical limitations of time-resolved phase-contrast magnetic resonance imaging (PC-MRI) that have hindered its widespread adoption in routine clinical practice. These limitations include low spatial resolution, k-space acquisition noise, non-velocity background phase errors, and velocity aliasing artifacts. To overcome these challenges, we developed a novel deep learning framework, termed the Input-Parameterized Physics-Informed Neural Network (IP-PINN), which integrates physical governing equations with learnable, input-driven representations of flow data. The IP-PINN framework was successfully implemented and evaluated across both major PC-MRI modalities: 4D Flow MRI and 2D PC-MRI. Our results demonstrate that IP-PINN effectively mitigates the aforementioned limitations, while also offering the potential to significantly reduce MRI acquisition time by eliminating the need for reference scans and enabling undersampling of velocity-encoded acquisitions. Additionally, the framework generates high-resolution spin-density magnitude images, which facilitate accurate segmentation of the vessel lumen—an essential prerequisite for reliable hemodynamic analysis in clinical applications.

Chapter 3 detailed the development of the Input-Parameterized Physics-Informed Neural Network (IP-PINN), a novel deep learning framework designed to overcome the limitations of traditional hemodynamic assessment techniques, including 4D Flow MRI and patient-specific CFD simulations. This chapter introduced the architectural foundations of IP-PINN, which combines convolutional neural networks (CNNs) and multilayer perceptrons (MLPs) with physics-informed neural network (PINN) principles to enable generalizable,

high-fidelity reconstruction of flow fields. The framework was designed to operate on input image data defined over localized spatio-temporal regions of interest (ROIs), and it outputs continuous volumetric velocity fields, pressure, and magnitude images. A key innovation introduced in this chapter is the parameterization of the PINN solutions using a CNN-based encoder that maps flow-encoded MRI data to a latent representation, enabling the model to generalize across varying anatomical and flow conditions without requiring ab initio retraining. The chapter also established a custom loss function that incorporates both data fidelity—computed in complex Cartesian image space to handle phase encoding physics—and a physics regularization term based on the incompressible Navier–Stokes equations. This hybrid loss formulation allows the network to produce physically consistent predictions while remaining robust to noise, aliasing, and phase offset artifacts. Collectively, the developments in Chapter 3 form the methodological backbone of this thesis and lay the groundwork for the application-focused evaluations presented in the subsequent chapters.

Chapter 4 rigorously evaluated the IP-PINN framework on synthetic yet physiologically realistic 4D-Flow MRI data. A complex-Cartesian data-fidelity term was introduced so that training could be performed directly on the acquired image data, thereby avoiding phase-unwrapping errors and making reference scans unnecessary. The network was first *pre-trained* on velocity-encoded images synthesized from a posterior cerebral-artery (PCA) aneurysm CFD model, and was subsequently *fine-tuned* on an unseen middle cerebral-artery (MCA) aneurysm dataset. Because the framework required only the three orthogonal velocity-encoded volumes—not the additional reference scan—the corresponding acquisition time would be reduced by roughly 25% in practice. Thanks to the input-parameterization strategy, fine-tuning each new region of interest took a mere  $\sim 1.5$  min, compared with the  $\sim 90$  min demanded by conventional PINNs, while still delivering super-resolved velocity fields, high-quality magnitude images, and robust suppression of noise, phase offsets, and velocity aliasing. These results confirm that IP-PINN generalizes effectively to previously unseen vascular anatomies while offering an order-of-magnitude acceleration over standard physics-

informed approaches.

Chapter 5 pushed the IP-PINN concept beyond volumetric (4D-Flow) imaging by showing that a single, slice-selective 2D PC-MRI acquisition can be turned into a high-fidelity four-dimensional dataset. We first introduced a minimally modified 2D PC-MRI sequence that cycles the bipolar gradients through the  $x$ ,  $y$  and,  $z$  directions across successive cardiac phases; because reference scans are omitted, this scheme already halves the nominal scan time. Feeding the three interleaved, complex  $x$ -space phase images into a pre-trained, input-parameterised PINN, We demonstrated that the network can infer super-resolved, noise- and alias-free 3D-3C velocity fields together with a 3D spin-density volume in the vicinity of the interrogation plane. Fine-tuning on an unseen MCA-aneurysm dataset required only 1.5 minutes. Crucially, the reconstructed 3D lumen geometry allowed computation of wall-shear-stress distributions that matched CFD ground-truth far better than conventional in-plane B-spline methods, underscoring the clinical value of volumetric information recovered from a fast, planar scan.

Chapter 6 proposes a *pseudo-1-point* 4D-Flow MRI sequence that acquires just a single velocity-encoded data set per slice and rotates the encoding axis slice-by-slice, thereby eliminating the reference scan and the two remaining orthogonal encodes. The raw data load therefore drops by roughly 75% compared with the conventional 4-point scheme, but the resulting inverse problem is extremely under-determined. To make the reconstruction tractable, we reformulated the IP-PINN so that its data-fidelity term is enforced directly in  $k$ -space rather than in complex Cartesian image  $x$ -space. A  $k$ -space loss gives several practical advantages: the generated  $k$ -space data by the IP-PINN’s predictions are compared with the exact quantities measured by the scanner and thus need not learn or invert the implicit filters that act after the inverse Fourier transform; local residuals remain local because each mismatch is confined to the Fourier coefficients that were actually sampled; the inevitable rect-window truncation, which would otherwise manifest as a global sinc-blur in image space, is handled transparently without forcing the network to perform an implicit deconvolution;

finally, physical effects that enter multiplicatively in  $k$ -space, such as  $T_2/T_2^*$  decay during read-out or coil-sensitivity weighting, can be incorporated in a straightforward, model-based fashion. Despite receiving just one directional encode per slice, fine-tuning the pre-trained revised IP-PINN for only two minutes on an unseen MCA aneurysm, the model produced super-resolved, denoised and alias-free three-component velocity fields together with a high-resolution magnitude image for the entire volume. Severe phase-wraps imposed by lower  $v_{enc}$  setting were completely resolved. The model also provided accurate lumen segmentation and temporal interpolation from 40 *ms* to 10 *ms*. Thus Chapter 6 demonstrates that a  $k$ -space IP-PINN can bridge the substantial gap between ultra-fast data acquisition and the volumetric, three-component flow information required for haemodynamic analysis.

## 7.2 Future Work

**Validation with in-vitro 4D-Flow MRI data.** All reconstructions in this dissertation were benchmarked against synthetic data generated from patient-specific CFD. The next logical step is to acquire time-resolved 4D-Flow MRI in a rigid, optically-clear flow phantom whose geometry and inlet waveforms replicate those of the simulated aneurysms. Because the ground-truth velocity can be measured independently with particle-image velocimetry, such experiments will provide an objective yard-stick for the absolute accuracy of IP-PINN under the full suite of scanner imperfections—gradient non-linearities, concomitant fields, eddy currents and coil sensitivity modulations—that synthetic data cannot fully emulate.

**Modelling and reconstructing turbulent flows.** The Navier–Stokes residual employed at present implicitly assumes laminar, newtonian fluid flows. Intracranial aneurysms at peak systole and stenotic aortic jets exhibit transitional or fully turbulent behaviour, for which the instantaneous velocity field contains sub-voxel eddies that cannot be resolved by MRI. One avenue is to augment the governing equations with an effective eddy-viscosity model, such as the Spalart–Allmaras or  $k$ – $\omega$  closure, and learn the turbulent viscosity as an auxiliary field. Alternatively, the ensemble-averaged Reynolds-averaged Navier–Stokes

(RANS) equations can be imposed, with the unresolved Reynolds stresses treated as latent variables inferred by the network. Either strategy will allow the reconstruction to remain faithful to the macroscopic MRI signal while respecting the altered momentum balance that turbulence introduces.

**Application to moving-boundary problems.** Cardiac chambers and large arteries undergo significant wall motion; yet all demonstrations herein assume a stationary domain. An interesting direction of research would include implementing IP-PINN and testing for moving boundaries of large arteries. Such a development would enable direct computation of wall shear stress and pressure on the moving endocardial surface and could obviate the need for separate segmentation and meshing pipelines.

**Extracting accurate lumen boundaries with level-set regularisation.** Although the network can already predict a high-resolution magnitude image, small intensity errors near the interface occasionally produce jagged or topology-inconsistent segmentations. Embedding a level-set evolution equation inside the training loop will provide an implicit, smooth representation of the boundary whose zero-level conforms to the predicted magnitude while being regularised by curvature or minimal-surface constraints. Coupling this signed-distance function to the velocity IP-PINN through a shared loss—e.g. penalising flux through the evolving surface—will ensure geometric consistency and furnish sub-voxel accurate normals for subsequent wall-shear calculations. The level-set could be updated by differentiable re-initialisation so that the entire pipeline remains end-to-end trainable.

## BIBLIOGRAPHY

- [1] R. J. Adrian and J. Westerweel. *Particle image velocimetry*. Number 30. Cambridge university press, 2011.
- [2] A. Aghaee and M. O. Khan. Performance of fourier-based activation function in physics-informed neural networks for patient-specific cardiovascular flows. *Computer Methods and Programs in Biomedicine*, 247:108081, 2024.
- [3] V. Alakbarzade and A. C. Pereira. Cerebral catheter angiography and its complications, 2018.
- [4] M. Aliakbari, M. Mahmoudi, P. Vadasz, and A. Arzani. Predicting high-fidelity multiphysics data from low-fidelity fluid flow and transport solvers using physics-informed neural networks. *International Journal of Heat and Fluid Flow*, 96:109002, 2022.
- [5] O. Amili, D. Schiavazzi, S. Moen, B. Jagadeesan, P.-F. Van de Moortele, and F. Coletti. Hemodynamics in a giant intracranial aneurysm characterized by in vitro 4d flow mri. *Plos one*, 13(1):e0188323, 2018.
- [6] N. S. Anavekar and J. K. Oh. Doppler echocardiography: a contemporary review. *Journal of cardiology*, 54(3):347–358, 2009.
- [7] G. Annio, R. Torii, B. Ariff, D. P. O’Regan, V. Muthurangu, A. Ducci, V. Tsang, and G. Burriesci. Enhancing magnetic resonance imaging with computational fluid dynamics. *Journal of Engineering and Science in Medical Diagnostics and Therapy*, 2(4):041010, 2019.
- [8] A. Arzani and S. C. Shadden. Characterizations and correlations of wall shear stress in aneurysmal flow. *Journal of Biomechanical Engineering*, 138(1):014503, 2016.

- [9] A. Bakhshinejad. *Novel Algorithms for Merging Computational Fluid Dynamics and 4D Flow MRI*. PhD thesis, The University of Wisconsin-Milwaukee, 2018.
- [10] A. Bakhshinejad, A. Baghaie, A. Vali, D. Saloner, V. L. Rayz, and R. M. D’Souza. Merging computational fluid dynamics and 4d flow mri using proper orthogonal decomposition and ridge regression. *Journal of biomechanics*, 58:162–173, 2017.
- [11] A. J. Barker, P. Van Ooij, K. Bandi, J. Garcia, M. Albaghdadi, P. McCarthy, R. O. Bonow, J. Carr, J. Collins, S. C. Malaisrie, et al. Viscous energy loss in the presence of abnormal aortic flow. *Magnetic resonance in medicine*, 72(3):620–628, 2014.
- [12] 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.
- [13] P. Berg, C. Roloff, O. Beuing, S. Voss, S.-I. Sugiyama, N. Aristokleous, A. S. Anayiotos, N. Ashton, A. Revell, N. W. Bressloff, et al. The computational fluid dynamics rupture challenge 2013—phase ii: variability of hemodynamic simulations in two intracranial aneurysms. *Journal of biomechanical engineering*, 137(12):121008, 2015.
- [14] H. Berhane, M. B. Scott, A. J. Barker, P. McCarthy, R. Avery, B. Allen, C. Malaisrie, J. D. Robinson, C. K. Rigsby, and M. Markl. Deep learning–based velocity antialiasing of 4d-flow mri. *Magnetic resonance in medicine*, 88(1):449–463, 2022.
- [15] R. B. Bird, R. C. Armstrong, and O. Hassager. *Dynamics of polymeric liquids. Vol. 1: Fluid mechanics*. John Wiley and Sons Inc., New York, NY, 1987.
- [16] M. M. Bissell, A. T. Hess, L. Biasioli, S. J. Glaze, M. Loudon, A. Pitcher, A. Davis, B. Prendergast, M. Markl, A. J. Barker, S. Neubauer, and S. G. Myerson. Aortic dilation in bicuspid aortic valve disease: flow pattern is a major contributor and differs with valve fusion type. *Circulation. Cardiovascular imaging*, 6(4):499–507, jul 2013.

- [17] K. T. Block, M. Uecker, and J. Frahm. Suppression of mri truncation artifacts using total variation constrained data extrapolation. *International journal of biomedical imaging*, 2008(1):184123, 2008.
- [18] L. Bohs, B. Geiman, M. Anderson, S. Gebhart, and G. Trahey. Speckle tracking for multi-dimensional flow estimation. *Ultrasonics*, 38(1-8):369–375, 2000.
- [19] E. Bollache, P. van Ooij, A. Powell, J. Carr, M. Markl, and A. J. Barker. Comparison of 4d flow and 2d velocity-encoded phase contrast mri sequences for the evaluation of aortic hemodynamics. *The international journal of cardiovascular imaging*, 32:1529–1541, 2016.
- [20] L. Boussel, V. Rayz, A. Martin, G. Acevedo-Bolton, M. T. Lawton, R. Higashida, W. S. Smith, W. L. Young, and D. Saloner. Phase-contrast magnetic resonance imaging measurements in intracranial aneurysms in vivo of flow patterns, velocity fields, and wall shear stress: comparison with computational fluid dynamics. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 61(2):409–417, 2009.
- [21] L. Boussel, V. Rayz, C. McCulloch, A. Martin, G. Acevedo-Bolton, M. Lawton, R. Higashida, W. S. Smith, W. L. Young, and D. Saloner. Aneurysm Growth Occurs at Region of Low Wall Shear Stress: Patient-Specific Correlation of Hemodynamics and Growth in a Longitudinal Study. *Stroke*, 39(11):2997–3002, nov 2008.
- [22] O. Brenner, J. Plogmann, P. Piroozmand, and P. Jenny. A variational data assimilation approach for sparse velocity reference data in coarse rans simulations through a corrective forcing term. *Computer Methods in Applied Mechanics and Engineering*, 427:117026, 2024.
- [23] L. Brix, S. Ringgaard, A. Rasmusson, T. S. Sørensen, and W. Y. Kim. Three dimensional three component whole heart cardiovascular magnetic resonance velocity

- mapping: comparison of flow measurements from 3d and 2d acquisitions. *Journal of Cardiovascular Magnetic Resonance*, 11(1):3, 2009.
- [24] J. Busch, D. Giese, and S. Kozerke. Image-based background phase error correction in 4d flow mri revisited. *Journal of Magnetic Resonance Imaging*, 46(5):1516–1525, 2017.
- [25] J. Busch, D. Giese, L. Wissmann, and S. Kozerke. Reconstruction of divergence-free velocity fields from cine 3d phase-contrast flow measurements. *Magnetic resonance in medicine*, 69(1):200–210, 2013.
- [26] F. M. Callaghan, R. Kozor, A. G. Sherrah, M. Vallely, D. Celermajer, G. A. Figtree, and S. M. Grieve. Use of multi-velocity encoding 4d flow mri to improve quantification of flow patterns in the aorta. *Journal of Magnetic Resonance Imaging*, 43(2):352–363, 2016.
- [27] M. Carlsson, J. Töger, M. Kanski, K. Markenroth Bloch, F. Ståhlberg, E. Heiberg, and H. Arheden. Quantification and visualization of cardiovascular 4d velocity mapping accelerated with parallel imaging or kt blast: head to head comparison and validation at 1.5 t and 3 t. *Journal of Cardiovascular Magnetic Resonance*, 13:1–7, 2011.
- [28] J. R. Cebral, C. M. Putman, M. T. Alley, T. Hope, R. Bammer, and F. Calamante. Hemodynamics in normal cerebral arteries: qualitative comparison of 4d phase-contrast magnetic resonance and image-based computational fluid dynamics. *Journal of engineering mathematics*, 64(4):367–378, 2009.
- [29] E. Cecchi, C. Giglioli, S. Valente, C. Lazzeri, G. F. Gensini, R. Abbate, and L. Manini. Role of hemodynamic shear stress in cardiovascular disease. *Atherosclerosis*, 214(2):249–256, 2011.
- [30] P. Chai and R. Mohiaddin. How we perform cardiovascular magnetic resonance flow assessment using phase-contrast velocity mapping. *Journal of Cardiovascular Magnetic Resonance*, 7(4):705–716, 2005.

- [31] S. Chakraborty. Transfer learning based multi-fidelity physics informed deep neural network. *Journal of Computational Physics*, 426:109942, 2021.
- [32] K. Y. Chan, B. Abu-Salih, R. Qaddoura, A.-Z. Ala'M, V. Palade, D.-S. Pham, J. Del Ser, and K. Muhammad. Deep neural networks in the cloud: Review, applications, challenges and research directions. *Neurocomputing*, 545:126327, 2023.
- [33] M. Cibis, K. Jarvis, M. Markl, M. Rose, C. Rigsby, A. J. Barker, and J. J. Wentzel. The effect of resolution on viscous dissipation measured with 4d flow mri in patients with fontan circulation: Evaluation using computational fluid dynamics. *Journal of biomechanics*, 48(12):2984–2989, 2015.
- [34] T. J. Colgan, D. Hernando, S. D. Sharma, and S. B. Reeder. The effects of concomitant gradients on chemical shift encoded mri. *Magnetic resonance in medicine*, 78(2):730–738, 2017.
- [35] D. Craiem, A. F. Pascaner, M. E. Casciaro, U. Gencer, J. Alcibar, G. Soulat, and E. Mousseaux. Automatic correction of background phase offset in 4d-flow of great vessels and of the heart in mri using a third-order surface model. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 32:629–642, 2019.
- [36] D. S. de Azevedo, A. S. M. Salinet, M. de Lima Oliveira, M. J. Teixeira, E. Bor-Seng-Shu, and R. de Carvalho Nogueira. Cerebral hemodynamics in sepsis assessed by transcranial doppler: a systematic review and meta-analysis. *Journal of clinical monitoring and computing*, 31(6):1123–1132, 2017.
- [37] M. Dehghani and H. Samet. Momentum search algorithm: A new meta-heuristic optimization algorithm inspired by momentum conservation law. *SN Applied Sciences*, 2(10):1720, 2020.
- [38] N. DeMarchi and C. White. Echo particle image velocimetry. *Journal of visualized experiments: JoVE*, (70), 2012.

- [39] F. Donati, C. A. Figueroa, N. P. Smith, P. Lamata, and D. A. Nordsletten. Non-invasive pressure difference estimation from pc-mri using the work-energy equation. *Medical image analysis*, 26(1):159–172, 2015.
- [40] M. M. Dua and R. L. Dalman. Hemodynamic influences on abdominal aortic aneurysm disease: Application of biomechanics to aneurysm pathophysiology. *Vascular pharmacology*, 53(1-2):11–21, 2010.
- [41] S. T. Duong, S. L. Phung, A. Bouzerdoun, H. G. Boyd Taylor, A. M. Puckett, and M. M. Schira. Anatomy-guided inverse-phase-encoding registration method for correcting susceptibility artifacts in sub-millimeter fmri. *bioRxiv*, page 779272, 2019.
- [42] M. Dushkoff and R. Ptucha. Adaptive activation functions for deep networks. *Electronic Imaging*, 28:1–5, 2016.
- [43] P. Dyverfeldt, M. Bissell, A. J. Barker, A. F. Bolger, C.-J. Carlhäll, T. Ebbers, C. J. Francios, A. Frydrychowicz, J. Geiger, D. Giese, et al. 4d flow cardiovascular magnetic resonance consensus statement. *Journal of Cardiovascular Magnetic Resonance*, 17(1):1–19, 2015.
- [44] M. D’Elia, L. Mirabella, T. Passerini, M. Perego, M. Piccinelli, C. Vergara, and A. Veneziani. Applications of variational data assimilation in computational hemodynamics. *Modeling of physiological flows*, pages 363–394, 2012.
- [45] T. Ebbers, L. Wigström, A. F. Bolger, J. Engvall, and M. Karlsson. Estimation of relative cardiovascular pressures using time-resolved three-dimensional phase contrast mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 45(5):872–879, 2001.
- [46] 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*, 12:1–10, 2010.

- [47] G. Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53:343–367, 2003.
- [48] G. Evensen. The ensemble kalman filter for combined state and parameter estimation. *IEEE Control Systems Magazine*, 29(3):83–104, 2009.
- [49] L. Fang, S. Li, R. P. McNabb, Q. Nie, A. N. Kuo, C. A. Toth, J. A. Izatt, and S. Farsiu. Fast acquisition and reconstruction of optical coherence tomography images via sparse representation. *IEEE transactions on medical imaging*, 32(11):2034–2049, 2013.
- [50] M. F. Fathi, A. Bakhshinejad, A. Baghaie, D. Saloner, R. H. Sacho, V. L. Rayz, and R. M. D’Souza. Denoising and spatial resolution enhancement of 4d flow mri using proper orthogonal decomposition and lasso regularization. *Computerized Medical Imaging and Graphics*, 70:165–172, 2018.
- [51] M. F. Fathi, A. Bakhshinejad, A. Baghaie, D. Saloner, R. H. Sacho, V. L. Rayz, and R. M. D’Souza. Denoising and spatial resolution enhancement of 4d flow mri using proper orthogonal decomposition and lasso regularization. *Computerized Medical Imaging and Graphics*, 70:165–172, 2018.
- [52] 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 nets. *Computer Methods and Programs in Biomedicine*, 197:105729, 2020.
- [53] 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 nets. *Computer Methods and Programs in Biomedicine*, 197:105729, 2020.
- [54] E. Ferdian, D. Marlevi, J. Schollenberger, M. Aristova, E. R. Edelman, S. Schnell, C. A. Figueroa, D. A. Nordsletten, and A. A. Young. Cerebrovascular super-resolution 4d flow mri—using deep learning to non-invasively quantify velocity, flow, and relative pressure. *bioRxiv*, 2021.

- [55] E. Ferdian, A. Suinesiaputra, D. J. Dubowitz, D. Zhao, A. Wang, B. Cowan, and A. A. Young. 4dflownet: super-resolution 4d flow mri using deep learning and computational fluid dynamics. *Frontiers in Physics*, page 138, 2020.
- [56] E. Ferdian, A. Suinesiaputra, D. J. Dubowitz, D. Zhao, A. Wang, B. Cowan, and A. A. Young. 4dflownet: super-resolution 4d flow mri using deep learning and computational fluid dynamics. *Frontiers in Physics*, 8:138, 2020.
- [57] S. Ferrari, S. Ambrogio, A. Walker, P. Verma, A. J. Narracott, I. Wilkinson, and J. W. Fenner. The ring vortex: concepts for a novel complex flow phantom for medical imaging. *Open Journal of Medical Imaging*, 7(1):28–41, 2017.
- [58] D. Firmin, P. Gatehouse, J. Konrad, G. Yang, P. Kilner, and D. Longmore. Rapid 7-dimensional imaging of pulsatile flow. In *Proceedings of computers in cardiology conference*, pages 353–356. IEEE, 1993.
- [59] A. Frydrychowicz, O. Wieben, E. Niespodzany, S. B. Reeder, K. M. Johnson, and C. J. François. Quantification of thoracic blood flow using volumetric magnetic resonance imaging with radial velocity encoding: in vivo validation. *Investigative radiology*, 48(12):819–825, 2013.
- [60] S. W. Funke, M. Nordaas, Ø. Evju, M. S. Alnæs, and K. A. Mardal. Variational data assimilation for transient blood flow simulations: cerebral aneurysms as an illustrative example. *International journal for numerical methods in biomedical engineering*, 35(1):e3152, 2019.
- [61] S. W. Funke, M. Nordaas, Ø. Evju, M. S. Alnæs, and K. A. Mardal. Variational data assimilation for transient blood flow simulations: cerebral aneurysms as an illustrative example. *International journal for numerical methods in biomedical engineering*, 35(1):e3152, 2019.

- [62] F. Gaidzik, S. Pathiraja, S. Saalfeld, D. Stucht, O. Speck, D. Thévenin, and G. Janiga. Hemodynamic data assimilation in a subject-specific circle of willis geometry. *Clinical Neuroradiology*, 31(3):643–651, 2021.
- [63] P. Garg, J. J. Westenberg, P. J. van den Boogaard, P. P. Swoboda, R. Aziz, J. R. Foley, G. J. Fent, F. Tyl, L. Coratella, M. S. ElBaz, et al. Comparison of fast acquisition strategies in whole-heart four-dimensional flow cardiac mr: Two-center, 1.5 tesla, phantom and in vivo validation study. *Journal of Magnetic Resonance Imaging*, 47(1):272–281, 2018.
- [64] A. Geers, I. Larrabide, A. G. Radaelli, H. Bogunovic, M. Kim, H. G. Van Andel, C. Majoie, E. VanBavel, and A. Frangi. Patient-specific computational hemodynamics of intracranial aneurysms from 3d rotational angiography and ct angiography: an in vivo reproducibility study. *American Journal of Neuroradiology*, 32(3):581–586, 2011.
- [65] H. Goldsmith and R. Skalak. Hemodynamics. *Annual Review of Fluid Mechanics*, 7(1):213–247, 1975.
- [66] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, et al. Recent advances in convolutional neural networks. *Pattern recognition*, 77:354–377, 2018.
- [67] H. Ha, G. B. Kim, J. Kweon, S. J. Lee, Y.-H. Kim, D. H. Lee, D. H. Yang, and N. Kim. Hemodynamic measurement using four-dimensional phase-contrast mri: quantification of hemodynamic parameters and clinical applications. *Korean journal of radiology*, 17(4):445–462, 2016.
- [68] M. Habibi, R. M. D’Souza, S. T. Dawson, and A. Arzani. Integrating multi-fidelity blood flow data with reduced-order data assimilation. *Computers in Biology and Medicine*, 135:104566, 2021.

- [69] A. Harloff, A. Nußbaumer, S. Bauer, A. F. Stalder, A. Frydrychowicz, C. Weiller, J. Hennig, and M. Markl. In vivo assessment of wall shear stress in the atherosclerotic aorta using flow-sensitive 4d mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 63(6):1529–1536, 2010.
- [70] L. Hatle, A. Brubakk, A. Tromsdal, and B. Angelsen. Noninvasive assessment of pressure drop in mitral stenosis by doppler ultrasound. *Heart*, 40(2):131–140, 1978.
- [71] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [72] W. Hong, H. Yu, J. Chen, J. Talamantes, D. M. Rollins, X. Fang, J. Long, C. Xu, and A. P. Sawchuk. A mock circulation loop to characterize in vitro hemodynamics in human systemic arteries with stenosis. *Fluids*, 8(7):198, 2023.
- [73] A. Hsiao, M. Lustig, M. T. Alley, M. J. Murphy, and S. S. Vasanaawala. Evaluation of valvular insufficiency and shunts with parallel-imaging compressed-sensing 4d phase-contrast mr imaging with stereoscopic 3d velocity-fusion volume-rendered visualization. *Radiology*, 265(1):87–95, 2012.
- [74] A. D. Jagtap, K. Kawaguchi, and G. E. Karniadakis. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, 404:109136, 2020.
- [75] K. M. Johnson and M. Markl. Improved snr in phase contrast velocimetry with five-point balanced flow encoding. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 63(2):349–355, 2010.
- [76] B. Jung, M. Honal, P. Ullmann, J. Hennig, and M. Markl. Highly k-t-space-accelerated

- phase-contrast mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 60(5):1169–1177, 2008.
- [77] B. Jung, C. Müller, W. Buchenberg, M. Ith, D. Reineke, F. Beyersdorf, and C. Benk. Investigation of hemodynamics in an in vitro system simulating left ventricular support through the right subclavian artery using 4-dimensional flow magnetic resonance imaging. *The Journal of thoracic and cardiovascular surgery*, 150(1):200–207, 2015.
- [78] D. F. Kallmes. Point: Cfd—computational fluid dynamics or confounding factor dissemination. *AJNR: American Journal of Neuroradiology*, 33(3):395, 2012.
- [79] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35, 1960.
- [80] C. Karmonik, R. Klucznik, and G. Benndorf. Comparison of velocity patterns in an acoma aneurysm measured with 2d phase contrast mri and simulated with cfd. *Technology and Health Care*, 16(2):119–128, 2008.
- [81] K. Kellaris. *Application of Data Assimilation Methods in Computational Fluid Dynamics*. PhD thesis, National Technical University of Athens, 2023.
- [82] G. Kerschen, J.-c. Golinval, A. F. Vakakis, and L. A. Bergman. The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: an overview. *Nonlinear dynamics*, 41(1):147–169, 2005.
- [83] A. Kheradvar, H. Houle, G. Pedrizzetti, G. Tonti, T. Belcik, M. Ashraf, J. R. Lindner, M. Gharib, and D. Sahn. Echocardiographic particle image velocimetry: a novel technique for quantification of left ventricular blood vorticity pattern. *Journal of the American Society of Echocardiography*, 23(1):86–94, 2010.
- [84] H.-B. Kim, J. R. Hertzberg, and R. Shandas. Development and validation of echo piv. *Experiments in fluids*, 36:455–462, 2004.

- [85] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [86] 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.
- [87] 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.
- [88] F. Klemens, S. Schuhmann, R. Balbierer, G. Guthausen, H. Nirschl, G. Thäter, and M. J. Krause. Noise reduction of flow mri measurements using a lattice boltzmann based topology optimisation approach. *Computers & Fluids*, 197:104391, 2020.
- [89] T. S. Koltukluoglu. *Data Assimilation in Computational Hemodynamics*. PhD thesis, ETH Zurich, 2019.
- [90] A. Kontogiannis and M. P. Juniper. Physics-informed compressed sensing for pc-mri: an inverse navier-stokes problem. *IEEE Transactions on Image Processing*, 32:281–294, 2022.
- [91] M. S. Kraus, A. C. Coblentz, V. S. Deshpande, J. M. Peeters, P. M. Itriago-Leon, and G. B. Chavhan. State-of-the-art magnetic resonance imaging sequences for pediatric body imaging. *Pediatric Radiology*, 53(7):1285–1299, 2023.
- [92] M. J. Krause, G. Thäter, and V. Heuveline. Adjoint-based fluid flow control and optimisation with lattice boltzmann methods. *Computers & Mathematics with Applications*, 65(6):945–960, 2013.
- [93] F. W. Kremkau. Diagnostic ultrasound: principles and instruments. (*No Title*), 1998.

- [94] Z. Kulcsár, A. Ugron, M. Marosfői, Z. Berentei, G. Paal, and I. Szikora. Hemodynamics of cerebral aneurysm initiation: the role of wall shear stress and spatial wall shear stress gradient. *American Journal of neuroradiology*, 32(3):587–594, 2011.
- [95] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [96] A. Leopaldi, R. Vismara, M. Lemma, L. Valerio, M. Cervo, A. Mangini, M. Contino, A. Redaelli, C. Antona, and G. B. Fiore. In vitro hemodynamics and valve imaging in passive beating hearts. *Journal of biomechanics*, 45(7):1133–1139, 2012.
- [97] B. Li, Y. Li, and X. Rong. The extreme learning machine learning algorithm with tunable activation function. *Neural Computing and Applications*, 22:531–539, 2013.
- [98] 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.
- [99] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998.
- [100] J. Lotz, C. Meier, A. Leppert, and M. Galanski. Cardiovascular flow measurement with phase-contrast mr imaging: basic facts and implementation. *Radiographics*, 22(3):651–671, 2002.
- [101] L. E. Ma, M. Markl, K. Chow, H. Huh, C. Forman, A. Vali, A. Greiser, J. Carr, S. Schnell, A. J. Barker, et al. Aortic 4d flow mri in 2 minutes using compressed sensing, respiratory controlled adaptive k-space reordering, and inline reconstruction. *Magnetic resonance in medicine*, 81(6):3675–3690, 2019.

- [102] S. Magder. Invasive hemodynamic monitoring. *Critical care clinics*, 31(1):67–87, 2015.
- [103] M. Markl, R. Bammer, M. Alley, C. Elkins, M. Draney, A. Barnett, M. Moseley, G. Glover, and N. Pelc. Generalized reconstruction of phase contrast mri: analysis and correction of the effect of gradient field distortions. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 50(4):791–801, 2003.
- [104] M. Markl, F. P. Chan, M. T. Alley, K. L. Wedding, M. T. Draney, C. J. Elkins, D. W. Parker, R. Wicker, C. A. Taylor, R. J. Herfkens, et al. Time-resolved three-dimensional phase-contrast mri. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 17(4):499–506, 2003.
- [105] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben. 4d flow mri. *Journal of Magnetic Resonance Imaging*, 36(5):1015–1036, 2012.
- [106] M. Markl, P. J. Kilner, and T. Ebbers. Comprehensive 4d velocity mapping of the heart and great vessels by cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 13:1–22, 2011.
- [107] M. Markl, S. Schnell, C. Wu, E. Bollache, K. Jarvis, A. Barker, J. Robinson, and C. Rigsby. Advanced flow mri: emerging techniques and applications. *Clinical radiology*, 71(8):779–795, 2016.
- [108] D. Marlevi, B. Ruijsink, M. Balmus, D. Dillon-Murphy, D. Fovargue, K. Pushparajah, C. Bertoglio, M. Colarieti-Tosti, M. Larsson, P. Lamata, et al. Estimation of cardiovascular relative pressure using virtual work-energy. *Scientific reports*, 9(1):1–16, 2019.
- [109] I. Marshall, S. Zhao, P. Papathanasopoulou, P. Hoskins, and X. Y. Xu. Mri and cfd studies of pulsatile flow in healthy and stenosed carotid bifurcation models. *Journal of Biomechanics*, 37(5):679–687, 2004.

- [110] L. Mejia, B. B. Ergin, and L. Rivera Lara. Neuro-icu: Cerebral hemodynamics and transcranial doppler (tcd/tccs) waveform interpretation in the most common neurocritical pathologies. In *Neurosonology in Critical Care*, pages 299–318. Springer, 2022.
- [111] H. Meng, V. Tutino, J. Xiang, and A. Siddiqui. High wss or low wss? complex interactions of hemodynamics with intracranial aneurysm initiation, growth, and rupture: toward a unifying hypothesis. *American Journal of Neuroradiology*, 2013.
- [112] J. Meunier and M. Bertrand. Ultrasonic texture motion analysis: theory and simulation. *IEEE transactions on medical imaging*, 14(2):293–300, 1995.
- [113] M. Michael, F. Alex, K. Sebastian, H. Mike, and W. Oliver. 4d flow mri. *Journal of Magnetic Resonance Imaging*, 36(5):1015–1036, 2012.
- [114] M. A. H. Mohd Adib, S. Ii, Y. Watanabe, and S. Wada. Minimizing the blood velocity differences between phase-contrast magnetic resonance imaging and computational fluid dynamics simulation in cerebral arteries and aneurysms. *Medical & biological engineering & computing*, 55(9):1605–1619, 2017.
- [115] J. Naqvi, K. H. Yap, G. Ahmad, J. Ghosh, et al. Transcranial doppler ultrasound: a review of the physical principles and major applications in critical care. *International journal of vascular medicine*, 2013, 2013.
- [116] K. S. Nayak, J.-F. Nielsen, M. A. Bernstein, M. Markl, P. D. Gatehouse, R. M. Botnar, D. Saloner, C. Lorenz, H. Wen, B. S. Hu, et al. Cardiovascular magnetic resonance phase contrast imaging. *Journal of Cardiovascular Magnetic Resonance*, 17(1):1–26, 2015.
- [117] R. A. NISHIMURA, F. A. MILLER Jr, M. J. CALLAHAN, R. C. BENASSI, J. B. SEWARD, and A. J. TAJIK. Doppler echocardiography: theory, instrumentation, technique, and application. In *Mayo Clinic Proceedings*, volume 60, pages 321–343. Elsevier, 1985.

- [118] L. Niu, M. Qian, L. Yan, W. Yu, B. Jiang, Q. Jin, Y. Wang, R. Shandas, X. Liu, and H. Zheng. Real-time texture analysis for identifying optimum microbubble concentration in 2-d ultrasonic particle image velocimetry. *Ultrasound in medicine & biology*, 37(8):1280–1291, 2011.
- [119] A. M. Nixon, M. Gunel, and B. E. Sumpio. The critical role of hemodynamics in the development of cerebral vascular disease: A review. *Journal of Neurosurgery JNS*, 112(6):1240 – 1253, 2010.
- [120] D. Nolte and C. Bertoglio. Inverse problems in blood flow modeling: A review. *International journal for numerical methods in biomedical engineering*, 38(8):e3613, 2022.
- [121] S. Ohtsuki and M. Tanaka. The flow velocity distribution from the doppler information on a plane in three-dimensional flow. *Journal of visualization*, 9(1):69–82, 2006.
- [122] C. W. Ong, I. Wee, N. Syn, S. Ng, H. L. Leo, A. M. Richards, and A. M. Choong. Computational fluid dynamics modeling of hemodynamic parameters in the human diseased aorta: A systematic review. *Annals of Vascular Surgery*, 63:336–381, 2020.
- [123] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanaawala, and M. Lustig. Robust 4d flow denoising using divergence-free wavelet transform. *Magnetic resonance in medicine*, 73(2):828–842, 2015.
- [124] W. H. Organization. Cardiovascular diseases (cvds) fact sheet, 2021. Accessed: 2024-08-22.
- [125] B. G. Osgood. *Lectures on Fourier transform and its applications*. SIAM, 2019.
- [126] S. Patankar. *Numerical Heat Transfer and Fluid Flow: CRC Press*. Taylor & Francis Group, LLC. Florida, United States, 1980.
- [127] N. J. Pelc, R. J. Herfkens, A. Shimakawa, D. R. Enzmann, et al. Phase contrast cine magnetic resonance imaging. *Magnetic resonance quarterly*, 7(4):229–254, 1991.

- [128] E. S. Peper, L. M. Gottwald, Q. Zhang, B. F. Coolen, P. van Ooij, A. J. Nederveen, and G. J. Strijkers. Highly accelerated 4d flow cardiovascular magnetic resonance using a pseudo-spiral cartesian acquisition and compressed sensing reconstruction for carotid flow and wall shear stress. *Journal of Cardiovascular Magnetic Resonance*, 22:1–15, 2020.
- [129] E. S. Peper, P. van Ooij, B. Jung, A. Huber, C. Gräni, and J. A. Bastiaansen. Advances in machine learning applications for cardiovascular 4d flow mri. *Frontiers in cardiovascular medicine*, 9:1052068, 2022.
- [130] V. M. Pereira, B. Delattre, O. Brina, P. Bouillot, and M. I. Vargas. 4d flow mri in neuroradiology: techniques and applications. *Topics in Magnetic Resonance Imaging*, 25(2):81–87, 2016.
- [131] V. M. Pereira, B. Delattre, O. Brina, P. Bouillot, and M. I. Vargas. 4d flow mri in neuroradiology: techniques and applications. *Topics in Magnetic Resonance Imaging*, 25(2):81–87, 2016.
- [132] I. Perez-Raya, M. F. Fathi, A. Baghaie, R. H. Sacho, K. M. Koch, and R. M. D’Souza. Towards multi-modal data fusion for super-resolution and denoising of 4d-flow mri. *International Journal for Numerical Methods in Biomedical Engineering*, 36(9):e3381, 2020.
- [133] C. Poelma. Ultrasound imaging velocimetry: a review. *Experiments in Fluids*, 58:1–28, 2017.
- [134] C. Poelma, J. Mari, N. Foin, M.-X. Tang, R. Krams, C. Caro, P. Weinberg, and J. Westerweel. 3d flow reconstruction using ultrasound piv. *Experiments in fluids*, 50:777–785, 2011.
- [135] W. V. Potters, H. A. Marquering, E. VanBavel, and A. J. Nederveen. Measuring

- wall shear stress using velocity-encoded mri. *Current Cardiovascular Imaging Reports*, 7:1–12, 2014.
- [136] W. V. Potters, P. van Ooij, H. Marquering, E. vanBavel, and A. J. Nederveen. Volumetric arterial wall shear stress calculation based on cine phase contrast mri. *Journal of Magnetic Resonance Imaging*, 41(2):505–516, 2015.
- [137] S. Prakash and C. R. Ethier. Requirements for mesh resolution in 3d computational hemodynamics. *J. Biomech. Eng.*, 123(2):134–144, 2001.
- [138] T. Puiseux. *Numerical simulations for phase-contrast magnetic resonance imaging*. PhD thesis, Université Montpellier, 2019.
- [139] S. Qian, H. Liu, C. Liu, S. Wu, and H. San Wong. Adaptive activation functions in convolutional neural networks. *Neurocomputing*, 272:204–212, 2018.
- [140] Q. Qin. Point spread functions of the t2 decay in k-space trajectories with long echo train. *Magnetic resonance imaging*, 30(8):1134–1142, 2012.
- [141] 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.
- [142] 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.
- [143] 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.
- [144] 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.
- [145] H. Rajabzadeh-Oghaz, P. van Ooij, S. S. Veeturi, V. M. Tutino, J. J. Zwanenburg, and H. Meng. Inter-patient variations in flow boundary conditions at middle cerebral artery from 7t pc-mri and influence on computational fluid dynamics of intracranial aneurysms. *Computers in biology and medicine*, 120:103759, 2020.
- [146] V. L. Rayz, L. Boussel, G. Acevedo-Bolton, A. J. Martin, W. L. Young, M. T. Lawton, R. Higashida, and D. Saloner. Numerical simulations of flow in cerebral aneurysms: comparison of cfd results and in vivo mri measurements. *Journal of biomechanical engineering*, 130(5), 2008.
- [147] Y. Ren, G.-Z. Chen, Z. Liu, Y. Cai, G.-M. Lu, and Z.-Y. Li. Reproducibility of image-based computational models of intracranial aneurysm: a comparison between 3d rotational angiography, ct angiography and mr angiography. *Biomedical engineering online*, 15(1):1–14, 2016.
- [148] V. C. Rispoli, J. F. Nielsen, K. S. Nayak, and J. L. Carvalho. Computational fluid dynamics simulations of blood flow regularized by 3d phase contrast mri. *Biomedical engineering online*, 14(1):1–23, 2015.
- [149] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [150] D. R. Rutkowski, A. Roldán-Alzate, and K. M. Johnson. Enhancement of cerebrovascular 4d flow mri velocity fields using machine learning and computational fluid dynamics simulation data. *Scientific reports*, 11(1):1–11, 2021.

- [151] D. R. Rutkowski, A. Roldán-Alzate, and K. M. Johnson. Enhancement of cerebrovascular 4d flow mri velocity fields using machine learning and computational fluid dynamics simulation data. *Scientific reports*, 11(1):1–11, 2021.
- [152] N. Saikrishnan, C.-H. Yap, N. C. Milligan, N. V. Vasilyev, and A. P. Yoganathan. In vitro characterization of bicuspid aortic valve hemodynamics using particle image velocimetry. *Annals of biomedical engineering*, 40:1760–1775, 2012.
- [153] A. Salari, M. B. Shafii, and S. Shirani. An experimental review on microbubble generation to be used in echo-particle image velocimetry method to determine the pipe flow velocity. *Journal of fluids engineering*, 135(3):034501, 2013.
- [154] L. M. Sangalli, P. Secchi, and S. Vantini. Aneurisk65: A dataset of three-dimensional cerebral vascular geometries. *Electronic Journal of Statistics*, 2014.
- [155] T. W. Secomb. Hemodynamics. *Comprehensive physiology*, 6(2):975, 2016.
- [156] P. P. Sengupta, G. Pedrizzetti, P. J. Kilner, A. Kheradvar, T. Ebbers, G. Tonti, A. G. Fraser, and J. Narula. Emerging trends in cv flow visualization. *JACC: Cardiovascular Imaging*, 5(3):305–316, 2012.
- [157] D. M. Sforza, C. M. Putman, and J. R. Cebal. Hemodynamics of cerebral aneurysms. *Annual review of fluid mechanics*, 41:91–107, 2009.
- [158] Y. Shen, B. Wang, F. Chen, and L. Cheng. A new multi-output neural model with tunable activation function and its applications. *Neural processing letters*, 20:85–104, 2004.
- [159] S. Shit, J. Zimmermann, I. Ezhov, J. C. Paetzold, A. F. Sanches, C. Pirkl, and B. H. Menze. Srflow: Deep learning based super-resolution of 4d-flow mri data. *Frontiers in Artificial Intelligence*, 5:928181, 2022.

- [160] F. Shone, N. Ravikumar, T. Lassila, M. MacRaid, Y. Wang, Z. A. Taylor, P. Jimack, E. Dall'Armellina, and A. F. Frangi. Deep physics-informed super-resolution of cardiac 4d-flow mri. In *International Conference on Information Processing in Medical Imaging*, pages 511–522. Springer, 2023.
- [161] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.
- [162] J. O. Smith. *Mathematics of the discrete Fourier transform (DFT): with audio applications*. Julius Smith, 2007.
- [163] G. Soulat, P. McCarthy, and M. Markl. 4d flow with mri. *Annual review of biomedical engineering*, 22:103–126, 2020.
- [164] A. F. Stalder, M. Russe, A. Frydrychowicz, J. Bock, J. Hennig, and M. Markl. Quantitative 2d and 3d phase contrast mri: optimized analysis of blood flow and vessel wall parameters. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 60(5):1218–1231, 2008.
- [165] Z. Stankovic, B. D. Allen, J. Garcia, K. B. Jarvis, and M. Markl. 4d flow imaging with mri. *Cardiovascular diagnosis and therapy*, 4(2):173, 2014.
- [166] T. L. Szabo. *Diagnostic ultrasound imaging: inside out*. Academic press, 2013.
- [167] T. L. Szabo and P. A. Lewin. Ultrasound transducer selection in clinical imaging practice. *Journal of Ultrasound in Medicine*, 32(4):573–582, 2013.
- [168] S. Tao, P. T. Weavers, J. D. Trzasko, Y. Shu, J. Huston III, S.-K. Lee, L. M. Frigo, and M. A. Bernstein. Gradient pre-emphasis to counteract first-order concomitant fields on asymmetric mri gradient systems. *Magnetic resonance in medicine*, 77(6):2250–2262, 2017.

- [169] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [170] T. Tieleman. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26, 2012.
- [171] A. N. Tikhonov, A. Goncharsky, V. Stepanov, and A. G. Yagola. *Numerical methods for the solution of ill-posed problems*, volume 328. Springer Science & Business Media, 1995.
- [172] J. Töger, M. J. Zahr, N. Aristokleous, K. Markenroth Bloch, M. Carlsson, and P.-O. Persson. Blood flow imaging by optimal matching of computational fluid dynamics to 4d-flow data. *Magnetic resonance in medicine*, 84(4):2231–2245, 2020.
- [173] J. Töger, M. J. Zahr, N. Aristokleous, K. Markenroth Bloch, M. Carlsson, and P.-O. Persson. Blood flow imaging by optimal matching of computational fluid dynamics to 4d-flow data. *Magnetic resonance in medicine*, 84(4):2231–2245, 2020.
- [174] F. Troger, C. Tiller, M. Reindl, I. Lechner, M. Holzknrecht, M. Pamminer, P. Poskaite, C. Kremser, H. Ulmer, E. R. Gizewski, et al. Slice positioning in phase-contrast mri impacts aortic stenosis assessment. *European Journal of Radiology*, 161:110722, 2023.
- [175] J. H. Tu. *Dynamic mode decomposition: Theory and applications*. PhD thesis, Princeton University, 2013.
- [176] C. Vagli, F. Fisticaro, L. Vinciguerra, V. Puglisi, M. S. Rodolico, A. Giordano, R. Ferri, G. Lanza, and R. Bella. Cerebral hemodynamic changes to transcranial doppler in asymptomatic patients with fabry’s disease. *Brain Sciences*, 10(8):546, 2020.
- [177] A. Vali, A. A. Abl, M. T. Lawton, D. Saloner, and V. L. Rayz. Computational fluid dynamics modeling of contrast transport in basilar aneurysms following flow-altering surgeries. *Journal of biomechanics*, 50:195–201, 2017.

- [178] J. Van Vaals and A. Bergman. Optimization of eddy-current compensation. *Journal of Magnetic Resonance (1969)*, 90(1):52–70, 1990.
- [179] U. Vovk, F. Pernus, and B. Likar. A review of methods for correction of intensity inhomogeneity in mri. *IEEE transactions on medical imaging*, 26(3):405–421, 2007.
- [180] A. K. Wake, J. N. Oshinski, A. R. Tannenbaum, and D. P. Giddens. Choice of in vivo versus idealized velocity boundary conditions influences physiologically relevant flow patterns in a subject-specific simulation of flow in the human carotid bifurcation. *Journal of biomechanical engineering*, 131(2), 2009.
- [181] P. G. Walker, G. B. Cranney, M. B. Scheidegger, G. Waseleski, G. M. Pohost, and A. P. Yoganathan. Semiautomated method for noise reduction and background phase error correction in mr phase velocity data. *Journal of Magnetic Resonance Imaging*, 3(3):521–530, 1993.
- [182] P. G. Walker, G. B. Cranney, M. B. Scheidegger, G. Waseleski, G. M. Pohost, and A. P. Yoganathan. Semiautomated method for noise reduction and background phase error correction in mr phase velocity data. *Journal of Magnetic Resonance Imaging*, 3(3):521–530, 1993.
- [183] B. Wang, Z. Sun, X. Jiang, J. Zeng, and R. Liu. Kalman filter and its application in data assimilation. *Atmosphere*, 14(8):1319, 2023.
- [184] L. Willats and F. Calamante. The 39 steps: evading error and deciphering the secrets for accurate dynamic susceptibility contrast mri. *NMR in Biomedicine*, 26(8):913–931, 2013.
- [185] K. K. L. Wong, S. C. P. Cheung, W. Yang, and J. Tu. Numerical simulation and experimental validation of swirling flow in spiral vortex ventricular assist device. *The International Journal of Artificial Organs*, 33(12):856–867, 2010.

- [186] D. T. Wymer, K. P. Patel, W. F. Burke III, and V. K. Bhatia. Phase-contrast mri: physics, techniques, and clinical applications. *Radiographics*, 40(1):122–140, 2020.
- [187] J. Xiang, V. Tutino, K. Snyder, and H. Meng. Cfd: computational fluid dynamics or confounding factor dissemination? the role of hemodynamics in intracranial aneurysm rupture risk assessment. *American Journal of Neuroradiology*, 35(10):1849–1857, 2014.
- [188] D. Yang, D. Cabral, E. N. Gaspard, R. B. Lipton, T. Rundek, and C. A. Derby. Cerebral hemodynamics in the elderly: a transcranial doppler study in the einstein aging study cohort. *Journal of Ultrasound in Medicine*, 35(9):1907–1914, 2016.
- [189] C.-C. Yu, Y.-C. Tang, and B.-D. Liu. An adaptive activation function for multilayer feedforward neural networks. In *2002 IEEE Region 10 Conference on Computers, Communications, Control and Power Engineering. TENCOM'02. Proceedings.*, volume 1, pages 645–650. IEEE, 2002.
- [190] J. Zhang, M. C. Brindise, S. M. Rothenberger, M. Markl, V. L. Rayz, and P. P. Vlachos. A multi-modality approach for enhancing 4d flow magnetic resonance imaging via sparse representation. *Journal of the Royal Society Interface*, 19(186):20210751, 2022.
- [191] M. Zhou, Y. Yu, R. Chen, X. Liu, Y. Hu, Z. Ma, L. Gao, W. Jian, and L. Wang. Wall shear stress and its role in atherosclerosis. *Frontiers in cardiovascular medicine*, 10:1083547, 2023.