

# Integrating Flow Imaging and Deep Learning into Patient-Specific Models Following Myocardial Infarction



Fergus Shone

University of Leeds  
Centre for Doctoral Training in Fluid Dynamics  
School of Computing

Submitted in accordance with the requirements for the integrated degree of

*Doctor of Philosophy and Master of Science*

January, 2025

## **Intellectual Property Statement**

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors' to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapter 4 includes work from the following publication: F. Shone, N. Ravikumar, T. Lassila, et al. "Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI", In: *Information Processing in Medical Imaging (IPMI) 2023, Springer Lecture Notes in Computer Science, vol. 13939* (2023), DOI: 10.1007/978-3-031-34048-2\_39 [205]. The candidate performed the literature review, prepared the original manuscript and revised the manuscript. Other authors contributed to the conceptualisation, supervision and review of the manuscript.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Fergus Shone to be identified as Author of this work has been asserted by Fergus Shone in accordance with the Copyright, Designs and Patents Act 1988.

© 2024

The University of Leeds and Fergus Shone.

## Acknowledgements

First and foremost, I would like to thank my supervisors Erica, Alex, Peter, and Zeike for their unwavering support and wisdom over the past four years. I am deeply grateful to my collaborators — Toni, Michael, Qiongyao, Ben, Rose, Cristina and Nishant — for sharing so much of their time, ideas and memories with me.

I am thankful to both David Higgins and Christel Kamani for their help in acquiring and processing much of the data used in the penultimate chapter of this thesis - and for answering my many queries about magnetic resonance imaging.

I extend my gratitude to the Fluid Dynamics Centre for Doctoral Training for their invaluable support throughout this journey, and to the many companions I have made along the way.

To my friends Otto and Sanjay: I am forever thankful for our shared passion for research and the fun we have had during our years together.

Thank you to my parents for their constant encouragement and affection, and to my sister Ellen and her beautiful family for always lifting my spirits.

Lastly, I would like to thank my dear partner Maia. She has been a constant presence throughout my academic journey, and I would not be where I am today without her daily love, support, and patience.

## Abstract

**Background** 4D-flow magnetic resonance imaging (MRI) provides non-intrusive blood flow reconstructions in the left ventricle (LV) and other cardiac chambers and has the potential to become a key tool in both research and clinic. However, low spatio-temporal resolution and the presence of significant noise artifacts hamper the accuracy of derived haemodynamic quantities and thus limit the effectiveness of the modality to establish links between haemodynamic abnormalities and pathologies. Furthermore, models that are constrained by boundary conditions are impacted by additional uncertainty which arises due to the low spatial resolution of the structural cine-MRI.

**Methods** 4D-flow MRI data corruption, introduced by low resolution and noise artefacts, may be alleviated through super-resolution and de-noising methods, which have been explored in the literature for haemodynamic flow in the vasculature. In this thesis, a physics-informed neural network (PINN) model is introduced to provide super-resolution and de-noising, specifically of cardiac 4D-flow MRI. The model is constrained through weak enforcement using the low-resolution 4D-flow MRI data, the no-slip boundary condition on the endocardium and the governing physical equations. Model components are compared and incorporated to address specific challenges introduced by modelling haemodynamic flow in the heart chambers, such as flow across a range of length and time scales within a heavily deforming domain. Validation of the model is performed across synthetic and *in vivo* studies, evaluating the robustness of the model to uncertainty in both the 4D-flow MRI data and the position of the deforming endocardium. Following this, the model is applied to a small cohort of LV remodelling patients.

**Results** It is demonstrated that the PINN model is able to effectively upsample and de-noise the velocity field across a range of spatio-temporal resolutions

and signal-to-noise ratios (SNR), and is robust to positional uncertainty in the deforming endocardium for flow variables measured away from the domain boundaries. Further, variables that are not directly measured, such as relative pressure and flow derivatives, are reconstructed to an acceptable degree of accuracy. In the dual-resolution *in vivo* validation study, it is shown the model is generally independent to the spatial resolution and SNR of the input 4D-flow MRI data.

**Conclusions** Through synthetic and *in vivo* validation studies, it is demonstrated that the PINN model introduced in this thesis is effective. It is concluded that the use of this type of model is feasible for super-resolution of cardiac 4D-flow MRI data, although certain limitations should be addressed and the model should be further validated using *in vivo* or *in vitro* data.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Motivation	4
1.2	Research Questions	5
1.3	Thesis Outline	6
<b>2</b>	<b>Literature Review</b>	<b>8</b>
2.1	Literature Search Approach	8
2.2	Haemodynamic Flow in the Left Ventricle	9
2.3	4D-Flow Magnetic Resonance Imaging	20
2.4	Physics-Informed Neural Networks	28
<b>3</b>	<b>Super-Resolution of Cardiac 4D-Flow MRI: An Idealised <i>in Silico</i> Study</b>	<b>55</b>
3.1	Introduction	55
3.2	Methodology	56
3.3	Results	74
3.4	Discussion	87
3.5	Conclusion	89
<b>4</b>	<b>Assessing Robustness to Positional Uncertainty in the Endocardium</b>	<b>91</b>
4.1	Introduction	91
4.2	Methodology	93
4.3	Results	103
4.4	Conclusion	140
<b>5</b>	<b>Super-Resolution of <i>in Vivo</i> Cardiac 4D-Flow MRI</b>	<b>142</b>
5.1	Introduction	142
5.2	Methodology	145

5.3	Results . . . . .	151
5.4	Discussion . . . . .	163
5.5	Conclusion . . . . .	167
<b>6</b>	<b>Conclusions and Outlook</b>	<b>169</b>
6.1	Summary . . . . .	169
6.2	Limitations . . . . .	171
6.3	Outlook . . . . .	174

# List of Figures

2.1	Peak systolic blood flow in the left ventricle [84] . . . . .	11
2.2	Example of 4D-flow MRI data . . . . .	15
2.3	The complex structures present on the endocardium [194] . . . . .	19
2.4	Physical principles of PC-MRI [233] . . . . .	20
2.5	4DFlowNet velocity results [75] . . . . .	25
2.6	SRFlow velocity results [203] . . . . .	26
2.7	PINN super-resolution model results [72] . . . . .	27
2.8	2D turbulent velocity field predicted using CycleGAN [126] . . . . .	31
2.9	A basic overview of the PINN workflow, using a simple 1D example. . . . .	35
2.10	Fixed activation functions . . . . .	39
2.11	Adaptive tanh activation function . . . . .	45
2.12	Siren network for image reconstruction [206] . . . . .	46
3.1	Our PINN model: an overview . . . . .	57
3.2	Outline of the mini-batch training process used when $N_d \neq N_b \neq N_p$ . . . . .	67
3.3	2D idealised ventricle case: mesh and domain . . . . .	69
3.4	Synthetic 4D-flow MRI data workflow . . . . .	71

3.5	2D idealised ventricle: velocity and pressure reconstructions . . . . .	75
3.6	2D idealised ventricle: velocity and pressure reconstructions . . . . .	78
3.7	[2D idealised ventricle: velocity and pressure reconstructions . . . . .	78
3.8	2D idealised ventricle: velocity and pressure reconstructions . . . . .	79
3.9	2D idealised ventricle: velocity and pressure reconstructions, comparing activation functions . . . . .	80
3.10	2D idealised ventricle: demonstration of over- and under-fitting . . . . .	84
3.11	2D idealised ventricle: velocity and pressure reconstructions . . . . .	85
3.12	2D idealised ventricle: velocity and pressure reconstructions . . . . .	86
4.1	endocardial surface reconstruction: MRI against CT . . . . .	94
4.2	Positional uncertainty study: workflow . . . . .	96
4.3	Segmented surface geometries from three patient-specific CT image sets, namely cases <i>A</i> , <i>B</i> and <i>C</i> . . . . .	97
4.4	The endocardial surfaces reconstructed using smoothing operations outlines in Tab. 4.1, for LV case C . . . . .	100
4.5	Mesh independence study . . . . .	102
4.6	Mesh independence study: velocity and pressure . . . . .	104
4.7	Optimisation study results . . . . .	105
4.8	Optimisation study results . . . . .	105
4.9	Case A: velocity magnitude . . . . .	107
4.10	Case B: velocity magnitude . . . . .	107
4.11	Case C: velocity magnitude . . . . .	109
4.12	Vorticity reconstruction at differing resolutions . . . . .	111

4.13 Case B: vorticity field . . . . .	112
4.14 Vortex core variables . . . . .	115
4.15 Case C: Vortex core reconstruction . . . . .	117
4.16 Kinetic energy figures . . . . .	120
4.17 Case C: Wall shear stress magnitude . . . . .	121
4.18 An explanation of the regions of the endocardium extracted during bull's eye plot analysis. . . . .	123
4.19 Mean wall shear stress . . . . .	124
4.20 Bull's eye plots of relative mean WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth. . . . .	124
4.21 Mean wall shear stress error . . . . .	125
4.22 Peak wall shear stress . . . . .	126
4.23 Bull's eye plots of relative peak WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth. . . . .	127
4.24 Peak wall shear stress error . . . . .	127
4.25 Velocity gradient fields and wall shear stress . . . . .	128
4.26 Wall shear stress magnitude . . . . .	130
4.27 Case B: relative pressure field . . . . .	132
4.28 Pressure drop curves . . . . .	134
4.29 Positional uncertainty: velocity and pressure . . . . .	136
4.30 Positional uncertainty: vorticity . . . . .	137
4.31 Positional uncertainty: kinetic energy . . . . .	138
4.32 Positional uncertainty: wall shear stress . . . . .	139

5.1	Dual-resolution study: velocity field	152
5.2	Dual-resolution study: pressure field	153
5.3	Dual-resolution study: haemodynamic forces	154
5.4	Dual-resolution study: kinetic energy	155
5.5	Exploratory study: relative pressure field	156
5.6	Exploratory study: pressure drop	157
5.7	Exploratory study: haemodynamic forces	157
5.8	Exploratory study: multiple variables	158
5.9	Exploratory study: multiple variables	159
5.10	Exploratory study: wall shear stress	159
5.11	Exploratory study: mean and peak wall shear stress	161
5.12	Exploratory study: kinetic energy	162
5.13	Exploratory study: kinetic energy	162

# List of Tables

2.1	An example automatic differentiation trace [13] . . . . .	37
3.1	PINN architecture and hyperparameter values used for the 2D idealised LV model. . . . .	75
3.2	2D idealised ventricle: velocity and pressure reconstruction errors . . . . .	77
3.3	Search space defined for the 2D hyperparameter optimisation study. . . . .	77
3.4	Optimal results for each activation function analysed in the ablation study. . . . .	80
3.5	Error metrics for different values of the initialisation frequency $\omega_0$ in the Siren network architecture. . . . .	81
3.6	2D idealised ventricle: Hyperparameter selections used for each configuration of loss weighting strategies, and quantitative results obtained from each loss weighting configuration. . . . .	83
3.7	2D idealised ventricle: Initialised weights, $\alpha_0$ and $\beta_0$ , and final weights (after $N$ weight updates), $\alpha_N$ and $\beta_N$ , for the LR annealing method. . . . .	83
4.1	Details of the smoothing parameters used to reconstruct surfaces $S_1^i$ , $S_2^i$ and $S_3^i$ for cases A, B and C. . . . .	100
4.2	hyperparameter search space . . . . .	101
4.3	The restricted search space identified for hyperparameter optimisation studies in cases B and C. . . . .	102

4.4	Relative error (from fine mesh) of selected variables for coarse and medium meshes at three lines of interest. . . . .	103
4.5	Optimal individual hyperparameter configurations for cases A, B and C, alongside a shared configuration to be used across all cases. . . . .	106
4.6	Velocity field errors . . . . .	109
4.7	A list of vorticity ( $\omega$ ) parameters analysed in this study with their respective descriptions. . . . .	110
4.8	Vorticity errors . . . . .	111
4.9	Vorticity errors . . . . .	113
4.10	Vorticity errors . . . . .	113
4.11	Peak and mean vorticity . . . . .	114
4.12	Vortex core errors . . . . .	117
4.13	Kinetic energy parameters . . . . .	118
4.14	Kinetic energy errors . . . . .	119
4.15	Point-wise, max-normalised RMS error of the relative pressure field. . . . .	134
5.1	Volunteer details (LV function). . . . .	147
5.2	Dual-resolution 4D-flow MRI acquisition details . . . . .	147
5.3	DTI INMI 4D-flow MRI acquisition details . . . . .	148

# Glossary

**4D-Flow MRI** Phase contrast magnetic resonance imaging technique used to visualise time-resolved, three-dimensional velocity fields in the cardiovascular system.

**aliasing** A phenomenon experienced in PC-MRI acquisition in which discontinuities appear in the constructed solution field due to velocity measurements exceeding the prescribed  $v_{enc}$ .

**background noise** Noise present in magnetic resonance imaging, arising from radiofrequency coil resistance, electronic noise in the preamplifier, and dielectric and inductive losses in the imaged object.

**cine-MRI** Any time-resolved MRI technique. Typically used to describe time-resolved MR images of tissue structure.

**diastole** The refilling period of the cardiac cycle, during which the ventricle myocytes relax to draw blood into the chamber.

**eccentric hypertrophy** In the LV, this describes an increase in the volume of the chamber through lengthening of the myocardium tissue.

**encoding velocity ( $v_{enc}$ )** Prescribed by the imaging expert during acquisition, this value must be higher than any expected velocities in the region of interest.

**endocardium** The innermost tissue of the myocardium, in contact with the ventricular chamber.

**haemodynamic forces (HDFs)** The force generated by the flow of intracardiac blood, calculated by integrating the pressure gradient field over the heart chamber.

**haemodynamics** The dynamics of blood flow within the cardiovascular system.

**heart failure (HF)** Occurs when the heart can no longer pump blood efficiently.

**hypertrophy** Increase in the size of muscular tissue, such as thickening of the myocardium in the LV.

**in silico** A process that takes place within a computer model/simulation.

**in vitro** A process that takes place outside of a living organism.

**in vivo** A process that takes place within a living organism.

**learning rate (LR)** A tuning parameter in machine learning that determines the step size at each iteration during optimisation.

**left ventricle (LV)** The largest chamber of the heart, tasked with supplying oxygenated blood to the body.

**left ventricular remodelling** Describes changes in the size, shape and function of the left ventricle as a result of myocardial infarction.

**magnetic resonance imaging (MRI)** A medical imaging technique used for visualising structure and processes within the body that is powered by magnetic fields.

**mitral valve** The heart valve between the left atrium and the left ventricle, through which blood is drawn in during diastole.

**myocardial infarction (MI)** Otherwise known as a heart attack, describes damage to the myocardium imparted by a loss of blood supply.

**myocardium** The muscle tissue in the heart.

**myocyte** A muscle cell.

**network generalisation** The capability of a neural network to be applied to data that lies outside of the training set.

**neural network (NN)** A type of machine learning algorithm, consisting of layers of connected nodes, that approximates functions by training the connecting weights between each node. In the case of *supervised* learning, a collection of labelled examples is used to train the network by minimising the *loss function*, which is the error between prediction and example.

**neural tangent kernel (NTK)** A kernel that may be used to describe the training dynamics of wide neural networks.

**phase contrast magnetic resonance imaging (PC-MRI)** Magnetic resonance imaging technique that utilises the magnetic phase shift rather than the signal magnitude, allowing quantification of fluid velocities.

**physics-informed neural network (PINN)** Similar to a neural network, but an additional loss function term is included, which contains any known physical laws of the system that is being predicted. The physics is included in the form of residual PDEs, and therefore minimising them ensures that the network prediction approximately satisfies them.

**registration** The process of transforming data sets onto a specified coordinate system.

**segmentation** Describes the generation of a computational domain from medical images.

**stroke volume** The volume of blood pumped from the left ventricle during systole.

**structured noise** Noise present in magnetic resonance imaging, arising due to subject motion, ghosting and reconstruction artefacts.

**super-resolution** The process of enhancing the resolution of an image.

**systole** The contraction period of the cardiac cycle, during which blood is pumped from the ventricle.

**trabeculae carneae** The complex muscular structures present on the inside of the myocardium in the left ventricle.

**wall shear stress (WSS)** The tangential force applied to a surface through viscous shearing of fluid in the boundary layer.

# Chapter 1

## Introduction

### 1.1 Motivation

The emergence of 4D-flow magnetic resonance imaging (MRI) has facilitated the quantification of a wide array of haemodynamic parameters in the left ventricle (LV), which has allowed for the discovery of novel clinical markers linked with the development of complex cardiac pathologies. It has been proposed that measurable disturbances to the haemodynamic field often precede those observed in the myocardium [174], and as such, quantification of certain flow parameters may also facilitate earlier-stage patient risk stratification. 4D-flow MRI provides reconstructions of time-resolved velocity fields in three spatial dimensions, allowing for comprehensive assessment of haemodynamic flow in the LV. Further, from the measured velocity field, it is possible to compute additional haemodynamic parameters, such as kinetic energy (KE), vorticity and pressure drops, vastly expanding the space of potential cardiac disease biomarkers when compared with studies utilising structural imaging techniques alone.

Despite the extensive use of 4D-flow MRI in a research capacity, its wider application in a clinical setting has been hampered by a variety of shortcomings. Low spatio-temporal resolution and noise artefacts reduce the accuracy of the measured velocity field, with quantities that are not directly measured highly susceptible to data corruption. As such, considerable uncertainty may be present in directly-measured and derived haemodynamic markers, reducing confidence in results. To address such shortcomings, a limited selection of super-resolution methods have been proposed to increase the resolution of the measured velocity fields whilst reducing the impact of noise [72, 75, 203, 195]. Super-resolution of

4D-flow MRI is an emerging field, with all publications so far dedicated to addressing the problem in vascular flow. Comparable studies in the cardiac chambers do not exist, which provides the underlying motivation for this thesis.

## 1.2 Research Questions

The primary research aims and corresponding objectives are outlined below:

1. Identification of a suitable computational method for super-resolution of cardiac 4D-flow MRI
  - Assess existing methods for super-resolution of vascular 4D-flow MRI
  - Identify requirements of the model relevant to the present application
    - Applicability to flow data within deforming domains
    - Ability to assimilate multiple sources of (incomplete) information
2. Adaptation of the selected model to the present application
  - Establish key sources of uncertainty/error in the data and model
    - Uncertainty present in the flow data (i.e. spatio-temporal resolution and noise)
    - Uncertainty in the boundary conditions (i.e. positional uncertainty in the segmented endocardium)
    - Model limitations (i.e. spectral bias and gradient imbalances)
  - Identification of methods to mitigate the above limitations
  - Assessment of the robustness of the model to the above limitations
3. Demonstration of the feasibility and validity of the super-resolution model
  - Synthetic validation studies in idealised and patient-specific flow domains
  - *In vivo* validations study in a single volunteer at two spatial resolutions
  - *In vivo* exploratory study in a small cohort of left ventricular remodelling patients

These aims and objectives are dispersed throughout the thesis, which is outlined in the next section.

## 1.3 Thesis Outline

The objective of this thesis is to develop and assess the suitability of a physics-informed neural network (PINN) model for the super-resolution of cardiac 4D-flow MRI, specifically in the LV. By constraining model outputs using the governing physical equations, PINNs are able to operate in the absence of high-resolution data labels, which are typically infeasible to acquire in cardiac 4D-flow MRI studies. Further, such physics-based regularisation allows for simple quantification of variables of interest that are not directly measured, such as relative pressure and velocity derivatives, avoiding the requirement for complex post-processing techniques. To the best of the authors knowledge, the work presented in this thesis represents the first attempt to super-resolve cardiac *in vivo* flow imaging data, obtained from either 4D-flow MRI or other modalities, and the first to model fluid flow in moving boundaries using PINNs.

### Literature review

In chapter 2, we provide a detailed review of established haemodynamic markers in cardiac disease progression, before outlining the principles and shortcomings of 4D-flow MRI. This is followed by an overview of current 4D-flow MRI super-resolution approaches, before providing a detailed introduction to physics-informed machine learning, its applications and key methodological contributions in the field.

### Robustness to 4D-flow MRI corruption

The image quality of 4D-flow MRI is degraded by three main sources of corruption, namely low spatial resolution, low temporal resolution and low signal-to-noise ratio (SNR), which impact the quality of the directly measured velocity field and reduce the accuracy of quantities that are derived from it. Inter-patient variability of these effects is also present, due to differences in heart rate, LV volume and patient discomfort. As such, our model must be robust across differing levels of data corruption to be applicable across patient cohorts. To this end, in chapter 3 we assess the robustness of our super-resolution model across a range of degradation configurations using two synthetic, idealised LV cases. We also establish a preferred model configuration by comparing a range of architecture choices, chosen specifically to negate the effects of spectral bias and loss component imbalances during training.

### Robustness to geometric uncertainty

Positional uncertainty poses a significant problem for MRI-based, boundary-constrained flow modelling approaches in the LV, where low spatial resolution of structural cine-MRI reduces the accuracy of the derived endocardium and other interior surfaces. This is a significant hindrance to CFD-based approaches, owing to their strong dependence on boundary conditions, but the impact such geometrical uncertainty may have on PINN-based models has not been explored. In Chapter 4, we obtain results in a cohort of synthetic cases in CT-derived, patient-specific LV geometries. For each case, we extract three distinct endocardial surfaces using different levels of smoothing to quantify the effects of geometrical uncertainty in the applied no-slip boundary condition. This is performed alongside detailed comparison of clinically-relevant variables derived from both PINN results and competing methods utilised in the literature.

### Application to real 4D-flow MRI data

In Chapter 5, the PINN model established in the preceding chapters is applied to a limited *in vivo* 4D-flow MRI validation data set and a small cohort of real LV cases with varying levels of LV remodelling. In the validation study, 4D-flow MRI were obtained at two distinct spatial resolutions, with PINN results compared using each data set. In the LV remodelling study, clinically-relevant variables of interest are obtained and compared, providing a brief comparison between disease states.

# Chapter 2

## Literature Review

### 2.1 Literature Search Approach

The literature discussed within this chapter was assembled using a variety of electronic databases (Pubmed, Google Scholar, Web of Science) using the keywords listed below, which are categorised based on the application. Given the limited amount of literature around the primary application (methods for super-resolution of 4D-flow MRI), a systematic approach to literature searching was deemed to not be necessary, although investigation into potential computational methods was limited to machine learning-based approaches only.

- **4D-flow MRI principles:** “4D-flow MRI/magnetic resonance imaging”, “cardio-vascular magnetic resonance”, “phase-contrast MRI/magnetic resonance imaging”, “compressed sensing”, “parallel imaging”, “encoding strategies”
- **4D-flow MRI super-resolution:** “4D-flow MRI/magnetic resonance imaging”, “super-resolution”, “acceleration”, “upsampling”, “de-noising”
- **Machine learning approaches for modelling fluid flow:** “machine learning” + “fluid dynamics/mechanics/flow”
- **Physics-informed machine learning:** “physics-informed machine learning”, “physics-informed neural network”, “scientific machine learning”
- **Haemodynamic markers evaluated in the cardiac chambers:** “left ventricle/-cardiac” + (“hemodynamic”, “velocity”, “4D-flow MRI/magnetic resonance imaging”, “pressure”, “forces”, “vorticity/vortex”, “kinetic energy”)

## 2.2 Haemodynamic Flow in the Left Ventricle

### 2.2.1 Clinical Motivation

Cardiovascular diseases (CVDs) are the leading cause of death worldwide [157], accounting for an estimated 17.9 million deaths per year [30]. CVDs also carry significant economical expense, where it is forecast that the direct medical costs of cardiovascular disease in the US alone will reach \$818 billion by 2030 [98]. In the UK, this figure currently stands at £7.4 billion, with further costs estimated to incur an extra £15.8 billion [97], highlighting the significant societal strain imparted by both the morbidity and mortality of CVDs.

The term CVD covers an array of pathologies across the cardiovascular system, however, those that develop within the heart chambers are the largest contributors to the global death toll. Ischaemic heart disease alone is responsible for more deaths per year than the next four most lethal CVDs combined [219], claiming an estimated 9.44 million lives each year, with heart failure (HF) expected to increase in prevalence by 46% by 2030 [16].

To reduce the impact of cardiac disease, improvements in early-stage patient risk stratification are required to provide targeted interventions. To do this, a greater understanding of how certain pathologies progress over time and what triggers their development is a necessity, leading researchers to look beyond traditional metrics that may not become measurable until lasting damage has been done to the heart. The identification of alternative metrics that are measurable at an earlier stage allows for targeted treatment prior to the onset of irreversible damage.

### 2.2.2 Intraventricular Flow

Flow fields in the left ventricle (LV) and other cardiac chambers have traditionally been considered as passive, with specific patterns and characteristics thought to impart no influence on the function of the beating heart. However, growing evidence suggests that this theory is untrue. Haemodynamic forces in the apico-basal and aortic valve direction have been shown to drive filling and ejection throughout the cardiac cycle, where minimal differences across healthy subjects and elite athletes when indexed to LV volume indicates that the forces are optimised for the individual heart [180]. Haemodynamic forces have also been demonstrated to stimulate healthy embryonic growth of the heart [172], while it is hypothesised that the LV shape is optimised during development to facilitate the proper

formation of the mitral vortex ring [7]. In the adult heart, the diastolic mitral vortex ring in diastole is a key indicator of healthy LV function, facilitating penetration of the inflow jet with minimal energy loss, whilst aiding the redirection of flow to the aortic valve in systole [156].

### 2.2.3 Haemodynamic Influence on Cardiac Pathology

Whilst the healthy LV may be characterised by certain haemodynamic properties, disturbances to such features may be indicative of maladaptive cardiac function. There exists an intimate relationship between the myocardium and the adjacent flow in the LV, and as such, subclinical mechanical dysfunction can impart measurable changes on haemodynamic parameters[198, 156, 174]. The emergence of *in vivo* flow imaging technologies, such as phase-contrast MRI (PC-MRI) and Doppler echocardiography, has enabled researchers to explore these connections, facilitating the discovery of novel markers for cardiac disease.

Although the directly-measured velocity field itself can provide key insights in cardiovascular analysis, it also facilitates the calculation of many additional haemodynamic parameters that are challenging to obtain otherwise. These include spatial flow derivatives, such as vorticity [156, 52, 198, 63, 131, 34, 100, 2, 33, 125] and stresses (wall shear stress, in particular) [27, 96, 220, 213, 41], and relative pressure [128, 173, 243, 104, 46, 46, 209, 70, 168, 95, 62, 77, 59, 155, 74], which have been extensively investigated in disease progression in both the vasculature and the cardiac chambers.

#### Velocity

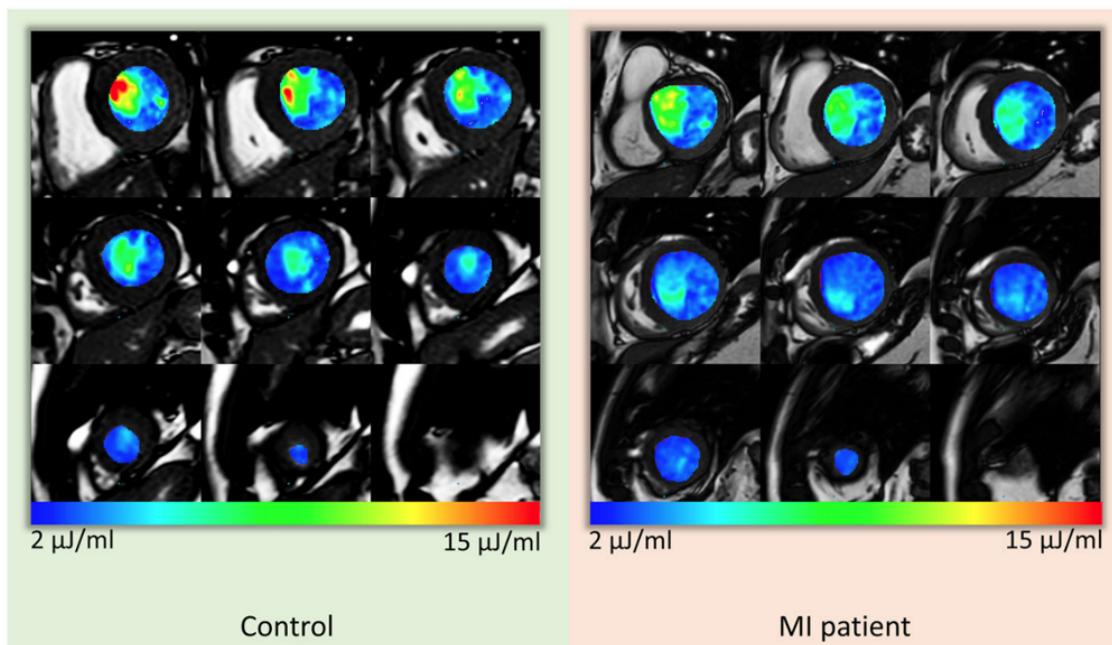
Metrics obtained directly from the velocity field, such as peak long-axis velocity, deceleration time and the E/A ratio (the ratio between peak velocities in E-wave and A-wave, respectively), have been explored extensively in relation to cardiac dysfunction [82, 53, 121]. Such parameters have been linked with diastolic HF [82], congestive HF [3] and LV remodelling [53], among more, although alternative quantities computed from the velocity field have gained more attention in recent years.

#### Kinetic Energy

Kinetic energy (KE), defined as:

$$KE = \frac{1}{2}m|\mathbf{u}|^2, \quad (2.1)$$

for fluid mass  $m$  and velocity  $\mathbf{u}$ , is an important variable of interest in cardiac studies, linked with myocardial infarction (MI) [84, 31, 15], LV remodelling [53, 15, 48], heart failure (HF) [119, 6] and ischaemic LV dysfunction [32]. KE in the LV is reflective of the work performed by the heart [84], and subtle changes in mechanical function may result in measurable changes in the KE patterns in the early stages of disease [6]. In [84], MI patients displayed a reduction in mean overall KE and systolic KE (indexed to end diastolic volume (EDV)), which can be observed in Fig. 2.1, whilst exhibiting an increase in the proportion of in-plane KE, which was independently associated with infarct size. Conversely, systolic KE was found to increase acutely post-MI, although this gradually decreased over the following 12 months [15]. Mean systolic KE was also linked with HF in [119], where HF patients showed a decrease in mean systolic KE when indexed to EDV. Using multivariate analysis, a reduction in peak A-wave diastolic KE was identified as an independent marker for LV remodelling at all levels of severity in [53]. HF, separated into subclinical left ventricular diastolic dysfunction (LVDD), HF with preserved ejection fraction (EF) and HF with reduced EF groups, was characterised by an increase in diastolic KE [6].



**Figure 2.1:** Peak systolic kinetic energy blood flow mapping: Comparison between healthy and infarcted patients, taken from [84].

### Vorticity

Vorticity, defined as:

$$\boldsymbol{\omega} = \nabla \times \mathbf{u}, \quad (2.2)$$

for velocity  $\mathbf{u}$ , is a vectorial quantity that describes the rotational motion of a continuum of fluid, and has emerged in recent years as a variable of interest for cardiac haemodynamic analyses. In the healthy heart, diastolic flow is characterised by the formation of a large vortex ring, alongside smaller vortices, induced by the formation of a shear layer on the mitral valve leaflets [174]. Previously considered a passive phenomenon, the formation of such flow structures is thought to facilitate efficient filling of the left ventricle (LV) [34], minimising viscous energy loss by a factor of  $2-4\times$  [63], whilst improving washout and aiding the redirection of flow to the aortic valve in systole [174, 202]. It is even hypothesised that the shape and volume of the LV is constructed during embryonic growth to optimise the formation and propagation of the mitral vortex ring [7].

Where the healthy heart can be characterised by the presence of the aforementioned vortical structures, measurable disturbances of such features could indicate cardiac dysfunction at a prior stage to the emergence of typical clinical markers [156]. To this end, a range of LV vorticity parameters have been investigated *in vivo* in the literature, using 4D-flow MRI [52, 198, 63, 131], planar PC-MRI [34] and echocardiographic particle imaging velocimetry (echo-PIV) [100, 2, 33]. Reduced vortex depth and pulsatility power in diastole were shown to contribute to apical thrombus formation in anterior MI patients in [125]. Similarly, mean and peak vorticity was shown to be lower in MI patients [52], with elevated mean and peak systolic vorticity at the basal region for MI patients with LV thrombus present. In [33], patients suffering from heart failure with reduced ejection fraction (HFrEF) exhibited fragmented vortex patterns, using maximum number of vortices in systole, early diastole and late diastole as a marker. It was also found that vortices had a reduced area (when indexed by LV end diastolic diameter), with a weaker systolic vortex. The connection between mitral vortex formation and viscous energy loss was investigated in [63], where it was found that abnormal vortex formation resulted in elevated viscous energy loss, particularly in cases where no vortex ring was formed. The suitability of 4D-flow MRI-derived vorticity as a marker for early LVDD in patients with mild-to-moderate chronic obstructive pulmonary disease (COPD) was assessed in [198]. COPD patients showed a significant reduction in early diastolic vorticity, displaying no or only mild signs of LVDD compared with controls, indicating that disturbances to the vorticity field precede mechanical changes. This finding is in agreement with [174], where it is postulated that mechanical dysfunction is preceded, and even driven, by unstable flow patterns, which are measurably reflected in the vorticity field.

## Pressure

There is a consensus that intraventricular pressure and pressure gradients play an important role in the healthy function of LV flow [173, 243], however their usage in a clinical setting has been restricted by the difficulties faced in acquiring accurate measurements [173]. The invasive nature of cardiac catheterisation to measure absolute pressure [104, 46] has motivated the development of image-based pressure quantification techniques for use in both cardiac and vascular studies.

Although coupled to the velocity field through the Navier-Stokes equations, reconstruction of cardiovascular pressure fields from velocity measurement data is not a trivial task. Traditional methods utilise velocity data obtained using Doppler-echocardiogram (Doppler-ECG) to approximate the pressure drop,  $\Delta p$ , between two distinct points in the flow region, using the Bernoulli theorem, in either a simplified or modified form [46, 209, 70, 168, 95, 62, 77], or the Euler equations [237, 20]. Recent methods have exploited advancements in imaging technology to derive relative pressure in the cardiovascular system from planar PC-MRI and 4D-flow MRI, with notable contributions including the use of a virtual work-energy function (vWERP) [59, 155, 74], pressure Poisson equations [242, 132, 69] and machine learning-based methods [204, 128].

Early approaches to quantify intraventricular pressure drop based on the simplified [209] and modified [77] Bernoulli equations have been shown to correlate poorly with catheter data [58], owing to the number of simplifications made about the intermediate flow field. The vWERP approach has been successfully adapted for use in the cardiac chambers [155], utilising the arbitrary Lagrangian-Eulerian form of the Navier-Stokes equations across dynamic sub-domains of the chamber. However, although this model is significantly more accurate than the aforementioned Bernoulli approaches, it also only provides pressure drop values between two regions in the LV. As with velocity, the pressure field in the LV is complex and spatially heterogeneous [69], and cannot be fully characterised by relative differences. However, despite the potential usefulness of reconstructing localised pressure differences, there has been limited investigation into suitable methods. In [69, 179], the pressure Poisson equations are used to reconstruct pressure distributions from 4D-flow MRI data, however the resulting field is highly dependent on the accuracy of flow gradients, which can be limited [155]. Alternatively, in [26], acceleration maps are integrated to reconstruct pressure distributions.

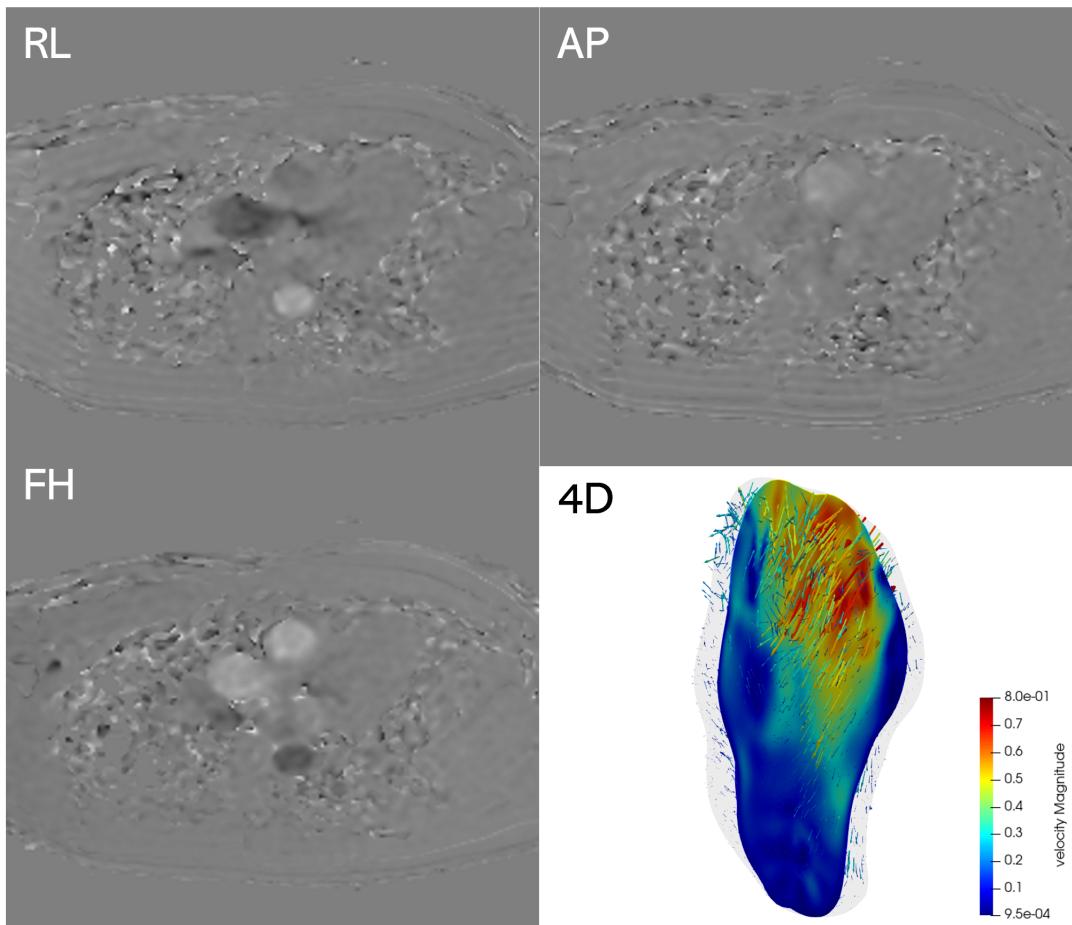
## Wall Shear Stress

Wall shear stress (WSS) is an important variable of interest in the vasculature, linked with aneurysm formation and rupture [27], atherosclerosis [96] and vessel remodelling [220]. It describes the tangential force applied to a surface through viscous shearing of the fluid in the boundary layer, and is defined as:

$$\mathbf{WSS} = 2\mu\tilde{\mathbf{S}} \cdot \hat{\mathbf{n}}, \quad (2.3)$$

for strain-rate tensor  $\tilde{\mathbf{S}}$  and inward wall-normal vector  $\hat{\mathbf{n}}$ . WSS is a highly sensitive variable that requires the accurate calculation of spatial velocity derivatives in the immediate vicinity of the wall, and therefore high spatial resolution in the near-wall region is desirable for acceptable results. This requirement proves problematic when attempting to approximate WSS using 4D-flow MRI, where low spatial resolution is a limiting factor [71]. Accurate reconstruction of the boundary layer in the near-wall region is crucial for calculating realistic velocity gradients, however, this flow feature typically occurs on length scales that are far smaller than the minimum voxel volume available to 4D-flow MRI, and thus are not properly resolved. This effect is further exaggerated at higher velocities, where the thickness of the boundary layer is reduced and voxel averaging becomes more problematic [42]. The result is that, while the reconstructed WSS distribution may be similar to reality, 4D-flow MRI persistently underestimates the magnitude of WSS by a significant factor [182]. To improve results, techniques such as parabolic curve fitting [208] and linear extrapolation [182], have been investigated to obtain sub-voxel velocity gradient approximations in the near-wall region [220, 213, 41]. When compared with Computational fluid dynamic (CFD) results in corresponding geometries, however, significant differences in WSS magnitude were still noted, although overall the WSS distributions were generally captured appropriately [214]. More recently, WSSNet, an ML-based model trained using CFD simulation data, was proposed [76]. WSSNet was shown to achieve significantly increased accuracy than the competing parabolic curve fitting method, particularly in attaining the correct WSS magnitude.

Despite the magnitude being under-predicted in *in vivo* imaging studies, an understanding of the role WSS plays in the progression of vascular disease is relatively well-established, where changes in the intensity or oscillatory characteristics of WSS have been shown to trigger a response from endothelial cells by means of mechanotransduction, which can cause disease [27, 96, 220]. The endocardium is similarly lined with endothelial cells, which have been shown to influence cardiac performance and remodelling [200], however, few studies have attempted to quantify the relationship with WSS. This is due to a combination



**Figure 2.2:** 4D-flow MRI: At each slice location, time-dependent phase images are reconstructed in three orthogonal directions, namely right-left (RL), anterior-posterior (AP) and foot-head (FH). Complete coverage of the LV cavity and cardiac cycle over the scan duration produces a time-varying, 3D volume of haemodynamic flow (shown as 4D).

of the aforementioned limitations of 4D-flow MRI, but more importantly, limitations in the structural MRI methods used to segment the moving endocardium. The complex structures that line the surface of the endocardium are not fully resolved, owing to low spatial resolution of structural cine-MRI, which means the surface on which WSS is being estimated is not anatomically correct. This is then coupled with extreme deformation of the boundary throughout the cardiac cycle, which is difficult to accurately capture, resulting in significant uncertainty in both the position and motion of the endocardium. This makes the task of deriving an accurate WSS approximation directly from 4D-flow MRI in the LV challenging, and further, given such significant differences in the boundary surface, the validity of such WSS measurements could also be questioned.

## 2.2.4 Quantification of Ventricular Flow

### 4D-Flow Magnetic Resonance Imaging

Given the inherent 3D nature of haemodynamic flow in the heart chambers, 4D-flow MRI has emerged as the gold standard *in vivo* flow imaging modality in LV studies, owing to its ability to reconstruct time-varying velocity fields in three spatial dimensions (3D) at a higher spatial resolution than competing methods (the title 4D refers to three spatial dimensions + time). This allows for full characterisation of the intraventricular velocity field, setting 4D-flow MRI apart from Doppler ECG and planar PC-MRI, which are only able to provide velocity in a single direction on a 2D plane. However, the acquisition of such a dense set of velocity measurements incurs large scan times, requiring 15-25 minutes to reconstruct a single cardiac cycle [152], and MRI machines are expensive and generally scarce within healthcare systems in low income countries. In Fig. 2.2, a simple demonstration of the phase images acquired and resulting velocity field is provided.

### Computational Fluid Dynamics

Computational fluid dynamics (CFD) describes the numerical simulation of fluid flows, and has seen extensive use in cardiac haemodynamic studies [39, 40, 222, 194, 135, 43, 134, 199, 60, 167, 9, 193]. CFD approximates the Navier-Stokes equations to a high degree of accuracy, allowing for the reconstruction of noise-free, high-resolution velocity and pressure fields in the cardiovascular system. The use of dense computational meshing procedures near flow boundaries also permits the calculation of WSS to a high degree of accuracy, which is challenging using *in vivo* flow imaging [213]. However, CFD approaches in the LV suffer from a number of significant drawbacks which reduce their clinical applicability. Firstly, the approximated solution fields are highly dependent on the accuracy of boundary conditions. This can be limited in cardiac studies, given uncertainty in the endocardium position and motion, and the inability to reconstruct the valve leaflets at the mitral inlet from structural imaging. The result of this is that, while they conform to the Navier-Stokes equations and the given boundary conditions, the reconstructed solution fields may not represent those truly present in the patient LV. This can be somewhat corrected using *in vivo* measurement data to calibrate the boundary conditions, but this requires the computation of multiple simulations. Secondly, CFD simulations in the LV are computationally expensive, and typically require many manual steps, from segmentation and registration to meshing [160]. This renders them challenging to apply to a large patient population.

Where CFD simulations can be particularly useful in this setting is for the generation of synthetic flow imaging data in real cardiovascular geometries [75, 71, 72, 203]. Although solution fields that have been simulated using CFD may not correspond directly to the true patient flow in a particular geometry, they are noise-free, of a high spatio-temporal resolution and satisfy the governing equations to an acceptable degree of accuracy. Therefore, it is possible to use the CFD results as reference ground truth data and artificially down-sample from this to produce synthetic image data. This type of study setup is helpful for validating 4D-flow MRI super-resolution methods, where validation *in vivo* is challenging given the absence of more accurate flow modalities. Further, using synthetic data allows for a controlled application of image corruption, which is helpful in determining the upper bounds of applicability of the model (in terms of levels of image degradation). Therefore, synthetic validation studies are common across all 4D-flow MRI super-resolution publications to date [75, 71, 72, 203].

### 2.2.5 The Complex Structures of the Endocardium

The interior surface of the LV cavity, called the *endocardium*, is lined with a matrix of complex, ridge-like structures called *trabeculae carneae*. From the endocardial surface, two *papillary muscles* protrude through the cavity and connect to the mitral valve by a web of fibrous tissue called *chordae tendineae*. The function of the papillary muscles and chordae tendineae in the LV is to prevent inversion of the mitral valve, (and subsequent regurgitation into the left atrium), during systole, where myocardial contraction generates extreme high pressures [162]. The functionality of trabeculae carneae, on the other hand, is less clear. Previously considered a passive byproduct of embryonic development [230], there is some evidence to suggest that they play a role in the haemodynamic function of the LV. It is hypothesised, for instance, that the trabeculae carneae help to squeeze blood from the apical region under contraction, improving washout in systole [170], while it is also postulated that their presence may reduce the diastolic load on the myocardium by slowing the inflow and thus dispersing the KE [25].

### Limitations of Structural Imaging

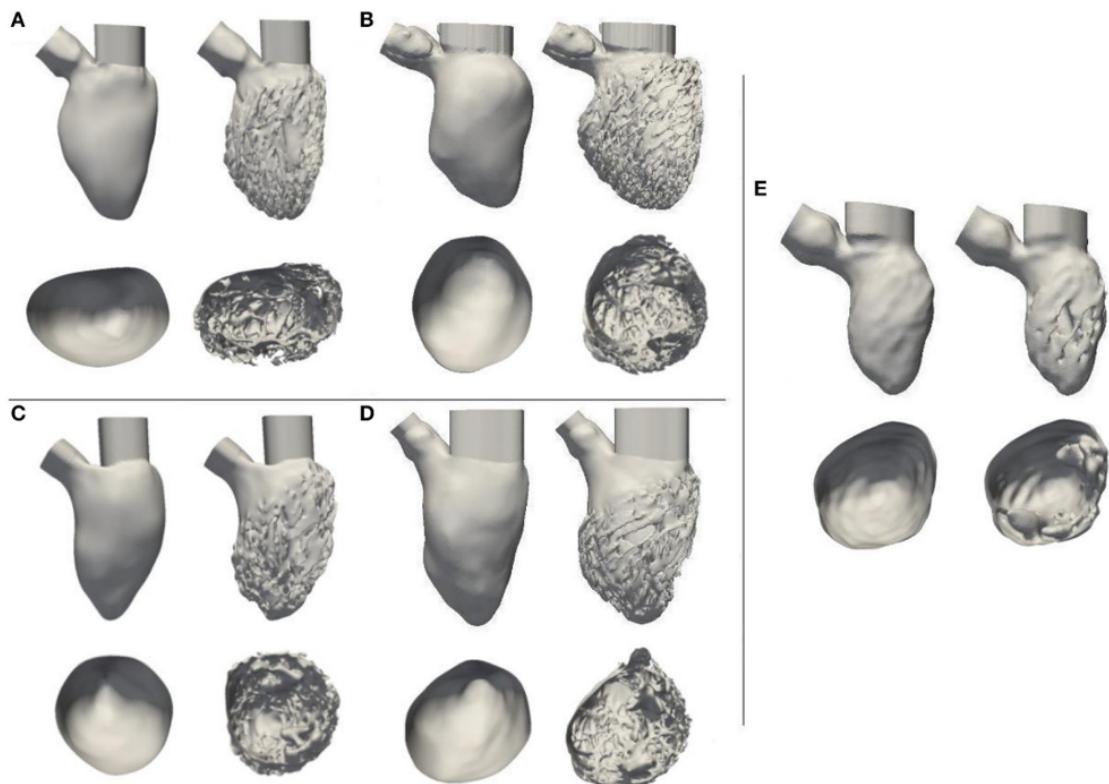
Regardless of their function, the presence of trabeculae carneae, papillary muscles and chordae tendineae has an impact on the observed intraventricular flow patterns due to their physical intrusion into the LV cavity [133]. However, most simulation-based studies of LV flow neglect to include these complex wall features, instead assuming that the

endocardium is a smooth surface with no protrusions [60, 39, 40, 28, 167, 90]. This simplification is primarily due to the low spatial resolution of *in vivo* structural imaging, in particular cine-MRI, but also due to the computational costs incurred by simulating such a complex, deformable mesh [194]. Cardiac cine-MRI has an in-plane spatial resolution of  $\sim 1.5\text{mm}^2$ , but a through-plane resolution of  $\sim 8\text{mm}$  [160], so while larger trabeculations may be visible within each planar slice, it is challenging to delineate these features in the apico-basal direction. Alternatively, cardiac computed tomography (CT) is able to achieve an isotropic spatial resolution of  $\sim 0.5\text{mm}^3$ , which is sufficient to reconstruct all but the smallest endocardial surface features. However, few publications have utilised this capability with CFD simulations, likely due to the computational challenges associated with prescribing motion to such a complex geometry [135].

There has been some investigation into the impact that trabeculae carneae, papillary muscles or chordae tendineae have on LV flow parameters [194, 222, 161, 164, 135, 133], with studies comparing CFD results with and without the presence of such features included on the endocardium. The influence of trabeculae carneae and papillary muscles on flow parameters and patterns is studied in [194], where CFD simulations are performed in patient-specific LV geometries using smoothed and trabeculated representations of the same endocardial surface, which can be seen in Fig. 2.3. It is observed that the trabeculations reduce WSS, increase pressure drop and alter the breakup pattern of the mitral vortex ring in diastole, in agreement with the theory discussed in [25]. It is also demonstrated that the inclusion of a thin, porous layer to represent the trabeculated region in the smoothed model can reproduce many of the flow characteristics observed in the trabeculated geometry, which may be a useful course of direction to emulate realistic intraventricular flow patterns without the computational cost of modelling the complex, anatomically-correct boundary. The authors use rigid walls, however, so the influence of contraction and relaxation of the different boundary representations is not considered. Such effects are investigated in [222], where a similar study is performed including motion of the lumen, again segmented from CT images. The authors found that the presence of trabeculations and papillary muscles produced more complex vortical features, inducing a deeper penetration of flow to the apex of the LV which may improve mixing and washout.

In [135], CT-derived CFD simulations are again performed in geometries with and without trabeculations and papillary muscles. The latter feature in particular was found to impact the simulated inflow due to blockage of the mitral jet, reducing flow to the apex, which is in contradiction with the findings in [222]. Further, regions of high residence

time were found to exist within the small spaces on the trabeculated surface, which may induce flow stagnation. However, it is unclear whether this phenomenon occurs in reality, or is a consequence of modelling simplifications that neglect contraction of the trabeculae during systole, which may act to squeeze the blood out [170]. A comparable study using segmentations from cine-MRI is performed in [161], in which CFD simulations are performed using three representations of the same endocardium, from smooth to trabeculated. It is found that the time-averaged WSS is elevated on ridged features near the mitral valve, with the trabeculated geometries exhibiting higher residence time in the apex due to decreased penetration of the mitral jet. This is in agreement with [135], although the endocardial segmentations used in this instance are of a poor quality, being derived from cine-MRI instead of CT, and do not appear anatomically accurate. The impact of chordae tendineae, which are comparatively less prominent than trabeculae carneae and papillary muscles, is analysed in [164]. The authors find that, although small, localised flow effects are induced behind the chordae tendineae, the effect on overall flow energetics and patterns is limited. Therefore, it is deemed that for most applications, the inclusion of chordae tendineae in CFD simulations is not required.

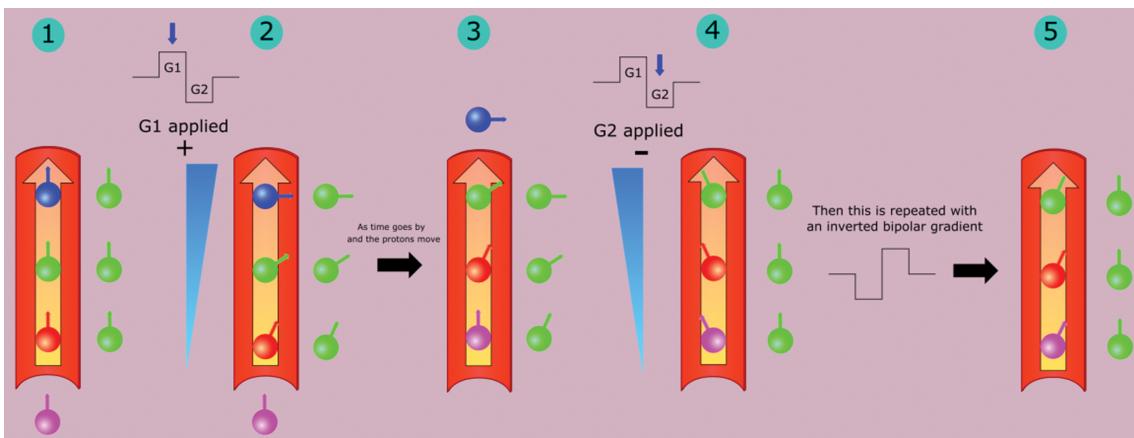


**Figure 2.3:** A demonstration of the level of detail present on the endocardial surface when segmented from computed tomography images (right-hand of each figure from A-E), displayed alongside the type of smoothed surface that cardiac blood flow modelling is typically performed in (left-hand of each figure from A-E) [194]

## 2.3 4D-Flow Magnetic Resonance Imaging

### 2.3.1 Physical Principles of Phase-Contrast Magnetic Resonance Imaging

Non-invasive quantification of blood flow in the cardiovascular system can be achieved using a variety of imaging techniques, however the increasingly dominant modality of choice for this application is Phase Contrast MRI (PC-MRI) [233]. Disregarded in structural MRI examinations, PC-MRI makes use of frequency and phase information recorded during acquisitions to reconstruct haemodynamic velocity fields in the cardiovascular system. Under the effect of a magnetic frequency gradient, the spin angle of a proton in a region of interest will change if observed for a period of time, leading to an altered phase shift. Applying two equal but opposite magnetic gradient pulses in sequence results in stationary protons having zero phase shift, whilst moving protons exhibit different degrees of phase shift dependent on their change in position relative to the gradient direction [233]. The phase shift experienced by a proton is directly proportional to its velocity in the gradient direction, and thus this information can be used to reconstruct a 2D slice of in-plane blood velocity. This process is outlined in Fig. 2.4.



**Figure 2.4:** This figure explains how a bipolar gradient pulse can be used to identify protons in motion, and thus calculate the related velocity. 1) Initially, each of the protons in both the moving frame (blue, green, red and purple protons) and stationary frame (remaining green protons) are in phase. 2) A gradient pulse is applied in the flow direction, producing a phase shift in each proton that is dependent on the position in the domain. 3) After the protons within the moving frame travel, their phase shift is different with respect to the neighbouring proton in the stationary frame. 4) An equal but opposite gradient pulse is applied. Protons in the stationary frame now exhibit zero net phase shift, whereas protons in the moving frame exhibit a constant phase shift, with protons travelling at the same velocity generating equal phase shift. 5) An inverted bipolar gradient is finally applied, resulting in an opposite phase shift. Values from the normal and inverted bipolar gradients are subtracted from one another, removing signal from the tissue regions [233].

The first application of this principle was realised in planar PC-MRI, in which through-plane velocity is measured on a single 2D slice [163]. This technique is well-suited to

studies in the vasculature, where analyses of cross-sectional flow can provide useful insights; However, the inherent 3D nature of flow in the heart chambers limits its usage in cardiac applications. The more recently developed 4D-flow MRI (3 spatial dimensions plus time) is far better suited for quantifying the haemodynamics of the heart [93, 84, 153, 218].

#### 4D-flow Magnetic Resonance Imaging

To construct 3D velocity fields for 4D-flow MRI, the principle of planar PC-MRI is extended by acquiring multiple orthogonal sequences to cover each spatial dimension at a particular point in time. Initially, to obtain velocity components in all three dimensions, six acquisition sequences were used to obtain the three dimensional velocity field (one positive and one negative in each spatial direction), resulting in the six-point method [233]. However, more recent encoding strategies including the four-point referenced method (one sequence in each spatial dimension, one sequence in some reference direction) [45] and the four-point balanced method (four equidistant points, with the initial used as a reference direction) [175] have become more popular. By acquiring two sequences fewer than the six-point method, the four-point methods reduce the overall scan time required, with the balanced approach improving velocity-to-noise ratio. The five-point balanced method [118] is one of the most recently developed encoding strategies, deploying an additional low-resolution encoding point to increase the velocity-to-noise ratio by 63% whilst increasing scan times by only 1% when compared to the four-point balanced method.

#### 2.3.2 Limitations of 4D-Flow MRI

Whilst providing unrivalled visualisation capabilities in the cardiovascular system, 4D-flow MRI is affected by a combination of hindrances. A balance between spatio-temporal resolution, velocity-to-noise ratio (VNR) and scan time must be realised to obtain acceptable results [72, 203, 75], where improved resolution and VNR can be achieved with scans of a longer duration, at the expense of more severe motion artefacts and increased flow averaging.

##### Scan Duration

The total scan time required to obtain adequate 4D-flow MRI results is universally large, when compared against Doppler-ECG and planar PC-MRI, although it differs depending on the location in the cardiovascular system. For cardiac imaging, acquisition times range from 5-25 minutes [61] (achieving a spatial resolution of  $\sim 2.5 - 3\text{mm}^3$ ), the highest of

any region in the body. Scan times of this duration or longer can result in the patient becoming uncomfortable and restless, particularly if they suffer from a debilitating illness, which can result in excessive motion artefacts. Further, scans of this duration may also increase variability in heart rate, leading to increased averaging in the measured flow and misalignment between cine-MR images and 4D-flow MRI [23]. Such flow averaging is inherent to 4D-flow MRI, potentially reducing peak velocities while removing beat-by-beat variations in the velocity field [210], but is exacerbated by longer scan times.

### Spatio-Temporal Resolution

For studies in the cardiac chambers, the recommended spatial resolution is  $2.5\text{mm}^3$ , up to a maximum of  $3\text{mm}^3$ , with scan duration a limiting factor [89]. At this resolution, small-scale flow features are missed and spatial gradients of the velocity field, which are required to calculate an array of clinically-relevant variables, can be of poor quality [182]. The temporal resolution of cardiac 4D-flow MRI is typically around 30–40ms [61], resulting in  $\sim 20$ –30 phases per cardiac cycle. Whilst this is not a limiting factor, it can result in a reduction in peak velocity measurements [61].

### Noise Corruption

Noise in MRI can be categorised into *structured noise*, like that occurring as a result of motion and ghosting, and *random noise*, which arises from the MRI machine itself [65]. Random noise in the measured *raw* MRI data, which consists of complex (real and imaginary) signals, is complex Gaussian in distribution [92, 29, 169, 106]. This is then skewed into a Rician distribution when converted to magnitude images [92, 29], whilst the phase information, from which velocity measurements are derived, can be characterised by a Gaussian distribution in high VNR regions, which is skewed in low VNR regions [106]. An important distinction to make here is that velocity-to-noise ratio (VNR) and signal-to-noise ratio (SNR) describe different parameters in the context of PC-MRI. SNR typically describes the signal-to-noise ratio in the magnitude images, which is independent of velocity values, whereas VNR is dependent on the measured velocity. The two parameters are related by:

$$\text{VNR} \simeq \frac{\pi}{2} \frac{|\mathbf{u}|}{v_{enc}} \text{SNR}, \quad (2.4)$$

for voxel speed  $|\mathbf{u}|$  and encoding velocity  $v_{enc}$  [138]. This equation explains why VNR is reduced when the  $v_{enc}$  is set excessively high.

In the synthetic studies in this thesis, we will consider data corruption in the form of re-

duced spatio-temporal resolution and increased levels of random noise. Whilst the impact of structured noise is an important consideration, realistic synthesis of such effects is challenging.

### 2.3.3 Advancing 4D-Flow MRI

Given the shortcomings discussed in the previous section, there has been a long-standing motivation in the field to accelerate acquisitions from as far back at the 1970s (for general MRI) [112], allowing for acquisitions of a higher quality. A range of approaches have been explored to achieve this goal, such as parallel imaging [57, 185, 217, 91, 102], compressed sensing [112, 149, 101, 207, 171, 85], novel encoding strategies [175, 118, 138, 45] and alternative  $k$ -space sampling patterns [178, 183, 12].

#### Parallel Imaging

Parallel imaging, utilised in almost all modern MRI machines, provides the framework to use multiple receiver coils in unison, which collectively acquire an undersampled set of  $k$ -space data. By acquiring a smaller amount of data from the  $k$ -space, significant acceleration is achieved [57]. This undersampling is made possible due to the fact that signals originating from coils at different spatial locations emit different information about spatial localisation [112]. MR images are then reconstructed from the data using techniques such as SENSE (Sensitivity encoding for fast MRI) [185, 217] and GRAPPA (Generalised Autocalibrating Partial Parallel Acquisition) [91, 102]. Although parallel imaging is a generalised MR method, it is similarly applicable 4D-flow MRI.

#### Compressed Sensing

Compressed sensing is a technique associated with a wide range of signal processing applications. It relies upon enforcing prior assumptions and conditions, primarily that the underlying signal is sparse (or, at least, can be represented sparsely), to reconstruct a representation of a signal using only a sparse set of measurements. Image compression, where raw images are condensed into much smaller files by utilising only a small set of frequency information without losing important details in the image, led to the concept of compressed sensing, where it was hypothesised that it may not be necessary to store such high-quality raw images since much of the information is discarded during compression [112]. This is of relevance in the field of MRI, where data acquisition is incredibly expensive, and thus in the last two decades this area of research has attained much focus.

Compressed sensing of MRI data is achieved through targeted undersampling of the  $k$ -space during acquisition [112]. Applied to cardiac 4D-flow MRI, an acceleration factor of 4.84 was achieved in [101] using Poisson disc  $k$ -space undersampling and a combined parallel imaging and compressed sensing algorithm (L1-SPIRiT). More recently, the respiratory controlled adaptive  $k$ -space reordering (ReCAR) has been introduced [149], for which an acceleration factor of 7.7 was achieved whilst maintaining relative errors of 13% for peak velocity and peak flow, compared with conventional 4D-flow MRI. This approach has been adopted in a number of recent publications across a range of cardiovascular applications [207, 171, 85].

### K-Space Sampling Strategies

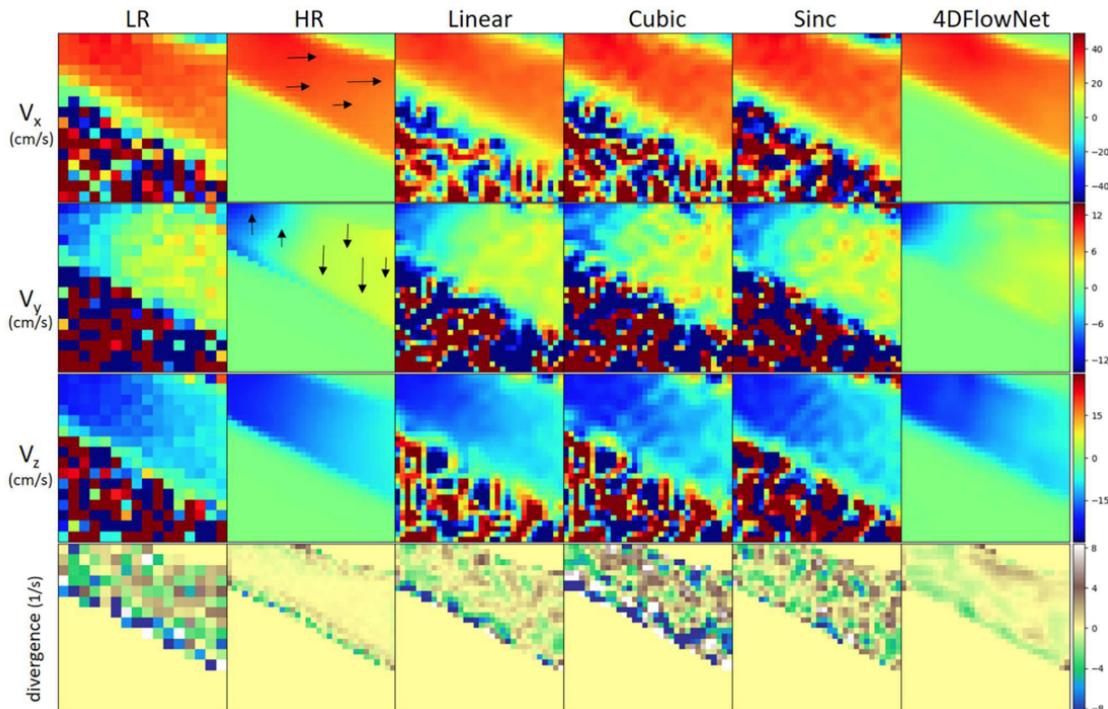
Whilst compressed sensing methods achieve acceleration through undersampling of  $k$ -space points, the grid on which points are sampled can also be modified to provide acceleration and improve robustness to motion artefacts [178]. Since the 1980s, Cartesian sampling strategies, in particular the spin-warp technique [78], have been the dominant choice, in which coverage is achieved by sampling lines across the  $k$ -space. The main benefit of this approach is that it improves computational efficiency, where the fast Fourier transform (FFT) is simple to apply to points on a Cartesian grid [68]. Radial sampling strategies, which actually out-date the more popular Cartesian methods, have since gained more attention with the emergence of the periodically rotated overlapping parallel lines with enhanced reconstruction (PROPELLER) approach [183], which has since been adopted by Siemens (BLADE) (Siemens Healthcare Limited, Camberley, UK) and Philips (MultiVane) (Philips Healthcare, Best, The Netherlands), among other manufacturers [67]. Radial sampling strategies are more robust to motion artefacts, owing to the natural over-sampling of the central region of the  $k$ -space image (within which the most important spatial information is contained) as the sampling lines rotate. However, re-gridding is required to convert the non-Cartesian sample points onto a regular grid to apply the FFT.

Specifically for use with 4D-flow MRI acquisitions, spiral sampling approaches have also been proposed [12], in which sampling occurs on a golden-angle spiral within the  $k$ -space. In a similar fashion, a Cartesian pseudo-spiral approach is deployed in addition to compressed sensing in [178], in which points are sampled in a spiral pattern, but on a Cartesian grid. In doing this, the fast reconstruction time of undersampled Cartesian sampling is maintained, however, the spiral pattern used mitigates eddy current artefacts by reducing the distance between sampling points.

### 2.3.4 Super-Resolution of 4D-Flow MRI

As mentioned in [203], the acceleration techniques discussed in section 2.3.3 can, rightly, be considered super-resolution techniques, as they allow for increased spatio-temporal resolution and VNR at the same scan duration. However, there is a small, but growing, community of researchers working on the development of direct super-resolution methods, which are applied directly to existing 4D-flow MRI velocity data. The main distinction between the two categories of acceleration-based super-resolution and direct super-resolution is that the former must be decided upon and applied prior to acquisition, whereas the second can be applied to existing data sets. It is also important to note that both approaches can be used in conjunction.

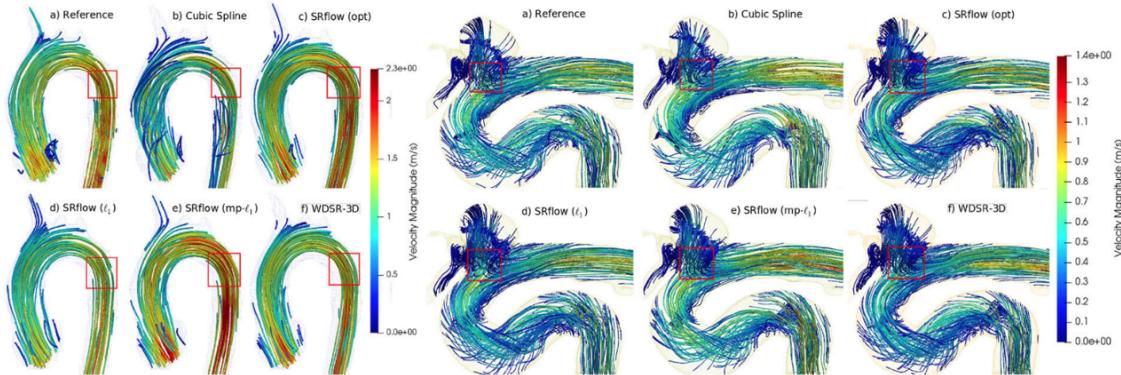
Direct super-resolution of 4D-flow MRI is an emerging field, and thus there are only a limited number of relevant publications to date. Early attempts utilised CFD simulations and proper orthogonal decomposition alongside ridge regression [9], Lasso regularisation [71] and dynamic mode decomposition [181]. Other approaches include [139], where an efficient super-resolution approach is introduced in which a Navier-Stokes-based  $L_2$ -penalised criterion is minimised, and [129], which uses a Fourier spectral dynamic method to assimilate 4D-flow MRI data into CFD studies.



**Figure 2.5:** Velocity component fields: Comparison between the 4DFlowNet model [75] and competing methods, taken from [75].

Given the groundbreaking success machine learning-based models have had when applied

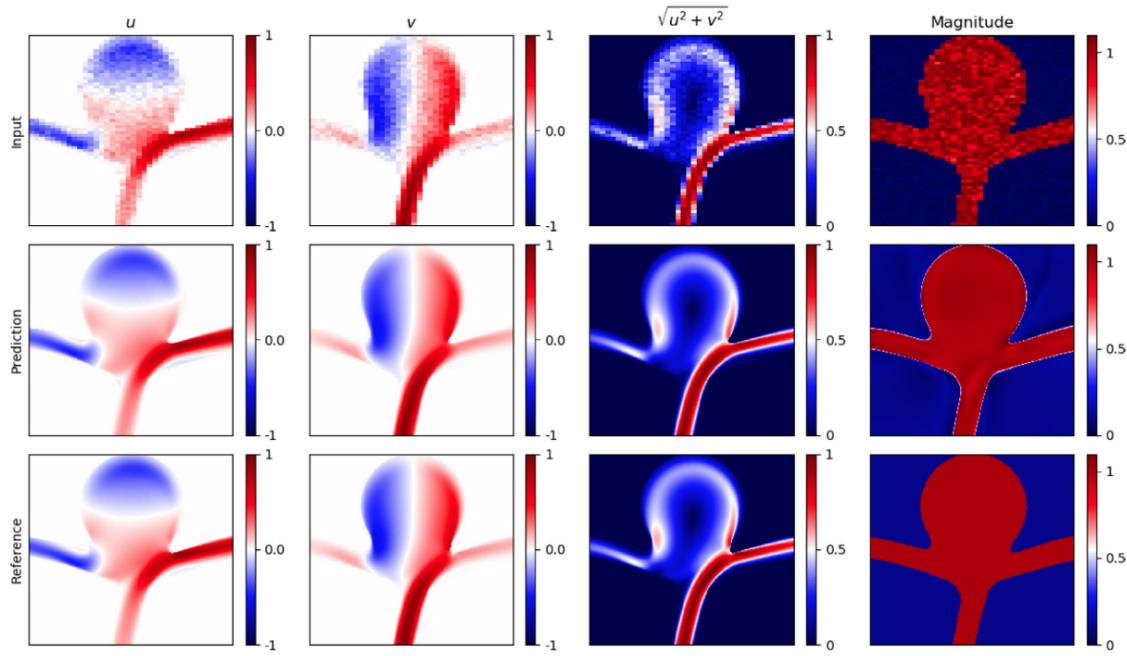
to super-resolution tasks, it is unsurprising that most recent methods in 4D-flow MRI super-resolution also deploy similar strategies. In [75], a residual network is adapted from the generator section of the super-resolution generative adversarial network (SRGAN) model [137], and is used to denoise vascular velocity fields whilst increasing spatial resolution by a factor of  $2\times$ , as visualised in Fig. 2.5.



**Figure 2.6:** Velocity streamlines: Comparison between the SRflow model [203] and competing methods, taken from [203].

A residual network is again deployed in [203] to perform spatial super-resolution, at an upsampling factor of  $4\times$ , and denoising in vascular flows, in which the authors introduce a mutually projected  $L_1$  loss designed to disentangle magnitude and direction of predicted velocities. Results are visualised in Fig. 2.6. Both approaches here use architectures designed for single image super-resolution, and therefore address only spatial upsampling. Not only does this not provide temporal super-resolution of 4D-flow MRI data, but the rich information stored in flow data at neighbouring time points is not utilised, which may also improve the spatial upsampling capabilities. Additionally, the methods rely upon the presence of paired low-resolution input data and high-resolution labels, which are not readily available in standard 4D-flow MRI studies.

In contrast, super-resolution in both space and time in vascular flows is achieved in [72], without high-resolution target data, using a physics-informed neural network (PINN), as seen in Fig. 2.7. Upsampling factors of  $100\times$  and  $5\times$  in space and time respectively are reported, and by utilising the often overlooked 4D-flow MRI magnitude images to estimate boundary location, their method avoids the need for complex manual steps such as segmentation and registration. However, whilst this is feasible for the simple rigid geometries presented in [72], it is unlikely that this approach would work in ventricular geometries due to the large deformation of the endocardium throughout the cardiac cycle. Further, the network architecture and internal components used are not capable of capturing complex, small-scale flow features, like those present in cardiac flow.



**Figure 2.7:** Velocity component and magnitude field: PINN super-resolution model reconstruction against reference data, taken from [72].

A key benefit of the PINN model [72], when compared to those based on residual networks [75, 203], is the inclusion of strict, physics-based model regularisation, enforced through loss constraints. This regularisation, the key feature of PINNs, ensures that predicted solution fields satisfy the underlying physical constraints that govern the haemodynamic flow. This provides confidence that the interpolated velocity field, and subsequent quantities that are derived from it, are physically consistent, which is of particular importance in cardiovascular studies where *in vivo* validation is challenging. This also links to the ill-posed nature of super-resolution, in that multiple (or even infinite) high-resolution predicted fields may correspond to the given low-resolution input, and thus having the assurance that the predicted fields at minimum satisfy the underlying physics is beneficial.

To date, there have been no publications addressing super-resolution of 4D-flow MRI in the cardiac chambers. There are a number of unique challenges posed in this setting, such as extreme deformation of the flow domain boundary and the presence of complex, multi-scale flow features, both of which are less prominent in vascular studies.

## 2.4 Physics-Informed Neural Networks

### 2.4.1 An Introduction to Machine Learning

The term *machine learning* (ML) covers a collection of methods from the field of artificial intelligence, in which models automatically learn relationships between input and output data. Model training strategies can be broadly divided into *supervised* or *unsupervised* learning (*self-supervised* and *semi-supervised* are also popular, but are omitted here). Supervised learning utilises data samples that are manually *labelled* prior to training. For instance, in supervised image recognition, an input-output data pair may consist of an input image, coupled with a text label output specifying the object in that image. Once trained on a set of paired data of sufficient size, the model will be able to classify new images based on the discrete set of labels defined beforehand. Typical use cases for supervised ML algorithms can be divided into *classification* and *regression* tasks. In the former, the outputs of an ML model are discrete, for instance in image recognition tasks. For regression tasks, an ML model can be used to fit a continuous function to a discrete set of data points, an example use case being super-resolution tasks [10]. Unsupervised learning is, conversely, training a model using unlabelled data. The classical use case for unsupervised learning is *clustering*, in which patterns and correlations in large sets of data are identified to generate natural groupings.

ML model training is facilitated through the minimisation of a *loss function* with respect to the *trainable parameters* of the network. For a fully connected network (otherwise labelled a multi-layer perceptron), the set of trainable parameters is a collection of weights and biases that connect a matrix of neurons, which apply a non-linear *activation function* to their output. Loss functions can take many forms depending on the application and available data type. As an example, for supervised regression tasks there will typically be a data loss function, which consists of the error (usually mean squared error) between predicted values and the corresponding data measurement at the given input point. During training, the gradient of the loss function with respect to the trainable parameters of the network is calculated using backpropagation, with these gradients then used to update the trainable parameters at each iteration using one of many optimisation algorithms (typically stochastic gradient descent and its variants).

In inference mode, a trained ML model is able to produce output predictions from a given input, regardless of whether that input was included in the training data set. The error

attained by testing the model using unseen input data is labelled *generalisation error*, which is typically the main metric of importance, and the most challenging to reduce satisfactorily. Generalisation error is closely tied to the problem of over-fitting, which describes the scenario in which an ML models begins to represent noise in the training data as the underlying approximated function itself. This occurs when a model is over-parameterised, or is allowed to train for too many iterations, to the detriment of the models generalisation capabilities.

### Machine Learning for Super-Resolution Tasks

Traditionally an imaging concept, the term *super-resolution* describes the task of recovering high-resolution (HR) images from low-resolution (LR) inputs. Due to the ill-posed nature of the problem, with multiple HR images corresponding to one LR image [229], exact recovery of HR images is not possible. However, there are a number of ways to achieve this objective to a satisfactory degree of accuracy, mainly through the use of filtering, statistics and prior knowledge. Given the challenging nature of the problem, it is difficult to formulate algorithms *a priori* that can perform this task effectively. Therefore, it is a field for which ML is well-suited.

There has been extensive exploration into the development of ML methods for super-resolution over the last decade, many of which can be found in within [229]. Similarly, these methods have been successfully applied to structural medical imaging data across a range of modalities. A review of these methods can be found in [140].

### Super-Resolution of Fluid Flows

Whilst super-resolution of 4D-flow MRI data is a field in its infancy, super-resolution of fluid flows in general is relatively well established. This is partly due to a long-standing demand for high-fidelity flow reconstructions calculated in a shorter time than those achievable with traditional methods such as direct numerical simulation (DNS) and large eddy simulation (LES), but also for improving the resolution of experimental flow measurements, which tend to be sparse and noisy [55]. As a consequence of the former, there is a notable interest in super-resolution of turbulent flows, particularly for industrial applications where DNS simulations are too costly [142]. Super-resolution techniques have been used for particle image velocimetry data for over two decades [123, 215, 211, 11, 236, 197]. A variety of approaches have been attempted, from Kalman-filtering and  $\chi^2$ -testing [215] to the integration of particle tracking velocimetry (PTV) algorithms [11, 211]. Further

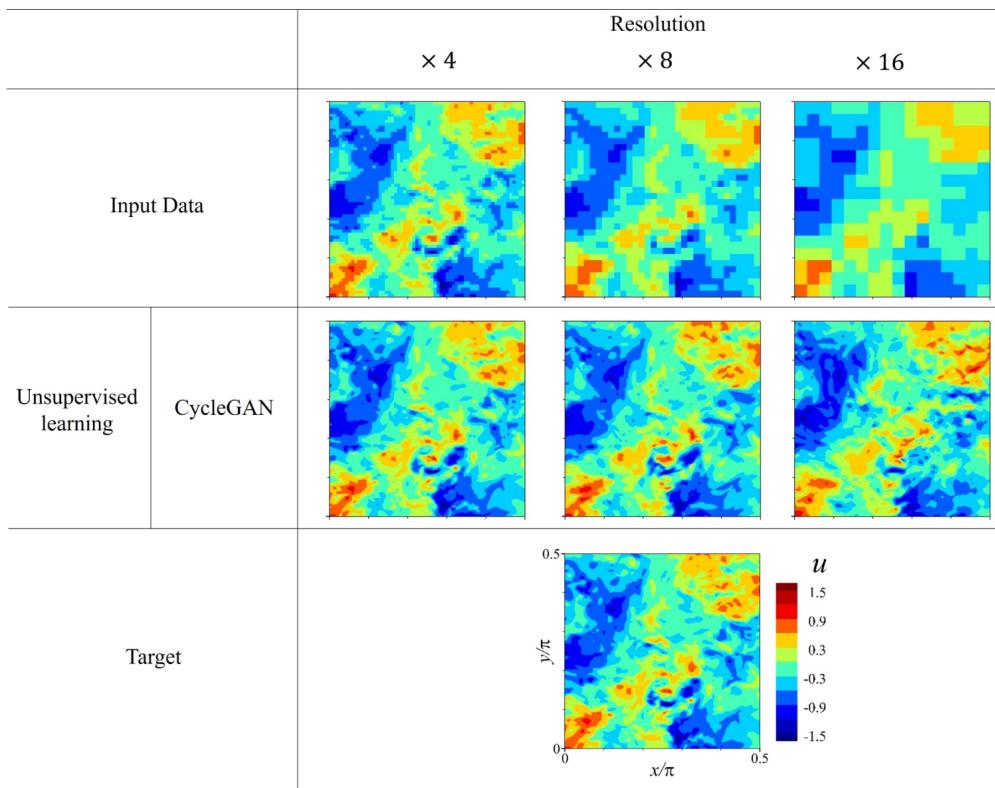
techniques include the successive abandonment method [236] and the redefinition of the resolution to correspond to either the mean particle spacing or mean particle displacement [123].

Well established in the imaging domain [229], machine learning is showing great promise in the field of fluid dynamics, especially in situations where measured data is readily available [24]. As with general machine learning-based super-resolution tasks, generative adversarial network (GAN) [55, 126, 234] and convolutional neural network (CNN) [80, 142, 79] architectures dominate more recent publications in this field, with low-resolution data used in these cases either downsampled from the high-resolution target [142, 126, 80, 55] or obtained using a turbulence closure model such as LES [126].

In [146], the authors deploy two CNNs to produce super-resolved reconstructions of turbulent flow, namely static CNNs (SCNNs) and multiple temporal paths CNNs (MTPC). The MTPC model is a novel contribution, leveraging not only the spatial LR data but also temporal data, unlike static models. The model is able to recover HR flow fields whilst preserving physical features such as kinetic energy spectra. With previous, static, methods, only the instantaneous spatial LR data is used at each time frame. Although this approach has been shown to be effective, data from other time frames is wasted. The MTPC method provides the network with a sequence of time frames, imparting further constraints whilst learning the temporally non-local properties of the turbulence. Further applications of CNNs to super-resolution are provided in [79, 80], who investigate the reconstruction of DNS-generated turbulent flows from both coarse input data and LES-generated data.

Generative adversarial networks (GANs) [88] are unsupervised learning algorithms that employ a two-player zero-sum game in which one player (the generator network) attempts to learn the distribution of input samples, and the other player (the discriminator network) learns to discriminate between real and generated samples [229]. The outcome is a generator network capable of producing new samples that are indistinguishable from the input data. These networks have been applied to the problem of single-image SR [229], and this has been successfully extended to flow SR. In [234], the TempoGAN method is proposed to provide SR of fluid flows, deploying a novel temporal discriminator which forces the generator to produce results displaying consistent temporal evolution. In [55], two methods previously used in image SR are adapted for use in fluid SR, super-resolution GAN (SRGAN) and enhanced SRGAN (ESRGAN). They use flow data acquired experimentally using PIV and are able to upsample by a factor of up to  $8^2$ . The drawback of this

approach is the requirement for paired HR data in training. This issue is circumvented in [126] by applying a cycle-consistent GAN (CycleGAN), capable of performing unsupervised SR with unpaired data. The results of their approach can be seen in figure (2.8), in which the authors use their model to reconstruct 2-dimensional turbulence (only the  $u$  velocity is displayed) from DNS-generated LR input data, displaying good capability even for downsampling by a factor of 16. Whilst this approach allows the use of unpaired data, it still assumes the existence of HR data, which are commonly not available. A drawback of GANs in general is their tendency to hallucinate features in the predicted fields. In tasks such as image super-resolution, this is not such a major issue. However, the presence of such hallucinations in predicted velocity fields would likely prove problematic when applied in practice. This effect can be seen in the right-hand column in Fig. 2.8. At first glance, the model has successfully super-resolved the velocity field, generating many of the small-scale flow features present in the target data. However, on closer inspection, the position and intensity of such features are misaligned. The underlying architectures



**Figure 2.8:** The  $u$  velocity of 2D turbulence reconstructed from LR input by cycleGAN [126]

of both CNNs and GANs were originally developed for single image super-resolution, and thus the input data typically have to be structured on a regular grid in order to perform the necessary convolutions. Therefore, these methods are difficult to apply to non-rectangular flow domains, which precludes their usage in a wide number of fluid applications. This issue is addressed in [83], in which the authors apply an elliptic coordinate transformation

to project a regular grid onto an irregularly shaped flow domain, permitting the use of a CNN. They are also able to address another limitation of the methods discussed above, namely the requirement for both low- and high-resolution data during training. They do this by using the Navier-Stokes equations and boundary conditions, enforced in the loss function, to constrain the predicted solutions. These physical constraints are imposed by convolving over the predicted velocity field using fixed finite difference filters, then minimising the resulting fields. This approach belongs to a class of machine learning methods, of which PINNs also belong, called *physics-informed machine learning*, which will be discussed later in this chapter.

#### 2.4.2 Physics-Informed Machine Learning

Around 2009, the emergence of highly parallel graphics processing unit (GPU) technology enabled significant acceleration of ML tasks, sparking interest in the development of complex models with millions of trainable parameters. This field of research became known as *deep learning*, owing to the high number of stacked layers within the neural networks used, i.e. *depth*. With the increased computational complexity of ML models came the capacity to learn features and correlations across vast, multi-dimensional input spaces [120]. Training of such models relies heavily on having large amounts of data across such input spaces, making deep learning models a natural fit for tasks such as image recognition and natural language processing where immense banks of publicly available data exist.

When working on scientific problems, comparatively large data sets are often unavailable, precluding the application of many deep learning approaches when modelling physical phenomena. This may be due to the cost of acquisition [189], difficulty in accessing the problem domain or inability to directly measure variables of interest [191]. However, in this setting we often have detailed *a priori* knowledge of the physical system at hand, not utilised by purely data-driven deep learning models, which can include known physical laws, boundary conditions, symmetries or empirically validated information [189]. This additional information can be utilised to heavily regularise model outputs by instilling physically-meaningful inductive biases, to compliment existing observational biases [120], which allows for effective model training in the *small-data* regime in which scientific problems commonly lie. Methods that utilise such *a priori* physical information form the field of *physics-informed machine learning* (PIML).

There exist numerous distinct realisations of this concept, with differences in how the

informed regularisation is applied. The two approaches can be broadly categorised into *soft* and *hard* constraint enforcement strategies. The former penalises outputs if they do not conform to a particular constraint, whereas the latter forces the model to follow some condition, usually via modifications of the architecture [148]. The majority of works in the PIML literature, including the popular PINNs [189], utilise soft constraints, which generally involve minimising physical laws or boundary conditions as additional loss terms.

### 2.4.3 An Introduction to PINNs

The term *PINN* is generally used to describe architectures like that used in the original publication [190], which consists of a fully-connected neural network that maps spatial and/or temporal coordinates to corresponding solution fields, constrained in a soft manner to obey some underlying physical condition via a *physics-informed* term within the loss function.

#### Overview

In Fig. 2.9, an overview of a standard PINN setup for a steady 1D problem with one output variable,  $u$ , is shown. The objective here is to predict a continuous solution to  $u$  across an entire domain in  $x$ , where the exact solution is shown in the upper figure. It is assumed that there is access to a sparse set of measurements of the variable at positions  $\bar{x}$ , given by  $\bar{u}$ , and it is known that the underlying partial differential equation (PDE) governing the dynamics of our system is given by

$$\partial_x u = G(u, x). \quad (2.5)$$

Given the sparsity of our measurements,  $\bar{u}$ , traditional data-driven machine learning tools to approximate  $u$  across the whole domain cannot be applied. However, using a PINN it is possible to exploit the *a priori* knowledge of the system, given by the PDE in Eq. 2.5.

The PDE is re-formulated as a residual, to give:

$$F(u, x) = \partial_x u - G(u, x), \quad (2.6)$$

which can be used as a soft constraint in the loss function. By including  $F(u, x)$  in the loss function, network predictions that violate the underlying PDE are penalised, which heavily restricts the space of possible solutions in  $u$ . The derivative terms to be used

within the PDE loss component are calculated using *automatic differentiation* [13], a standard tool included in all major machine learning codes which is used primarily for calculating gradients during backpropagation. This powerful algorithm is exploited in PINNs to calculate the gradients of the output variable, here  $u$ , with respect to the input variable, here  $x$ , to machine precision. This PDE constraint is used in conjunction with a data loss term, which penalises predictions that do not comply with the sparse set of existing measurements. The total loss is therefore taken to be the sum of the individual loss components, i.e.

$$\mathcal{L} = \mathcal{L}_{PDE} + \mathcal{L}_{data}, \quad (2.7)$$

where

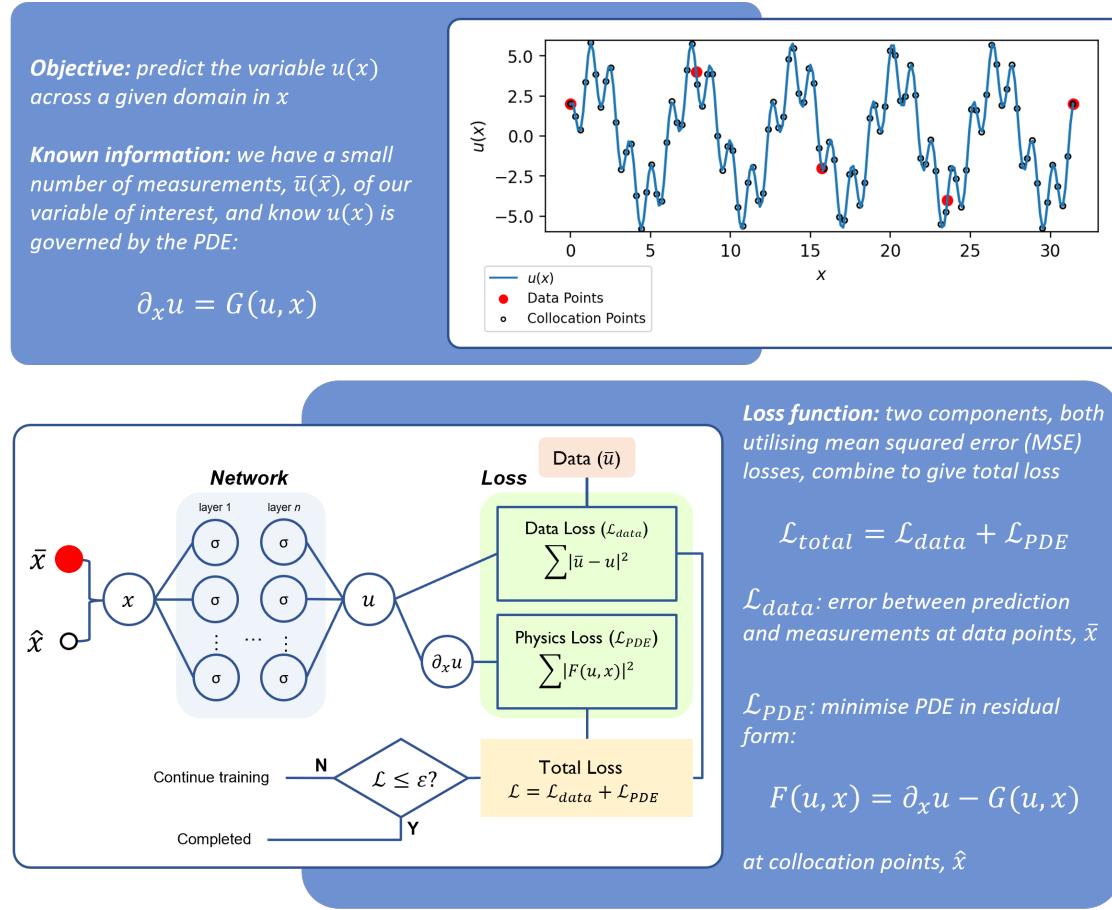
$$\mathcal{L}_{PDE} = \frac{1}{N_p} \sum_{i=0}^{N_p} |F(u(x_i), x_i)|^2, \quad (2.8)$$

$$\mathcal{L}_{data} = \frac{1}{N_d} \sum_{i=0}^{N_d} |u(x_i) - \bar{u}_i|^2, \quad (2.9)$$

for collocation (PDE) set size  $N_p$  and data set size  $N_d$ . By minimising both loss terms in unison, the PINN is trained to fit a continuous approximation of  $u$  to the underlying data whilst also satisfying the governing equations.

A distinct set of coordinates are used to interrogate both data and PDE loss terms, as seen in the upper figure in Fig. 2.9, where, conventionally, the PDE loss points are named ‘*collocation*’ points. It is typical to sample the collocation point set far more densely than the data point set, as it is computationally inexpensive to do so in comparison with acquiring more measurement data. In an analogy to numerical methods, the collocation point cloud can be viewed as the computational mesh, although there is no concern for connectivity here. Sampling more densely generally provides more accurate solutions, but unlike with conventional numerical methods, the parallelisation afforded by GPU architectures means increasing the density of the collocation point set is not as costly as increasing the density of a computational mesh. However, there exists a number of adaptive collocation cloud refinement strategies to further reduce computational overhead, which will be discussed later in this chapter.

Whilst they are able to operate without measurements of the system at hand, arguably the greatest benefit PINNs have over numerical and other such methods is the ease with which data is assimilated, and therefore a large number of application-focused publications utilise PINNs in this manner. In a practical setting, this is advantageous as data measurements



**Figure 2.9:** A basic overview of the PINN workflow, using a simple 1D example.

are rarely constrained to a boundary or initial state, and rather distributed randomly across a solution domain.

### Network Architecture

As defined in [190], the internal architecture of a PINN takes the form of a fully connected neural network (FCNN), also referred to as a multi-layer perceptron (MLP). A FCNN is built by stacking multiple layers of *neurons*, specifically, an input layer,  $m - 2$  hidden layers, and an output layer. Each hidden layer,  $\mathcal{L}_i$  for  $i \in 2, \dots, m - 1$ , contains  $m_i$  neurons (which is generally fixed across layers but may vary), with each neuron in layer  $i$  connecting to each neuron in layer  $i + 1$  and vice versa for  $i \in 2, \dots, m - 1$ . Hidden layer  $\mathcal{L}_i$  of the network receives the output from the previous layer, given by  $\mathbf{x}^{i-1} \in \mathbb{R}^{N_{i-1}}$ , and applies a transformation of the form:

$$\mathcal{L}_i(x^{i-1}) = \mathbf{W}^i \mathbf{x}^{i-1} + \mathbf{b}^i. \quad (2.10)$$

Here,  $\mathbf{W}^i \in \mathbb{R}^{N_i \times N_{i-1}}$  corresponds to the weight matrix of layer  $i$ , with  $\mathbf{b}^i \in \mathbb{R}^{N_i}$  the bias vector. The matrix  $\mathbf{W}^i$  is composed of all the weights connecting each neuron in layer

$i - 1$  to layer  $i$ , where the vector  $\mathbf{b}^i$  contains individual bias values for each neuron in layer  $i$ . A non-linear activation function,  $\sigma_i(\cdot)$ , is then applied to each neuron in layer  $i$ , and this setup is repeated for each layer in the network. Finally, the complete network is composed as:

$$\mathbf{y}_\theta(\mathbf{x}_0) = (\mathcal{L}_m \circ \sigma_{m-1} \circ \mathcal{L}_{m-1} \circ \dots \circ \sigma_2 \circ \mathcal{L}_2)(\mathbf{x}_0), \quad (2.11)$$

for network input  $\mathbf{x}_0$  and complete set of trainable parameters  $\theta$ , where

$$\theta = \sum_{i=2}^m \mathbf{W}_i + \mathbf{b}_i. \quad (2.12)$$

The input to a PINN typically consists of spatial and temporal coordinates within the solution domain in a *low-dimensional* manner, meaning one training sample is given by a single point in space and time. This is distinct from many common machine learning approaches, where *high-dimensional* input spaces are used to encode large input arrays, such as pixels structured on an image. Similarly, the output to a PINN is also low-dimensional, where predicted variables at each given spatio-temporal location are produced. This structure is referred to as an *implicit neural representation* [206], where the neural network itself can be viewed as a continuous parameterisation of the underlying function being approximated on the domain of interest. There are notable benefits to using such a structure. Firstly, the network approximates a continuous solution field during training, meaning predictions are agnostic to the grid resolution. This allows simple evaluation on grids of differing resolution if required. Secondly, and most importantly in the context of PINNs, by structuring the network in such a way we are able to calculate the gradients required for the physics-based loss constraints to machine precision, using automatic differentiation.

## Automatic Differentiation

Automatic differentiation refers to a class of methods used to compute function derivatives, and is used across a wide range of applications from machine learning [13] and optimisation [223], to fluid dynamics [86, 21]. It is algorithmically distinct to both numerical (finite difference, finite element, finite volume etc.) and symbolic differentiation (as found in Mathematica (Wolfram Research, Inc., 2024), Maple (Maplesoft, a division of Waterloo Maple Inc., Waterloo, Ontario) etc.), and alleviates many of the issues associated with these techniques. Numerical differentiation is prone to round-off and truncation errors [114], but more importantly, scales poorly for gradients. Given that machine learning requires gradients with respect to millions of parameters, this effectively rules out the utilisation of numerical differentiation for this task. Conversely, symbolic differentiation

is not so prone to the same errors as numerical differentiation, but requires converting computational code into a mathematical expression. For large-scale, complex tasks, such as machine learning, this is almost intractable.

Automatic differentiation suffers from none of the aforementioned shortcomings, allowing for computation of derivatives at machine precision with minimal computational overhead [13]. The cornerstone of automatic differentiation is the decomposition of differentials, facilitated by the rules of differential calculus [223]. Every numerical computation can be expressed as the composition of a finite set of operations, for which derivatives are known, such as arithmetic operations and trigonometric functions [13]. It is therefore possible to construct a set of operations and the corresponding derivative compositions, from which a trace of elementary operations can be constructed to form a graph. An example of *forward mode* automatic differentiation, published originally in [13], can be found in Tab. 2.1. Here, the function  $y$  is decomposed into a set of operations, given by  $v_i$ , to form the primal trace, found in the left side of the table. Then, for each  $v_i$ , the corresponding derivatives,  $v'_i$ , are listed in the derivative (or tangent) trace, on the right hand side of the table. By evaluating both traces in unison at a particular set of values for  $x_1$  and  $x_2$ , we are able to calculate the specified derivative. A review of this topic and further applications of automatic differentiation can be found in [13], with an introduction to the topic found in [223].

**Table 2.1:** An example taken from [13], in which the forward mode automatic differentiation is used to evaluate the derivative,  $\frac{\partial y}{\partial x_1}$ , of the function  $y = \ln(x_1) + x_1 x_2 - \sin(x_2)$  at  $(x_1, x_2) = (2, 5)$ . With  $x'_1 = 1$ , we specify that we intend to calculate the derivative with respect to  $x_1$ .

Forward primal trace	Forward derivative trace
$v_0 = x_1 = 2$	$v'_0 = x'_1 = 1$
$v_1 = x_2 = 5$	$v'_1 = x'_2 = 0$
$v_2 = \ln v_0 = \ln 2$	$v'_2 = v'_0/v_0 = 1/2$
$v_3 = v_0 \times v_1 = 2 \times 5$	$v'_3 = v'_0 \times v_1 + v'_1 \times v_0 = 1 \times 5 + 0 \times 2$
$v_4 = \sin v_0 = \sin 5$	$v'_4 = v'_1 \times \cos v_1 = 0 \times \cos 5$
$v_5 = v_2 + v_3 = 0.693 + 10$	$v'_5 = v'_2 + v'_3 = 0.5 + 5$
$v_6 = v_5 - v_4 = 10.693 + 0.959$	$v'_6 = v'_5 - v'_4 = 5.5 - 0$
$y = v_5 = 11.625$	$y' = v'_6 = 5.5$

Backpropagation, which is ubiquitous in the field of machine learning and optimisation, is a specific instance of *reverse mode* automatic differentiation. Therefore, it is natural that other, more general forms, of automatic differentiation are implemented in modern machine learning packages such as TensorFlow [1], which provide ease of usage. This functionality is exploited in PINNs to generate the necessary gradients for use in the loss function, where, in practical terms, one simply has to specify an input and output variable over which to calculate the gradient, provided the two are connected through the

computational graph. For instance, returning to the example in Fig. 2.9, assume we would like to calculate the first and second derivatives of our output variable,  $u$ , with respect to the input variable,  $x$ , (which is possible because  $x$  and  $u$  are connected through the neural network). Written in TensorFlow, this operation is simply carried out using the code:

```
def gen_gradients(x):
    u = PINN(x)
    u_x = tensorflow.gradients(u, x)[0]
    u_xx = tensorflow.gradients(u_x, x)[0]
    return u_x, u_xx
```

Note that here  $\text{PINN}()$  represents the neural network itself, and thus  $u = \text{PINN}(x)$  generates the network output with respect to  $x$ .

## Activation Function

The term *activation function*, denoted  $\sigma$ , defines the non-linear function applied to the output of a neuron within a neural network. The use of non-linear activation functions is necessary for a neural network to approximate non-linear functions, since a composition of linear functions is itself a linear function. Common activation functions include sigmoid (or logistic) ( $S(x)$ ), hyperbolic tangent ( $\tanh(x)$ ) and rectified linear unit ( $\text{ReLU}(x)$ ) [81].

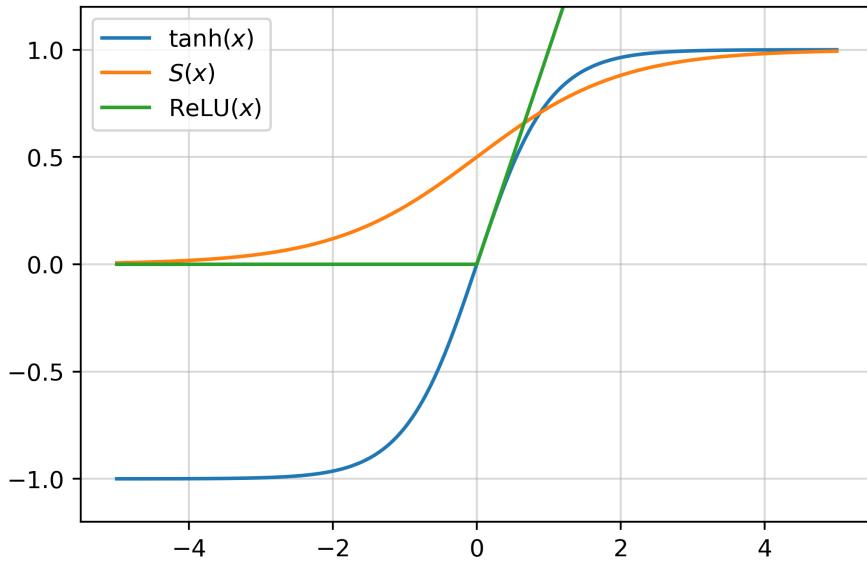
These functions are defined as:

$$S(x) = \frac{1}{1 + e^{-x}}, \quad (2.13)$$

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad (2.14)$$

$$\text{ReLU}(x) = \max(0, x). \quad (2.15)$$

In Fig. 3.9, we plot curves for these three activation functions, displaying the characteristics of each function. Across all machine learning applications in the last decade, the ReLU activation function is by far the most dominant choice owing to several useful properties. Firstly, neurons receiving inputs that are less than zero produce no output which promotes sparsity, reducing model complexity. Secondly, the derivative of a ReLU is always constant, alleviating issues around vanishing gradients which are present with other activation functions. However, it has been shown theoretically and empirically that a PINN model operating with non-smooth activation functions, such as ReLU, does not converge to the exact solution under ideal circumstances [151], owing to the fact that



**Figure 2.10:** Curves for three of the most common fixed activation functions,  $S(x)$  (sigmoid),  $\tanh(x)$  and  $\text{ReLU}(x)$ .

non-smooth functions are only Lipschitz continuous [159]. Therefore, smooth activation functions such as  $S(x)$  and  $\tanh(x)$  are used as standard in the PINN literature, alongside more advanced activation functions which will be discussed later in this section.

### Loss Function

For the purposes of this section, it is to be assumed that the loss function is of the form in Eqn. 2.7, which consists of one measurement error constraint and one physics regularisation constraint. This basic form encompasses much of the wider field of PINNs, in which measurement data is assimilated into physics-based models. There are PINN architectures and applications that do not use measurement data, in which model outputs are constrained by physical constraints and boundary conditions only; however, these are in the minority across the wider literature and are less relevant for the present application.

As introduced in 2.7, the total loss function,  $\mathcal{L}$ , in a PINN is formed as the summation of the individual loss constraints, in this case  $\mathcal{L}_{PDE}$  and  $\mathcal{L}_{data}$ . For the individual loss components, the mean-squared error (MSE) loss is used almost ubiquitously throughout the PINN literature, taking the form as shown in Eqn. 2.8. The key benefit to the MSE function is its ability to suppress outliers present within data, although models incorporating the MSE in the loss function may be sensitive to these outliers if significant levels of high-magnitude noise are present. Since the function is squared, higher-magnitude errors between measurement and prediction contribute more significantly to the loss value.

However, this benefit can also become a hindrance if a large quantity of significant outliers (i.e. high levels of noise) are present within a training data set, where outliers impart too strong an influence on model training at the expense of learning the underlying function.

As is perhaps intuitively obvious, combining two or more distinct loss components in the form of Eqn. 2.7 can produce imbalances in training convergence if the magnitude of individual loss components is imbalanced. To account for this, static and dynamic weighting strategies have been proposed, which will be discussed in detail later in this chapter.

#### 2.4.4 The Limitations of PINNs

##### Spectral Bias

The term *spectral bias* characterises the tendency of deep neural networks, including PINNs, to favour low-frequency solutions, with high-frequency solution components requiring significant training time to learn without the implementation of necessary countermeasures [187]. A consequence of spectral bias is prohibitively slow convergence to functions that exhibit multiple scales or dominant high-frequency components, which are frequently encountered in the field of fluid dynamics, for instance. In [187], this phenomenon is demonstrated theoretically and empirically, analysing the rate of learning across different frequency ranges in both synthetic and real experiments. In approximating complex sinusoidal functions, there is a clear hierarchy of learning throughout the training cycle, with low-frequencies learned first, followed by higher-frequencies. An interesting discovery is the apparent attenuation of spectral bias when embedding data that lies on a low-dimensional manifold into a higher-dimensional input space, dependent on the shape of the data manifold. It is shown that the learning of higher-frequency components commences earlier when data is embedded on manifolds of increasingly complex shapes.

Loosely related to this work is the Fourier feature network, introduced in [216]. Here, the authors focus specifically on implicit neural representation networks (labelled ‘coordinate-based’ MLPs here), and find that mapping input coordinates onto a vector of Fourier features [188], seen in equation 2.16, greatly improves the capability of such networks to capture high-frequency solution modes. In this formulation, input coordinates  $\mathbf{x}$  are mapped to:

$$\gamma(\mathbf{x}) = [a_1 \cos(2\pi \mathbf{b}_1^T \mathbf{x}), a_1 \sin(2\pi \mathbf{b}_1^T \mathbf{x}), \dots, a_m \cos(2\pi \mathbf{b}_m^T \mathbf{x}), a_m \sin(2\pi \mathbf{b}_m^T \mathbf{x})]^T \quad (2.16)$$

before being passed to the hidden layers of the network, for frequency vectors  $\mathbf{b}_i$ .

The problem of spectral bias is framed in the context of neural tangent kernel (NTK) theory [107], which relates the training of neural networks to kernel methods by using a kernel (the NTK) applied to wide neural network architectures. The NTK describes the dynamics of a neural network during training, where it has been demonstrated in [5] and other publications that convergence rate during training is governed by eigenvalues of the NTK matrix, where eigenvectors corresponding to higher eigenvalues are learned at a faster rate. In relation to spectral bias, it is shown in [216] that the eigenvalue spectrum of an NTK decays as a function of frequency. With a standard neural network, low-frequency function components correspond to NTK eigenvectors with larger eigenvalues, and thus will be learned at a faster rate. However, the eigenvalues of the NTK corresponding to higher-frequency components typically decay rapidly, which means that learning function components that relate to such eigenvectors is slow [216]. Thus, by transforming the NTK to increase the magnitude of eigenvectors corresponding to higher-frequency solution components, it can be expected that effects associated with spectral bias may be alleviated. The inclusion of a Fourier feature mapping achieves this, whilst making the NTK stationary (or shift-invariant), a desirable feature for image processing applications. There still remains a problem-specific tuning required to select the correct frequency vectors,  $\mathbf{b}_i$ , which is discussed at length in [216].

There are similarities between Fourier feature [216] and Siren [206] networks, in which sinusoidal activation functions are used in conjunction with a specific initialisation strategy, which are highlighted in [239]. Through Thm. 1 in [239], it is shown that both network compositions encode signals in a similar manner, and share a similar expressive power (the space of functions that can be represented with these networks).

Where the frequency vectors,  $\mathbf{b}_i$ , are tuned to attenuate spectral bias in Fourier feature networks, tuning an initialisation frequency,  $\omega_0$ , has a similar effect within Siren networks. When  $\mathbf{b}_i$  are generated by drawing samples from  $\mathcal{N}(0, \sigma^2)$  (which is shown to be the most effective technique in [216]), selecting a  $\sigma$  value that is too low can lead to underfitting, whereas setting this value too high leads to overfitting and noisy solution fields. An almost identical effect is seen when altering  $\omega_0$  in the Siren network. The relationship between these two parameters becomes apparent when comparing the first layer of the Siren network with the Fourier mapping of the Fourier feature network. Firstly, consider a Fourier mapping that consists solely of sinusoidal functions with  $a_i = 1$ , and the elements

of  $\mathbf{b}_i$  drawn from  $\mathcal{N}(0, \sigma^2)$ , given by:

$$\gamma(\mathbf{x}) = \sin(\mathbf{B}\mathbf{x}), \quad (2.17)$$

where  $\mathbf{B} = [\mathbf{b}_1^T, \dots, \mathbf{b}_m^T]$ . Then, the first layer of a Siren network, after the application of the activation function, is given by:

$$\mathbf{z}^1(\mathbf{x}) = \sin(\omega_0(\mathbf{W}^0\mathbf{x} + \mathbf{b}^0)). \quad (2.18)$$

At initialisation, the weights in the Siren network,  $w_i$ , are randomly drawn from a uniform distribution, given by  $\mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$ , where  $n$  is the input shape (of  $\mathbf{x}$ ). It is shown in [206] that under this initialisation strategy, the input to each sinusoidal activation function is normally distributed with a standard deviation of 1, and thus, in Eq. 2.18, elements in the array  $\mathbf{z}_0 = \mathbf{W}^0\mathbf{x} + \mathbf{b}^0$  will follow this normal distribution. Therefore, multiplication of  $\mathbf{z}_0$  by  $\omega_0$  acts to transform the standard deviation from 1 to  $\omega_0$ , which is analogous to altering  $\sigma$  directly in the Fourier feature network. It is apparent from Eqs. 2.17 and 2.18 that the input layer in the Siren network is essentially a special case of the Fourier feature mapping, in which only sinusoidal features are considered. Therefore, it should perhaps be expected that both network architectures possess similar capabilities in modulating spectral bias.

In fluid flow applications, high-frequency solution components correspond to flow features that occur on small length and time scales. Given the multi-scale nature of LV flow, the modulation of spectral bias is important to fully resolve the corresponding velocity field.

### Generalisation of Physics-Informed Neural Networks

Generalisation of an ML model refers to its applicability to data not seen in the training set. In the context of PINNs, generalisation can be interpreted in two ways: Either generalisation across unseen spatio-temporal locations for a single computational domain, or generalisation across differing computational domains (which can be defined by varying the geometry or boundary conditions, for example). PINNs are effective at generalisation under the former definition, where they have demonstrated exceptional performance in interpolation and, to a lesser degree, extrapolation tasks. However, they are incapable of generalisation under the latter definition, which we will define as *inter-domain generalisation* from this point onward. This is due to the formulation of PINNs, at least in the form introduced in [190], in that the network takes as input the spatio-temporal coordinates only. Therefore, a single PINN model is constructed such that it is only able to approxi-

mate solution fields across one configuration of a physical system, i.e. one set of governing physical laws, geometry and boundary conditions [122, 176].

There have been two main courses of action to address the issue of inter-domain generalisation. Firstly, acceleration methods can circumvent the problem by providing significant speed-ups with regards to training convergence. Techniques that have been applied here include transfer learning [235, 176, 56], which pre-trains the hidden layers of a shared PINN on a family of differential equations before fine-tuning for specific configurations.

The second approach to address inter-domain generalisation has been to re-formulate the architecture of PINNs to allow domain-specific information to be learnt. Examples here include the utilisation of PointNets [122, 186], in which the geometric features of a computational domain are learnt alongside solution approximation, and hypernetworks [14, 94], in which one network, the *hypernetwork*, learns to generate weights for another network, in this case the PINN, based on parameterisations of the physical system.

The ultimate goal for our present application is to develop a model that is capable of performing flow super-resolution across large cohorts of 4D-flow MRI data, in a reasonable time and with minimal manual input required. Therefore, in future work it would be highly beneficial to develop a PINN model configuration that has either a fast convergence rate, without the requirement for case-specific hyperparameter tuning, or some natural inter-domain generalisation capabilities. To this end, we will discuss future opportunities in this regard in the final chapter of this thesis.

### Multi-Component Loss Imbalance

NTK theory has been used to analyse and diagnose common failings during PINN training, beyond issues of spectral bias. When analysing the individual eigenvalues associated with both the boundary condition loss and the PDE residual loss, it is found in [228] that, in some cases, there is a large discrepancy in eigenvalue magnitude between the two individual kernel matrices. The consequence of this finding is a discrepancy in convergence between the two loss terms, where the term with the largest eigenvalue magnitude in its kernel matrix converges faster, to the detriment of the competing loss term. This results in the trained PINN satisfying one constraint whilst violating the other – an undesirable property that frequently results in the PINN converging to the trivial solution to the Navier-Stokes equations given by zero-values everywhere in the domain. By simply using a weighted sum approach, attaching manually-selected weights to each of the loss components, the authors

are able to balance the magnitudes of competing eigenvalues and reduce the relative  $L_2$  error by two orders of magnitude.

The importance of addressing loss imbalances is increased when noise is present in the measurement data. In this instance, not only is the magnitude of individual loss components a concern, but loss components may also produce conflicting gradients, where outliers present in the measurement data may not conform to underlying physics. This is observed empirically in chapter 4.

The issue of balancing distinct constraints is not unique to PINNs, and there is a well-established collection of literature in the broader field of multi-objective (MO) optimisation. Most work in this setting is framed in the context of identifying the Pareto front – a set of ‘trade-off’ optimal solutions given by different model configurations in the face of conflicting objectives [51]. To identify the Pareto front of a particular MO model, the weighted sum approach is used extensively [154], alongside more advanced methods such as evolutionary MO optimisation [51]. While similarities exist, a key distinction between general MO optimisation and deep learning with MO loss functions is the considerable difference in convergence time for a single model configuration. Deep learning models can take days to fully converge for a single hyperparameter configuration, and thus identifying the Pareto front using the weighted sum method and other approaches is typically infeasible. Further, PINN models specifically must be trained on a case-by-case basis, and the optimal configuration of loss weights is not common across cases. While fixed weights can be optimised using a hyperparameter ablation study, doing so alongside all other hyperparameters can significantly expand the size of the search space, leading to excessively large parameterisation studies. Therefore, there has been a motivation to develop dynamic loss weighting strategies in the MO deep learning literature, as with the PINN literature, which will be discussed in the following section.

#### 2.4.5 Key Methodological Contributions

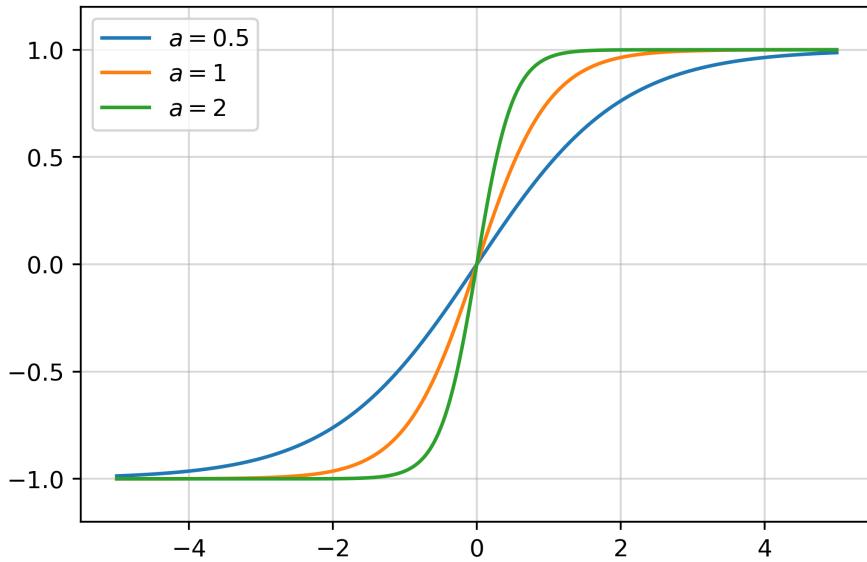
##### Novel Activation Functions

Given the unique nature of PINNs and other implicit neural representation models, there has been considerable effort afforded to the development of tailored activation functions. These include functions that are locally and globally adaptive [110, 111, 109], or that express derivatives to a high degree of accuracy [206].

Whilst the use of adaptive activation functions is not unique to PINNs [105], their first application in this context can be found in [110], where globally adaptive activation functions are applied to a set of benchmark PINN problems. Adaptivity is introduced to fixed activation functions through a trainable slope parameter,  $a$ . For example, the tanh activation function is reformulated as:

$$\tanh(ax) = \frac{e^{ax} - e^{-ax}}{e^{ax} + e^{-ax}}. \quad (2.19)$$

In Fig. 2.11 This method is considered *globally* adaptive since  $a$  is a single scalar that



**Figure 2.11:** Curves displaying the effect of varying the slope parameter  $a$  in the adaptive activation function  $\tanh(ax)$ .

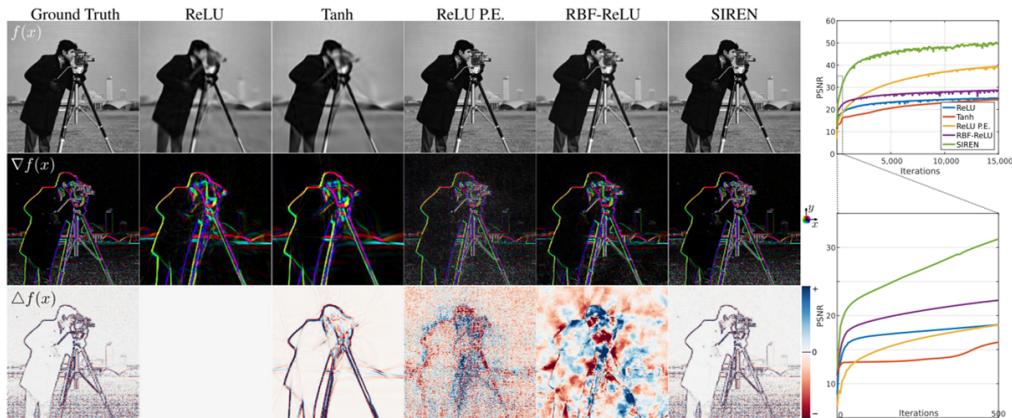
controls the slope of every activation function in the network. The introduction of a trainable slope parameter is shown to reduce training time across all studies investigated, with decreased  $L_2$  errors in each case. The authors investigate performance through frequency analysis, where it is shown that the adaptive activation functions somewhat mitigate the impact of spectral bias, where the authors use the F-principle [117] to demonstrate the capability of their models to capture high-frequency solution components in fewer training iterations.

*Locally* adaptive activation functions for PINNs are introduced by the same author in [109], in which the principle of globally adaptive activation functions is extended to allow for distinct layer-wise or neuron-wise scalable parameters. The authors demonstrate improved convergence and accuracy over both fixed activation functions and the globally adaptive

predecessor, whilst incurring minimal additional overhead. An additional slope recovery term is introduced to the loss function, which has the effect of increasing the magnitude of the activation slope quickly, accelerating overall training convergence.

Alongside adaptive activation functions, there has been interest in finding alternative activation functions that are designed specifically for use in implicit neural representation networks. Introduced in the previous section, the Siren network [206] uses periodic sine activation functions alongside a specific initialisation strategy.

The suitability of the sine activation function for implicit neural representation tasks arises from the fact that the derivative of a sine function is itself a phase-shifted sine function, and thus derivatives of a Siren network inherit the properties of the Siren. This ensures that derivatives that are calculated for loss constraints, such as those used in physics-based constraints in PINNs, are well behaved, which is an important feature for any model being used to represent signals. In Fig. 2.12, multiple implicit network architectures are compared alongside the Siren network for an image reconstruction task. Qualitatively, it is apparent that the Siren network captures the function derivatives most accurately, whilst also achieving the highest peak signal-to-noise-ratio.



**Figure 2.12:** An example case using the Siren implicit network to reconstruct an image,  $f(x)$  (top row), with comparison against other activation functions. The network is supervised by the image alone, however the first- and second-order derivatives are displayed (middle and lower respectively). The Siren network is able to represent both derivatives to a high degree of accuracy, whilst achieving significant improvements in peak signal-to-noise ratio (PSNR) [206].

Whilst there is limited discussion on the topic of spectral bias in [206], further publications have highlighted some properties of Siren networks that negate some of the associated effects [239].

## Loss Weighting Strategies

To mitigate the effects of gradient imbalances in MO applications, dynamic loss weighting strategies have been proposed. For general MO deep learning tasks, methods such as GradNorm [37], SoftAdapt [99], dynamic weight averaging (DWA) [145], uncertainty-based weighting schemes [124, 141] and gradient surgery [238] have been proposed. GradNorm [37] acts to balance gradient contributions between different loss terms by modifying the magnitudes of individual gradient values passed to the network weights during training, dependent on historical convergence rate. The update step for weight modification incurs an additional loss component, and thus an additional optimisation round, which, combined with storing and accessing gradient terms from previous iterations, makes this approach computationally expensive. Inspired by GradNorm, DWA [145] dynamically adjusts the loss term weights using only the rate of change of individual task losses, which improves efficiency by avoiding costly operations directly on the internal gradients of the model. The underlying algorithm of SoftAdapt [99] is similar to DWA, with both methods updating the weights based on the rate of change of loss components, however, the equation for updating the weights in SoftAdapt is formed as a Softmax function. In [124], a weighting scheme is proposed that utilises information on the homoscedastic (or task-dependent) uncertainty of each learning task, followed by a revised (and generally improved) formulation in [141].

A variety of PINN-specific dynamic weighting strategies have been proposed, where there is some methodological crossover with previous methods for general MO deep learning applications. Returning to [228], the authors not only analyse PINN performance across a range of manually-selected weights, but also introduce a dynamic weighting scheme derived from eigenvalues of the individual kernel matrices. The same authors also introduce a ‘learning rate annealing’ (LRA) method [225], which updates weights based on the ratio between gradient contribution magnitudes. Given a loss function containing a PDE loss term,  $\mathcal{L}_{PDE}$ , and data loss term,  $\mathcal{L}_{data}$ , the total loss at iteration  $k$  is given by:

$$\mathcal{L} = \mathcal{L}_{PDE} + \alpha^k \mathcal{L}_{data}, \quad (2.20)$$

for dynamic loss weight  $\alpha_k$ , where:

$$\mathcal{L}_{PDE} = \frac{1}{N_p} \sum_{i=0}^{N_p} |F(u(x_i), x_i)|^2, \quad (2.21)$$

$$\mathcal{L}_{data} = \frac{1}{N_d} \sum_{i=0}^{N_d} |u(x_i) - \bar{u}_i|^2. \quad (2.22)$$

The weighting term  $\alpha^k$  is calculated in moving-average form as:

$$\alpha^k = (1 - \lambda)\alpha^{k-1} + \lambda\hat{\alpha}^k, \quad (2.23)$$

for constant  $\lambda$ , with a recommended value of 0.1 suggested in [225], previous weight  $\alpha^{k-1}$  and weight update term:

$$\hat{\alpha}^k = \frac{\max_{\theta} \{|\nabla_{\theta} \mathcal{L}_{PDE}|\}}{|\nabla_{\theta} \alpha^{k-1} \mathcal{L}_{data}|}, \quad (2.24)$$

where  $\max_{\theta} \{|\nabla_{\theta} \mathcal{L}_i|\}$  and  $|\nabla_{\theta} \mathcal{L}_i|$  are the maximum and mean, respectively, of the gradient of loss component  $i$  w.r.t. the network weights,  $\theta$ . There are similarities between this approach and GradNorm [36] in that gradient information is required to perform the weight updates, and therefore an additional back-propagation step is required. However, GradNorm requires an additional loss constraint and optimisation step to update the weights, whereas LRA simply uses gradient statistics [158]. In [22], the LRA, GradNorm and SoftAdapt schemes are compared for both forward and inverse PINN problems. The authors also introduce another dynamic loss weighting scheme titled ReLoBRaLo (relative loss balancing with random lookback), which incorporates some of the most beneficial features from the previous approaches. As with SoftAdapt, the ReLoBRaLo operates on loss statistics instead of gradient information, and thus is less computationally expensive than GradNorm or LRA. Of the four methods compared, the LRA and ReLoBRaLo techniques generally give the best performance for forward and inverse approximation across the examples considered, with the ReLoBRaLo approach achieving slightly lower train and validation errors. Further, ReLoBRaLo is less computationally intensive than LRA, with a lower standard deviation of error across different training runs. Aside from generally poor performance, GradNorm and, to some degree, SoftAdapt are inadequate for accurate estimation of inverse parameters, and thus not recommended for use with PINNs in their current formulation. The LRA method introduced in [225] is modified slightly in [116], specifically for use in fluid flow problems. Equation 2.24 is replaced with:

$$\hat{\alpha}^k = \frac{|\nabla_{\theta} \mathcal{L}_{PDE}|}{|\nabla_{\theta} \mathcal{L}_{data}|}. \quad (2.25)$$

Following this, equation 2.23 is calculated as before. Note that the  $\max_{\theta}$  operators in the numerators are replaced by the mean value, and the contributions from previous weights  $\alpha^k$  and  $\beta^k$  in the gradient calculations in the denominators are removed. Results therein indicate minimal differences in performance between the formulation given in equations 2.24 and 2.25.

An entirely distinct approach is introduced in [158]. Inspired by advances in deep learning architectures, specifically attention mechanisms, for natural language processing [221], the authors propose a self-adaptive soft attention mechanism for PINNs, titled SA-PINNs. Expanding on the idea of dynamic weighting of distinct loss terms, the SA-PINN model attaches a trainable weight to each individual training sample. Doing so allows the model to not only balance contributions between loss terms, but also between training samples, allowing for greater ‘attention’ on samples that incur higher losses. This typically translates to larger weights attached to samples in regions with higher solution gradients, such as in the immediate vicinity of a boundary. As with loss weighting, increasing the magnitude of weights attached to the samples in the SA-PINN scheme improves the rate of convergence at these points, while decreasing the rate of convergence at points with relatively lower weights. Assuming a two-component loss function as in the previous paragraph, using the mean-squared error loss function, the SA-PINN is formulated as:

$$\mathcal{L} = \mathcal{L}_{PDE} + \mathcal{L}_{data}, \quad (2.26)$$

where

$$\mathcal{L}_{PDE} = \frac{1}{N_p} \sum_{i=0}^{N_p} m(\lambda_p^k) |F(u(x_i), x_i)|^2, \quad (2.27)$$

$$\mathcal{L}_{data} = \frac{1}{N_d} \sum_{i=0}^{N_d} m(\lambda_d^k) |u(x_i) - \bar{u}_i|^2. \quad (2.28)$$

Here,  $N_p$  and  $N_d$  correspond to the number of collocation, data and wall training points respectively, where vectors  $\lambda_p$  and  $\lambda_d$  correspond to the self-adaptive weight vectors, with each value  $\lambda_i^k$  within  $\lambda_i = \{\lambda_i^1, \dots, \lambda_i^{N_i}\}$  corresponding to a single weight. The mask function,  $m(\lambda_i)$ , is defined to be a non-negative, differentiable and strictly increasing function of  $\lambda_i$ , with polynomial and logistic maps suggested in [158]. Whilst the network weights,  $\mathbf{w}$ , are minimised during training, the trainable weight vectors  $\lambda_p$  and  $\lambda_d$  are instead *maximised*, transforming the optimisation problem from strictly minimisation to a min-max formulation, given by:

$$\min_{\mathbf{w}} \max_{\lambda_p, \lambda_d} = \mathcal{L}(\mathbf{w}, \lambda_p, \lambda_d). \quad (2.29)$$

The maximisation of  $\lambda_p$  and  $\lambda_d$  is carried out using distinct optimisers for each loss term, which allows for some additional tuning to *a priori* improve the balance between loss components, such as modifying the starting weights or the learning rates of respective optimisers. These can be tuned as part of a hyperparameter ablation study, as we will

show later in this thesis. Whilst this provides some user-defined flexibility, it is not strictly necessary for effective function of the model, since the weights  $\lambda_p$  and  $\lambda_d$  already act to balance gradient contributions based on both spatio-temporal location *and* associated loss terms.

An extension to the SA-PINNs formulation is proposed in [241], titled differentiable adversarial SA-PINNs (DASA-PINNs). The DASA-PINN approach shares the same underlying logic of SA-PINNs, which is that dynamically adjusting only loss term weights ignores the interplay between individual training samples, and thus neglects the heterogeneity of the solution fields across the spatio-temporal domain. An additional sub-network is introduced to predict the attention weights, which is again maximised using a distinct optimiser, although much of the approach is closely related to SA-PINNs. The authors provide a comparison between their method, DASA-PINNs, and SA-PINNs, finding consistently improved performance over the latter on a small collection of problems.

A residual-based attention (RBA) scheme is introduced in [4], providing similar functionality to both SA-PINNs and DASA-PINNs by applying weights to each training sample individually (but applied only to the collocation point cloud here, giving single weight vector  $\lambda$ ). A key distinction in the RBA scheme is the use of PDE residual values at the corresponding collocation points to compute the attention weights, instead of learnable parameters, reducing the computational overhead by removing additional backpropagation steps. Further, the incorporation of the PDE residual values directly relates the attention weights to the local performance of the network in capturing the physics. At iteration  $k$ , the attention weight  $i$  is updated as:

$$\lambda_i^{k+1} = \gamma \lambda_i^k + \eta^* \frac{|F(u(x_i), x_i)|}{\max_i (|F(u(x_i), x_i)|)}, \quad i \in \{0, 1, \dots, N_p\}, \quad (2.30)$$

for decay parameter  $\gamma$  and RBA learning rate  $\eta^*$  (distinct from the network learning rate,  $\eta$ ).

### Adaptive Collocation Cloud Refinement

Three of the approaches discussed above, SA-PINNs [158], DASA-PINNs [241] and RBA [4], can increase attention in regions of the solution field that are challenging to learn by increasing multiplicative weights attached to samples in those regions. However, the usefulness of these approaches diminishes if the density of the collocation point cloud, at which the physics-based regularisation is applied, is too low. For instance, in areas

containing steep solution gradients, if the collocation point cloud is too sparse to obtain multiple samples within the high-gradient region, sharp solution information will not be accurately reconstructed, regardless of increased attention. Therefore, the collocation point cloud used for PINN applications is typically required to be dense [47], particularly when approximating complex solution fields with high-frequency components. However, when using PINNs to approximate three-dimensional, time-dependent solution fields, to capture all solution scales to an acceptable degree of accuracy can require a collocation point cloud with a prohibitively large cardinality, incurring significantly increased computational costs. Drawing inspiration from adaptive meshing techniques for numerical [18] and meshless methods [177], adaptive collocation cloud refinement techniques have been proposed for use with PINN methods. The underlying assumption is that increased density is required in regions where the solution is challenging to approximate, such as areas containing high solution gradients, whereas regions of relatively simple solution complexity do not require as many points to accurately reconstruct.

In [147], the residual-based adaptive refinement (RAR) method is introduced – the first method proposed of its kind in the context of PINNs. Starting from a coarse collocation point cloud, this approach iteratively increases the density of collocation points in regions where the PDE residual loss component is largest, which typically occurs near steep solution gradients. As such, the implementation of this approach can significantly enhance the performance of a PINN when used to approximate solution fields that are governed by PDEs exhibiting such behaviour [231]. Monte Carlo integration is used to probe the PDE residual values at each point in some randomly sampled dense set  $\mathcal{S}$ , where the  $m$  points with the greatest residual value are added to the collocation point cloud (here  $m$  is a user-defined scalar parameter). This simplified selection criterion can result in regions with smaller residual values being disregarded during training [231], and thus modifications to this strategy have been investigated. An alternative *importance sampling* approach is proposed in [165], where, at each update step, an entirely new set of collocation points is selected based on a probability density function (PDF). The theoretically ideal PDF is based on the  $L_2$ -norm of the PDE loss gradient at each new point, however, this is computationally expensive, requiring a backpropagation pass for each new collocation point. Therefore, the PDF is instead constructed from the actual PDE loss values, which are correlated to the gradient norm values [165], and thus the PDF is proportional to PDE residual value at a candidate collocation point. Using the PDF to draw samples instead of a strict in/out criterion defined by a scalar value, as with RAR, helps to nullify the issue of neglecting regions with low PDE residual values. Further, since the collocation point

cloud is resampled each time, the cardinality of the collocation point set does not increase to an unmanageable size, as with RAR.

The issue of growing collocation point set size for the RAR method is addressed in [240] and [50]. In [240], the authors introduce residual-based adaptive exchange (RAE). This is similar to RAR in that the  $m$  points with the highest PDE residual values are added to the collocation point set at each update, however, the RAE method additionally removes the  $m$  points with the *lowest* PDE residual values, maintaining the same collocation point cloud density throughout training. When compared with RAR, it is shown that this approach does not reduce the accuracy of the PINN model, but reduces the computational overhead by  $\sim 25\%$ . In [50], an evolutionary sampling (Evo) technique is proposed, which, as with RAE, removes points with the lowest PDE residual values, based on a threshold value. However, instead of directly generating new points where the PDE residuals are highest, additional points are simply sampled from a uniform distribution instead. This approach is more closely related to evolutionary algorithms in optimisation, in that the most important aspect is the ‘survival’ of samples, rather than a targeted increase of density in problem regions. Of course, redistributing collocation points in such a manner will still, over time, increase the density of points in regions with high PDE residuals, but it may not be the most efficient strategy for doing so. Although not discussed in the related literature, it would be an intuitive assumption that RAE, Evo and other similar methods exacerbate the issue of ‘neglect’ in regions with low PDE residual values, particularly in the latter stages of training, where the redistribution may result in regions with very sparse collocation point cloud density.

In [231], the RAR [147] and importance sampling [165] strategies are compared, and two new approaches are introduced, namely residual-based adaptive distribution (RAD) and residual-based adaptive refinement with distribution (RAR-D). The RAD approach draws inspiration from the importance sampling technique, however the sampling PDF is modified by raising the PDE loss value contained to a power  $k \geq 0$ , and introducing an additive constant  $c \geq 0$ , which generalises the importance sampling strategy (which itself is equivalent to RAD with  $k = 1$  and  $c = 0$ ). The constant  $k$  acts to increase the concentration of new points where the highest PDE loss values are located, whereas  $c$  acts to uniformly spread points across the domain. The RAR-D method combines functionality from both the importance sampling and RAR techniques, in that a smaller set of points are sampled from the same PDF as with RAD, however, these are added to the existing collocation point cloud instead of replacing it. Across a range of problems,

the RAD formulation is shown to consistently outperform non-adaptive and competing adaptive sampling strategies [231], where the inclusion of parameters  $k$  and  $c$  allows for problem-specific tuning, dependent on the governing PDE and boundary conditions.

In [212], it is argued that excessive attention on local optimisation of the solution fields may result in convergence to local minima at the expense of global convergence. To alleviate this issue, a cosine annealing strategy is deployed, which periodically alters the ratio of uniformly-sampled points to adaptively sampled points. The PDF for the adaptive sampling is either based directly on the PDE loss constraint (ADAPTIVE-R) or the spatial gradient of the PDE loss constraint (ADAPTIVE-G).

### Improving Inter-Case Generalisation

As discussed earlier in this section, inter-case generalisation is a present problem for PINNs and has hampered their applicability to real-life problems. In this context, transfer learning has been applied successfully to significantly accelerate the training convergence of PINNs when deployed across similar domains [184, 56, 176, 235]. The main theory behind this approach is that the majority of trainable parameters learnt in a model are not entirely specific to one individual task. Therefore, it is possible to use a shared weight configuration for an array of distinct, but similar, tasks, training only a small number of task-specific weights (typically within the output layer) when applying the model to individual tasks. An example in the domain of medical imaging is seen in [150], in which an image classification model is pre-trained using ImageNet [44], a dataset of over 1.2 million natural images, before fine-tuning to be utilised for classification and detection of Alzheimer’s disease from brain MRI. Although a model trained entirely on brain MR images would likely perform better, such a large database of labelled images does not exist. However, the convolutional filters for image classification models, particularly those closer to the input layer, learn to extract generic features such as edges, whereas the domain-specific features are typically only extracted in the final few layers of the model [150]. Therefore, regardless of the image type used for training, the majority of the learned weights in such a model will be similar. It is therefore more beneficial to utilise a far larger, but generalised, dataset for pre-training, before fine-tuning the final few layers on a smaller, domain-specific dataset.

This concept holds true for PINN applications, where we may want to apply a model to distinct computational domains that share close similarities, such as similar domain geometries and underlying physical laws. This is certainly the case for the present application, where the variance across LV morphologies is small (relative to the wider context

of fluid modelling applications), the governing equations and boundary conditions remain fixed and there is limited variability in the temporal scales (affected only by heart rate). In [56], transfer learning is applied to a selection of ordinary differential equations (ODEs) and PDEs. The hidden layers of a PINN model are pre-trained across a family of related differential equations, with domain-specific fine-tuning of a single linear output layer performed when applying the model to specific cases in inference mode. Training of this single linear layer in inference mode for an unseen parameterisation of the Poisson equation is completed in only 33.2 seconds, achieving a test accuracy of  $3.6 \times 10^{-5}$ . This approach is elaborated on in [176], in which a multi-head output architecture is deployed to facilitate knowledge transfer of distinct initial conditions or potential field parameterisations for a given non-linear ODE. Further applications of transfer learning with PINN models can be found in [184, 235].

## Chapter 3

# Super-Resolution of Cardiac 4D-Flow MRI: An Idealised *in Silico* Study

Much of the content in this chapter has been published in F. Shone, N. Ravikumar, T. Lassila, et al. “Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI”, In: *Information Processing in Medical Imaging (IPMI) 2023, Springer Lecture Notes in Computer Science, vol. 13939* (2023), DOI: 10.1007/978-3-031-34048-2\_39 [205].

### 3.1 Introduction

In this chapter, we introduce the PINN model to be used for super-resolution of left ventricular 4D-flow MRI, and demonstrate its capabilities using synthetic, idealised LV cases in 2D and 3D. We investigate the robustness of the PINN model across varying spatio-temporal resolutions and SNRs, which is of importance in real studies, and provide a comparison between competing methods used to mitigate the effects of spectral bias and gradient imbalances during training.

As discussed in chapter 2, corruption of 4D-flow MRI can be characterised by low spatial resolution, low temporal resolution, background noise (influencing SNR and thus VNR) and structured noise. The presence of such effects has hampered the application of 4D-flow MRI in clinic due to uncertainty in the measured velocity field. Moreover, clinically-

relevant quantities like pressure, vorticity and WSS are not directly measured, and are instead obtained through calculation of the spatial derivatives of the velocity field, which are particularly susceptible to effects of corruption [42]. In this chapter, we consider the impact of low spatio-temporal resolution and background noise, whilst assuming perfect reconstruction of the endocardium and its motion. We do not consider the effects of structured noise, as it is challenging to accurately synthesise and is generally mitigated by other approaches during acquisition [183].

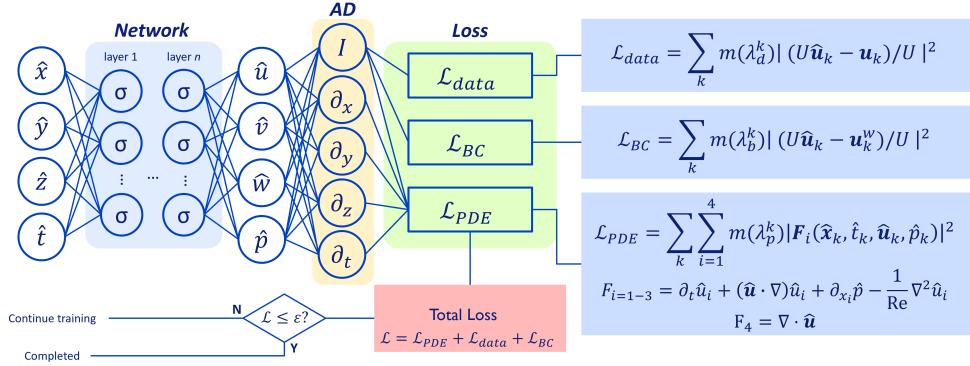
Spatio-temporal super-resolution and de-noising of 4D-flow MRI using ML has previously been addressed in vascular flows [75, 74, 72, 203, 195]. Fixed spatial upsampling factors of 2 [75, 74] and 4 [203] have been achieved, utilising residual networks, alongside arbitrary spatial upsampling rates [72, 195] using implicit neural representation networks (to which PINNs belong). De-noising of background noise has been demonstrated for SNR values ranging from 2 to 20, and temporal upsampling of arbitrary rates has been demonstrated [72, 195] using implicit neural representation networks.

In this work, we utilise a PINN architecture [190] to achieve spatio-temporal super-resolution and de-noising of ventricular 4D-flow MRI. We opt to use this approach for a few reasons. Firstly, the physics regularisation used in the loss function ensures that the calculated fields and quantities derived from them conform to the underlying physics of the system, providing explainability. Secondly, PINNs do not require paired low- and high-resolution training data, improving their applicability to real 4D-flow MRI datasets where paired data is challenging to obtain. PINNs have no natural inter-domain generalisation capabilities, and thus must be re-trained for each case, placing them at a disadvantage against the approaches introduced in [75, 203]. However, it is the author’s opinion that, through the application of appropriate acceleration methods, PINNs may be able to provide state-of-the-art results in this domain.

## 3.2 Methodology

In this section, we describe the PINN model used and the procedures followed to generate synthetic 4D-flow MRI data sets. PINNs belong to a class of ML methods called implicit neural representation networks, in which a fully-connected neural network is deployed to produce a continuous approximation to a function on a given spatio-temporal domain. In this context, the PINN model learns to approximate continuous velocity and pressure fields at the spatio-temporal coordinates defined for the given LV geometries. These predictions

are regularised by the available low-fidelity 4D-flow MRI data, alongside physics-based constraints. Once trained, the combination of these loss components produces a network capable of making predictions that both correspond to the existing data and obey the underlying physics.



**Figure 3.1:** The physics-informed neural network (PINN) used in this thesis approximates a mapping from spatio-temporal coordinates to the corresponding velocity and pressure fields in the flow domain. The network is fully connected, utilising adaptive Siren activation functions [206]. The predictions are constrained by a multi-component loss function, consisting of a data loss term to minimise velocity prediction-measurement error, a boundary loss term to enforce the no-slip condition on the domain walls, and a physics loss term that enforces the Navier-Stokes equations in the flow domain. A self-adaptive weighting strategy is used in the loss function to balance the gradient contributions of each sample point during training [158]. Automatic differentiation is used to generate the relevant derivatives for the physics loss term.

### 3.2.1 PINN Model

The PINN approach used in this thesis deploys a fully connected deep neural network to approximate a function mapping spatio-temporal coordinates within the domain to corresponding velocity and pressure fields. The predicted solutions are constrained through a multi-component loss function using velocity measurements and known physical laws, namely the Navier-Stokes equations and no-slip boundary condition. Unlike purely data-driven approaches, the presence of these physics-based constraints in the loss function restricts the space of possible solutions by penalising non-physical predictions, facilitating efficient training with sparse data. The schematic for the PINN (shown here for 3D applications, using a self-adaptive loss weighting strategy [158]) can be found in Fig. 3.1.

The loss function consists of three components: *Data loss*,  $\mathcal{L}_{data}$ , in which the error between velocity predictions and 4D-flow MRI measurement data is minimised, *PDE (partial differential equation) loss*,  $\mathcal{L}_{PDE}$ , containing residuals of the Navier-Stokes equations, and *BC (boundary condition) loss*,  $\mathcal{L}_{BC}$ , where predictions on the domain wall are constrained

to obey the no-slip condition. In all cases, mean-squared error (MSE) loss is used.

### 3.2.1.1 Governing Physical Laws

In large vessels and cardiac chambers, blood is considered to be a Newtonian, incompressible fluid. Therefore, the flow physics are governed by the incompressible Navier-Stokes equations, given by

$$\rho(\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u}) = -\nabla p + \mu \nabla^2 \mathbf{u} \quad (3.1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (3.2)$$

for velocity  $\mathbf{u}$ , density  $\rho$ , pressure  $p$  and dynamic viscosity  $\mu$ . For all cases discussed in this chapter, it is assumed that  $\rho = 1066 \text{ kg m}^{-3}$  and  $\mu = 0.0035 \text{ Pa s}$  [196].

The no-slip boundary condition, which is applied on the chamber walls, states that no fluid may penetrate or slip tangentially at a given boundary, and is thus formulated as:

$$\mathbf{u} = \mathbf{u}_w, \quad (3.3)$$

for fluid velocity  $\mathbf{u}$  and wall velocity  $\mathbf{u}_w$ . If such a wall is rigid, this condition reduces to  $\mathbf{u} = 0$ , enforcing zero flow. However, in our use case,  $\mathbf{u}_w$  is non-zero due to the myocardial deformation across the cardiac cycle.

### 3.2.1.2 Nondimensionalisation and standardisation

The importance of nondimensionalising input and output variables when working with PINNs, as with neural networks in general, has been highlighted [128, 72, 195]. Network training is typically ineffective when the scales of input and output variables are significantly different [87], as is commonly the case in fluid dynamics problems. To address this issue, we nondimensionalise each input and output variable using characteristic scales to ensure each is of an approximate order of magnitude of 1. Whilst physiology-based characteristic scales for velocity and length are used in [128], we have found that simply using the maximum velocity magnitude,  $U$ , and maximum distance,  $L$ , encountered in the training set provides adequate results. The pressure and time scales are then selected to be  $\rho U^2$  and  $L/U$  respectively, for fluid density  $\rho$ . This dynamic pressure scale is appropriate for ventricular flow as, for the most important phases of the cardiac cycle, inertial flow

effects dominate. Thus, our dimensionless variables are given by:

$$\mathbf{x}^* = \frac{\mathbf{x}}{L} \quad (3.4)$$

$$t^* = \frac{Ut}{L} \quad (3.5)$$

$$\hat{\mathbf{u}} = \frac{\mathbf{u}}{U} \quad (3.6)$$

$$\hat{p} = \frac{p}{\rho U^2}. \quad (3.7)$$

To conform with best practices for ML applications [136], we follow this by standardising our network inputs to have zero mean and unit variance, as per [128], giving dimensionless and standardised spatio-temporal coordinates:

$$\hat{x}_i = \frac{x_i^* - \mu_{x_i^*}}{\sigma_{x_i^*}} \quad (3.8)$$

$$\hat{t} = \frac{t^* - \mu_{t^*}}{\sigma_{t^*}}, \quad (3.9)$$

for  $i \in \{1, 2, 3\}$ , means  $\mu$  and standard deviations  $\sigma$ . The dimensionless and standardised variables are used as inputs and outputs in the network and so in inference mode we must nondimensionalise and standardise the network inputs, and multiply the network outputs by their respective scales.

Using the characteristic scales, we are able to convert the Navier-Stokes equations to a dimensionless form, given by:

$$\mathbf{F}_{1-3} = \frac{1}{\sigma_{t^*}} \partial_{\hat{t}} \hat{\mathbf{u}} + (\hat{\mathbf{u}} \cdot \hat{\nabla}) \hat{\mathbf{u}} + \hat{\nabla} \hat{p} - \frac{1}{\text{Re}} \hat{\nabla}^2 \hat{\mathbf{u}} = 0 \quad (3.10)$$

$$F_4 = \hat{\nabla} \cdot \hat{\mathbf{u}} = 0, \quad (3.11)$$

where

$$\hat{\nabla} = \left( \frac{1}{\sigma_{x_1^*}} \partial_{\hat{x}_1}, \frac{1}{\sigma_{x_2^*}} \partial_{\hat{x}_2}, \frac{1}{\sigma_{x_3^*}} \partial_{\hat{x}_3} \right)^T, \quad (3.12)$$

$$\hat{\nabla}^2 = \frac{1}{\sigma_{x_1^*}^2} \partial_{\hat{x}_1}^2 + \frac{1}{\sigma_{x_2^*}^2} \partial_{\hat{x}_2}^2 + \frac{1}{\sigma_{x_3^*}^2} \partial_{\hat{x}_3}^2, \quad (3.13)$$

for Reynolds number  $\text{Re} = \rho U L / \mu$ . Note that we must account for the multiplicative effects of the standard deviations used in standardising our input variables.

### 3.2.1.3 Loss Function

The total loss function,  $\mathcal{L}$ , is given (in unweighted form) by:

$$\mathcal{L} = \mathcal{L}_{PDE} + \mathcal{L}_{BC} + \mathcal{L}_{data}. \quad (3.14)$$

Each loss term utilises a distinct set of coordinates, since in the present application there is sparsity in available measurement data.

Let the spatio-temporal LV flow domain be denoted  $\Omega$ , enclosed by boundary  $\partial\Omega$ .  $\mathcal{L}_{data}$  is evaluated at the voxel centres of the 4D-flow MRI data, masked to exist only within  $\Omega$ , and considered coarse in space and time owing to its low spatio-temporal resolution. Conversely,  $\mathcal{L}_{PDE}$  is evaluated at a dense set of *collocation* points across  $\Omega$  in space and time, which are distinct to the 4D-flow MRI voxel coordinates. The point cloud for  $\mathcal{L}_{BC}$  is sampled densely on  $\partial\Omega$ , where nodal correspondence is required between phases to calculate the associated wall velocity. As previously mentioned, all loss components assume the MSE form.

#### PDE Loss Term

The available 4D-flow MRI data in cardiac studies is sparse and corrupted by noise, precluding the use of purely data-driven ML approaches in this setting (assuming no availability of high-resolution data labels). However, we have knowledge of the physical system at hand, in the form of the governing PDEs. Therefore, it is possible to constrain the outputs of an ML model using such information, which heavily restricts the space of possible solutions.

To this end, we ensure solutions fields approximated by our PINN model conform to the underlying physics using a PDE loss term that minimises residuals of the Navier-Stokes equations, as defined in Eq. 3.10. In vanilla form, the PDE loss component is given by:

$$\mathcal{L}_{PDE} = \frac{1}{N_p} \sum_{i=1}^4 \sum_{k=1}^{N_p} \|F_i(\hat{\mathbf{x}}_k, \hat{t}_k, \hat{\mathbf{u}}_k, \hat{p}_k)\|^2, \quad (3.15)$$

$$(\hat{\mathbf{x}}_k, \hat{t}_k) \in \hat{\Omega}, \quad (3.16)$$

for collocation set of cardinality  $N_p$ , and dimensionless and standardised domain  $\hat{\Omega}$ .

### Data Loss Term

The data loss term assumes the form of a typical supervised loss term, in which (mean-squared) errors between model outputs and available data are minimised. The data loss term, in vanilla form, is given by:

$$\mathcal{L}_{data} = \frac{1}{N_d} \sum_{k=1}^{N_d} \|\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k/U\|^2, \quad (3.17)$$

$$(\hat{\mathbf{x}}_k, \hat{t}_k) \in \hat{\Omega}, \quad (3.18)$$

for total number of 4D-flow MRI voxels in space and time  $N_d$  and associated velocity data  $\mathbf{u}_k$ .

In [72], the data loss term is reformulated into Cartesian images, which are able to negate effects of aliasing (or phase-wrapping), which occurs when velocity values in the domain exceed the pre-defined  $v_{enc}$ . However, in practice, the  $v_{enc}$  is typically selected prior to acquisition to exceed the maximum expected velocity, and thus such modifications to the loss function are generally not necessary. Further, the authors introduce a voxel averaging step, in which the velocity in each voxel is spatio-temporally averaged using Gaussian quadrature. It was found in this work that this additional step had no impact on the accuracy of results, whilst incurring additional computational overhead. It was also noted that the synthetic data used in [72] did not account for such temporal averaging effects, although this component may have been beneficial in their phantom 4D-flow MRI case.

### BC Loss Term

When using PINNs in a setting for which we have sparse and noisy measurement data, it is sensible to utilise as many sources of *a priori* domain knowledge as possible to constrain model outputs. In real cardiac 4D-flow MRI studies, we typically do not have access to accurate inflow and outflow boundary conditions, but we know that the no-slip boundary condition must apply to all interior surfaces in the domain, which we are able to segment and track in time. Therefore, we are able to apply this knowledge as an additional constraint in our loss function, and thus we enforce the no-slip boundary condition on the endocardium. In all studies in this thesis, we apply the boundary condition on the endocardium only, with the valve regions and much of the LV base excluded. This is done to ensure the model can operate using a partial boundary condition, as is typically the case with real 4D-flow MRI data sets, where it is challenging to accurately reconstruct the basal and valvular region of the LV using structural cine-MRI.

In real 4D-flow MRI studies, using a registration technique with nodal correspondence allows temporal interpolation of the endocardium, from which the wall velocity,  $\mathbf{u}_w$ , can be calculated using either a numerical central difference scheme in time, or by taking the first derivative of the interpolation spline. In the synthetic case discussed in this chapter, however, we assume that we have complete knowledge of the boundary motion, and thus the CFD data at the boundaries is used to enforce the no-slip condition.

In vanilla form, the BC loss component is given by:

$$\mathcal{L}_{BC} = \frac{1}{N_b} \sum_{k=1}^{N_b} \|\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k^w/U\|^2, \quad (3.19)$$

$$(\hat{\mathbf{x}}_k, \hat{t}_k) \in \partial\hat{\Omega}, \quad (3.20)$$

for BC point set of cardinality  $N_b$ , wall velocity data point  $\mathbf{u}_k^w$  and dimensionless and standardised boundary  $\partial\hat{\Omega}$ .

Note that the velocity scaling,  $U$ , is included in Eqs. 3.19 and 3.17. Our model output,  $\hat{\mathbf{u}}$ , is dimensionless, and therefore both  $\mathbf{u}_k$  (interior velocity data) and  $\mathbf{u}_k^w$  (wall velocity data) are divided by  $U$  to match. Division of the velocity data by  $U$ , as opposed to multiplication of the model output, is done to ensure that  $\mathcal{L}_{BC}$  and  $\mathcal{L}_{data}$  are both dimensionless, as with  $\mathcal{L}_{PDE}$ , removing any dependency of the loss terms on  $U$ .

Additionally, since pressure data and associated boundary conditions are not readily available in real 4D-flow MRI studies, no pressure constraints beyond the derivative terms that appear in the Navier-Stokes equations are enforced. However, the PDE and BC loss constraints are sufficient to allow the PINN to fully reconstruct to the relative pressure field in  $\Omega$ .

## Loss Weighting

Much attention in the PINN literature has been focused on the development of loss weighting schemes to balance gradient contributions of each loss component during training, where unweighted approaches commonly result in strong imbalances in the rate of descent of individual loss components. This imbalance tends to produce sub-optimal solutions, where one or more loss constraints may not be satisfied to an acceptable level. This issue is further exacerbated when using noisy measurement data, as the loss components may produce conflicting gradients.

Fixed, non-unity weights have been shown to address these issues [227, 226, 72], however the selection of such weights is heavily dependent on the problem at hand. Therefore, dynamic weighting schemes have received much attention. In our previous work [205], the learning rate (LR) annealing loss weighting technique, introduced in [116], was utilised. However, results published in [205] were found to be heavily influenced by noise corruption at higher spatial downsampling rates, with the PINN model seemingly over-fitting to noise in the training data. Therefore, alternative strategies have since been explored to mitigate these issues. In this chapter, we provide a comparison between three loss weighting approaches, namely fixed non-unity weights, the LR annealing scheme introduced in [116], and the self-adaptive soft attention (SA-PINN) scheme introduced in [158]. Although more dynamic weighting strategies have been proposed in the PINN literature [22, 228, 226, 241], there is significant overlap between competing methods, where the three schemes compared here represent three completely distinct approaches, which led to their inclusion.

For both the fixed weight configuration and the LR annealing scheme, the total loss function is given by:

$$\mathcal{L} = \mathcal{L}_{PDE} + \alpha \mathcal{L}_{data} + \beta \mathcal{L}_{BC}, \quad (3.21)$$

where  $\alpha$  and  $\beta$  are weights associated with the data loss and BC loss terms respectively. These weights can be fixed or dynamically updated. For the LR annealing scheme, at iteration  $k$  we calculate:

$$\hat{\alpha}^{k+1} = \frac{\overline{|\nabla_{\theta} \mathcal{L}_{PDE}|}}{\overline{|\nabla_{\theta} \mathcal{L}_{data}|}}, \quad \hat{\beta}^{k+1} = \frac{\overline{|\nabla_{\theta} \mathcal{L}_{PDE}|}}{\overline{|\nabla_{\theta} \mathcal{L}_{BC}|}}, \quad (3.22)$$

where  $\overline{|\nabla_{\theta} \mathcal{L}_i|}$  is the mean of the gradient of loss component  $i$  w.r.t. the network weights,  $\theta$ . Following this, the weights for iteration  $k + 1$  are updated as:

$$\alpha^{k+1} = (1 - \lambda) \alpha^k + \lambda \hat{\alpha}^{k+1} \quad (3.23)$$

$$\beta^{k+1} = (1 - \lambda) \beta^k + \lambda \hat{\beta}^{k+1}, \quad (3.24)$$

for constant  $\lambda$ , chosen to be 0.1 as per [115], and previous weights  $\alpha^k$  and  $\beta^k$ .

Conversely, for the SA-PINN model [158], the total loss to be minimised during training is given by

$$\mathcal{L} = \mathcal{L}_{PDE}^{SA} + \mathcal{L}_{data}^{SA} + \mathcal{L}_{BC}^{SA}, \quad (3.25)$$

where

$$\mathcal{L}_{PDE}^{SA} = \frac{1}{N_p} \sum_{i=1}^4 \sum_{k=1}^{N_p} m(\lambda_p^k) \|F_i(\hat{\mathbf{x}}_k, \hat{t}_k, \hat{\mathbf{u}}_k, \hat{p}_k)\|^2 \quad (3.26)$$

$$\mathcal{L}_{data}^{SA} = \frac{1}{N_d} \sum_{k=1}^{N_d} m(\lambda_d^k) \|\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k/U\|^2 \quad (3.27)$$

$$\mathcal{L}_{BC}^{SA} = \frac{1}{N_b} \sum_{k=1}^{N_b} m(\lambda_b^k) \|\hat{\mathbf{u}}(\hat{\mathbf{x}}_k, \hat{t}_k) - \mathbf{u}_k^w/U\|^2. \quad (3.28)$$

The vectors  $\boldsymbol{\lambda}_p$ ,  $\boldsymbol{\lambda}_d$  and  $\boldsymbol{\lambda}_b$  correspond to the self-adaptive weight vectors, which are non-negative, trainable values to be updated during network training. Each value  $\lambda_i^k$  within  $\boldsymbol{\lambda}_i = \{\lambda_i^1, \dots, \lambda_i^{N_i}\}$  corresponds to a single weight, which is applied specifically to training point  $k$ . During training, this allows the network to balance gradient contributions from each loss component, as with the previous approaches, but also balances contributions based on local solution complexity. In practice, regions of the solution field that are difficult to learn or have a significant impact on global convergence (for instance, regions near a boundary in space or near an initial state in time) are given larger weights, increasing the gradient magnitude from these points. Finally, the mask function,  $m(\boldsymbol{\lambda}_i)$ , is defined to be a non-negative, differentiable and strictly increasing function of  $\boldsymbol{\lambda}_i$ . In [158], polynomial and logistic maps are proposed, however we simply use the identity mask,  $m(\boldsymbol{\lambda}_i) = \boldsymbol{\lambda}_i$ , in line with the experiments used in the original publication [158]. Whilst the network weights,  $\mathbf{w}$ , are minimised during training, the trainable weight vectors  $\boldsymbol{\lambda}_p$ ,  $\boldsymbol{\lambda}_d$  and  $\boldsymbol{\lambda}_b$  are instead *maximised*, using individual optimisers, transforming the optimisation problem from strictly minimisation to a min-max formulation, given by:

$$\min_{\mathbf{w}} \max_{\boldsymbol{\lambda}_p, \boldsymbol{\lambda}_d, \boldsymbol{\lambda}_b} = \mathcal{L}(\mathbf{w}, \boldsymbol{\lambda}_p, \boldsymbol{\lambda}_d, \boldsymbol{\lambda}_b). \quad (3.29)$$

We make a minor alteration to the scheme used in [158] by allocating individual learning rates for the physics, data and boundary condition SA weight optimisers, given respectively by  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ . It has been observed empirically that there is, generally, a negative correlation between the number of trainable parameters in a neural network and the optimal learning rate of the optimiser used. Therefore, it is reasonable to assume that such a correlation may exist for the self-adaptive loss weighting strategy. Given the large differences between  $N_p$ ,  $N_d$  and  $N_b$ , and thus the cardinality of  $\boldsymbol{\lambda}_p$ ,  $\boldsymbol{\lambda}_d$  and  $\boldsymbol{\lambda}_b$ , it was stipulated that a shared learning rate may not provide the best performance for each optimiser. Further, individual learning rates allow for accelerated learning of particular loss components, which may be beneficial under certain circumstances. This hypothesis

was supported by results obtained during optimisation studies in this thesis, with the optimal configurations in each case exhibiting  $\eta_d, \eta_b > \eta_p$ , where  $N_d < N_b < N_p$ .

### 3.2.1.4 Activation Function

As with loss weighting schemes, the selection of an appropriate activation function within the hidden layers of a PINN model has been shown to have a significant effect on the rate and stability of convergence during training [108, 111, 206]. This is primarily related to the issue of spectral bias, which describes the tendency of PINNs (and other deep learning models) to favour low-frequency solution components, particularly during the earlier stages of training. This results in vanilla PINN architectures not fully resolving high-frequency components until very late in the training cycle, if at all. This is of particular importance to this study, where ventricular flow is complex and occurs across a range of length and time scales.

The network architecture used in this chapter and throughout this thesis is a fully connected Siren network, in which each hidden layer uses a periodic sine activation function [206], except for the output layer which uses a linear activation. For hidden layer  $i$ , the activation is formulated as

$$\phi_i(x_i) = \sin(\mathbf{W}_i x_i + \mathbf{b}_i), \quad (3.30)$$

for input  $x_i$ , weights  $\mathbf{W}_i$  and biases  $\mathbf{b}_i$ . Weights in each hidden layer,  $w_i$ , are drawn from  $w_i \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$  at initialisation, as per [206]. Additionally, the input layer to the network is constructed as:

$$\phi_0(x_0) = \sin(\omega_0 \mathbf{W}_0 x_0 + \mathbf{b}_0), \quad (3.31)$$

for initialisation frequency  $\omega_0$ , and is instead initialised by drawing weights from  $w_0 \sim \mathcal{U}(-1, 1)$ .

Siren networks were proposed for implicit neural representation tasks for their natural signal representation properties and their ability to reconstruct high-order derivatives accurately [206]. However, in a similar fashion to the popular Fourier feature network [216], it has since been demonstrated that the initialisation frequency,  $\omega_0$ , can be used to modulate the effects of spectral bias, where the value of  $\omega_0$  dictates the dominant frequency modes of predicted solutions [239]. For this reason, Siren networks have previously been

applied to super-resolve 4D-flow MRI data [195], without the use of a PDE constraint, to good effect. Their ability to modulate spectral bias and thus resolve high-frequency solution components makes them an ideal candidate for approximating complex, multi-scale flow fields, such as those encountered in the cardiac chambers.

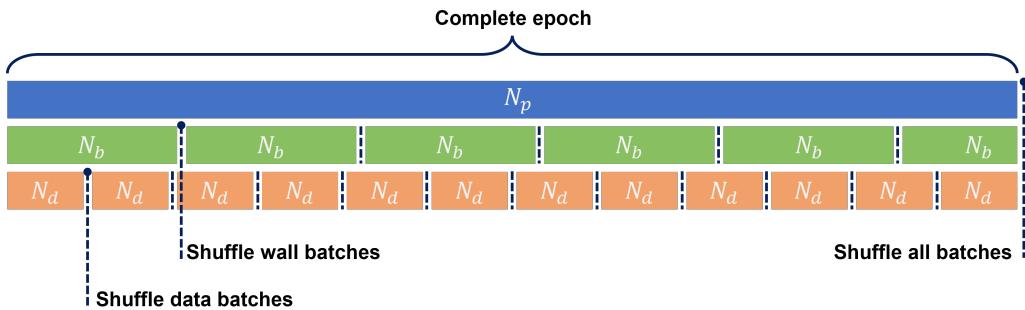
The choice of activation was made after testing with a variety of functions, including hyperbolic tangent (as per [189]) and Swish [192], finding consistently improved qualitative and quantitative performance with Siren network architectures, owing to their ability to capture complex, sub-grid flow features. Details of this study can be found in section 3.3.

### 3.2.1.5 Network Specifics

PINNs can be trained using a dual-optimiser strategy, composed of an initial training phase using the stochastic ADAM optimiser [127], followed by a fine-tuning phase with the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm [144]. The L-BFGS algorithm is a second-order, quasi-Newton method which can provide improved accuracy when used in a dual-optimiser setup. However, it is only applicable as a full-batch method since the descent direction requires an accurate approximation of the entire Hessian matrix. Given the large data set sizes used in this paper, it is infeasible to use a full-batch algorithm due to memory limitations, and thus we use the ADAM optimiser alone. Whilst multi-batch L-BFGS approaches have been explored in the literature [17], they have not yet been applied to PINNs and their reliability is relatively unknown. However, given the performance improvements offered, the development and application of multi-batch, quasi-Newton optimisation algorithms for PINNs could be a fruitful avenue of investigation in future work.

In the present study, an initial maximum learning rate of  $1 \times 10^{-4}$  is used, which was selected to be the highest possible value without divergence of loss values. If divergence of loss values occurs during the training cycle, the learning rate is divided by 4 and training restarted from the previously completed epoch. Additionally, plateau-based learning rate annealing is used, where the learning rate decays by a factor of 0.1 if the validation loss plateaus, within a tolerance of  $10^{-6}$ . The annealing scheme initiates after 5 epochs, and has a cool down period of 3 epochs following a decay action. All code is written in Python 3.7, primarily using the package TensorFlow 2 [1].

As discussed previously,  $N_d$ ,  $N_b$  and  $N_p$  are very large, and thus a mini-batch approach



**Figure 3.2:** Outline of the mini-batch training process used when  $N_d \neq N_b \neq N_p$ .

must be used to reduce the memory overhead during network training. Unlike with the L-BFGS algorithm, this is not an issue with stochastic gradient descent (SGD) algorithms, such as ADAM. In fact, it is commonly recommended that mini-batch schemes are used as opposed to full-batching to avoid convergence to local minima during training [136]. Based on empirical findings, we selected the largest batch size possible, which was restricted by the 4D-flow data size,  $N_d$ . This resulted in batch sizes of 500 for the 2D case, and 5000 for the 3D cases. Our findings are consistent with those presented in [35], where it was demonstrated that larger batch sizes, between  $\sim 1000 - 16000$  reduced prediction error. This is at odds with common practice in the machine learning field, where batch sizes between 2 and 32 have been shown to improve both generalisation and convergence in many applications. However, PINNs are a very distinct sub-field of machine learning, and likely require special treatment.

A final consideration when using mini-batching with  $N_d \neq N_b \neq N_p$  is how to structure the training cycle most efficiently. To maintain a consistent batch size across each data type, we loop continuously through the batches of the smaller data sets until the full epoch is completed, shuffling the data after each complete pass through the respective data type. This ensures that the model weights are continuously updated based on gradients from all three loss components during training, whilst using a common batch size. A visual explanation of this process can be found in Fig. 3.2

### 3.2.1.6 Hyperparameter Optimisation Studies

A critical stage of ML model construction is the optimisation of the network architecture and hyperparameter values using ablation studies, with the function approximation capabilities of ML models highly sensitive to these choices. This holds true for PINNs, where it has been consistently shown in publications across a range of applications that problem-specific hyperparameter tuning is vital to ensure acceptable results. Since PINNs have

no natural inter-domain generalisation capabilities, problem-specific tuning is required to achieve the best results. However, this becomes intractable when required for very large studies, which has, in part, motivated the search for adaptive and dynamic methods for PINN applications [111, 158, 116, 227, 226].

In this chapter, we wish to establish a PINN model configuration that provides consistent results across a range of data degradation levels. This is important for this application since spatio-temporal resolution and SNR are not fixed across patients, due to variations in LV morphology and heart rate.

There are a wealth of optimisation strategies available for ML applications, from grid searching to Bayesian optimisation strategies. To perform our studies we used the Bayesian tree-structured parzen estimator algorithm from the Python package Hyperopt [19]. Bayesian optimisation methods are typically more efficient than exhaustive methods such as grid searching, since the total number of model evaluations required is reduced by selective choice of updated model parameters, decided upon by the particular Bayesian algorithm used.

### 3.2.2 Synthetic Data Generation

As outlined in chapter 2, the utilisation of synthetic data in 4D-flow MRI super-resolution studies is critical [75, 203, 72, 195], since validation of results *in vivo* poses significant challenges. CFD-generated synthetic data provide clearer insight into model performance by allowing control over the severity of artefacts such as noise and resolution, whilst permitting access to high-fidelity ground truth results.

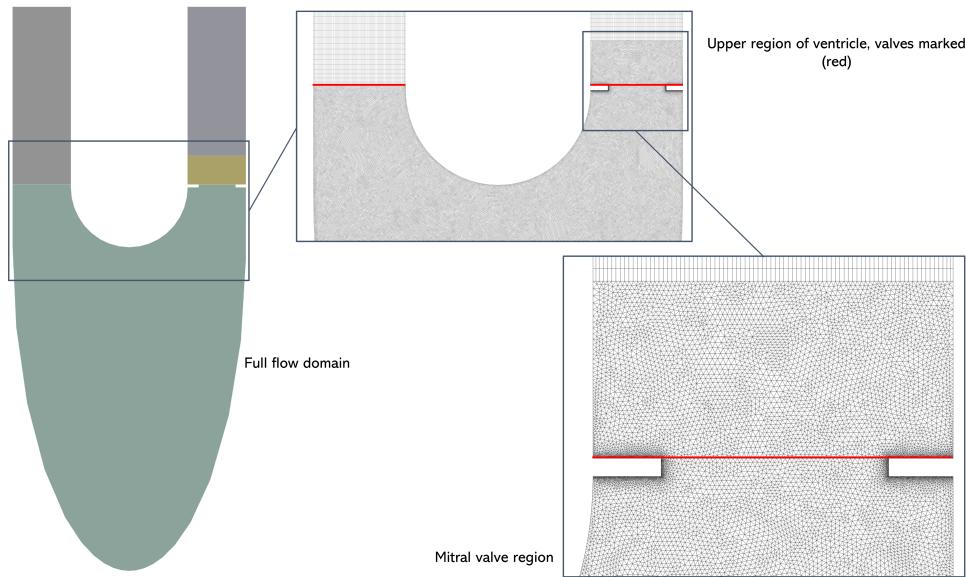
#### 3.2.2.1 Computational Fluid Dynamics

All model generation, meshing and CFD modelling was undertaken using ANSYS SpaceClaim, Mechanical and Fluent (ANSYS Inc., Canonsburg, PA).

#### 2D Idealised Left Ventricle

The synthetic case used in this chapter consists of a 2D idealised ventricle in which a one-way fluid-structure interaction (FSI) coupling drives the fluid flow in both cardiac phases. The primary focus for this experiment is to test various PINN components, such

as activation functions and loss weighting schemes, without using large, 3D data sets, which would incur much greater computational costs. To facilitate this, we aimed to replicate conditions found in typical LV flow in 2D, such as moving boundaries, valve leaflets and flow features across a range of spatial and temporal scales. For reference, the domain and mesh detail for the 2D idealised ventricle can be found in Fig. 3.3.



**Figure 3.3:** 2D idealised ventricle. Flow domain in end diastole (left) with mesh details (middle and right). In end diastole, the mesh consisted of 330,000 elements.

Boundary motion was prescribed analytically across the lower edge of the domain using a sinusoidal function, with node position on the boundary updated at time step  $i$  as:

$$f^i = -0.04 - 0.03 \cos(4\pi t_i) \quad (3.32)$$

$$x^i = x^{i-1} \quad (3.33)$$

$$y^i = f^i \sqrt{1 - (x^i)^2 / 0.000625}. \quad (3.34)$$

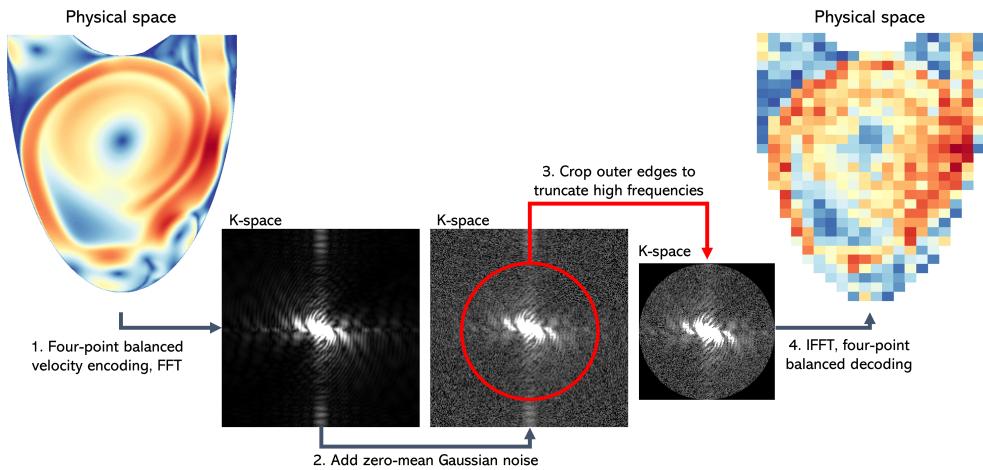
for simulation time  $t_i$ . This resulted in a complete cardiac cycle of 0.5 seconds, producing a peak Reynolds number of 5600, in line with expectations for true LV flow. In diastole, the inlet was opened and the outlet closed, and vice versa in systole. Two small protrusions were placed at the inlet to represent the open leaflets of the mitral valve, designed to produce vortex shedding in diastole. Zero-normal stress conditions were applied to the inlet and outlet when opened, with the no-slip condition applied on the walls and closed openings. We used a highly-resolved mesh of  $\sim 330k$  elements (in end-diastole) and error-based adaptive time-stepping to capture flow features across a range of scales and emulate similar flow patterns to those observed in a real LV. The minimum time step was  $10^{-6}$

seconds. Adaptive meshing was used with diffusion-based smoothing, remeshing every 3 time steps based on a maximum cell skewness of 0.6, and minimum and maximum length scales based on the initial mesh sizing.

The Navier-Stokes equations were solved using direct numerical simulation (DNS), neglecting to include a sub-grid turbulence model. Given the time- and space-averaged nature of cardiac 4D-flow MRI data, the inclusion of turbulent effects in the training data was considered to be of minimal benefit. It would also be unrealistic to assume that, in practice, the PINN model would be able to reconstruct accurate turbulent features of the flow, and it is unlikely that such features would provide beneficial insights into cardiac disease. Further, blood flow in the heart is not fully turbulent, with a wide range of velocities experienced throughout the cardiac cycle including low-velocity flow between systole and diastole. The use of a turbulence model in this setting may introduce non-physical effects which would likely not be possible to capture using the PINN model.

### 3.2.2.2 Synthetic 4D-Flow MRI

Synthetic 4D-flow MRI data were generated from CFD results following a modality-specific downsampling procedure, modified from that found in [72] and [75]. Whilst temporal downsampling was performed by simply dropping a certain number of frames from the cardiac cycle, spatial downsampling and noise synthesis was applied in the  $k$ -space, matching the characteristics of real 4D-flow MRI acquisitions. We have included the equivalent procedure for generating both 2D and 3D synthetic 4D-flow MRI data, which will be used throughout this chapter and the following chapter. The process used was as follows: 1) velocity data at each time step were interpolated onto a high-resolution Cartesian grid, and then a square (in 2D) or cubic (in 3D) region of interest (ROI) was extracted, removing the inflow and outflow channels; 2) the encoding velocity,  $v_{enc}$ , was determined to be the maximum velocity experienced across all time steps, plus 10%; 3) a four-point balanced encoding [175, 118] was then applied to the velocity data to form four phase images; 4) complex signals were generated and then converted to  $k$ -space images using the fast Fourier transform (FFT); 5) spatial downsampling was achieved by cropping the outer edges of the  $k$ -space images using a circular (in 2D) or spherical (in 3D) mask, which has the effect of truncating high-frequency modes [75]. To maintain the square (in 2D) or cubic (in 3D) image shape, zero-padding was applied to the corner regions outside the mask; 6) zero-mean Gaussian noise was added in the frequency domain to both real and imaginary signals, matching the true noise distribution of 4D-flow MRI data [75, 72]. A



**Figure 3.4:** Shown here is the workflow used to process the synthetic 4D-flow MRI data, where spatial downsampling and noise synthesis are applied in the  $k$ -space.

target signal-to-noise ratio (SNR) was achieved by controlling the standard deviation of the synthetic noise based on the signal power in the particular phase image; 6) finally, the inverse FFT (IFFT), followed by the balanced four-point decoding, was applied to recover the downsampled solution data in the spatial domain. The outline of this process (in 2D) can be seen in Fig. 3.4, with the equations used for generating the phase images and complex signals shown below. In 3D, the four phase images are given by:

$$\phi_1 = \frac{\pi}{v_{enc}}(-u - v - w), \quad (3.35)$$

$$\phi_2 = \frac{\pi}{v_{enc}}(u + v - w), \quad (3.36)$$

$$\phi_3 = \frac{\pi}{v_{enc}}(u - v + w), \quad (3.37)$$

$$\phi_4 = \frac{\pi}{v_{enc}}(-u + v + w), \quad (3.38)$$

where  $v_{enc}$  is the encoding velocity. For the 2D procedure, we construct four phase images as above, with  $w = 0$  for each image. For phase image  $j$ , the complex signals are constructed as:

$$S_j = M \exp(i\phi_j), \quad (3.39)$$

where  $M$  is the magnitude image and  $i$  is the imaginary unit. The magnitude image is constructed so that each voxel inside the flow domain is set to 1 and each voxel outside the flow domain set to 0, which ensures that there are distinct noise distributions in the flow and non-flow regions, namely Gaussian noise in the flow domain and uniform noise in the non-flow domain. Since we only consider flow voxels when downsampling in this study, however, we can neglect  $M$ .

## Noise Synthesis

To synthesise the correct distribution and intensity for MRI background noise, we follow the process outlined in [75], in which the appropriate standard deviation of zero-mean Gaussian noise is determined inversely from a target SNR. First, a target SNR (in decibels) is selected *a priori*, which is converted to an ‘effective’ SNR,  $\text{SNR}^{10}$ , given by:

$$\text{SNR}^{10} = 10^{\text{SNR}/10}. \quad (3.40)$$

Then, at time step  $j$ , the signal power is approximated as:

$$P_s(S_j) = \frac{1}{N} \sum_{n=0}^N |S_j^n|^2, \quad (3.41)$$

for total number of voxels  $N$ . By taking the absolute value of the complex signal,  $S_j$ , this equation is calculating the signal power based on the magnitude image, which is related to the VNR as described in chapter 2 and [138]. Next, the noise power in the complex image,  $P_n(S_j)$ , which corresponds to the variance,  $\sigma_j^2$ , is calculated as:

$$P_n(S_j) = \frac{P_s(S_j)}{\text{SNR}^{10}}. \quad (3.42)$$

We add the noise independently to the real and imaginary components of  $S_j$ , and thus the standard deviation of zero-mean Gaussian noise added to each signal component is given by:

$$\sigma_j = \sqrt{\frac{P_n(S_j)}{2}}. \quad (3.43)$$

Under this procedure, the SNR of the synthetic data remains fixed, whilst  $\sigma_j$  is variable dependent on the signal power at time step  $j$ . The SNR synthesised here is translated to the visible VNR [138] as:

$$\text{VNR} \simeq \frac{\pi}{2} \frac{|\mathbf{u}|}{v_{enc}} \text{SNR}. \quad (3.44)$$

Here, it becomes apparent why visible noise in 4D-flow MRI is exacerbated in low-flow regions, or when the  $v_{enc}$  is set too high. The selection of the  $v_{enc}$  should therefore correspond to the lowest possible value without causing aliasing, or phase-wrapping, of the velocity reconstruction.

### 3.2.2.3 Boundary Conditions and Physics Interrogation

The final consideration is how best to apply the boundary conditions and generate the collocation point cloud, at which the physics regularisation is probed.

#### Boundary Conditions

The location and motion of the endocardium in real 4D-flow MRI studies is determined through segmentation and registration of structural MRI scans, which are acquired alongside the 4D-flow MRI data. Therefore, synthetic boundary velocity data do not require the same downsampling procedure as synthetic 4D-flow MRI. For both cases discussed in this chapter, we use the clean CFD data at the boundary nodes as our boundary condition. By doing this, we assume that we have complete knowledge of the applied boundary condition, allowing us to focus efforts solely on uncertainty in the flow data. Positional uncertainty is present in real MRI-based cardiac studies, although these effects will be investigated in chapter 4

#### Collocation Point Cloud

The final consideration is the generation of the collocation point clouds, on which we apply the PDE loss during PINN training. In both cases in this chapter, the point cloud is chosen to simply correspond to the nodes of the CFD mesh used, since this is already highly resolved and evenly distributed in space and time. As has been covered extensively in the literature, the collocation point cloud must be spatio-temporally dense to achieve satisfactory results. In studies with real 4D-flow MRI data, uniform sampling of points across the segmented LV domain can provide the collocation point cloud. When downsampling synthetic 4D-flow MRI data temporally, we assume that we maintain access to both the boundary velocities and collocation point cloud at all time steps. This approach is consistent with studies using real 4D-flow MRI data later in this thesis, where we are able to interpolate the endocardial surface in time to generate both point clouds and associated boundary velocity data using point-wise splines and their derivatives.

### 3.2.3 Metrics

In this chapter, we assess the quantitative accuracy of the PINN model with respect to both the velocity and pressure fields. We defer analysis of clinically-relevant derived quantities, such as WSS and vorticity, to Chapter 4. We evaluate performance using max-normalised

velocity root mean square error (RMS), for velocity ( $\text{RMS}_u$ ) and pressure ( $\text{RMS}_p$ ), with additional velocity metrics in the form of alignment similarity index (ASI) and magnitude similarity index (MSI)[72]. These metrics are defined as:

$$\text{RMS}_u := \frac{1}{\max|\mathbf{u}|} \sqrt{\frac{1}{N} \sum_{k=1}^N (\mathbf{u}_{pred} - \mathbf{u})_k^2}, \quad (3.45)$$

$$\text{RMS}_p := \frac{1}{\max|p|} \sqrt{\frac{1}{N} \sum_{k=1}^N (p_{pred} - p)_k^2}, \quad (3.46)$$

$$\text{ASI} := \frac{1}{2} \left( 1 + \frac{\mathbf{u}_{pred} \cdot \mathbf{u}}{|\mathbf{u}_{pred}| |\mathbf{u}|} \right), \quad (3.47)$$

$$\text{MSI} := 1 - \left| \frac{|\mathbf{u}|}{\max|\mathbf{u}|} - \frac{|\mathbf{u}_{pred}|}{\max|\mathbf{u}_{pred}|} \right|, \quad (3.48)$$

where  $\mathbf{u}_{pred} = U\hat{\mathbf{u}}$  and  $p_{pred} = \rho U^2 \hat{p}$ , with  $\mathbf{u}$  and  $p$  the ground truth velocity and pressure, respectively, at  $N$  interrogation points. Given that the pressure prediction is accurate only up to a constant, we calculate  $\text{RMS}_p$  by zero-centring the mean value for prediction and ground truth data at each time step, in essence extracting the relative pressure. ASI and MSI can be calculated at each interrogation point, in which case the form above is used, or the mean value across  $N$  points can be used to provide a single value.

### 3.3 Results

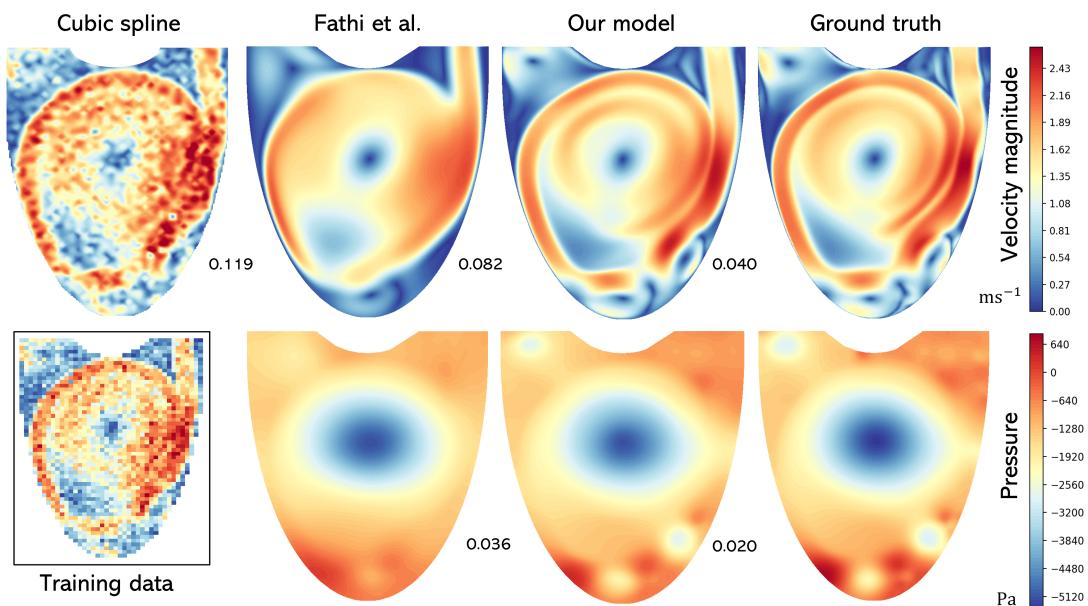
This results section is structured as follows: first, the optimised model configuration is introduced, with results presented across a range of data degradation levels (in terms of resolution and SNR); following this, competing architecture choices, such as activation function and loss weighting strategies, are compared and assessed with regards to performance with low-resolution and noisy data.

#### 3.3.1 2D Idealised Ventricle

The PINN model hyperparameter configuration used in this section can be found in Tab. 3.1, alongside descriptions of the hyperparameters specified. A hyperparameter optimisation study was used to tune network depth and width, dropout rate and activation function

**Table 3.1:** PINN architecture and hyperparameter values used for the 2D idealised LV model.

Components		
Activation function	Periodic sine (Siren) [206]	
Loss weighting	Self-adaptive PINN (SA-PINN) [158]	
Hyperparameters		
Depth	7	Number of hidden layers in the network
Width	800	Number of neurons per hidden layer
Initial learning rate	$10^{-4}$	Learning rate defined at the start of training
Dropout rate	0.55	Governs the proportion of neurons randomly removed from specified layers in the network
Batch size	3,200	Size of data, PDE and BC mini-batches
Initialisation frequency ( $\omega_0$ )	10	Parameters used to modulate spectral bias in the Siren network [206]
Data loss learning rate ( $\eta_d$ )	0.08	Learning rate for the ADAM optimiser used to update data loss weights in the SA-PINN scheme [158]
PDE loss learning rate ( $\eta_p$ )	0.0001	Learning rate for the ADAM optimiser used to update PDE loss weights in the SA-PINN scheme [158]
BC loss learning rate ( $\eta_b$ )	0.03	Learning rate for the ADAM optimiser used to update BC loss weights in the SA-PINN scheme [158]

**Figure 3.5:** 2D idealised ventricle: velocity magnitude (top row) and pressure (bottom row) contour plots in late diastole, comparing our model (centre right) with that used in [72] (centre left) and cubic spline (left). Here we have used training data with temporal and spatial downsampling rates of 5 and 4 respectively, with a SNR of 6.6. Insert displays the resolution of training data. Reported values are RMS for the respective fields.

choice, following the technique outlined in Section 3.2.1.6. The search space defined for this hyperparameter study can be found in Tab. 3.3, where we performed 140 training runs in total. Further details about this study can be found in section 3.3.2. A grid search was used to tune the SA-PINN learning rates, rather than using Bayesian optimisation. This was done because the specific values of  $\eta_d$ ,  $\eta_p$  and  $\eta_b$  are dependent on the level of corruption present in the synthetic 4D-flow MRI data, and as such, extensive optimisation using a single data configuration would reduce the ability of the model to generalise across different configurations of resolution and noise. Further details for this study can be found in section 3.3.4.

In Fig. 3.5, we compare the performance of our PINN model configuration given in Tab. 3.1 against two competing methods: cubic spline interpolation and the PINN configuration used in Section 4.2 of [72]. Qualitatively, our model shows robustness to noise whilst still being able to capture most of the solution features present in the ground truth data. Conversely, the cubic spline algorithm is heavily corrupted by the presence of noise and is clearly unsuitable in this context, while the PINN model used in [72] is able to remove the data noise at the expense of removing the finer flow details, which is also shown quantitatively with increased RMS values. The architecture and components used in [72] correspond to the ‘vanilla’ formulation of PINNs, as used in the original publication [190]. As such, no techniques were implemented to mitigate the effects of spectral bias, and fixed weights were utilised in the loss function. Further, the network used was relatively narrow, at a width of 25. Therefore, the model did not have the requisite capacity to represent all length and time scales present in the underlying velocity and pressure fields.

The impact of increasing spatial resolution is assessed in Figs. 3.6 (for noise-free data) and 3.7 (SNR of 6.6), where we compare PINN predictions across three spatial downsampling rates. For both noisy and noise-free data, the PINN model is able to effectively reconstruct both velocity and pressure fields effectively across all downsampling rates analysed. In the presence of noise at the spatial downsampling rate of 16, some small-scale flow features are not captured and peak flow values are reduced, although this downsampling rate corresponds to a spatial resolution of  $5\text{mm}^2$ , below what would be expected in real 4D-flow MRI studies. An animated version of these results can be found at: <https://drive.google.com/drive/folders/1pxuEz3VfdZSJNxUtvj0o41oDvkgHAAZ8?usp=sharing>. In this animation, PINN results, training data and the ground truth data are shown, with data obtained at spatial and temporal downsampling rates of 8 and 5, respectively, at a SNR of 6.6.

In Tab. 3.2, results are compared across two temporal downsampling rates, three spatial downsampling rates and two SNR levels. Across all configurations of data degradation,  $\text{RMS}_u$  ranges from a minimum of 0.023 to a maximum of 0.079, whilst  $\text{RMS}_p$  ranges from 0.011 to 0.031. For reference, the competing method displayed in Fig. 3.5 exceeds the maximum  $\text{RMS}_u$  and  $\text{RMS}_p$  values attained across all data degradation configurations.

**Table 3.2:** 2D idealised ventricle: Errors in predicted velocity and pressure fields across all configurations of data degradation types, using the model configuration introduced in Tab. 3.1.

SNR	DS <sub>T</sub>	DS <sub>S</sub>	RMS <sub>u</sub>	ASI	MSI	RMS <sub>p</sub>
$\infty$	0	4	0.023	0.990	0.982	0.011
$\infty$	0	8	0.038	0.978	0.971	0.015
$\infty$	0	16	0.054	0.964	0.957	0.021
$\infty$	5	4	0.039	0.981	0.972	0.019
$\infty$	5	8	0.051	0.970	0.962	0.022
$\infty$	5	16	0.065	0.956	0.949	0.027
6.6	0	4	0.029	0.985	0.976	0.014
6.6	0	8	0.041	0.973	0.965	0.018
6.6	0	16	0.066	0.950	0.945	0.027
6.6	5	4	0.040	0.975	0.968	0.020
6.6	5	8	0.052	0.963	0.960	0.023
6.6	5	16	0.079	0.936	0.936	0.031

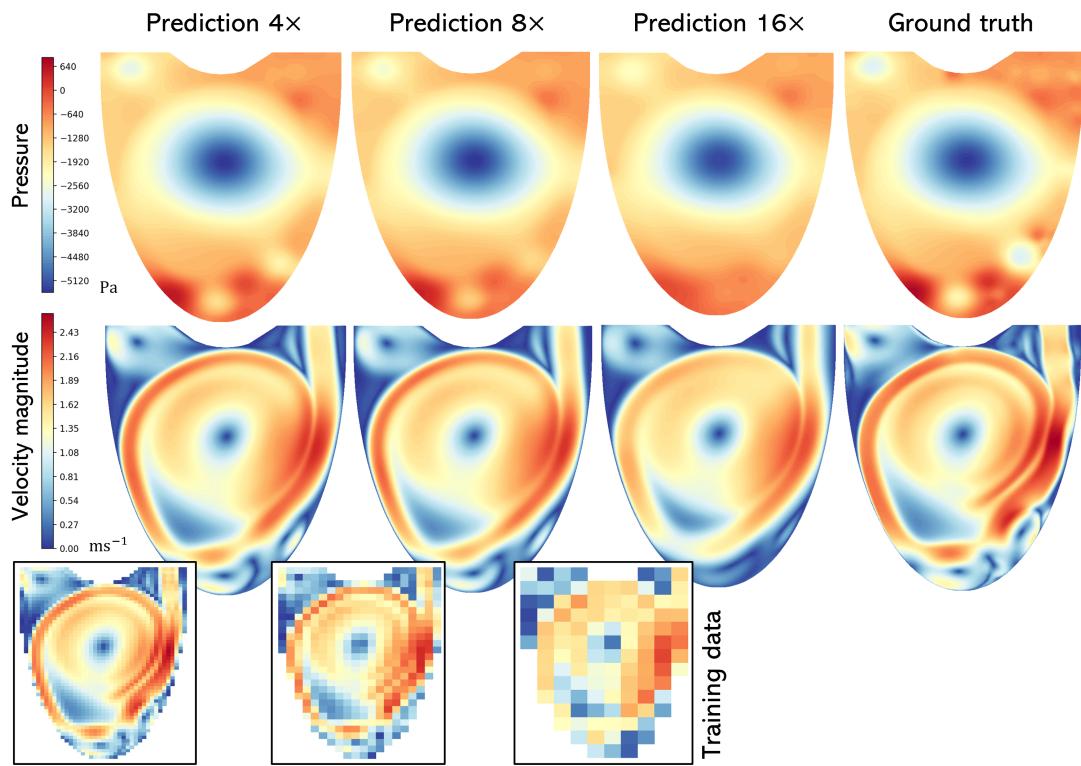
### 3.3.2 Activation Function

The selection of an appropriate activation function can have a significant impact on PINN model convergence rates and reconstruction quality. Most early PINN publications utilised fixed tanh activation functions, alongside ReLU [81], Swish [192] and, less frequently, Sigmoid functions. However, as has been demonstrated for PINNs and other implicit neural representation networks [109, 110, 216, 206], the use of such activation functions can result in a model that is unable to resolve high-frequency solution components within a reasonable training time, due to spectral bias. Since then, focus has been placed on the development of activation functions that are able to modulate spectral bias, utilising adaptive components [110, 109] or input feature projection [206, 216, 239].

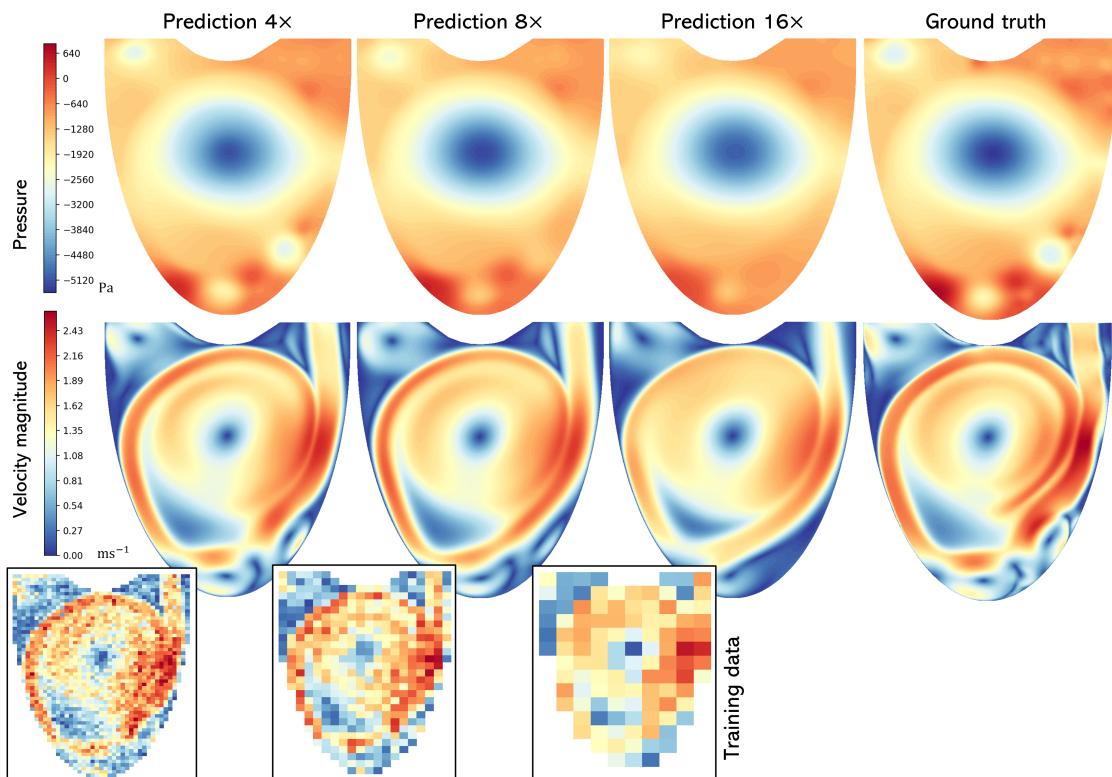
In this section, we compare model performance using three fixed activation functions,

**Table 3.3:** Search space defined for the 2D hyperparameter optimisation study.

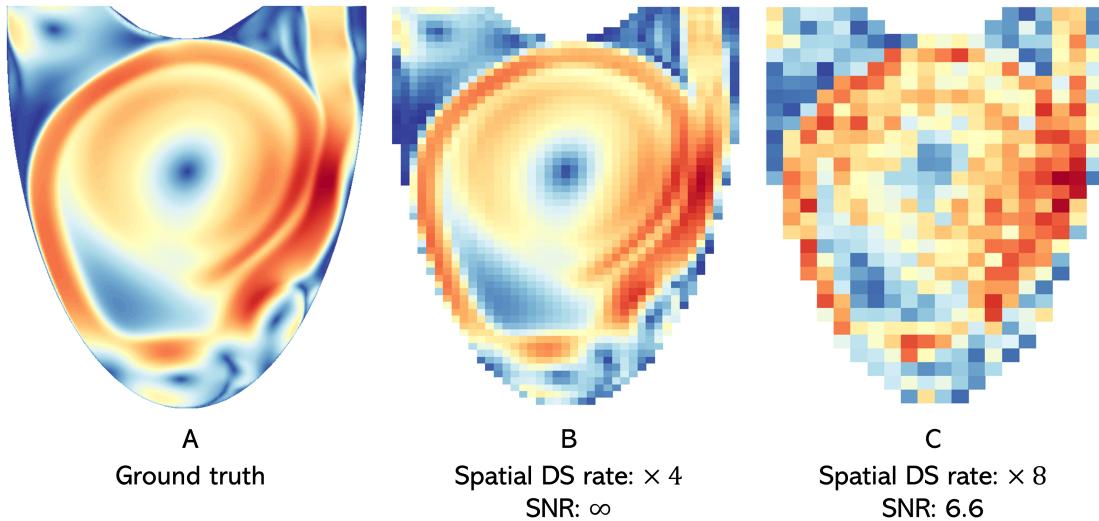
Search space: activation function study	
Hidden layers	5-12
Neurons per layer	100-1000
Dropout rate	0-0.8
Activation Function	tanh, Swish, Siren



**Figure 3.6:** 2D idealised ventricle: pressure (top row) and velocity magnitude (second row) contour plots in late diastole, using noise-free training data with a temporal downsampling rate of 5, and spatial downsampling rates specified at the top of the figure, alongside the ground truth fields (right). Inserts show the resolution of training data.



**Figure 3.7:** 2D idealised ventricle: pressure (top row) and velocity magnitude (second row) contour plots in late diastole, using training data with an SNR of 6.6, with a temporal downsampling rate of 5, spatial downsampling rates specified at the top of the figure, alongside the ground truth fields (right). Inserts show the resolution of training data.



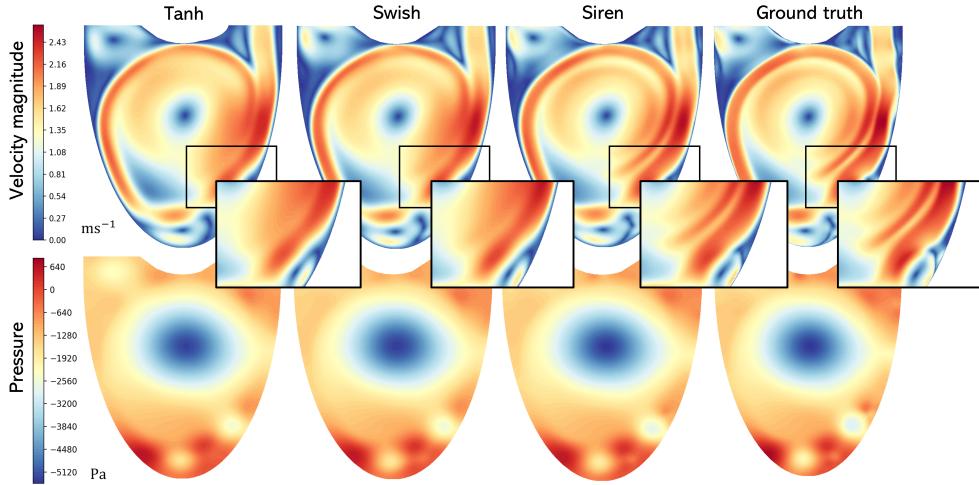
**Figure 3.8:** 2D idealised ventricle: Here we show contour plots of the velocity magnitude at one time step for ground truth CFD data (A) and synthetic 4D-flow MRI data using different levels of data degradation (B and C). The data used in to compare activation functions (B) is noise-free, with a spatial downsampling (DS) rate of 4, whereas the data used to compare loss weighting strategies (C) has a SNR of 6.6 and a spatial DS rate of 8.

namely tanh, Swish [192] and periodic sine functions (Siren network) [206]. Under the assumption that network size and dropout rate interact with each activation, we performed an ablation study across activation function choice, network depth and width and dropout rate. The complete search space for this study can be found in Tab. 3.1. The data used in this study was noise-free and downsampled by factors of 5 and 4 in time and space. The main qualitative marker for performance here is the *capacity* of each activation function to resolve small-scale flow features, and thus we used noise-free data to remove the potential risk of over-fitting. Further, using noise-free data removes the potential for conflict between loss components, reducing the impact of sub-optimal loss weight configurations. Thus, we were able to use the weight configuration proposed in our previous publication [205], which was selected to provide the best performance for this data configuration. A visualisation of the synthetic training data used can be found in Fig. 3.8B.

In Fig. 3.9, we compare velocity and pressure results for the best performing configurations found for each activation function. From the magnified region, we see the periodic sine activation is superior in its ability to resolve small-scale flow features, which is consistent with the results presented in [206]. From Tab. 3.4, we find that the Siren also outperformed the other activation functions quantitatively, achieving the best scores across all error metrics. Given that the Siren network architecture incurs minimal additional computational overhead, the results shown here present it as the obvious candidate for this application. The collection of activation functions studied here is not exhaustive, and additional exploration into adaptive activation functions may be beneficial.

**Table 3.4:** Optimal results for each activation function analysed in the ablation study.

Activation	RMS <sub>u</sub>	ASI	MSI	RMS <sub>p</sub>
tanh	0.0582	0.963	0.958	0.027
Swish	0.0551	0.968	0.961	0.0253
Siren	<b>0.0437</b>	<b>0.98</b>	<b>0.972</b>	<b>0.0209</b>

**Figure 3.9:** 2D idealised ventricle: contour plots of velocity magnitude (top) and pressure (bottom) at one time step, comparing three activation functions: Tanh, Swish and Siren. Qualitatively, the Siren activation provides the greatest similarity to the ground truth data (right), preserving more of the high-frequency solution details, as shown in the insert.

### 3.3.3 Initialisation Frequency

An additional consideration when using the Siren network is the tuning of the initialisation frequency,  $\omega_0$ . This parameter dictates the dominant frequency spectra of predicted solutions, which enables the architecture to modulate the effects of spectral bias [239]. It was found empirically that, for our application, the default value used in the original publication [206],  $\omega_0 = 30$ , did not provide satisfactory results. To this end, we reduced the value to  $\omega_0 = 10$  for the Siren architectures used throughout this chapter. Here, we demonstrate the quantitative effect of changing  $\omega_0$ . Using noise-free data, the appropriate selection of  $\omega_0$  is not so impactful, since there is no risk of over-fitting. Therefore, to test the effects of this value we use noisy data, utilising the configuration found in Fig. 3.8C. In Tab. 3.5, we observe results for this study, where  $\omega_0 = 10$  provides considerably better performance than other values. The optimal value of  $\omega_0$  is case-specific, since it is dependent on the relevant frequency modes of that are of importance in the underlying solution fields.

**Table 3.5:** Error metrics for different values of the initialisation frequency  $\omega_0$  in the Siren network architecture.

$\omega_0$	RMS <sub>u</sub>	ASI	MSI	RMS <sub>p</sub>
1	0.129	0.906	0.887	0.054
10	<b>0.052</b>	<b>0.963</b>	<b>0.960</b>	<b>0.023</b>
30	0.064	0.949	0.950	0.027
50	0.076	0.949	0.933	0.036

### 3.3.4 Loss Weighting Strategies

Following the publication of [205], it was observed that the PINN configuration used was adversely impacted by noise corruption at higher spatial downsampling rates. Although a dynamic weighting strategy was used in [205], the influence of the initial weights,  $\alpha_0$  and  $\beta_0$ , was not fully characterised. Therefore, it was determined that the over-fitting was related to sub-optimal weighting of the individual loss components, in particular over-weighting of  $\mathcal{L}_{data}$ , and thus alternative strategies were investigated. To this end, we compare PINN results using three distinct loss weighting strategies, namely 1) fixed loss weights, 2) LR annealing scheme as proposed in [116] and 3) SA-PINN scheme [158], exploring different parameter choices for each method.

There is a key distinction between scheme 3) and the other two that should be reiterated. Where methods 1) and 2) apply a single weight to each loss component, given by  $\alpha$  and  $\beta$ , the SA-PINN approach applies a weight to each individual training sample, with each loss component using a distinct optimiser to apply the necessary gradients (for instance,  $\mathcal{L}_{PDE}$  has its own optimiser to maximise the weights associated with the collocation points). This allows the network to balance gradient contributions from each loss component, as with the previous approaches, but also balances contributions based on local solution complexity. In practice, regions of the solution field that are difficult to learn or have a significant impact on global convergence (for instance, regions near a boundary in space or near an initial state in time) are given larger weights, increasing the magnitude of gradients from these points.

The motivation for this study is to establish a weighting strategy that both performs well under significant data degradation, but is also insensitive to initial parameter choices. There is an emphasis on the latter point, since we require our PINN model to be robust across different levels of data degradation and would like the model configuration to be independent of this. Further, we also wish to reduce the dependency of PINN performance on initial parameter choices, since the model must be applicable to large cohorts of patients

for which case-specific tuning is infeasible.

When comparing competing activation functions earlier in this chapter, we used noise-free data at a relatively low spatial downsampling rate, since the aim was to assess the capacity of the selected activation functions to resolve high-frequency solution components without concern over artefacts due to noise over-fitting. However, here we are motivated by improving model robustness to noise at high spatial downsampling rates, for which we have conflicting loss objectives which are more challenging to weight correctly. Therefore, the flow data here are downsampled by factors of 5 and 8 in time and space, respectively, with a SNR of 6.6. This combination of parameters results in heavily corrupted training data, which can be seen visually in Fig. 3.8C.

For weighting strategies 1) and 2), we produce results with weights initialised (or fixed) across four distinct orders of magnitude. With strategy 3), we instead look at distinct values for the respective learning rates (LRs), given by  $\eta_p$  (PDE loss LR),  $\eta_d$  (data loss LR) and  $\eta_b$  (BC loss LR). As stated in [158], it is possible to also initialise the self-adaptive weights using non-unity values, however, altering the respective learning rates should have an equivalent effect and provide a solution that is less problem-dependent. For strategies 1) and 2), the weights  $\alpha$  and  $\beta$  are applied to  $\mathcal{L}_{data}$  and  $\mathcal{L}_{BC}$ , respectively, where  $\mathcal{L}_{PDE}$  maintains a fixed weight of 1. Therefore, adjusting  $\alpha$  and  $\beta$  modifies the *relative* balance of weights in the loss function, where the absolute value of the weights is less important. This does not hold true for the equivalent parameters,  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ , in the SA-PINN configuration, where the absolute magnitude impacts the sensitivity of each optimiser. Improper configuration of these values may result in training divergence or convergence to local minima. Therefore, we compare SA-PINN configurations using a fixed relative balance between the respective learning rates (configurations SA-PINN<sub>D</sub>, SA-PINN<sub>E</sub> and SA-PINN<sub>F</sub>) to establish the effect of absolute values, alongside a comparison quantifying the effect of relative imbalances (configurations SA-PINN<sub>opt</sub>, SA-PINN<sub>A</sub>, SA-PINN<sub>B</sub> and SA-PINN<sub>C</sub>). In chapter 4, we explore this area further, where we establish that the key determinant for acceptable results is identifying the correct ratio between  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ . The loss strategy configurations used can be found in Tab. 3.6.

The performance of each loss weighting configuration using the data defined in Fig. 3.8C can be found in Tab. 3.6. In the LR annealing scheme [116], it was found that model performance was highly dependent on the selection of appropriate initial weights,  $\alpha_0$  and  $\beta_0$ , which is unconventional for a dynamic method. Whilst the weights update dynamically throughout training, the magnitude of weights selected at initialisation persisted.

**Table 3.6:** 2D idealised ventricle: Hyperparameter selections used for each configuration of loss weighting strategies, and quantitative results obtained from each loss weighting configuration.

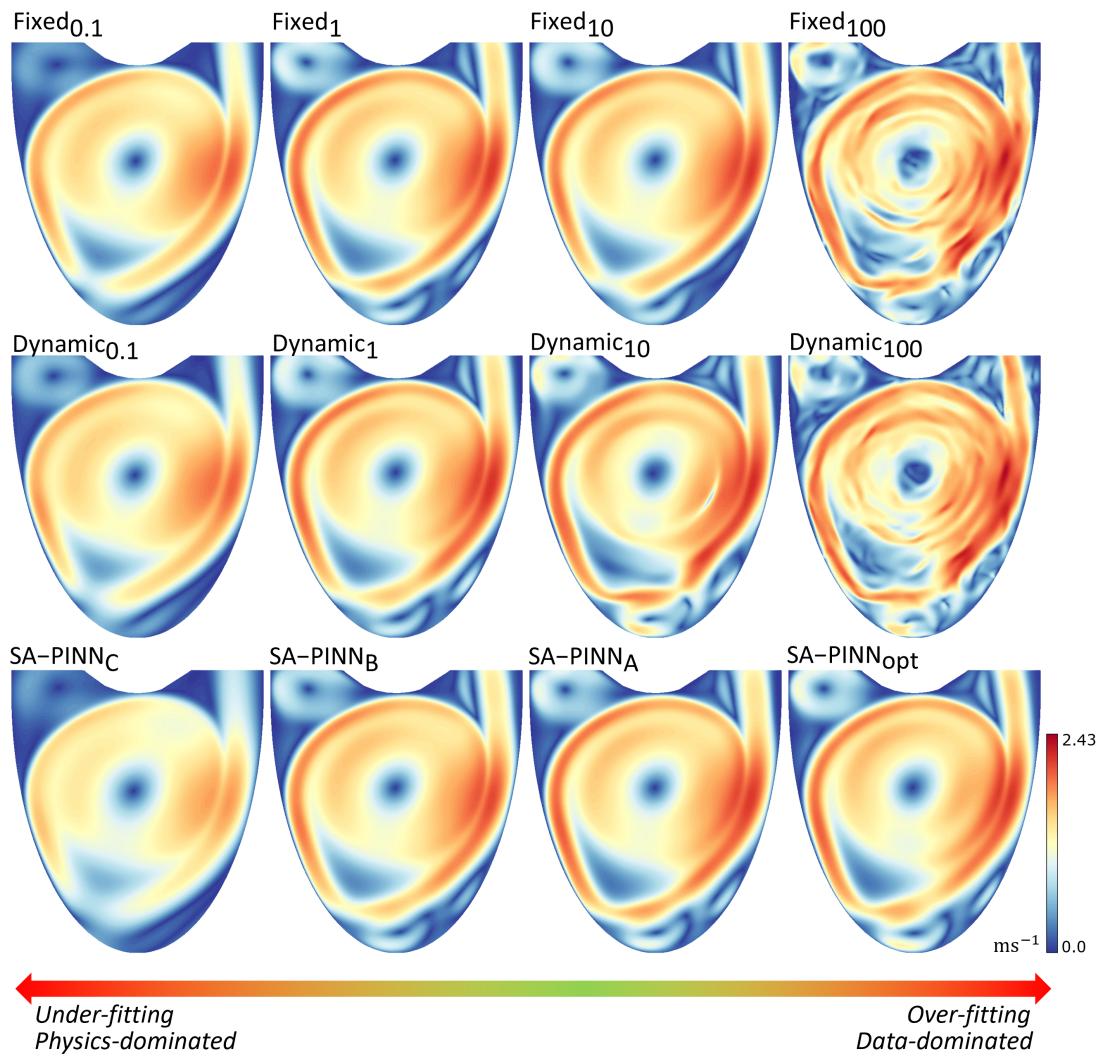
Config.	$\alpha^0$	$\beta^0$	$\eta_p$	$\eta_d$	$\eta_b$	RMS <sub>u</sub>	ASI	MSI	RMS <sub>p</sub>
Fixed <sub>0.1</sub>	0.1	0.1	-	-	-	0.086	0.935	0.929	0.039
Fixed <sub>1</sub>	1	1	-	-	-	0.058	0.957	0.956	0.026
Fixed <sub>10</sub>	10	10	-	-	-	0.059	0.958	0.955	0.026
Fixed <sub>100</sub>	100	100	-	-	-	0.072	0.939	0.940	0.028
Dynamic <sub>0.1</sub>	0.1	0.1	-	-	-	0.088	0.933	0.927	0.039
Dynamic <sub>1</sub>	1	1	-	-	-	0.056	0.960	0.957	0.025
Dynamic <sub>10</sub>	10	10	-	-	-	0.055	0.956	0.957	<b>0.023</b>
Dynamic <sub>100</sub>	100	100	-	-	-	0.072	0.939	0.941	0.028
SA-PINN <sub>opt</sub>	-	-	0.0001	0.08	0.03	<b>0.052</b>	<b>0.963</b>	<b>0.960</b>	<b>0.023</b>
SA-PINN <sub>A</sub>	-	-	0.001	0.08	0.03	0.055	0.961	0.958	0.024
SA-PINN <sub>B</sub>	-	-	0.01	0.08	0.03	0.058	0.958	0.954	0.026
SA-PINN <sub>C</sub>	-	-	0.1	0.08	0.03	0.114	0.919	0.897	0.053
SA-PINN <sub>D</sub>	-	-	0.0001	0.0001	0.0001	0.057	0.959	0.957	0.025
SA-PINN <sub>E</sub>	-	-	0.001	0.001	0.001	0.057	0.960	0.956	0.025
SA-PINN <sub>F</sub>	-	-	0.01	0.01	0.01	0.066	0.954	0.946	0.03

This can be seen in Tab. 3.7, where we display the initial and final loss weights for the four configurations studied. Similarly, the individual learning rates,  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ , affected prediction accuracy with the SA-PINN method. The relative balance between these parameter values was shown to have a greater impact than the absolute values, although this did play a role, as can be seen in SA-PINN configurations D-F in Tab. 3.6. We find the best performing configuration, overall, to be SA-PINN<sub>opt</sub>, which achieves the best scores across all metrics.

**Table 3.7:** 2D idealised ventricle: Initialised weights,  $\alpha_0$  and  $\beta_0$ , and final weights (after  $N$  weight updates),  $\alpha_N$  and  $\beta_N$ , for the LR annealing method.

Config.	$\alpha^0$	$\beta^0$	$\alpha^N$	$\beta^N$
Dynamic <sub>0.1</sub>	0.1	0.1	0.35	0.38
Dynamic <sub>1</sub>	1	1	2.56	6.19
Dynamic <sub>10</sub>	10	10	17.0	39.2
Dynamic <sub>100</sub>	100	100	235	598

Based on the findings in Tab. 3.6, the relative balance between  $\mathcal{L}_{PDE}$  and the other two loss constraints appears to be the determining factor for acceptable performance. The regularisation provided by the PDE loss constraint acts to modulate learning of the velocity data, where configurations balanced too heavily in favour of  $\mathcal{L}_{PDE}$  lead to under-fitting. Under these conditions, solutions are too heavily regularised by the physics constraint, resulting in overly smoothed solution fields. This effect can be seen qualitatively in Fig. 3.10, where each loss weighting scheme displays under-fitting in the ‘physics-dominated’ regime, for which  $\mathcal{L}_{PDE}$  dominates the training cycle. Conversely, configurations balanced

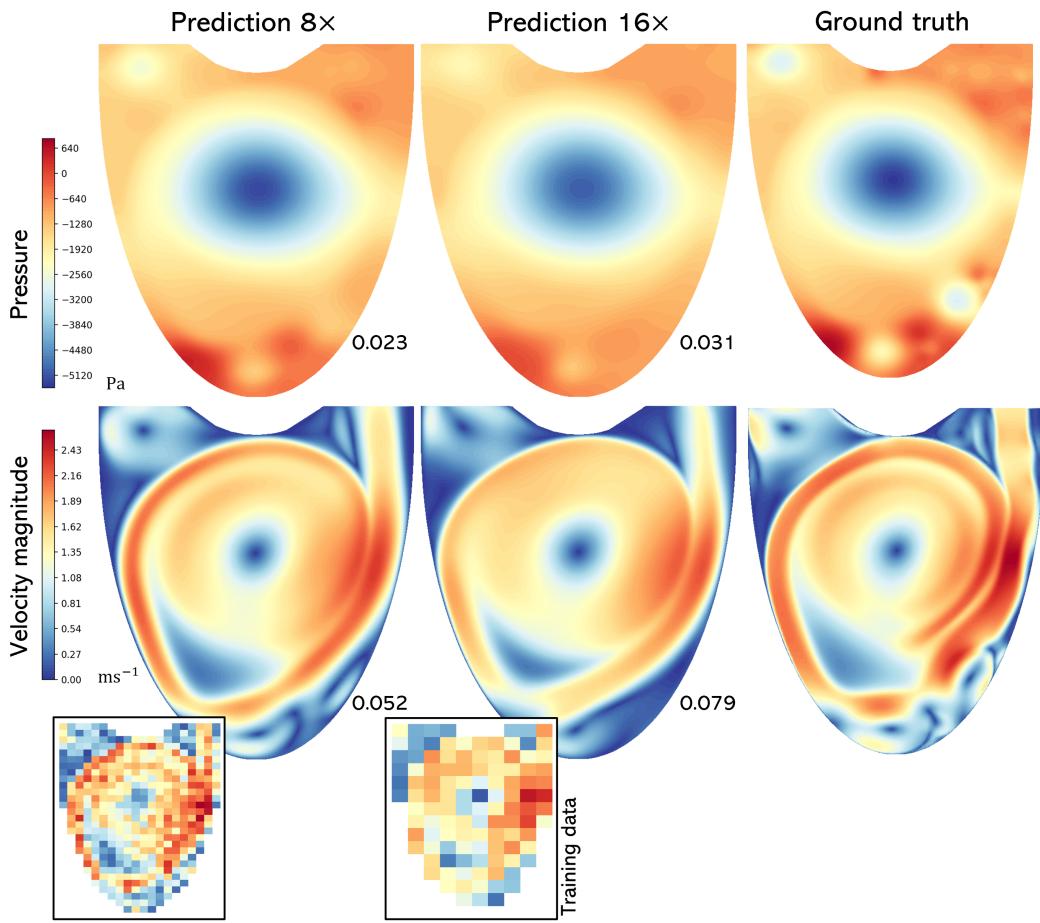


**Figure 3.10:** 2D idealised ventricle: A demonstration of under-fitting and over-fitting induced by poorly configured weights. Velocity magnitude contour plots in late diastole, predicted using different configurations of fixed weights, LR annealing scheme initial weights (Dynamic) and SA-PINN learning rates.

too heavily in favour of  $\mathcal{L}_{data}$  lead to over-fitting in the presence of data noise.

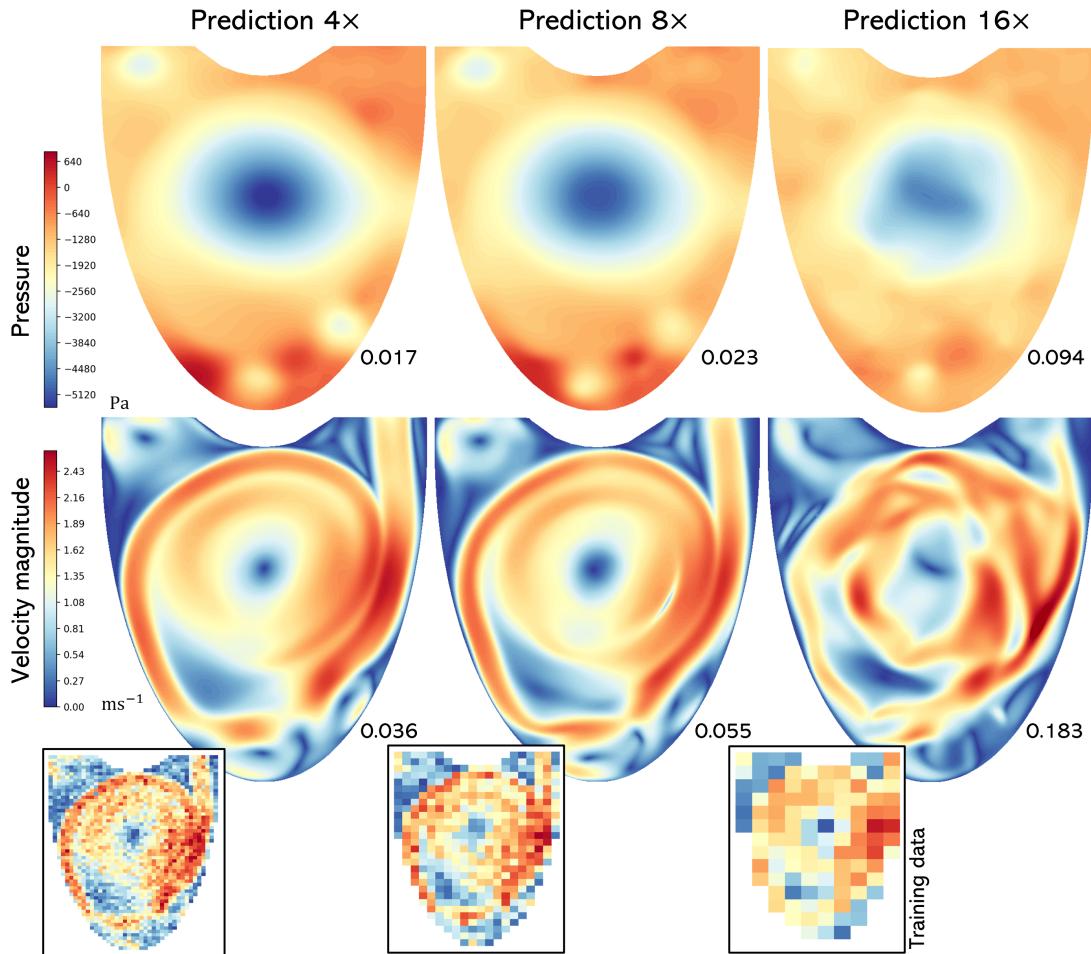
We observe that the fixed and LR annealing schemes are severely impacted by noise artefacts for configurations  $Fixed_{100}$  and  $Dynamic_{100}$ . The SA-PINN scheme appears less prone to over-fitting, which is corroborated in Fig. 3.10 and corresponding configurations in Tab. 3.6. Across SA-PINN configurations D-F, it can also be observed that the SA-PINN model is only slightly affected by changes in the absolute values of the learning rates  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ . This suggests a reduced sensitivity of the SA-PINN configuration to initial parameters, which is a desirable property for the present application, in which we wish to restrict the requirement for hyperparameter tuning between distinct cases.

Beyond comparing performance in a single data configuration, we also wish to assess the



**Figure 3.11:** 2D idealised ventricle: Predicted solution fields using a fixed configuration of the SA-PINN scheme. Pressure (top row) and velocity magnitude (bottom row) contour plots in late diastole, using training data with an SNR of 6.6, with a temporal downsampling rate of 5, spatial downsampling rates specified at the top of the figure. Insert displays the resolution of training data. Reported values are RMS for the respective fields.

applicability of each scheme across multiple spatial resolutions, without additional hyper-parameter tuning. To this end, the best performing configurations from loss weighting schemes 2) and 3) are deployed across three distinct data degradation configurations, each of which is downsampled by a factor of 5 in time with a SNR of 6.6, across spatial downsampling rates of 4, 8 and 16. In Fig. 3.11, the SA-PINN<sub>opt</sub> configuration is used to produce results across the data ranges specified above. Although not all features are reconstructed at the spatial downsampling rate of 16, there are no visible issues with over-fitting. Further, reconstruction of the pressure field is consistent across all spatial resolutions, with the dominant solution features approximated well. As a comparison, in Fig. 3.12, the best performing LR annealing configuration, Dynamic<sub>10</sub>, is applied to the same collection of data. It is observed that this model actually outperforms SA-PINN<sub>opt</sub> for the data with a spatial downsampling rate of 4, quantitatively. However, for data at the spatial downsampling rate of 16, severe over-fitting occurs, significantly degrading solution accuracy.



**Figure 3.12:** 2D idealised ventricle: Predicted solution fields using a fixed configuration of the LR annealing (dynamic) scheme. Pressure (top row) and velocity magnitude (bottom row) contour plots in late diastole, using training data with an SNR of 6.6, with a temporal downsampling rate of 5, spatial downsampling rates specified at the top of the figure. Insert displays the resolution of training data. Reported values are RMS for the respective fields.

## 3.4 Discussion

### 3.4.1 Spectral Bias

The selection of an appropriate activation function is critical in the current study, where ventricular flow is characterised by features across a range of length and time scales. In vanilla form, PINNs are generally able to capture large-scale (or low-frequency) solution components well, but spectral bias reduces the capacity of the model to reconstruct small-scale (or high-frequency) solution components. The effects of spectral bias can be mitigated by modifying certain network components, including the activation function used in the hidden layers. It has been demonstrated qualitatively and quantitatively in this chapter that the Siren network, which utilises periodic sine activation functions, is an effective tool in alleviating issues of spectral bias, in line with the content discussed in section 2.4.4.

### 3.4.2 Gradient Imbalances

PINNs are typically able to regularise network outputs to mitigate the effects of measurement uncertainty through the imposition of the physical loss terms,  $\mathcal{L}_{PDE}$  and  $\mathcal{L}_{BC}$ . However, under the influence of significant levels of measurement noise, appropriate weighting of these terms is more important, since  $\mathcal{L}_{data}$ ,  $\mathcal{L}_{PDE}$  and  $\mathcal{L}_{BC}$  provide conflicting objectives. This makes the learning task more challenging, since the backpropagated gradients from each loss term may counter-act each other, and excessive weighting of one term in this scenario promotes sub-optimal solutions that satisfy one term too heavily, to the detriment of the other objectives. As seen in Fig. 3.12, increased degradation of the important low-frequency solution information at highest downsampling rate (16) results in increased uncertainty in the velocity data, and thus increased disagreement between the PINN loss terms. Therefore, the initial weight used in the LR annealing scheme for  $\mathcal{L}_{data}$ , given by  $\alpha_0 = 10$ , leads to excessively large gradient contributions from  $\mathcal{L}_{data}$  while neglecting the conflicting gradients from  $\mathcal{L}_{PDE}$  (since  $\mathcal{L}_{BC}$  is only applied to points on the boundary, its impact on flow away from the wall is minimal). This reduces the regularising effects of the physics-based loss term and induces over-fitting. Whilst the LR annealing scheme is intended to be dynamic, results presented in Tab. 3.7 demonstrate that the respective weights remain fixed around the order of magnitude selected at initialisation. As such, strong gradient imbalances are perpetuated throughout the training cycle.

In Fig. 3.12, we observe the impact of increased spatial downsampling rates with an SNR of 6.6 for LR annealing scheme. Whilst the model is able to effectively de-noise the velocity fields and uncover the pressure fields to an acceptable quality for spatial downsampling rates of 4 and 8, severe over-fitting can be observed at a rate of 16. This distinction between results using different spatial downsampling rates in the presence of noise can be understood by re-examining the method used to construct the synthetic noise in the  $k$ -space. With the target SNR (and thus  $\text{SNR}^{10}$ ) fixed, the power of the noise at time point  $j$  is determined by the equation:

$$P_n(S_j) = \frac{P_s(S_j)}{\text{SNR}^{10}}, \quad (3.49)$$

for signal power  $P_s(S_j)$ , which is calculated by taking the mean  $k$ -space signal value. As seen in Fig. 3.4, before downsampling, much of the  $k$ -space (away from the centre of the image) is populated by low-magnitude signal intensities, which will contribute to a relatively small overall mean signal power, and thus noise power. When this noise is applied to the  $k$ -space image, the low-signal regions in the higher-frequency bands become populated by comparatively large amounts of noise, which helps to explain why the background noise is visually characterised by high-frequency patterns in the spatial domain at higher resolutions. However, the comparative intensity of this noise in relation to the high-intensity region in the centre of the image is small. This central region corresponds to the lower-frequency spatial information, within which the most important structural details of the solution are contained. Therefore, since the intensity of the noise added to this region is comparably small, most of the important structural information remains intact. Conversely, at higher downsampling rates (in which more of the outer region of the  $k$ -space image is cropped), a higher proportion of the image is populated by the higher-intensity signals in the central region, thus producing a greater noise power. Since the intensity of this noise is then larger in comparison to signals in the central region of the  $k$ -space image, the important structural information in the spatial domain is more heavily degraded. To summarise, increasing the spatial downsampling rate in this synthetic setup promotes the generation of noise with a power that scales with the signal power of decreasing frequency, and thus increasing intensity.

With real MRI acquisitions, increasing the voxel size has the effect of *increasing* SNR, since the total signal received within a voxel is greater [66]. However, in this synthetic study, the point of interest is assessing model performance at a particular target SNR. Therefore, this is the quantity that remains fixed, which implies that the noise power must increase with decreasing spatial resolution. Such a exaggerated SNR values at very low spatial

resolutions is unrealistic, however, it demonstrates the effectiveness of the PINN model in extreme conditions.

Results presented in this chapter suggest that the SA-PINN scheme is less sensitive to initial hyperparameter choices, as seen when comparing Figs. 3.11 and 3.12, which is a desirable quality for the present application. In real 4D-flow MRI studies, variability in LV size and heart rate across subjects results in different spatio-temporal resolution and SNR values, and as such, optimising loss weighting parameters on a case-by-case basis is infeasible. Whilst the LR annealing scheme outperforms the SA-PINN scheme at the spatial downsampling rate of 4 in Figs. 3.11 and 3.12, this is less important than the results obtained at the higher downsampling rates. Data obtained at rates of 8 and 16 correspond to spatial resolutions of 2.5 and 5mm<sup>2</sup>, which are closer to the recommended values in clinical practice in LV studies [23]. Therefore, acceptable reconstruction at these resolutions is of greater importance than at the spatial downsampling rate of 4, which corresponds to a spatial resolution of 1.25mm<sup>2</sup>, far higher than is plausible in the LV.

### 3.4.3 Robustness to Data Corruption

In Figs. 3.6 and 3.7, we illustrate velocity and pressure reconstructions at distinct spatial downsampling rates, with and without noise added. The similarity in reconstruction quality between noise-free and noisy data, without re-configuring hyperparameter values, demonstrates robustness of the PINN model to changes in SNR, where an SNR value of 6.6 exceeds what would be expected in real 4D-flow MRI studies. Similarly, there is only a small degradation in quality across spatial downsampling rates of 4, 8 and 16, with no noise artefacts present at higher downsampling rates, unlike those observed with the LR annealing scheme.

A comprehensive assessment across all data degradation levels can be found in Tab. 3.2. As expected, model accuracy increases with increasing spatial resolution, temporal resolution and SNR.

## 3.5 Conclusion

In this chapter, the PINN model used to perform cardiac 4D-flow MRI super-resolution was introduced, and competing architecture choices were compared to mitigate the im-

pacts of spectral bias and gradient imbalances. An extensive study was performed to establish a model configuration able to produce accurate results using data of a low SNR at high spatio-temporal downsampling rates. Model selection was based on both performance at a fixed level of data degradation, but also insensitivity to initial hyperparameter choices. The selected model configuration, which utilised the SA-PINN weighting scheme [158], was shown to perform best on both metrics. It was then demonstrated that a fixed model configuration,  $\text{SA-PINN}_{opt}$ , was able to operate successfully across a range of spatio-temporal downsampling rates and SNR values without case-specific tuning. It was also demonstrated that the Siren network architecture (with periodic sine activation functions) [206] provided the best results when compared with competing activation functions. The ability to modulate the effects of spectral bias using the initialisation frequency,  $\omega_0$ , presented a clear advantage, where the accurate reconstruction of small-scale flow features is of importance in the current study.

## Chapter 4

# Assessing Robustness to Positional Uncertainty in the Endocardium

### 4.1 Introduction

Fluid flow in the LV is subject to a complex fluid-structure interaction (FSI), where systolic and diastolic flow are driven by contraction and relaxation of the myocardium. The resulting haemodynamic flow field is therefore highly dependent on the structure and motion of the endocardium with which it interacts. For MRI-based haemodynamic simulation studies in the LV that require the imposition of boundary conditions, this dependency on the position and motion of the boundary is problematic, as low spatial resolution, particularly in the apico-basal direction, results in poorly reconstructed endocardial surfaces. Much of the wall detail that is present on the real endocardial surface, such as trabeculae carneae and papillary muscles, is not included in resulting segmentations, and through-plane motion in the apico-basal direction is not appropriately resolved. This results in significant *positional uncertainty*, which likely has a detrimental effect on the simulation of patient-specific flow in the LV. This is particularly problematic in the absence of intraventricular velocity measurements, such as with CFD-based studies, as the simulated solution fields are entirely dictated by the boundary conditions. However, such issues also likely persist when using PINN-based models that are only partially constrained by the boundary condition, even given the availability of flow measurement data within the domain.

The impact of positional uncertainty on CFD simulations in the LV has been covered in the

literature, with studies characterising the fluid response to trabeculae carneae [194, 222, 135, 161], papillary muscles [194, 222, 135, 161] and chordae tendineae [164]. Whilst the function of trabeculae carneae is still unknown, the presence of these features along with papillary muscles were observed to significantly impact fluid flow reconstruction within the LV, reducing WSS, increasing pressure drop and altering the breakup of vortical structures in diastole [222, 194].

The primary motivation behind these publications has been to better understand the haemodynamic function of trabeculae carneae and papillary muscles. However, in the present study we are primarily concerned with the accurate reconstruction of *patient-specific* haemodynamic solution fields. As these previous publications demonstrate, the presence of complex endocardial features can exert a significant influence, quantitatively and qualitatively, on the observed LV flow. Therefore, simulation studies that assume the endocardium to be smooth, (or otherwise not representative of the true topography), produce solution fields that are likely not patient-specific, and generate quantitative values that do not accurately reflect the underlying flow.

Mitigation of positional uncertainty with PDE-constrained optimisation methods is addressed in [166], where flow fields in synthetic 3D domains in the vasculature are reconstructed from planar PC-MRI. Such optimisation methods require segmentation of the vascular lumen, however, positional uncertainty can restrict the accuracy of such methods. Instead of enforcing the no-slip condition, an alternative, parameterised slip/transpiration boundary condition is proposed on the segmented lumen, under the assumption that the approximated boundary lies within the flow domain. The parameters of this boundary condition are then estimated as part of the optimisation process. The assumption that the segmented lumen lies within the flow domain precludes the usage of this approach in the current study, where segmented LV boundaries typically lie outside the physical flow domain along the compact endocardium.

Given the utilisation of existing intraventricular velocity measurements, which are derived directly from the patient-specific velocity field, the PINN-based model introduced in this thesis is only partially constrained by the boundary condition. This reduced dependency on the boundary constraint should somewhat mitigate the impact of positional uncertainty, when compared with CFD-based methods, but this area has not yet been explored. Given the uncertainty present in the endocardial position, it may appear sensible to simply remove the boundary condition constraint from the PINN loss function. However, doing so removes the capability to accurately resolve the near-wall flow, and thus WSS, which

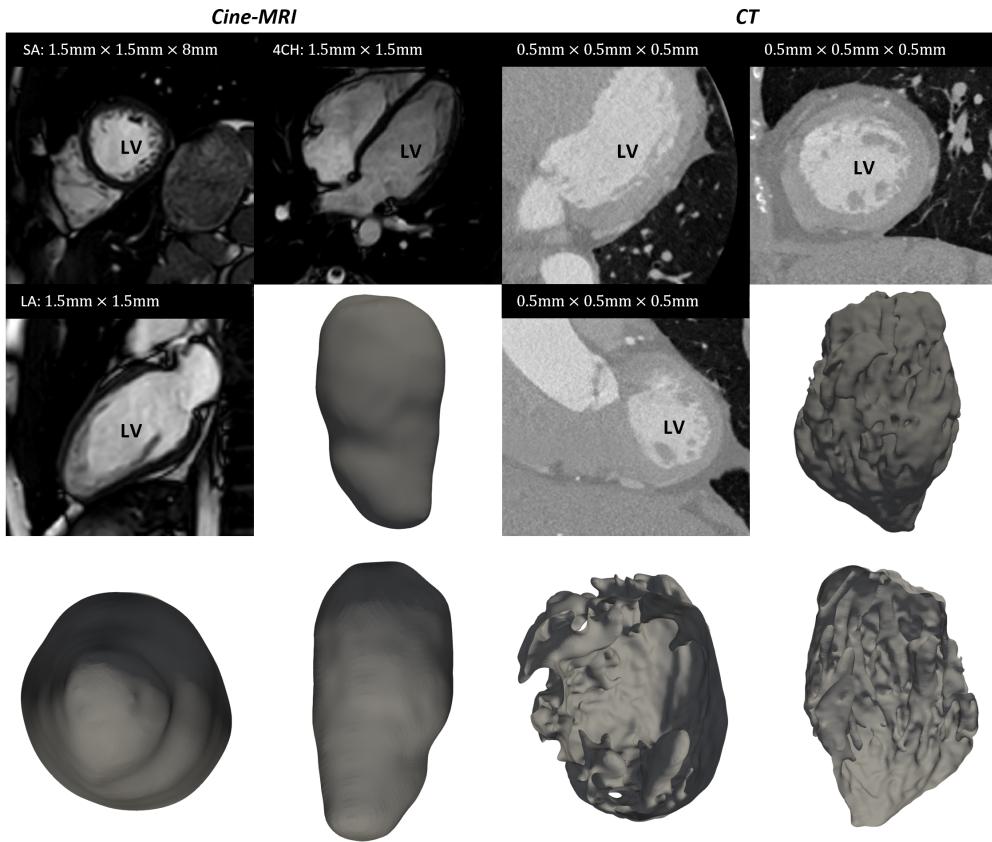
is a desirable variable of interest in the LV.

In this chapter, errors arising from positional uncertainty are characterised. Synthetic 4D-flow MRI data are generated in three patient-specific LV geometries, generated from CFD simulations using end-diastolic CT segmentations. The use of CT imaging data allows for the reconstruction of all but the smallest details present on the endocardial surface, and thus the generated CFD data are dependent on such features. From these high-resolution segmentations, three approximated boundary representations are produced, each using different quantities of smoothing. The boundary generated using the highest level of smoothing is intended to represent an expected segmentation from cine-MRI. We isolate the effects of positional uncertainty by assuming we have complete knowledge of the ground truth boundary motion, omitting the characterisation of uncertainty arising due to inaccurate approximation of the myocardial deformation. Super-resolved solution fields are obtained with each boundary representation and results are compared, so as to quantify the impact of positional uncertainty. Beyond the error metrics introduced in chapter 3, we also derive clinically-relevant flow variables that have been explored in the literature, comparing the PINN results with those produced directly from synthetic 4D-flow MRI and obtained using competing methods. These variables include vorticity, relative pressure, kinetic energy and WSS.

## 4.2 Methodology

### 4.2.1 Study Setup

The focus of this study is to investigate the impact of uncertainty in the position of the endocardial surface, examining the effect that different levels of smoothing impart on model performance. In standard 4D-flow MRI studies, the position and velocity of the endocardium are extracted from a sequence of structural cine-MR images. At each time point in the cardiac cycle, these typically consist of a short-axis stack of 2D slices in the apico-basal direction, with a slice thickness of 6 – 10mm and in-plane spatial resolution of 1 – 2mm<sup>2</sup> [130]. Although trabeculations and papillary muscles may be partially discernible within each slice, the low spatial resolution in the apico-basal direction makes it highly challenging to extract such complex endocardial features across an entire ventricular volume, which would require spatial interpolation between slices. Even still, details present between slices would likely be highly inaccurate. Trabeculae carneae range in



**Figure 4.1:** A comparison of endocardial surfaces segmented from cine-MRI and CT (in different patients). For each modality, three figures are included to illustrate the segmented endocardial surface, which demonstrates the significant difference in level of detail available between cine-MRI and CT. Alongside the short-axis stack (as seen for one slice in the upper left figure), with cine-MRI, single slice acquisitions are also produced in four-chamber (4CH) and long-axis (LA) views, at an in-plane resolution of  $1 - 2\text{mm}^2$ . Comparatively, CT is capable of achieving an isotropic spatial resolution of  $\sim 0.5\text{mm}^3$ , which is sufficient to reconstruct most of the endocardial wall features. However, CT is only obtained at a fixed phase in the cardiac cycle, typically end-diastole.

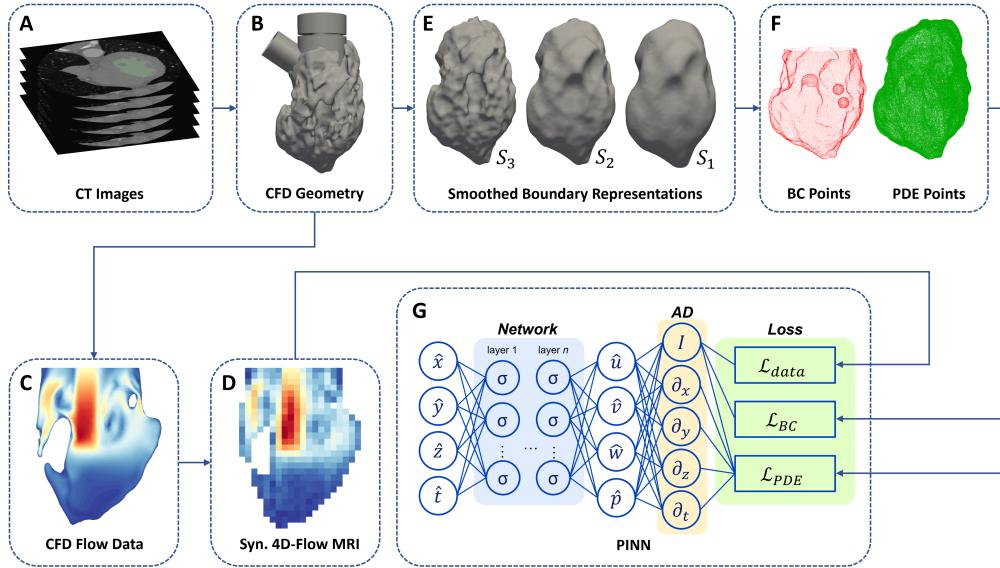
length from around 0.5 – 3mm, and posses a diameter of around 50 – 500 $\mu$ m. Therefore, they are challenging to reconstruct, even using higher-resolution structural imaging modalities such as CT. Similarly, they exist in the sub-grid size in relation to the 4D-flow MRI resolution.

Additional effects, such as through-plane motion and the presence of motion artefacts imparted by long, potentially uncomfortable, scan times, further degrade the quality of segmentations [160]. Therefore, high levels of uncertainty arise in both the position and velocity of the endocardium when segmenting and registering from MR images. In Fig. 4.1, typical segmentations obtained from both cine-MRI and CT are compared. The comparable lack of detail in the cine-MRI segmentation, and the resulting source of positional uncertainty, is apparent.

Standard practice in MRI-based CFD studies in the LV is to produce a smooth, approximated endocardial boundary, usually corresponding to the location of the compact endocardium [199]. Inter-slice interpolation of such a boundary is a simpler task, resulting in a uniformly smooth surface across the entire volume. However, it has been shown that the presence of complex endocardial features significantly impact the resulting flow both quantitatively and qualitatively [194, 222, 135, 161]. The effect of endocardial smoothing, or positional uncertainty in general, has not been quantified for PINN-based models such as ours, where the presence of flow data in the chamber alongside the boundary condition is likely to pose unique challenges. In CFD simulations, flow results are highly dependent on the imposed boundary conditions, where small inaccuracies in position or velocity may significantly impact results. However, with PINN-based models, predictions are constrained primarily by the available 4D-flow MRI data, with the boundary condition imparting less influence, at least far from the wall.

In this study, we quantify the impact of endocardial smoothing on predicted solution fields, assessing results using a number of distinct boundary approximations across increased levels of smoothing. Three patient-specific geometries are used to obtain three distinct cases of ground truth flow data, each including a high level of endocardial detail such as trabeculations and papillary muscles. The resulting CFD data are therefore dependent on these complex wall structures, and the purpose of this study is to determine whether the PINN super-resolution model can still reasonably predict all necessary variables without full knowledge of the endocardial structure. For simplicity, we assume that the motion, and thus velocity, of the boundary, which is again prescribed using an analytical function, is known across all boundary representations, and thus the task is restricted to quantifying

sensitivity to error in position. A schematic for this experiment can be found in Fig. 4.2.



**Figure 4.2:** Assessing the impact of positional uncertainty – experimental setup. (A-C) CFD data are generated in patient-specific LV geometries, segmented from end-diastolic CT imaging data with trabeculae carneae and papillary muscles present. (D) Synthetic 4D-flow MRI data (denoted syn. 4D-flow MRI) are generated from the CFD results, to be used as flow measurement training data. (E) From the high-resolution CT segmentation, three distinct boundary representations are produced, namely  $S_1$ ,  $S_2$  and  $S_3$ , each generated using different quantities of smoothing. (F) From these, three sets of boundary condition and collocation point clouds, and associated velocities, are generated. (G) Individual PINN models are then used to produce results across each of the boundary representations for comparative analysis.

## 4.2.2 Synthetic Data Generation

### 4.2.2.1 Ground Truth Flow Data

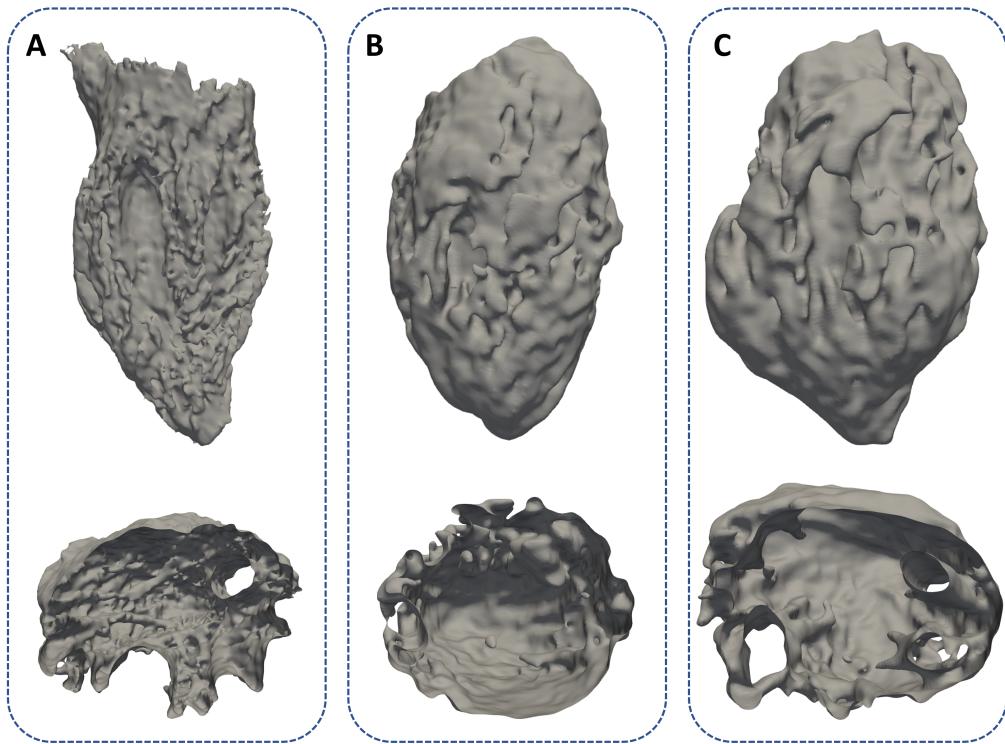
#### Acquisition and Segmentation

The image data used in this study were provided by the Multi-Modality Whole Heart Segmentation (MM-WHS) challenge [244], from which we randomly selected three LV volumes acquired using computed tomography (CT) imaging. Each volume was acquired in end-diastole, with an in-plane resolution of  $0.78\text{mm}^2$  and a slice thickness of 1.6mm, capable of accurately extracting many of the complex features on the endocardium. The MM-WHS challenge contains a mixture of healthy and impaired subjects, suffering from a variety of cardiovascular diseases, including cardiac function insufficiency (NYHA II), cardiac edema, hypertension (III), sick sinus syndrome, arrhythmia, atrial flutter, atrial fibrillation, artery plaque, coronary atherosclerosis, aortic aneurysm (dilated aorta), Tetralogy of Fallot (right ventricle hypertrophy), dilated cardiomyopathy (left ventricle), aortic stenosis, pulmonary artery stenosis or a combination of the above. This produces geometries with a range of shapes and sizes. The specific condition of the cases used in this study is

unknown, however, the cases were selected to cover a range of ventricular volumes.

Ground truth LV volumes, with which CFD simulations were performed, were segmented using the open source medical image processing software 3D Slicer (<http://www.slicer.org>) [73], where the grow from seeds region growing algorithm was used. This resulted in three surface geometries that displayed a realistic level of wall detail, including trabeculae carneae and protrusions from papillary muscles. The ground truth LV volumes produced for this study were compared with the reference endocardium segmentations included within the MM-WHS data set to ensure consistency, with good agreement shown between our segmentations and those produced by the authors of the data set.

The three segmented geometries are illustrated in Fig. 4.3.



**Figure 4.3:** Segmented surface geometries from three patient-specific CT image sets, namely cases A, B and C.

### CFD Setup

The CT image data provided in the MM-WHS challenge set were obtained at one fixed time point, in end diastole, and thus motion could not be derived directly from a sequence of volumes as with functional MRI-based studies. Therefore, we utilised an analytical function to prescribe the node-wise motion of the endocardium throughout the cardiac cycle. A sinusoidal function was used to define the motion in time, controlling radial and axial contraction and dilation, and, distinct from the approach used in chapter 3, torsional

motion. Under torsional motion, the apex of the LV was rotated by approximately  $15^\circ - 25^\circ$ , based on values suggested in [201], whilst the base remained fixed in space. The range of values explored was constrained by numerical stability of the surface during CFD simulation. The equations governing boundary node displacement at time step  $i + 1$  are given by:

$$x^{i+1} = r^i(1 - RS^i\delta_t^i z^i) \cos(\phi^i + TS^i\delta_t^i z^i), \quad (4.1)$$

$$y^{i+1} = r^i(1 - RS^i\delta_t^i z^i) \sin(\phi^i + TS^i\delta_t^i z^i), \quad (4.2)$$

$$z^{i+1} = z^i(1 + AS^i\delta_t^i), \quad (4.3)$$

where

$$S^i = -\sin(2\pi t_i), \quad (4.4)$$

$$r(x, y) = \sqrt{x^2 + y^2}, \quad (4.5)$$

and

$$\phi(x, y) = \begin{cases} \tan^{-1}(y/x), & \text{if } x > 0 \text{ and } y > 0 \\ \tan^{-1}(y/x) + 2\pi, & \text{if } x > 0 \text{ and } y < 0 \\ \tan^{-1}(y/x) + \pi. & \text{otherwise/} \end{cases}$$

Constants R, A and T are used to control the amount of radial, axial and torsional motion respectively, whilst  $\delta_t^i$  represents the time step size at time step  $i$ . The current time is given by  $t_i$ . For cases B and C, we used  $R = 45, T = 75$  and  $A = 0.3$ , whereas for case A we used  $R = 22.5, T = 75$  and  $A = 0.3$ . These values were selected empirically for numerical stability during the CFD simulation, where case A required a lower amount of radial compression since the lower initial chamber volume resulted in challenging deformation of the mesh.

The mitral valves used in this study transition directly from fully open to fully closed, and mitral valve leaflets have not been included (with no fluid-structure interaction modelled on the leaflets). Further, the boundary condition on the papillary muscles does not account for contraction and relaxation throughout the cardiac cycle, and these features do not directly attach (via chordae tendineae) to the mitral valve. Although this presents a simplification, it was deemed that the impact on the fluid flow characteristics within the modelled LV would be limited, and therefore this detail was omitted. It is also worth noting that the ground truth data, alongside the boundary condition as used by the PINN model, is produced using this condition.

The use of a sinusoidal function to dictate temporal motion of the endocardium is an over-simplification of myocardial deformation, where the use of realistic LV volume curves to produce distinct A-wave and E-wave phases would be more accurate. However, the focus of this study is the recovery of high-resolution ground truth solution fields from synthetic 4D-flow MRI data, and thus the use of realistic myocardial deformation is not necessarily required, so long as the solution complexity in the ground truth flow data is sufficiently similar to real LV flow.

As with chapter 3, the Navier-Stokes equations were solved via DNS.

#### 4.2.2.2 Synthetic 4D-flow MRI

Synthetic 4D-flow MRI data were obtained following the same workflow as outlined in section 3.2.2.2. To replicate realistic acquisition conditions, the 4D-flow MRI data should be masked using the approximated LV segmentations, which would include some flow voxels that correspond to tissue regions. In the magnitude images, these voxels contain zero values since there is no signal present. This produces uniform noise, which scales with the encoding velocity ( $v_{enc}$ ), in the corresponding regions in the phase images [75], which can therefore induce severe corruption. Since we are primarily concerned with positional uncertainty in this study, we neglect to include such tissue regions in our synthetic 4D-flow MRI data.

For each case in this section, the synthetic 4D-flow MRI data were generated to achieve a spatial resolution of  $3\text{mm}^3$ , a temporal resolution of 48ms and a SNR of 10. The spatio-temporal resolution was selected to align with expected values in typical cardiac 4D-flow MRI studies [23], which corresponds to spatial and temporal downsampling rates of 4 and 5, respectively. The particular SNR value used can be considered relatively low when compared to synthetic 4D-flow MRI studies in the literature [155, 72, 75], but was selected to demonstrate the effectiveness of our model under challenging conditions.

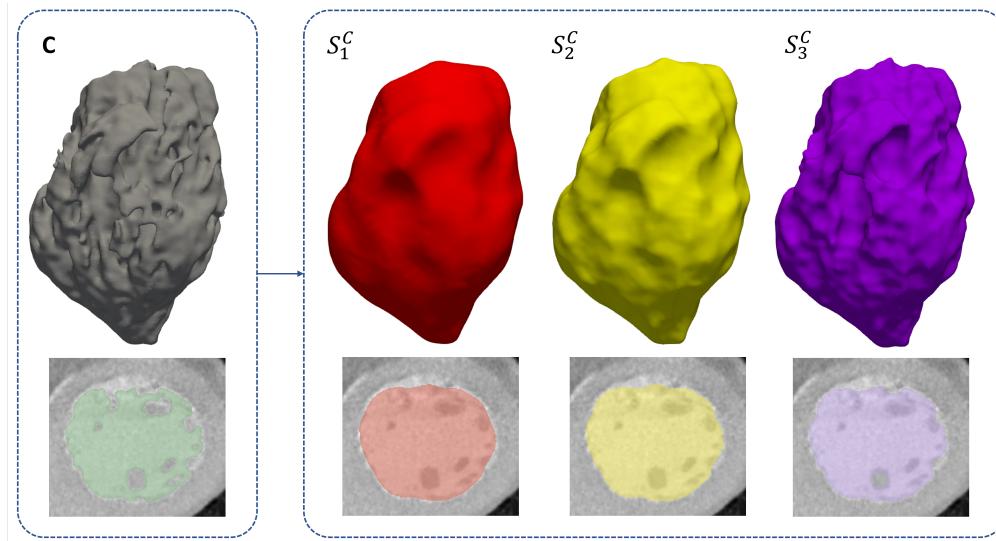
#### 4.2.2.3 Boundary Representations

For each case,  $A$ ,  $B$  and  $C$ , we generated three boundary approximations,  $S_1^i$ ,  $S_2^i$  and  $S_3^i$  for  $i \in \{A, B, C\}$ , where smoothing operations were applied sequentially to remove progressively larger features from the surface. The surface smoothing was performed in 3D Slicer (<http://www.slicer.org>) [73] using the ‘closing holes’ and ‘Gaussian’ smoothing

operations. The ‘closing holes’ operation fills sharp corners and holes smaller than the specified kernel size (in mm), and does not remove any material from the segmentation. This operation is used to add material to the segmentation, which has the effect of dilating the endocardial surface. The ‘Gaussian’ operation smooths all details with a specified kernel size (in mm), which typically shrinks the segmentation. This operation is used to smooth the smaller details that remain on the segmentation after the ‘closing holes’ operation has been applied. The surfaces with the highest level of smoothing applied,  $S_1^i$ , were designed to replicate segmentation of the compact endocardium, in which all wall features and protrusions are removed. Details of the parameters used for the smoothing operations used can be found in Tab. 4.1, with the corresponding surfaces produced (for case C) shown in Fig. 4.4.

**Table 4.1:** Details of the smoothing parameters used to reconstruct surfaces  $S_1^i$ ,  $S_2^i$  and  $S_3^i$  for cases A, B and C.

Surface	Closing holes (mm <sup>3</sup> )	Gaussian (mm <sup>3</sup> )
$S_3^i$	5	1
$S_2^i$	10	1.5
$S_1^i$	20	2



**Figure 4.4:** The endocardial surfaces reconstructed using smoothing operations outlined in Tab. 4.1, for LV case C.

A limitation in this study is the procedure used to downsample the LV volumes, where smoothing operations are applied directly to the surface. A more appropriate workflow may be provided by artificial downsampling of the CT images, both in-plane and through-plane, to match the typical resolution and slice thickness of cine-MRI before segmenting the low-resolution LV volumes. However, it was deemed that smoothing the surfaces directly would not significantly affect the objectives of this study, although this may be a beneficial avenue for further research.

### 4.2.3 PINN Model

The model used in this chapter, and associated components, is that established in chapter 3. As such, the PINN model architecture is composed of a Siren network [206] and utilises the self-adaptive PINN (SA-PINN) loss weighting scheme [158].

#### 4.2.3.1 Hyperparameter Optimisation Studies

As in chapter 3, the establishment of a common model configuration for the three cases introduced in this chapter is important. Here, we keep spatio-temporal resolution and SNR fixed across cases, however the geometry, boundary motion and simulated flow are different. As in chapter 3, the Bayesian tree-structured Parzen estimator algorithm from the Python package Hyperopt [19] is again utilised here to perform optimisation studies. Since three distinct cases are considered here, the goal is not to establish an optimal model configuration for each geometry, but instead identify a range of common hyperparameter values that perform adequately in each case. In each case, the ground truth boundary condition was used so that an effective model configuration could be identified without considering the effects of positional uncertainty. Data were downsampled by a factor of 2 in both space and time, with a SNR of 10. An initial, exhaustive study was performed using case A, the search space of which can be found in Tab. 4.2. Given the similarities in problem setup between cases, it was assumed that a more restricted search space could be used for cases B and C, based on the exclusion of model configurations that performed poorly in the initial optimisation study. As such, the restricted search space can be found in Tab. 4.3. For the parameter  $\eta_p$ , the lower bound is also extended slightly, as is the upper bound of  $\eta_d$ .

**Table 4.2:** The complete search space identified for hyperparameter optimisation studies in case A. The parameter  $N_p$  here represents the percentage of the total collocation cloud used, with a maximum value of 160M points. The upper bound for  $N_p$ , selected to be 30%, was required to reduce the overall training time. Training runs performed above this value were prohibitively slow.

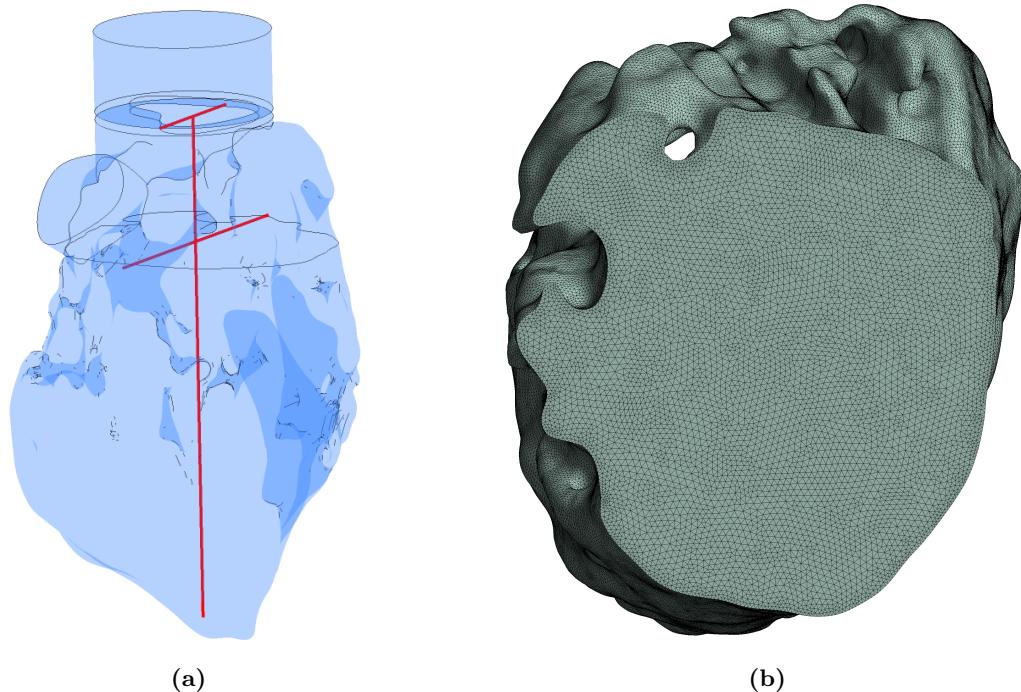
Complete search space	
Hidden layers	[5 .. 9]
Neurons per layer	[400 .. 800]
Batch size	[500 .. 8000]
Dropout rate	[0.1, 0.9]
Siren initialisation frequency ( $\omega_0$ )	[1 .. 30]
PDE loss SA-PINN learning rate ( $\eta_p$ )	[0.0001, 0.05]
Data loss SA-PINN learning rate ( $\eta_d$ )	[0.001, 0.1]
BC loss SA-PINN learning rate ( $\eta_b$ )	[0.0001, 0.05]
$N_p$ (%)	[1 .. 30]

**Table 4.3:** The restricted search space identified for hyperparameter optimisation studies in cases B and C.

Restricted search space	
Hidden layers	[8 .. 9]
Neurons per layer	[500 .. 800]
Batch size	[4000 .. 8000]
Dropout rate	[0.1, 0.9]
Siren initialisation frequency ( $\omega_0$ )	[1 .. 30]
PDE loss SA-PINN learning rate ( $\eta_p$ )	[0.00005, 0.01]
Data loss SA-PINN learning rate ( $\eta_d$ )	[0.05, 0.15]
BC loss SA-PINN learning rate ( $\eta_b$ )	[0.01, 0.05]
$N_p$ (%)	[20 .. 30]

#### 4.2.4 Mesh Independence Study

We evaluate mesh convergence in one case, namely case A, by analysing velocity and pressure errors at one time step along three lines of interest, as seen in Fig. 4.5a. The time step was chosen to be in late-diastole in the final completely simulated cardiac cycle. The most complex flow features are observed in late-diastole, and thus we would expect to find the greatest level of deviation here. Further, by selecting the final cardiac cycle, we eradicate any effects from poor initialisation. We produce results at three resolutions, given by fine (12,892,639 elements), medium (8,375,301) and coarse (3,771,507 elements) meshes. Given the similarity between cases, and by using an identical meshing procedure across each, it is sufficient to perform the mesh independence study once.



**Figure 4.5:** Lines of interest used to compare mesh results (left) and top-down slice displaying the medium mesh resolution (right).

In Tab. 4.4, we compare relative errors of both pressure and velocity for the coarse and medium meshes (against the fine mesh). The medium mesh obtains mean relative errors of 3.4% and 0.16 % for velocity and pressure, respectively, which falls within an acceptable tolerance of accuracy. In Fig. 4.6, we plot velocity and pressure profiles for each mesh across our three lines of interest. We find the profiles for medium and fine meshes match well across each line investigated. Given the similarity in results using the medium and fine meshes, we conclude that the solution obtained with the medium mesh is satisfactorily independent of the mesh resolution, and therefore will be used across all studies.

**Table 4.4:** Relative error (from fine mesh) of selected variables for coarse and medium meshes at three lines of interest.

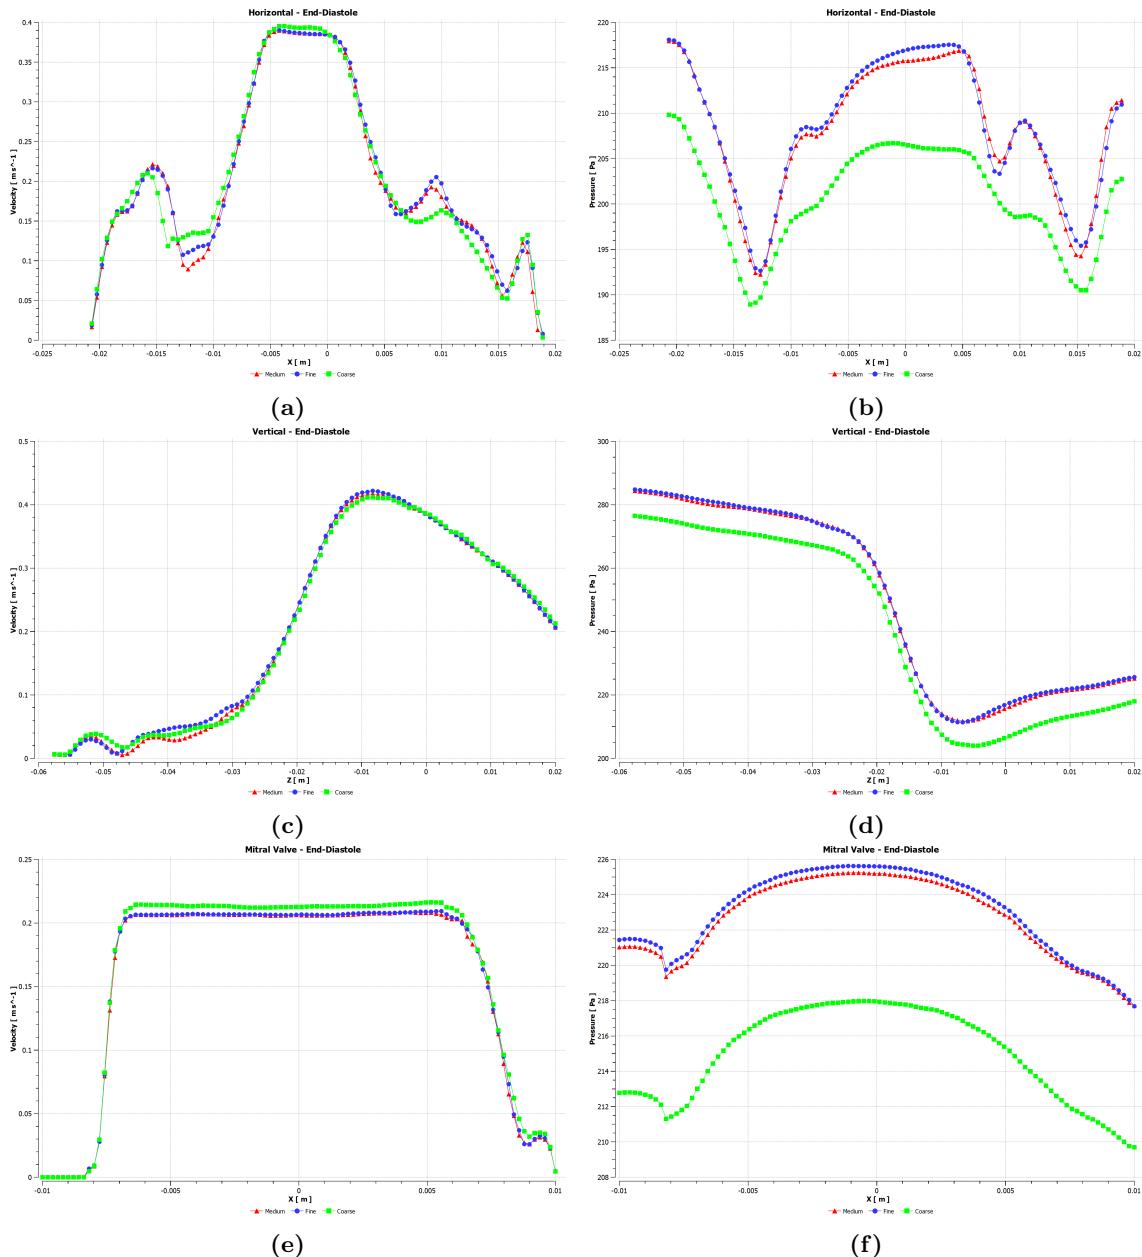
Variable	Mesh	Horizontal line	Vertical line	Mitral valve line	Mean
Velocity	Coarse	0.1485	0.0516	0.0384	0.0795
	Medium	0.0594	0.0077	0.0348	<b>0.034</b>
Pressure	Coarse	0.0381	0.0324	0.0358	0.0354
	Medium	0.0019	0.0016	0.0014	<b>0.0016</b>

## 4.3 Results

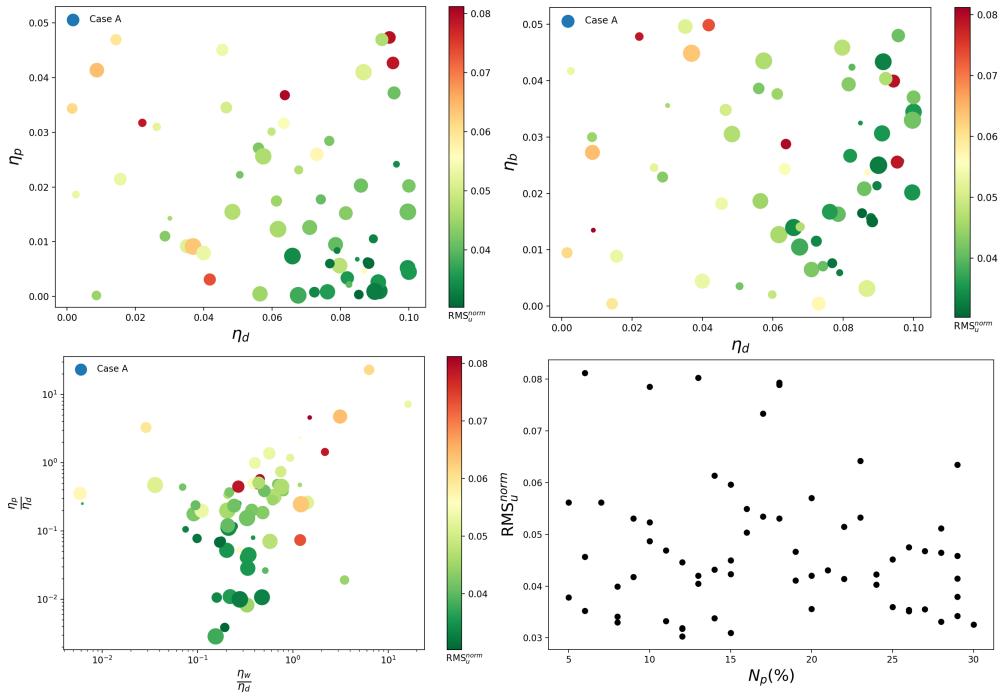
### 4.3.1 Optimisation Studies

Results for the extended hyperparameter optimisation study, using only case A, can be found in Fig. 4.7. As in chapter 3, the respective learning rates for each loss term optimiser in the SA-PINN scheme, given by  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ , played an important role in the accuracy and convergence of PINN model configurations. In Fig. 4.7a, a clear relationship between  $\eta_d$  and  $\eta_p$  emerges, where runs performed using high values of  $\eta_d$  and low values of  $\eta_p$  produced the most accurate results. This is corroborated in Fig. 4.7c, where it can be observed that  $\text{RMS}_u$  decreases as the ratio  $\eta_p/\eta_d$  decreases. The ratio between BC loss and data loss learning rates,  $\eta_b/\eta_d$ , also appears to play some role, although from Fig. 4.7b, the main influence is seemingly exerted by  $\eta_d$  as opposed  $\eta_b$ , given more pronounced variation in the horizontal axis than the vertical axis. There is no obvious trend when considering the total number of collocation points used, as seen in Fig. 4.7. The expected result would be that a greater number of collocation points produces the most accurate results, but that does not transpire, where the best three configurations use  $N_p \approx 10 - 15\%$ .

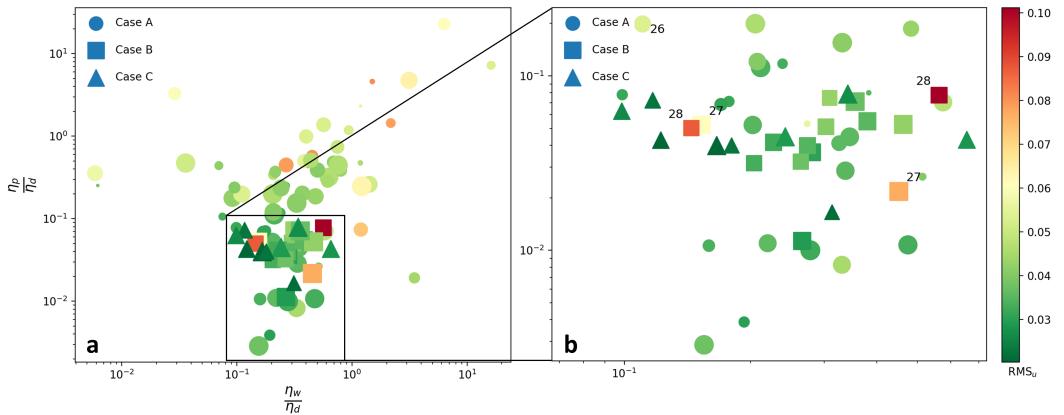
Results from the restricted optimisation study, where cases A, B and C are included, can be found in Fig. 4.8. Ratios between  $\eta_p$ ,  $\eta_d$  and  $\eta_b$  are again compared. As demonstrated by the reported  $\text{RMS}_u$  values across the three cases, the search space identified for the



**Figure 4.6:** Comparing velocity (left column) and pressure (right column) profiles of three mesh resolutions across each line of interest, defined in Fig. 4.5a.



**Figure 4.7:** Optimisation study results for case A, primarily examining the impact of the respective SA-PINN learning rates,  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ . Comparisons between  $\eta_d$  and  $\eta_p$  (a), and  $\eta_d$  and  $\eta_b$  (b) are shown. Analysis of the ratios between learning rates is performed (c), alongside the impact of collocation point set size,  $N_p$ . In (a-c), marker colour and size is defined by  $RMS_u$  and  $N_p$ , respectively.



**Figure 4.8:** Restricted optimisation study results for all cases, examining the impact of ratios between the respective SA-PINN learning rates,  $\eta_p$ ,  $\eta_d$  and  $\eta_b$ . Marker colour and size is defined by  $RMS_u$  and  $N_p$ , respectively, with reported values in (b) referring to  $\omega_0$ .

restricted study (as shown by the box in Fig. 4.8a) appears to define an appropriate range of  $\eta_p$ ,  $\eta_d$  and  $\eta_b$  values. There are some outliers for case B, however, the reported  $\omega_0$  values (which were only included for configurations with  $\omega_0 > 25$ ) explain this, where  $\omega_0$  values close to 30 appear to negatively affect results.

Above values of  $\sim 1000$ , batch size was found have minimal impact on performance, similarly for depth (above  $\sim 6$  layers) and width (above  $\sim 600$  neurons per layer).

The optimal configurations uncovered for individual cases A, B and C can be found in Tab. 4.5, alongside the shared configuration to be used across all three cases. The shared configuration was selected based on aggregating the five best performing runs in each individual study, instead of basing this choice entirely on the optimal configurations in each case.

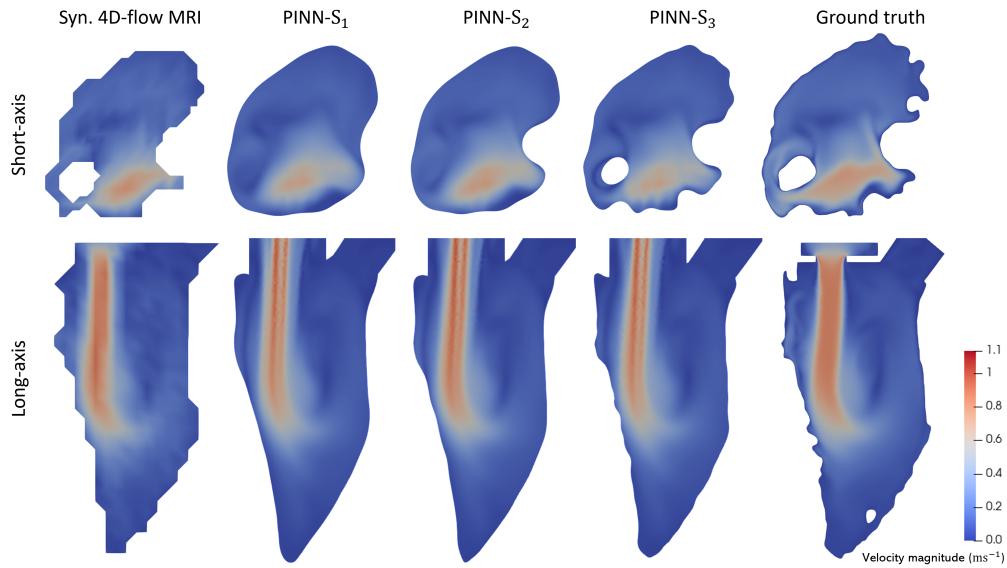
**Table 4.5:** Optimal individual hyperparameter configurations for cases A, B and C, alongside a shared configuration to be used across all cases.

Hyperparameter	Case A	Case B	Case C	Shared configuration
Hidden layers	6	8	8	6
Neurons per layer	687	746	785	750
Batch size	2885	5309	5180	4000
Dropout rate	0.71	0.47	0.776	0.65
$\omega_0$	19	7	13	12
$\eta_p$	0.0003	0.0015	0.0044	0.0045
$\eta_d$	0.0853	0.1319	0.1033	0.113
$\eta_b$	0.0165	0.0350	0.0126	0.024
$N_p$ (%)	12	25	24	20

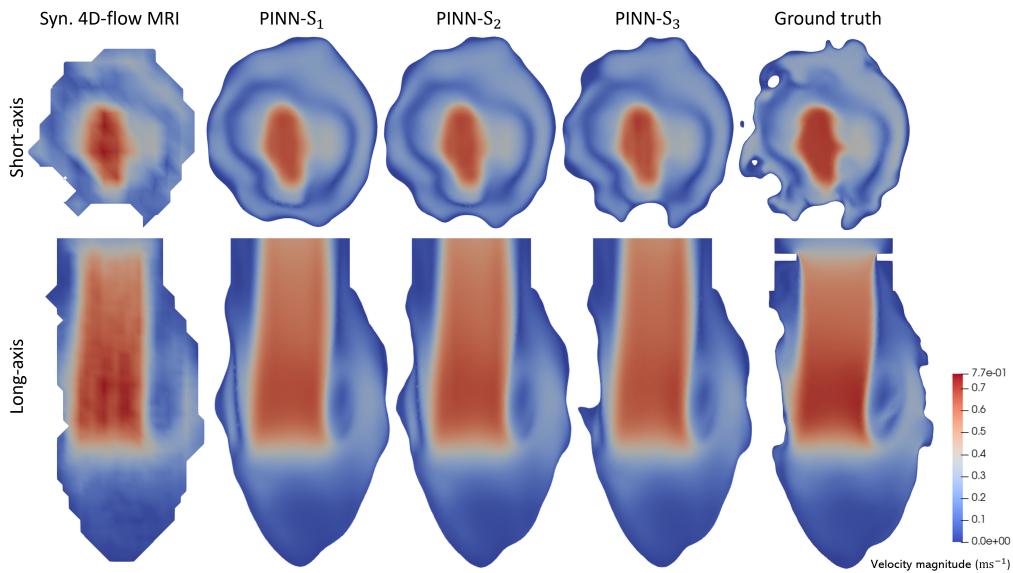
### 4.3.2 Clinically-Relevant Quantities

#### 4.3.2.1 Velocity Field

To assess the performance of our PINN model, we first analyse its ability to reconstruct the velocity field, as per chapter 3. We use  $\text{RMS}_u$ , ASI and MSI to assess results across the three boundary representations, and compare these against the synthetic 4D-flow MRI data. In Tab. 4.6, velocity errors are obtained using each PINN configuration for each case, alongside those derived directly from the synthetic 4D-flow MRI data. Error values for the synthetic 4D-flow MRI data are calculated through nearest neighbour interpolation of the high-resolution ground truth data onto the low-resolution grid, whereas PINN error metrics are obtained directly on the ground truth grid. Additionally, the PINN error metrics are obtained at all time points in the ground truth data, where the synthetic



**Figure 4.9:** Case A: Velocity magnitude contour plots in late-diastole at short-axis and long-axis planes. Fields derived from synthetic 4D-flow MRI (left), our three PINN configurations (centre) and the ground truth velocity data (right).



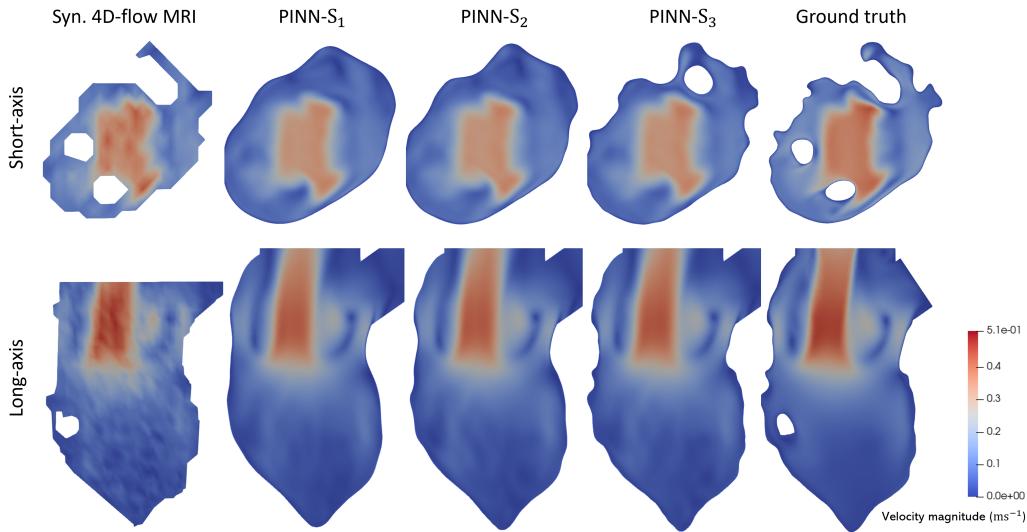
**Figure 4.10:** Case B: Velocity magnitude contour plots in late-diastole at short-axis and long-axis planes. Fields derived from synthetic 4D-flow MRI (left), our three PINN configurations (centre) and the ground truth velocity data (right).

4D-flow MRI data are downsampled by a factor of 4 in time. As such, 75% of the spatio-temporal coordinates at which PINN error metrics are evaluated are obtained at time points for which no flow data is present. This helps in explaining why the synthetic 4D-flow MRI data attains better scores for both similarity indices, although all PINN configurations achieve lower  $\text{RMS}_u$  values across the three cases.

In Figs. 4.9, 4.10 and 4.11, we compare PINN-derived velocity fields with the synthetic 4D-flow MRI data and the ground truth field. The PINN is able to effectively de-noise the velocity fields, although regions of peak velocity in the mitral jet are not fully resolved across all three cases. Flow near the valve region is also not accurately reconstructed by the PINN, since knowledge of the valve structure is assumed to be unavailable and thus the no-slip boundary condition is not applied to the valve leaflets. However, given that the flow reconstruction within the chamber is generally unaffected, such discrepancies around the valve are not important for the current study as the clinically-relevant quantities that are discussed in this chapter are measured deeper in the ventricle. In case A, PINN predictions across all boundary configurations exhibit streaking artefacts in the mitral jet.

As a consequence of the simplified endocardial motion, the mitral jet penetrates deeply into the ventricular chamber, particularly in Cases A and B. This arises due to the relatively slow expansion of the chamber, as prescribed by the sinusoidal function in this synthetic study. This causes an abrupt stalling of the inflow jet towards the apex of the ventricle. In reality, diastolic filling occurs in two distinct phases - E-wave and A-wave - which both occur, individually, over short time scales. During E-wave filling, the myocardium expands rapidly through relaxation to induce filling, with deceleration of the mitral jet occurring closer to the base of the ventricle. Similarly, during A-wave filling the left atrium contracts rapidly to force blood into the ventricle, again inducing deceleration closer to the base.

It is perhaps important to note that, while the ground truth flow data may not entirely reflect what is observed within the true LV, this likely does not invalidate the effectiveness of the PINN model. The key challenges for the model in this setting are super-resolution of flow across a range of length and time scales within heavily deforming domains, for which artificial positional uncertainty is introduced. These challenges are present within the ground truth flow data used here, and it is unlikely that the PINN model would not be applicable to data that were more realistic. Similar penetration of the mitral jet is observed, for instance, in experimental studies [232], which are still considered to be realistic.



**Figure 4.11:** Case C: Velocity magnitude contour plots in late-diastole at short-axis and long-axis planes. Fields derived from synthetic 4D-flow MRI (left), our three PINN configurations (centre) and the ground truth velocity data (right).

In Fig. 4.11, it is apparent that a number of holes are present through the ground truth slices. These are produced by the papillary muscles penetrating through the blood pool, and large-scale trabeculations present on the endocardial surface. This case exhibited a number of large trabeculations on the endocardial surface, which are reproduced in the reference segmentations as provided within the MM-WHS data set. The cause for such features is not apparent, however, it can be assumed that the impact on the ground truth flow data is similar to that of the papillary muscles.

**Table 4.6:** Error metrics for predicted velocity and pressure fields, alongside values derived directly from synthetic 4D-flow MRI. Reported RMS values are  $\text{RMS} \pm \text{one standard deviation of the relative error}$ , and reported SI values are  $\text{SI} \pm \text{one standard deviation of SI}$ .

Case	$\text{RMS}_u$	ASI	MSI
A	PINN- $S_1$	$0.068 \pm 0.053$	$0.881 \pm 0.192$
	PINN- $S_2$	$0.068 \pm 0.054$	$0.881 \pm 0.194$
	PINN- $S_3$	<b><math>0.065 \pm 0.051</math></b>	$0.886 \pm 0.190$
	Syn. 4D-flow MRI	$0.072 \pm 0.049$	<b><math>0.895 \pm 0.180</math></b>
B	PINN- $S_1$	$0.073 \pm 0.055$	$0.905 \pm 0.179$
	PINN- $S_2$	<b><math>0.072 \pm 0.054</math></b>	$0.907 \pm 0.175$
	PINN- $S_3$	$0.072 \pm 0.055$	$0.906 \pm 0.176$
	Syn. 4D-flow MRI	$0.089 \pm 0.057$	<b><math>0.918 \pm 0.153</math></b>
C	PINN- $S_1$	$0.073 \pm 0.052$	$0.909 \pm 0.176$
	PINN- $S_2$	$0.070 \pm 0.051$	$0.916 \pm 0.166$
	PINN- $S_3$	<b><math>0.068 \pm 0.050</math></b>	$0.916 \pm 0.169$
	Syn. 4D-flow MRI	$0.081 \pm 0.043$	<b><math>0.909 \pm 0.157</math></b>

### 4.3.2.2 Vorticity

Vorticity, defined as:

$$\boldsymbol{\omega} = \nabla \times \mathbf{u}, \quad (4.6)$$

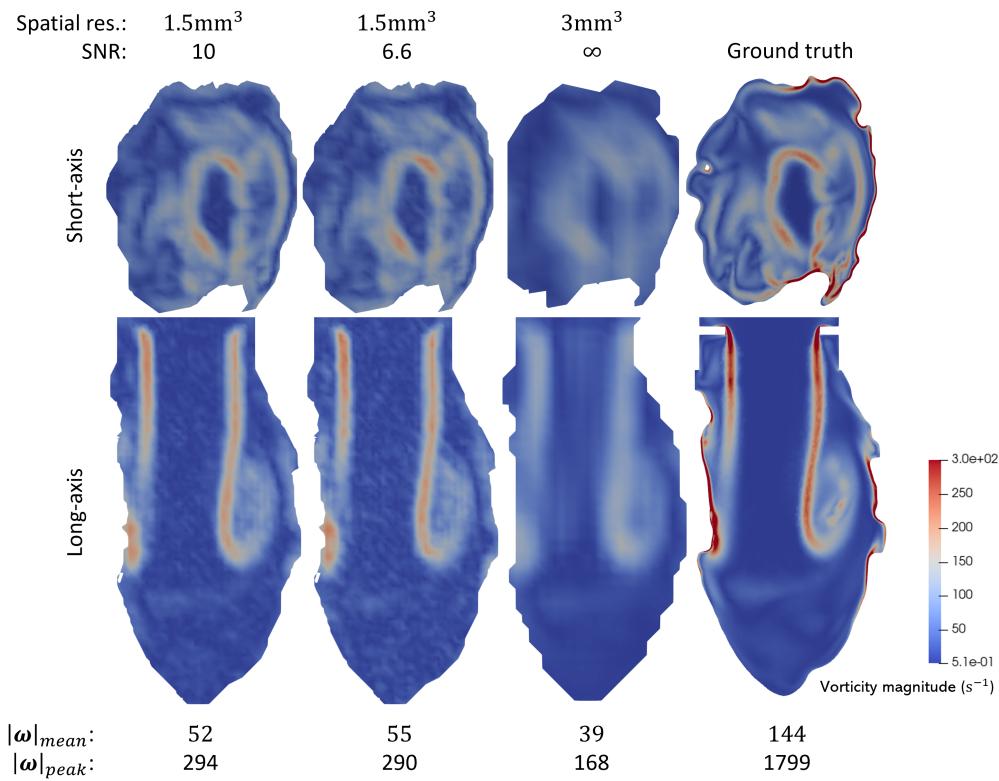
is a vectorial quantity that describes the rotational motion of a continuum of fluid, and has emerged in recent years as a variable of interest for cardiac haemodynamic analyses. Guided by studies in the literature, in particular those that utilised 4D-flow MRI, a selection of vorticity parameters at relevant time points were identified, which can be found in Tab. 4.7. Details on such studies and associated vorticity parameters can be found in section 2.2.3 in chapter 2.

**Table 4.7:** A list of vorticity ( $\boldsymbol{\omega}$ ) parameters analysed in this study with their respective descriptions.

$\text{SSIM}^t$	Structural similarity index against ground truth at time $t$
$\text{MS-SSIM}^t$	Multi-scale structural similarity index against ground truth at time $t$
$\text{MI}^t$	Mutual information against ground truth at time $t$
$\text{RMS}_u^t$	Max-normalised root mean square error against ground truth at time $t$
$ \boldsymbol{\omega} _{mean}^t$	Mean vorticity at time $t$
$ \boldsymbol{\omega} _{peak}^t$	Peak vorticity at time $t$
$\varepsilon_C$	Absolute error of vortex core centre point from ground truth
$\varepsilon_\theta$	Absolute error of vortex core orientation from ground truth

As vorticity is calculated by taking spatial derivatives of the velocity field, it is highly susceptible to low spatial resolution and noise corruption in the measured velocity field [42]. As a consequence of the former, averaging of flow values across voxels results in less sharp changes in velocity, and thus lower magnitude spatial gradients. As such, vorticity values derived directly from 4D-flow MRI data are likely to be under-predicted. This effect is visualised in Fig. 4.12, where it can be observed that the synthetic 4D-flow MRI data with the lowest spatial resolution produces reconstructions of the lowest quality, with significant under-prediction of vorticity. This is apparent qualitatively, but is also reflected in the quantitative values. The influence of noise is less pronounced, as observed across the other two data configurations. Reduction of the SNR to 6.6 from 10 results in slightly increased mean vorticity, but coincides with a decrease in peak vorticity.

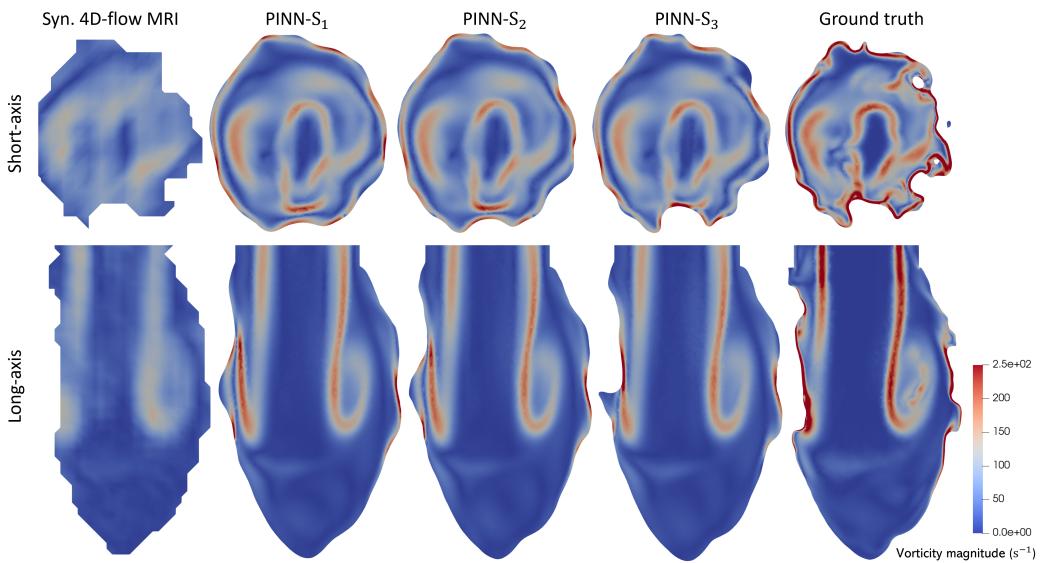
In Tabs. 4.8 and 4.9, vorticity errors across all cases are quantified at four distinct phases of the cardiac cycle for the three PINN configurations and the synthetic 4D-flow MRI data, the latter of which is linearly interpolated onto the ground truth mesh for comparison. All



**Figure 4.12:** Case B: Demonstrating the impact of SNR and spatial resolution on vorticity reconstruction. Contour plots of vorticity magnitude in late-diastole at short-axis and long-axis slices, derived from synthetic 4D-flow MRI data at three distinct downsampling configurations, alongside ground truth vorticity fields.

**Table 4.8:** Max-normalised vorticity RMSE at 4 time points, namely mid systole ( $T^{MS}$ ), early diastole ( $T^{ED}$ ), mid diastole ( $T^{MD}$ ), and late diastole ( $T^{LD}$ ), comparing performance of each PINN configuration against the synthetic 4D-flow MRI data, which here is interpolated onto the ground truth mesh to provide error values.

Case	$RMS_v^{MS}$	$RMS_v^{ED}$	$RMS_v^{MD}$	$RMS_v^{LD}$
A	PINN- $S_1$	0.0504	0.0479	0.0518
	PINN- $S_2$	0.0507	0.0479	0.0517
	PINN- $S_3$	<b>0.0498</b>	<b>0.0478</b>	<b>0.0515</b>
	Syn. 4D-flow MRI	0.0531	0.0492	0.0537
B	PINN- $S_1$	0.0681	0.0514	0.0501
	PINN- $S_2$	0.0674	<b>0.0513</b>	0.0499
	PINN- $S_3$	<b>0.0671</b>	0.0514	<b>0.0498</b>
	Syn. 4D-flow MRI	0.0708	0.0547	0.0525
C	PINN- $S_1$	0.0675	0.0477	0.0556
	PINN- $S_2$	0.0667	0.0475	0.055
	PINN- $S_3$	<b>0.0665</b>	<b>0.0474</b>	<b>0.0546</b>
	Syn. 4D-flow MRI	0.0681	0.0511	0.0577



**Figure 4.13:** Case B: Contour plots of vorticity magnitude in late-diastole at short-axis and long-axis slices, derived directly from the synthetic 4D-flow MRI data (left), alongside the ground truth vorticity data (right) and PINN predictions across three boundary configurations (centre).

PINN configurations provide more accurate reconstructions than direct measurement from synthetic 4D-flow MRI, as is demonstrated with each error metric. The most pronounced difference can be observed for the MS-SSIM metrics, where significant improvement over the synthetic data is observed. Both of these metrics assess global differences between structures within each field, and therefore, slight misalignment of the vortices are not as heavily penalised as with point-wise metrics like RMS. Of the four phases of the cardiac cycle investigated, the highest errors are attained in late-diastole, for both PINN predictions and synthetic 4D-flow MRI. This is, perhaps, to be expected, where breakup and dissipation of the mitral vortex ring produces small-scale coherent structures that are challenging to reconstruct to the level of detail present in CFD results, as observed in Fig. 4.13.

In Tab. 4.10, results from Tabs. 4.8 and 4.9 are condensed to provide mean and standard deviation values across all time steps and cases. Using T-tests, shared p-values for each error metric, across all PINN configurations, range from 0.27 to 0.92, demonstrating minimal variation in accuracy across each boundary representation. Similar analysis comparing collated PINN results with synthetic 4D-flow MRI produces p-values well under 0.05 for all metrics except  $RMS_v$ , which produces a p-value of 0.43. As discussed above, point-wise error quantification on gradient fields with sharp features and large peak values heavily penalises misalignment of features, and thus the PINN configurations generally perform poorly with  $RMS_v$ . These results demonstrate that significantly improved vorticity results can be obtained using the PINN model. Further, minor variation in results across PINN

boundary configurations suggests that positional uncertainty in the endocardium plays a minimal role in vorticity reconstruction.

Qualitatively, the PINN predictions are of a higher quality than direct calculation of vorticity from the synthetic 4D-flow MRI data, as can be seen in Fig. 4.13. Although peak vorticity values are visually not captured, the vortical structures present in the flow field are more refined and are present in the correct regions of the LV. Such under-estimation of peak vorticity is reflected quantitatively in Tab. 4.11, where peak and mean vorticity values are produced at three phases of the cardiac cycle. Across all PINN configurations in each case, both peak and mean vorticity are under-predicted (except PINN predictions of peak vorticity for case C in mid-systole and early-diastole), with significant variation in peak vorticity across PINN boundary configurations. Mean vorticity is predicted more consistently across PINN boundary representations, with relative errors ranging from 0.6% to 36%. For comparison, results obtained through direct calculation of vorticity from synthetic 4D-flow MRI are under-predicted by as much as two orders of magnitude.

**Table 4.9:** Vorticity MS-SSIM at 4 time points, namely mid systole ( $T^{MS}$ ), early diastole ( $T^{ED}$ ), mid diastole ( $T^{MD}$ ), and late diastole ( $T^{LD}$ ), comparing performance of each PINN configuration against the synthetic 4D-flow MRI data, which here is interpolated onto the ground truth mesh to provide error values.

Case		MS-SSIM <sup>MS</sup>	MS-SSIM <sup>ED</sup>	MS-SSIM <sup>MD</sup>	MS-SSIM <sup>LD</sup>
A	PINN- $S_1$	0.769	0.781	0.78	0.731
	PINN- $S_2$	0.778	0.793	0.813	0.736
	PINN- $S_3$	<b>0.793</b>	<b>0.816</b>	<b>0.822</b>	<b>0.754</b>
	Syn. 4D-flow MRI	0.638	0.695	0.729	0.613
B	PINN- $S_1$	0.821	0.86	0.857	<b>0.861</b>
	PINN- $S_2$	0.821	0.864	<b>0.858</b>	0.859
	PINN- $S_3$	<b>0.822</b>	<b>0.871</b>	<b>0.858</b>	0.853
	Syn. 4D-flow MRI	0.682	0.744	0.763	0.717
C	PINN- $S_1$	0.871	0.806	0.824	0.798
	PINN- $S_2$	0.871	0.811	0.832	0.804
	PINN- $S_3$	<b>0.877</b>	<b>0.839</b>	<b>0.851</b>	<b>0.824</b>
	Syn. 4D-flow MRI	0.762	0.678	0.754	0.695

**Table 4.10:** Vorticity error metrics across all evaluated time points, calculated for individual configurations across all three cases, alongside collated results for all configurations and those derived directly from synthetic 4D-flow MRI. Reported values are mean  $\pm$  one standard deviation, accounting for variation across cases and time steps.

Configuration	RMS <sub>v</sub>	MI	SSIM	MS-SSIM
PINN- $S_1$	$0.059 \pm 0.011$	$1.225 \pm 0.069$	$0.739 \pm 0.047$	$0.813 \pm 0.042$
PINN- $S_2$	$0.058 \pm 0.010$	$1.234 \pm 0.055$	$0.745 \pm 0.044$	$0.820 \pm 0.038$
PINN- $S_3$	$0.058 \pm 0.010$	$1.230 \pm 0.055$	$0.752 \pm 0.040$	$0.832 \pm 0.033$
PINN <sub>all</sub>	$0.058 \pm 0.010$	$1.230 \pm 0.060$	$0.746 \pm 0.044$	$0.822 \pm 0.039$
Syn. 4D-flow MRI	$0.061 \pm 0.011$	$1.147 \pm 0.062$	$0.623 \pm 0.047$	$0.706 \pm 0.046$

Such large discrepancies between predicted and ground truth mean and peak vorticity

**Table 4.11:** Peak and mean vorticity values in mid systole ( $T^{MS}$ ), early diastole ( $T^{ED}$ ) and mid diastole ( $T^{MD}$ ), for each PINN configuration, synthetic 4D-flow MRI and ground truth data.

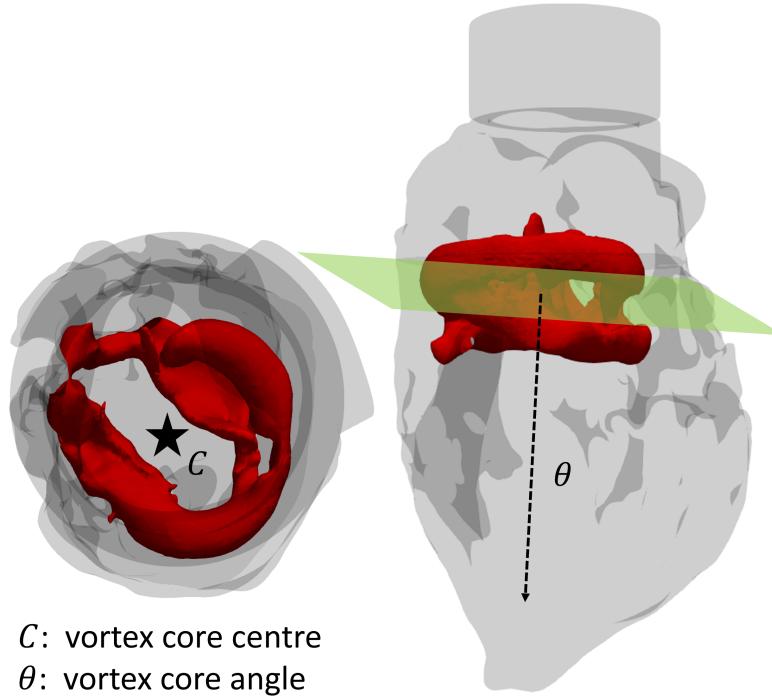
Case		$ \omega _{peak}^{MS}$	$ \omega _{mean}^{MS}$	$ \omega _{peak}^{ED}$	$ \omega _{mean}^{ED}$	$ \omega _{peak}^{MD}$	$ \omega _{mean}^{MD}$
A	PINN- $S_1$	519.7	34.7	244.7	22.2	519.3	48.3
	PINN- $S_2$	<b>970.5</b>	<b>39.4</b>	<b>418</b>	24.1	<b>1905.3</b>	59.5
	PINN- $S_3$	391.8	38.3	309.3	<b>24.8</b>	865.2	<b>61.4</b>
	Syn. 4D-flow MRI	98	23.6	122.4	6.6	216.3	43.5
	Ground truth	1638.6	75.6	659.4	43.9	2747.7	131.1
B	PINN- $S_1$	<b>865.8</b>	42.4	<b>595.8</b>	<b>30.2</b>	<b>869.6</b>	<b>52.2</b>
	PINN- $S_2$	455.5	<b>58.8</b>	365.1	28.8	791.2	49.3
	PINN- $S_3$	313	37.3	218.8	28	777.2	49
	Syn. 4D-flow MRI	143.2	25.6	124.5	19.4	163.4	38.2
	Ground truth	1107.1	73.1	686.5	50.7	1463.1	112.2
C	PINN- $S_1$	478.5	<b>20.5</b>	634.1	<b>28.4</b>	<b>1349.5</b>	<b>38.8</b>
	PINN- $S_2$	413.5	18.2	370.2	24.6	992.9	32.9
	PINN- $S_3$	<b>394.8</b>	17.4	<b>474</b>	26.7	966.7	33.9
	Syn. 4D-flow MRI	85	6.6	49.8	6	150.7	11.9
	Ground truth	350.6	29.4	448	51.2	1661.2	71.4

values can be partially explained by analysing the differences in the endocardial boundary. As can be seen in the short-axis plots in Fig. 4.13, vorticity values with the greatest magnitude are attained near the flow boundaries. As such, peak and mean vorticity are highly dependent on the position of the boundary and, in particular, the tangential velocity of the fluid in adjacent regions. The ground truth CFD data generated with the real endocardial boundaries experience interactions between high-velocity flow and the wall, generating large velocity gradients. Conversely, the approximated boundaries used by the PINN models typically lie outside of the masked flow domain, and thus no synthetic 4D-flow MRI measurements exist in close proximity to the wall. As such, near-wall flow is generally of a lower velocity, producing smaller gradients. There are exceptions to this, which can be seen in near-wall regions that share some topology with the ground truth endocardium in Fig. 4.13.

The centre point and orientation of the mitral vortex ring in diastole has been explored previously as a marker for cardiac disease [64, 63, 54]. Here, we investigate errors in the vortex core centre point,  $\varepsilon_c$ , and angle of penetration,  $\varepsilon_\theta$ , across PINN predictions and measurements made directly from synthetic 4D-flow MRI data. The vortex cores are formed using the Q-criterion, which is defined as

$$Q_{crit} = \frac{1}{2} (||\bar{\Omega}||^2 - ||S||^2), \quad (4.7)$$

for vorticity tensor  $\bar{\Omega}$  and strain rate tensor  $S$ . It is used in the field of fluid dynamics to visualise coherent structures, which can be considered vortices [113], since it isolates regions of flow where either rotational or viscous effects dominate. To isolate the vortex ring, an appropriate  $Q$ -criterion value is first selected so as to maximise the vortex ring volume whilst removing any trailing features and maintaining a fully connected ring. Whilst the volume of the vortex ring is sensitive to the choice of  $Q_{crit}$ ,  $c$  and  $\theta$  are less so. Further, since we are considering the error in  $c$  and  $\theta$  between model predictions and ground truth data, the selected  $Q_{crit}$  values are consistent across all prediction types. For cases A and B, a value of  $Q_{crit} = 1000$  was used, whereas for case C, a value of  $Q_{crit} = 500$  was used. Then, all coherent structures in the domain except the largest connected region are discarded, leaving only the vortex ring. For ground truth isosurfaces, some manual clipping is also required to remove trailing features that intersect the vortex ring. Once the ring is isolated, the centre point,  $C$ , is identified as the centre of mass of the structure. The angle of penetration,  $\theta$ , is defined by the normal vector direction to the best fit plane through the vortex ring, which is constructed using singular value decomposition. In Fig. 4.14, a visual explanation of the parameters  $C$  and  $\theta$  is provided.

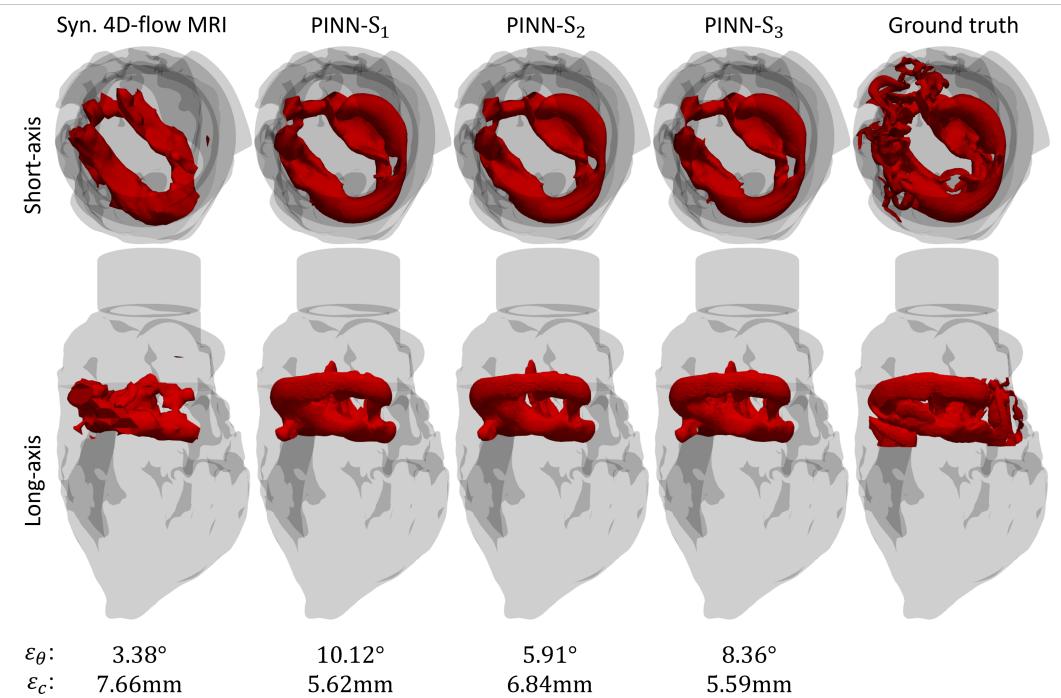


**Figure 4.14:** Two variables of interest in vortex core analysis, vortex core centre position and angle of penetration. Instead of calculating values for  $C$  and  $\theta$ , we instead only analyse errors in both measurements.

In the literature [64, 63, 54],  $C$  and  $\theta$  are defined relative to a cylindrical cardiac coordinate system, allowing for comparison and analysis across cases. However, here we are only interested in the accuracy of such measurements, and thus we only analyse the error

between the measured vortex core parameters and the ground truth values, given by  $\varepsilon_c$  and  $\varepsilon_\theta$ . In Fig. 4.15, qualitative and quantitative assessment of vortex ring formation is performed for synthetic 4D-flow MRI data, PINN predictions and ground truth flow data in case C in mid-diastole. Visually, delineation of the vortex ring directly from synthetic 4D-flow MRI data appears a feasible approach. However, from  $\varepsilon^{GT}$  values in Tab. 4.12, errors in the vortex centre position exceed those of all PINN predictions in each case. Additionally, errors in the angle of penetration are reduced using the PINN model in cases A and B (except for PINN- $S_3$  in case A), although they are increased in case C. Vortex ring orientation can vary across subjects by as much as  $83^\circ$  [63], and as such, errors of  $\sim 10^\circ$ , although not insignificant, are relatively small. Further, extraction of an isolated vortex ring from the ground truth data is challenging, since there exist many smaller-scale coherent structures that intersect the main vortex core. As such, manual clipping of the isosurfaces is required, which cannot completely remove all trailing features, imparting some influence on the measured vortex centre position and angle of penetration. When isosurfaces of the Q-criterion are derived from *in vivo* 4D-flow MRI data, using either direct measurement or a data assimilation model, the interference of small-scale coherent structures is typically less problematic. Such structures arise in CFD data as a consequence of its high spatial resolution, which allows for the reconstruction of flow features on very small scales. Voxel averaging effects, in both space and time, remove most of the small-scale flow features present in the velocity field with *in vivo* flow imaging, and although some recovery of these features is possible with the PINN model, most of the smaller coherent structures are not reconstructed. As such, discrepancies in vortex position and angle arising due to the presence of trailing coherent structures in the ground truth data are isolated to studies comparing synthetic 4D-flow MRI data with CFD data, such as this.

To assess variability of results across the three boundary representations, errors in the position and angle of the mitral vortex ring are also derived from the mean values obtained across each PINN configuration, which can be found in the  $\varepsilon^{mean}$  values in Tab. 4.12. Good agreement in vortex ring position is achieved across the three PINN configurations used in each case, with more significant positional errors arising with results derived from the synthetic 4D-flow MRI data. Further, errors in the angle of penetration are reduced across all PINN configurations when compared with the synthetic 4D-flow MRI data. As such, there is minimal variation across the PINN boundary configurations used, suggesting that positional uncertainty in the endocardial boundary condition does not significantly influence reconstruction of the mitral vortex ring.



**Figure 4.15:** Case C: Vortex core formation in mid-diastole. Isosurfaces of  $Q_{crit} = 500$  delineate the mitral vortex ring in the synthetic (syn.) 4D-flow MRI (left), PINN predictions (centre) and ground truth data (right). Absolute errors in vortex core centre point,  $\varepsilon_c$ , and angle of penetration,  $\varepsilon_\theta$ , are displayed underneath each figure.

**Table 4.12:** Absolute errors in the position and angle of penetration of the mitral vortex ring, calculated from the ground truth ( $\varepsilon^{GT}$ ) mean values across PINN configurations ( $\varepsilon^{mean}$ ).

Case		$\varepsilon_c^{GT}$ (mm)	$\varepsilon_\theta^{GT}$ (°)	$\varepsilon_c^{mean}$ (mm)	$\varepsilon_\theta^{mean}$ (°)
A	PINN-S <sub>1</sub>	1.53	3.67	0.29	2.49
	PINN-S <sub>2</sub>	1.70	5.73	0.52	6.50
	PINN-S <sub>3</sub>	1.85	11.0	0.35	8.99
	Syn. 4D-flow MRI	2.67	10.2	2.89	13.1
B	PINN-S <sub>1</sub>	2.66	4.60	0.86	1.95
	PINN-S <sub>2</sub>	1.67	4.80	1.19	1.23
	PINN-S <sub>3</sub>	3.14	2.61	0.81	2.72
	Syn. 4D-flow MRI	4.36	11.5	4.28	8.34
C	PINN-S <sub>1</sub>	5.62	10.12	0.51	2.48
	PINN-S <sub>2</sub>	6.84	5.91	1.11	4.10
	PINN-S <sub>3</sub>	5.59	8.36	0.61	2.41
	Syn. 4D-flow MRI	7.66	3.38	4.56	4.59

### 4.3.2.3 Kinetic Energy

Kinetic energy (KE) is defined as:

$$\text{KE} = \frac{1}{2} \int_v \rho |\mathbf{u}|^2, \quad (4.8)$$

for fluid of density  $\rho$  in volume  $v$ . It has been explored extensively in the literature in relation to cardiac disease, details of which can be found in section 2.2.3 of chapter 2.

Based on findings from the literature, a collection of KE parameters were identified for study in this chapter, given in Tab. 4.13. Note that instead of peak E-wave (early diastolic filling) and A-wave (atrial contraction) KE, which have been investigated in previous publications, only the peak diastolic KE ( $KE_{peak}^D$ ) is included here. This is because the simplified LV motion in our synthetic cases does not include distinct E-wave and A-wave phases. In-plane KE (IPKE) is defined as the proportion of KE that is directed in the short-axis direction (defined as a percentage), either at a particular cardiac phase or across the whole cardiac cycle [84, 53, 48]. This is an important variable of interest as KE should be primarily directed in the apico-basal direction for efficient pumping, with deviations from this arrangement indicative of pathological flow.

**Table 4.13:** A list of kinetic energy (KE) parameters analysed in this study with their respective descriptions. For simplicity, all parameters are normalised by the fluid mass.

$KE_{min}$	Minimum KE value across cardiac cycle
$KE_{mean}$	Mean KE value across cardiac cycle
$IPKE$	Proportion of in-plane KE to total KE
$KE_{mean}^S$	Mean systolic KE
$KE_{mean}^D$	Mean diastolic KE
$KE_{peak}^D$	Peak diastolic KE

In Tab. 4.14, predicted KE parameter values for each PINN boundary configuration are compared against those extracted directly from synthetic 4D-flow MRI data. The PINN model reconstructs each parameter to a higher level of accuracy than synthetic 4D-flow MRI data for all KE parameters, except  $KE_{peak}^D$  for case C. Aside from  $KE_{mean}^D$  in case C, the values extracted directly from the 4D-flow MRI data significantly overestimate the mean KE. This irregularity can be attributed to the presence of noise in the synthetic 4D-flow MRI data, where near-zero flow regions in the ground truth data are corrupted by non-zero velocity noise when downsampled to generate the 4D-flow MRI data. This effect can also be seen in Fig. 4.16, where overall (a-c) and mean KE (d-f) are far higher

for the synthetic 4D-flow MRI data than PINN predictions and ground truth. In this context, the de-noising capabilities of the PINN model are demonstrated by the improved accuracy of mean KE parameters.

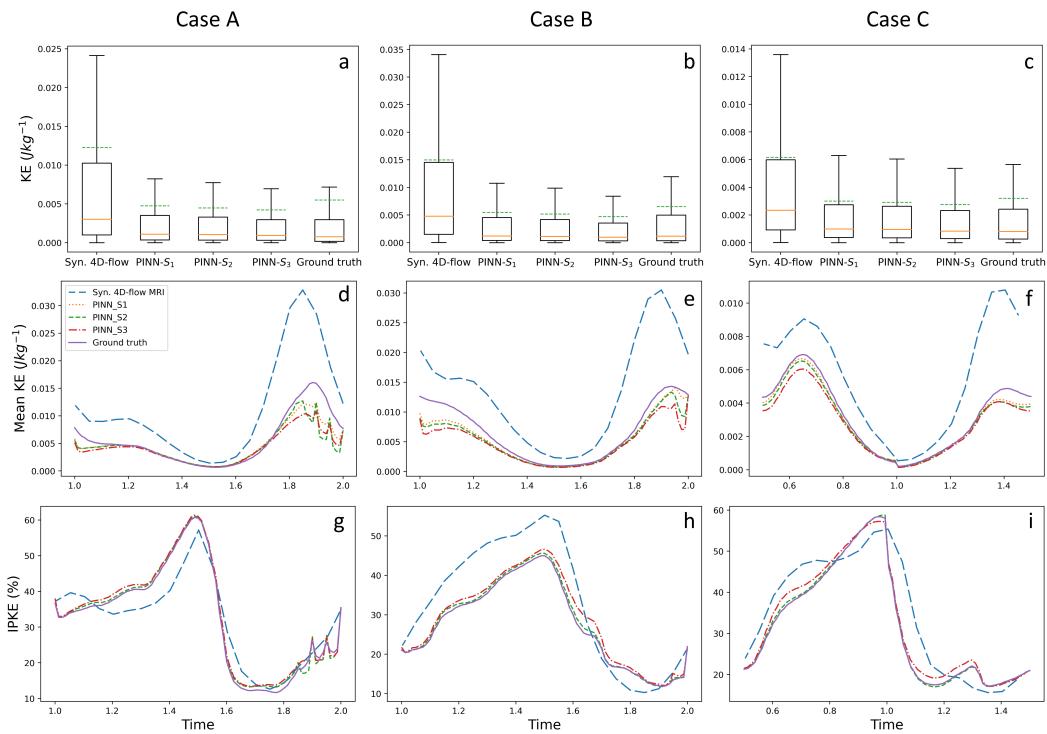
In Fig. 4.16, we see that the PINN model provides significant improvement in overall KE characteristics (a-c), mean KE throughout the cardiac cycle (d-f) and IPKE (g-i) compared to direct quantification from synthetic 4D-flow MRI. The plots shown in Fig. 4.16 demonstrate only small variations between the three PINN boundary configurations, particularly when compared to the difference between PINN results and synthetic 4D-flow MRI data. Across all three cases, mean, peak and median KE decreases slightly from PINN  $S_1$ - $S_3$ , with the inverse trend demonstrated for in-plane KE. However, differences are small, and it can be deduced that positional uncertainty in the endocardium does not heavily influence KE quantification. Conversely, KE parameters derived directly from synthetic 4D-flow MRI are highly susceptible to noise corruption. It is unclear exactly what relationship governs the interplay between SNR and KE in 4D-flow MRI data, but removing this dependency using the PINN model improves the repeatability of results across cases.

**Table 4.14:** Quantification of kinetic energy (KE) parameters derived using the three PINN configurations, synthetic 4D-flow MRI and ground truth data.

Case		KE <sub>min</sub>	KE <sub>mean</sub>	IPKE	KE <sup>S</sup> <sub>mean</sub>	KE <sup>D</sup> <sub>mean</sub>	KE <sup>D</sup> <sub>peak</sub>
A	PINN- $S_1$	$2.7 \times 10^{-9}$	<b>0.00475</b>	0.514	<b>0.00439</b>	<b>0.00592</b>	0.404
	PINN- $S_2$	$1.004 \times 10^{-8}$	0.00448	0.525	0.00425	0.0052	0.4
	PINN- $S_3$	<b><math>1.12 \times 10^{-9}</math></b>	0.00423	<b>0.538</b>	0.00391	0.00528	<b>0.397</b>
	Syn. 4D-flow MRI	$9.151 \times 10^{-7}$	0.0114	0.5	0.00743	0.0132	0.417
	Ground truth	$1.29 \times 10^{-14}$	0.0055	0.613	0.00506	0.00692	0.322
B	PINN- $S_1$	$4.231 \times 10^{-9}$	<b>0.00547</b>	0.444	<b>0.00508</b>	<b>0.00677</b>	0.233
	PINN- $S_2$	$1.014 \times 10^{-8}$	0.00518	0.457	0.00487	0.00619	0.238
	PINN- $S_3$	<b><math>8.07 \times 10^{-10}</math></b>	0.00472	<b>0.482</b>	0.00451	0.00537	<b>0.248</b>
	Syn. 4D-flow MRI	$2.765 \times 10^{-6}$	0.0133	0.468	0.0121	0.0126	0.295
	Ground truth	$6.199 \times 10^{-11}$	0.00654	0.496	0.00636	0.00713	0.245
C	PINN- $S_1$	$5.68 \times 10^{-9}$	<b>0.003</b>	0.432	<b>0.00274</b>	<b>0.00388</b>	0.117
	PINN- $S_2$	$1.91 \times 10^{-8}$	0.00292	0.434	0.00266	0.0038	0.117
	PINN- $S_3$	<b><math>2.25 \times 10^{-9}</math></b>	0.00275	<b>0.46</b>	0.00252	0.00355	0.116
	Syn. 4D-flow MRI	$1.837 \times 10^{-6}$	0.0056	0.436	0.00615	0.0038	<b>0.135</b>
	Ground truth	$1.812 \times 10^{-12}$	0.0032	0.476	0.00294	0.00407	0.132

#### 4.3.2.4 Wall Shear Stress

Wall shear stress (WSS) describes the tangential force applied to a surface through viscous shearing of the fluid in the boundary layer, and is defined for a Newtonian fluid with

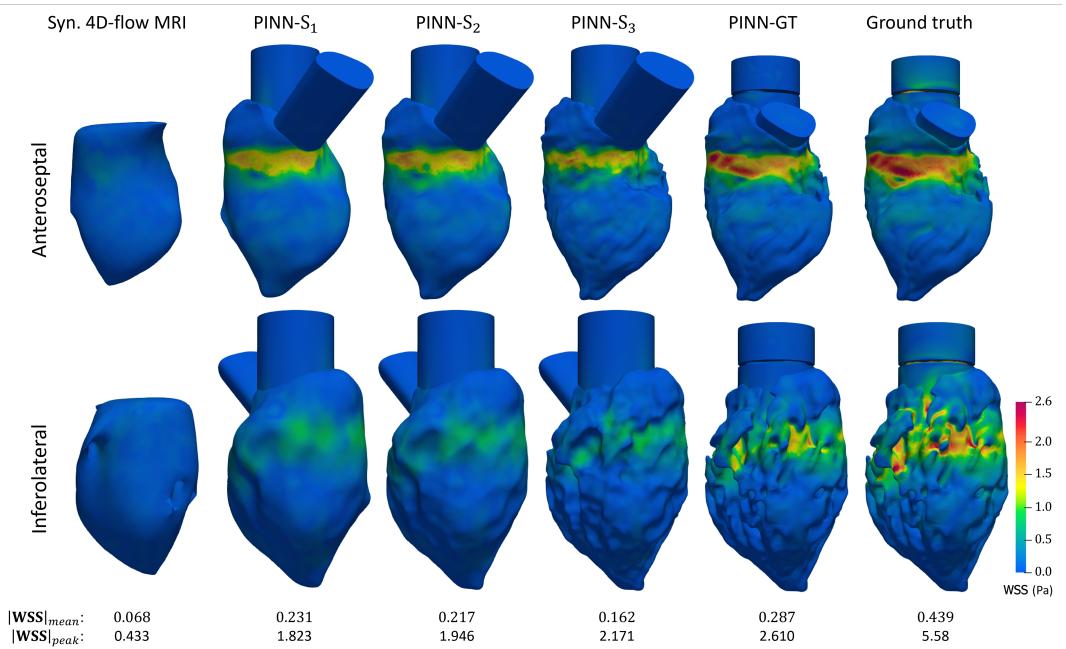


**Figure 4.16:** Comparison of kinetic energy (KE) characteristics of synthetic 4D-flow MRI data, PINN results and ground truth data for cases A, B and C, from left to right. (a - c) box and whisker plots for point-wise KE across whole cardiac cycle, where the boxes represent the first and third quartiles, whiskers display maximum and minimum values (with outliers removed), the solid orange line represents the median value and the green dashed line represents the mean value. (d - f) mean KE throughout the cardiac cycle. (g - i) in-plane KE throughout the cardiac cycle.

viscosity  $\mu$  as:

$$\mathbf{WSS} = 2\mu\tilde{\mathbf{S}} \cdot \hat{\mathbf{n}}, \quad (4.9)$$

for strain-rate tensor  $\tilde{\mathbf{S}}$  and inward wall-normal vector  $\hat{\mathbf{n}}$ . WSS is a highly sensitive variable that requires the accurate calculation of spatial velocity derivatives in the immediate vicinity of the wall region, and therefore high spatial resolution in the near-wall region is desirable for accurate reconstruction. This requirement proves problematic when attempting to approximate WSS using 4D-flow MRI, where low spatial resolution can significantly hamper accuracy [71]. Accurate reconstruction of the boundary layer in the near-wall region is crucial in calculating realistic velocity gradients, however, this flow feature is typically far smaller than the minimum voxel volume available to 4D-flow MRI, and thus not properly resolved. This effect is further exaggerated at higher velocities, where the thickness of the boundary layer is further reduced and voxel averaging becomes more problematic [42]. To negate the impact of low spatial resolution, interpolation techniques, such as parabolic curve fitting, have been investigated to obtain sub-voxel velocity gradient approximations in the near-wall region [220, 213, 41]. When compared with CFD results in corresponding geometries, however, significant differences in WSS magnitude were still noted, although overall the WSS distributions were generally captured appropriately [214].



**Figure 4.17:** Case C: WSS magnitude on the endocardium in late-diastole, derived directly from synthetic 4D-flow MRI data (left), from PINN results using four distinct boundary representations (centre) and CFD ground truth (right). PINN-GT utilises the exact ground truth boundary during training.

Despite the magnitude of WSS being under-predicted in *in vivo* imaging studies, an understanding of the role it plays in the progression of vascular disease is relatively well-

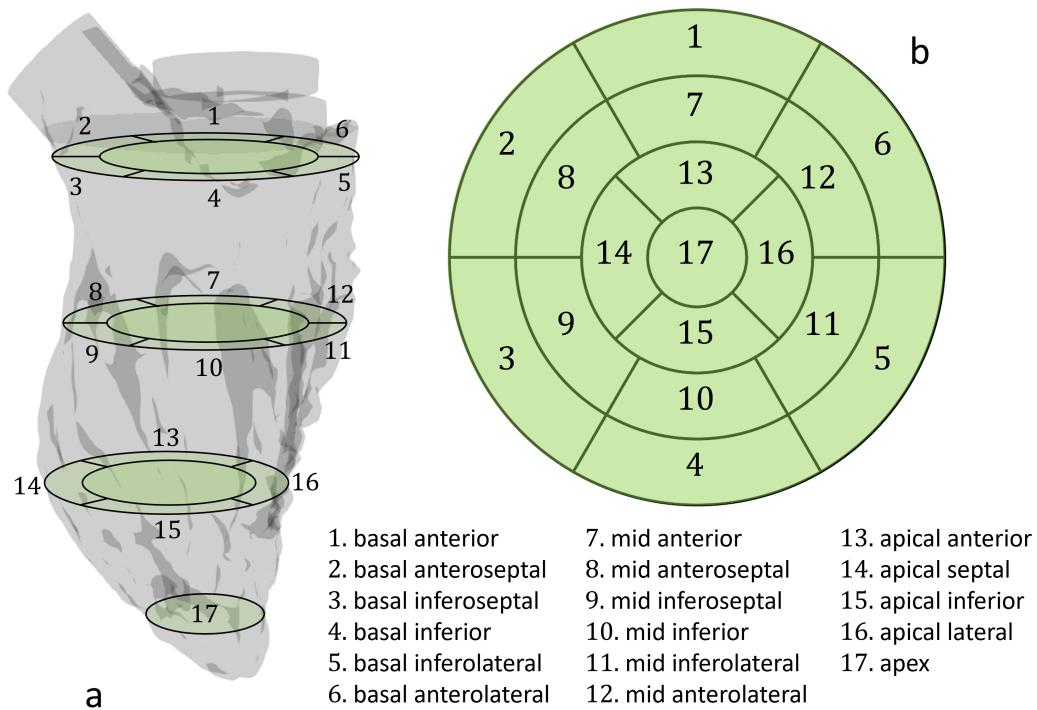
established, where changes in the intensity or oscillatory characteristics of WSS have been shown to trigger a response from endothelial cells by means of mechanotransduction, which can lead to disease [27, 96, 220]. The endocardium is similarly lined with endothelial cells, which have been shown to influence cardiac performance and remodelling [200], however, few studies have attempted to quantify the relationship with WSS. This is due to a combination of the aforementioned limitations of 4D-flow MRI but, equally importantly, limitations in the structural MRI methods used to segment the moving endocardium. As discussed previously, the complex structures on the surface of the endocardium are not fully resolved, owing to low spatial resolution of structural cine-MRI, which means the surface on which WSS is being estimated is not anatomically correct. This is then coupled with extreme motion of the boundary throughout the cardiac cycle, which is difficult to accurately capture, resulting in significant uncertainty in both the position and motion of the endocardium. This makes the task of deriving an accurate WSS approximation in the LV challenging, and further, given such significant differences in the boundary surface, the validity of such WSS measurements could also be questioned.

The problem setup outlined in this chapter allows for the characterisation of the relationship between WSS obtained directly from the anatomically-correct endocardium and that calculated on approximated boundary surfaces. To this end, we compare WSS distributions across each boundary representation,  $S_1^i - S_3^i$ , for each case, and compare against PINN predictions using the ground truth boundary and the ground truth WSS itself.

In Fig. 4.17, contour plots of WSS magnitude on the endocardium are compared in case C in late-diastole. We analyse results across the three PINN boundary configurations,  $S_1$ ,  $S_2$  and  $S_3$ , and also compare these against PINN results obtained using the ground truth boundary, labelled PINN-GT. On the anteroseptal side, the WSS distribution is approximately correct across PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$ , although there is a clear discrepancy in magnitude compared to the ground truth result. This is somewhat repeated on the inferolateral side of the LV, although there is a greater drop-off in WSS magnitude. This side of the LV experiences the greatest difference in boundary surface detail, as is visible in the figure. As such, it is expected that WSS reconstructions are of the lowest accuracy here. Quantitatively, peak WSS values increase from PINN- $S_1$  to PINN- $S_3$ , which can be visualised in the anteroseptal view. Conversely, mean WSS decreases, which is a surprising result. With increasing detail on the approximated endocardium, it would be expected that WSS accuracy increases too. This hypothesis is supported by the PINN-GT model, which attains the best qualitative and quantitative results, although there is

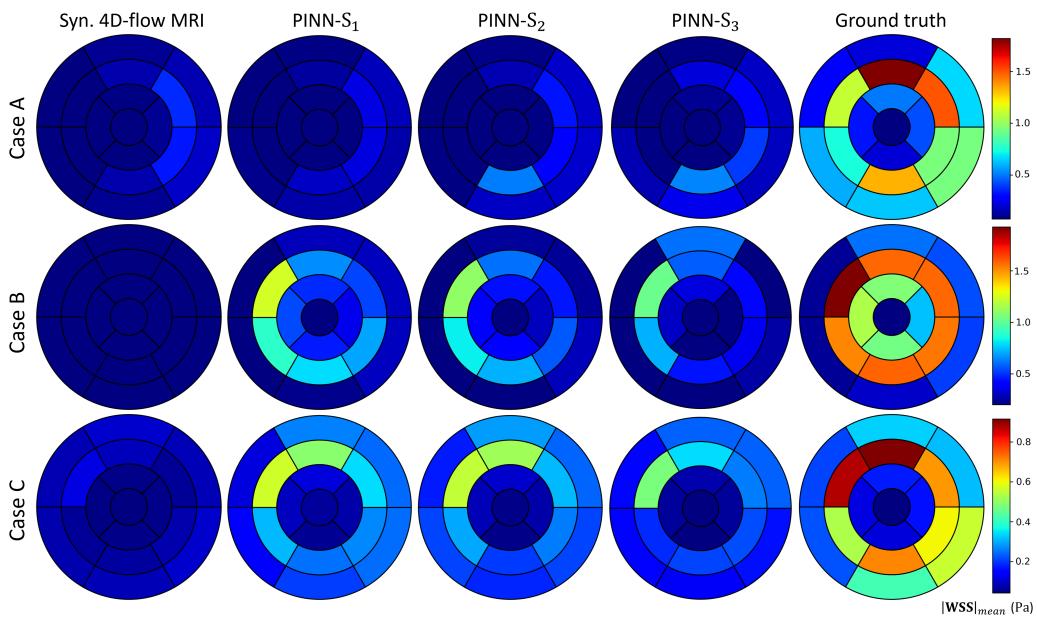
still some under-prediction in WSS magnitude.

Direct point-wise quantification of WSS error is not possible in the current study, since the surface topology is different between boundary approximation types. Instead, bull's eye plots are utilised to quantify errors in mean and peak WSS across seventeen distinct regions of the endocardium, as implemented in [194]. By subdividing the myocardium into 17 distinct regions, these plots are useful for comparing surface variables in the LV across distinct topologies, and have been used extensively in the literature to classify regional parameters such as scar extent [179] and myocardial strain [143]. In Fig. 4.18, the regions of the endocardium extracted for the bull's eye plots are displayed. The exact regions in Fig. 4.18a extend beneath the numbered sections to the next slices, so that the entire endocardium is dissected.

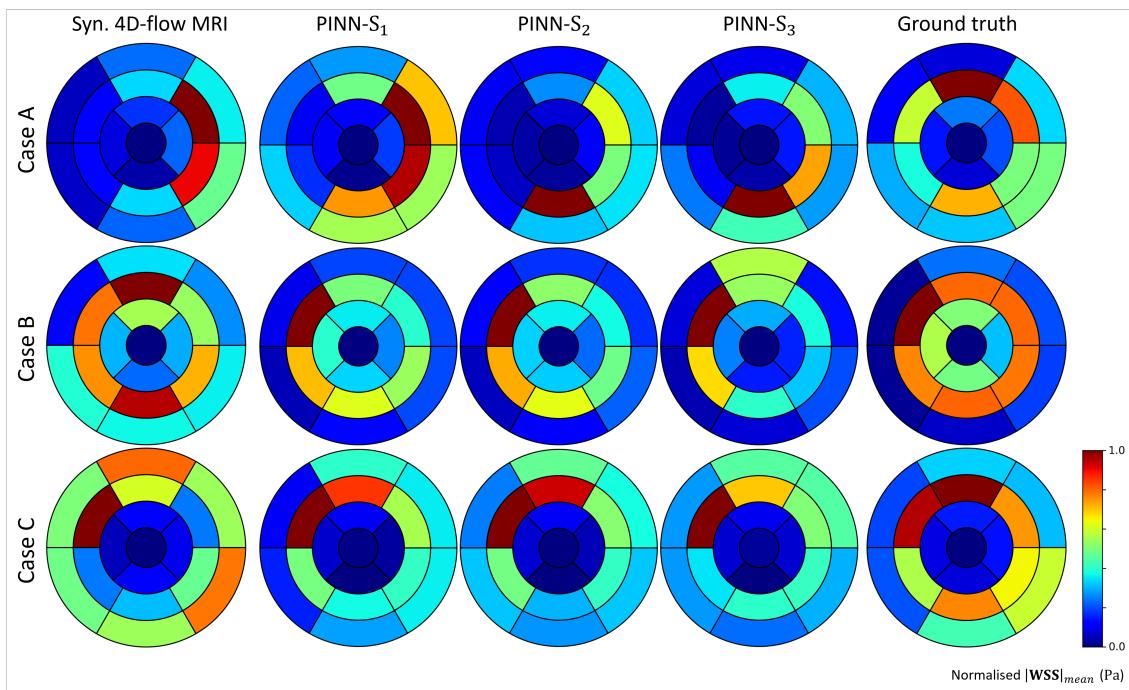


**Figure 4.18:** An explanation of the regions of the endocardium extracted during bull's eye plot analysis.

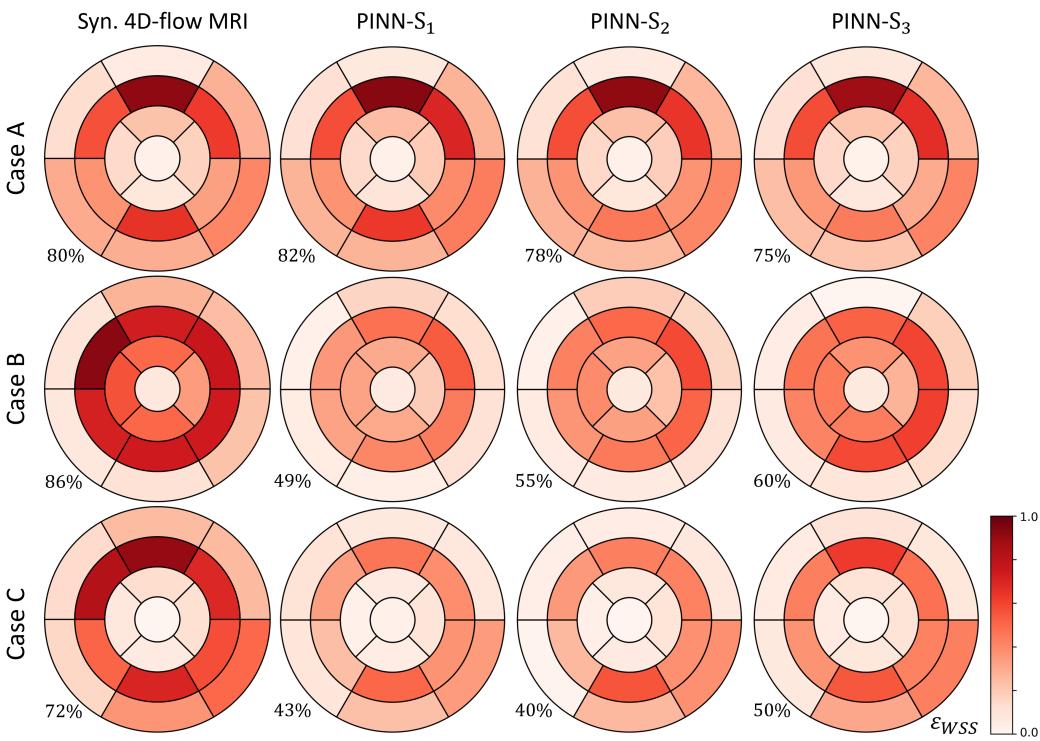
In Fig. 4.19, absolute mean WSS values are collated in bull's eye plots across all configurations in each case. As noted in Fig. 4.17, WSS is uniformly under-predicted by all PINN configurations. In cases B and C, there is little variation across the WSS patterns obtained using the three PINN boundary configurations, although the PINN model predictions are uniformly poor when compared against the ground truth. In case B, elevated mean WSS in the mid anteroseptal region is correctly identified, with PINN- $S_1$  and PINN- $S_2$  also reconstructing the slightly higher mean WSS around the mid septal and inferior



**Figure 4.19:** Bull's eye plots of absolute mean WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth.



**Figure 4.20:** Bull's eye plots of relative mean WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth.



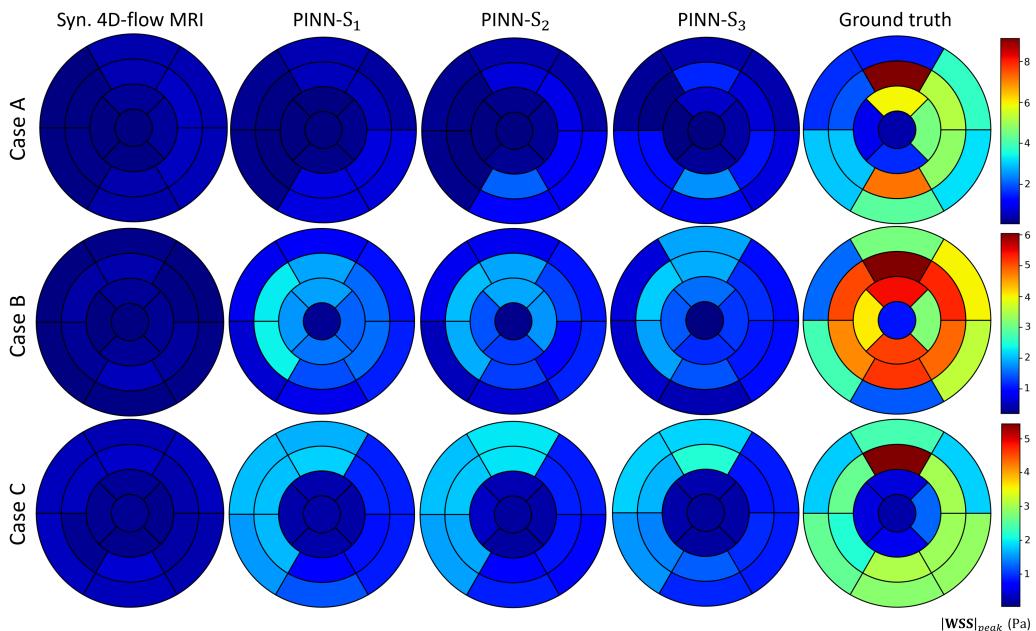
**Figure 4.21:** Bull's eye plots of relative mean WSS error (against ground truth) in late-diastole. Reported values are mean relative error across all segments.

regions. The relatively lower mean WSS in the basal region is also well characterised. In case C, the elevated mean WSS in the mid anterior and mid anteroseptal regions is well reconstructed, alongside relatively lower mean WSS in the apical region. Mean WSS is poorly reconstructed by all three PINN boundary configurations in case A. In Fig. 4.21, relative mean WSS errors are displayed across all cases. For case A, errors across all PINN boundary configurations are comparable with the synthetic 4D-flow MRI data. However, in cases B and C, the PINN offers a pronounced reduction in relative error, although errors still remain high.

In Fig. 4.20, bull's eye plots are used again, characterising max-normalised mean WSS, where normalisation is performed individually for each configuration. This allows for exposure of the mean WSS distribution, without accounting for differences in magnitude. Given that WSS is derived on distinct boundary representations, some discrepancy in magnitude should be expected, but the ability to reconstruct a similar distribution of WSS could also prove beneficial. For case B, the distribution of mean WSS is generally good for all three PINN boundary configurations, with significantly elevated mean WSS in the mid anteroseptal region, and reduced mean WSS in the basal and apex regions. relative mean WSS is slightly under-predicted across the entire mid region, and with PINN-S<sub>3</sub> there is a slight over-prediction in the basal anterior region. For case C, the two regions

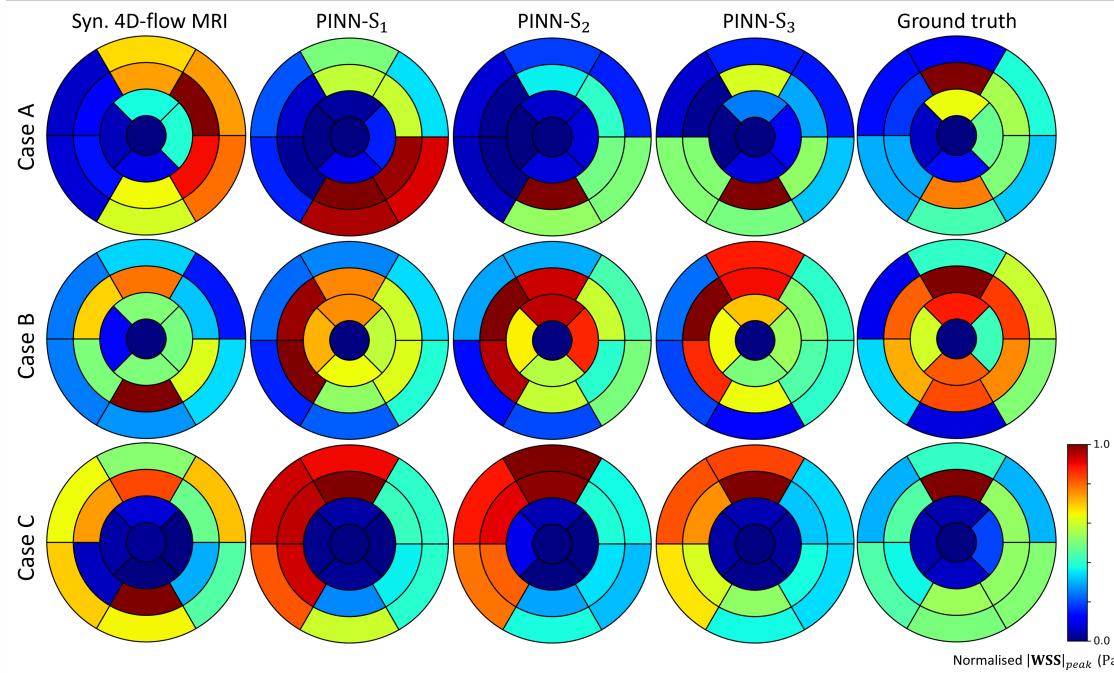
that experience the greatest mean WSS values, the mid anteroseptal and mid anterior regions, are correctly characterised, alongside comparatively reduced mean WSS in the apical region. PINN- $S_1$  also correctly reconstructs the reduced mean WSS in the basal septal region, although PINN- $S_2$  and PINN- $S_3$  fail to do so. Whilst direct calculation from the synthetic 4D-flow MRI data does correctly provide the elevated mean WSS in the mid anteroseptal region, relative mean WSS in the basal region is significantly over-predicted. As is evident from the absolute mean WSS bull's eye plots, WSS reconstruction in case A is again poor, with little consistency across PINN configurations.

Bull's eye plots of absolute and max-normalised peak WSS are shown in Figs. 4.22 and 4.23, respectively, with peak WSS error shown in 4.24. Absolute and max-normalised peak WSS in cases *A* and *C* is generally poorly reconstructed, although in case *C* all PINN configurations are able to reconstruct the maximum peak WSS in the correct region, namely the mid anterior, with relatively low peak WSS in the apical region. Elevated peak WSS in the mid LV region, with relatively low peak WSS in the basal and apex regions, is demonstrated for case *B* across ground truth, PINN predictions and synthetic 4D-flow MRI.

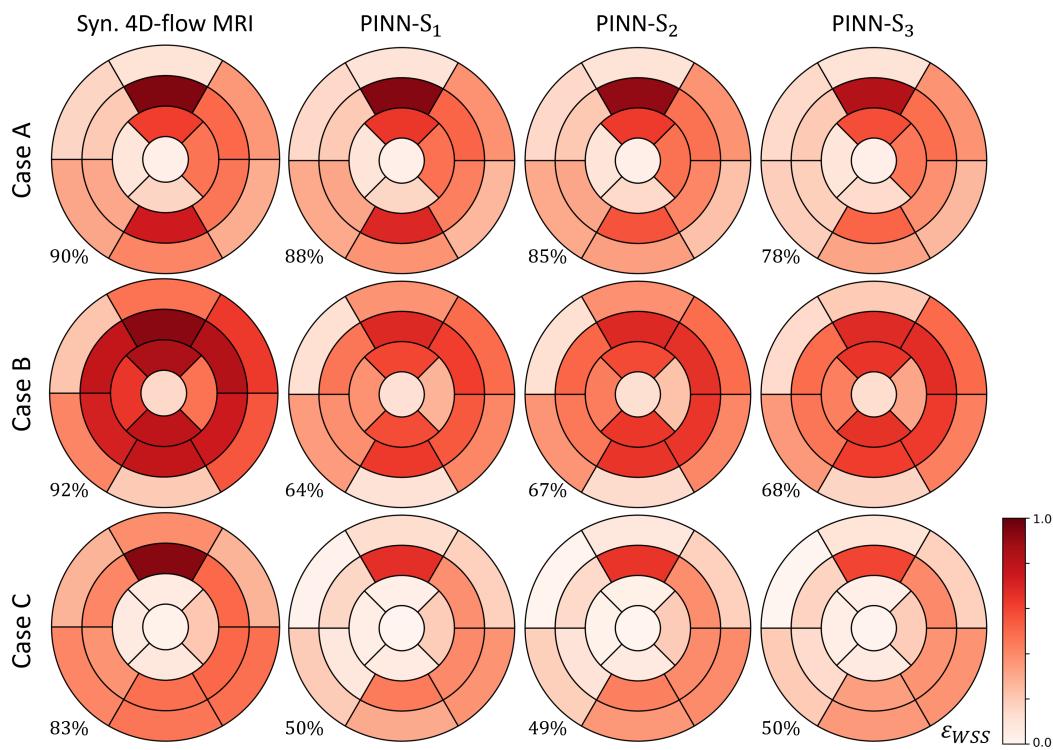


**Figure 4.22:** Bull's eye plots of absolute peak WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth.

As has been well documented [76], direct quantification of WSS from 4D-flow MRI is unreliable, which can be seen qualitatively and quantitatively in Figs. 4.17, 4.19, 4.21, 4.22, 4.24. In vascular studies, linear or parabolic curve fitting can be utilised to interpolate the velocity field in the near-wall region, improving the accuracy of spatial derivative calculations.

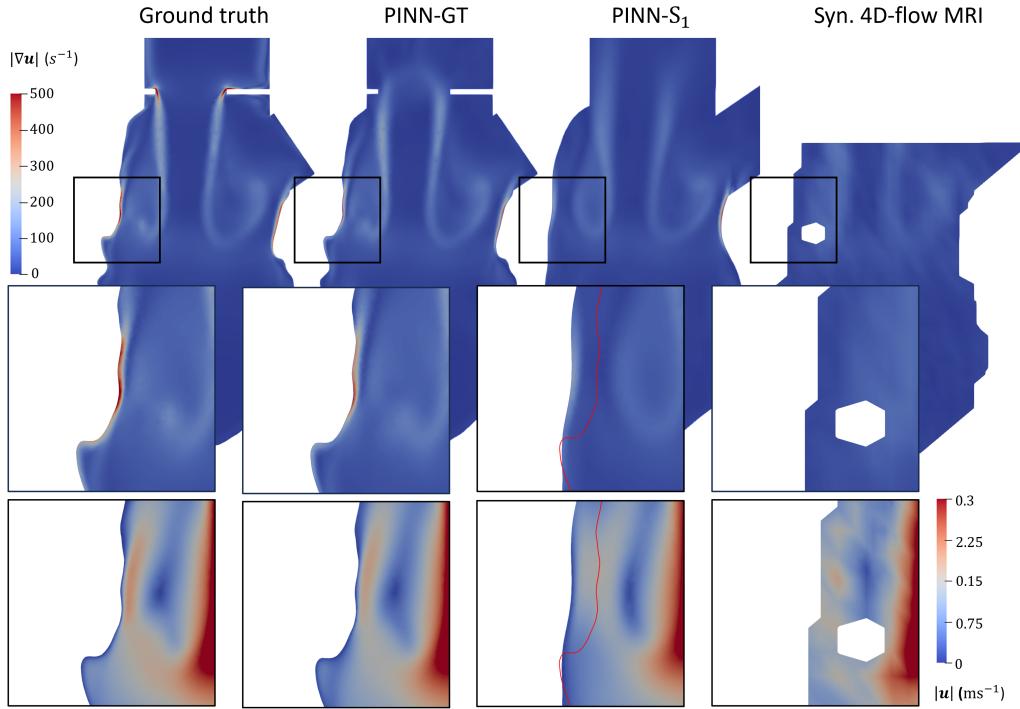


**Figure 4.23:** Bull's eye plots of relative peak WSS in mid-diastole, comparing synthetic 4D-flow MRI, PINN predictions and ground truth.



**Figure 4.24:** Bull's eye plots of relative peak WSS error (against ground truth) in late-diastole. Reported values are mean relative error across all segments.

However, such approaches assume axisymmetric flow with parabolic profiles in the near-wall region [38], which is not a valid assumption in the cardiac chambers. As such, they are not viable methods for comparison in this setting.



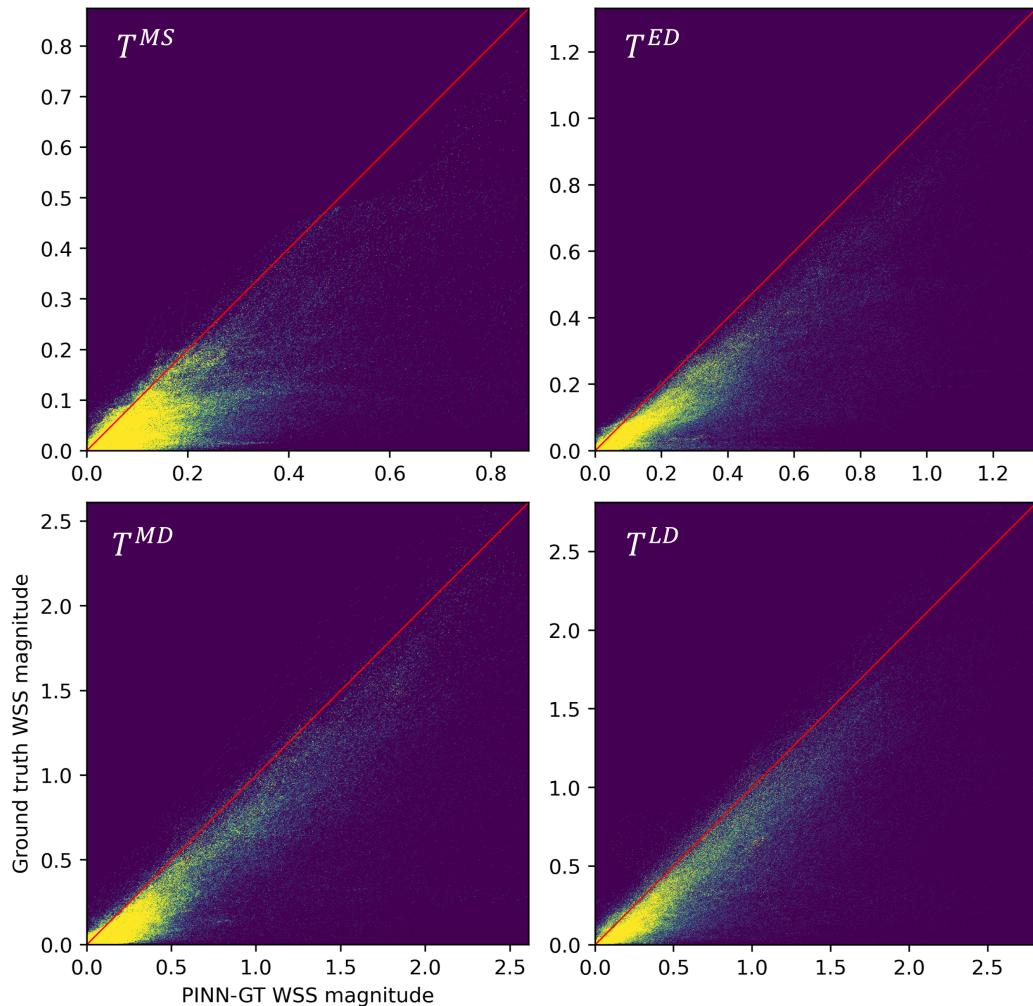
**Figure 4.25:** Case C, mid-diastole: Contour plots of velocity gradient magnitude (large figures and top row of inserts) and velocity magnitude (bottom inserts) for ground truth solution fields (left), PINN predictions using the ground truth boundary (centre left), predictions using PINN- $S_1$  (centre right) and synthetic 4D-flow MRI data (right). The red line displayed in the PINN- $S_1$  inserts corresponds to the ground truth boundary position.

In general, results presented in this section suggest that WSS is under-predicted by the PINN model when using approximated endocardial boundary representations. Given the discrepancy in boundary position, differences in magnitude are to be expected. Surfaces  $S_1$ ,  $S_2$  and  $S_3$  are typically positioned to the exterior of the true flow domain, and as such, there is no flow data present in the immediate vicinity of the boundary. In the absence of measurement data, velocities predicted by the PINN in the near-wall region are likely to be low, which in turn generates lower magnitude velocity gradients. In Fig. 4.25, velocity gradient magnitude in the near-wall region is compared for ground truth, PINN predictions in both the ground truth boundary and the  $S_1$  approximated boundary, alongside synthetic 4D-flow MRI data. The region for analysis, identified by the black box in the top row of figures, was selected to be an area with high ground truth WSS magnitude. This corresponds to an area containing high local velocity gradients in the near-wall region, which can be observed in the ground truth, and, to a lesser extent, the PINN-GT fields in the top inserts. As demonstrated by the red line in the PINN- $S_1$  inserts, the approximated  $S_1$  boundary is positioned to the exterior of the true endocardium, and

therefore does not intersect the higher-velocity flow near the incoming mitral jet to the same extent. This results in lower velocity in the near-wall region, and thus velocity gradients of a lower magnitude.

The enforcement of the no-slip boundary condition on the endocardium is key in facilitating the quantification of WSS using the PINN model, since it allows for the accurate calculation of velocity gradients in the near-wall region (relative to the boundary position and adjacent flow). As such, it is more beneficial to use an approximated boundary condition than no boundary condition at all, even in the presence of positional uncertainty. Overlooking the under-prediction of WSS magnitudes, the similarities in WSS distributions between ground truth data and PINN configurations  $S_1$ ,  $S_2$  and  $S_3$ , at least for cases B and C, suggest that the calculation of WSS on approximated boundaries may still provide useful information. Further, the variation in WSS magnitude and distribution across the three PINN boundary configurations, particularly in isolating the peak and minimum WSS regions, is quite low, demonstrating some level of robustness to small positional differences in segmented surfaces.

Although the greatest discrepancy in WSS magnitude can be accounted for by positional uncertainty, the PINN results obtained using the ground truth boundary, PINN-GT, suggest that under-prediction of WSS magnitude remains an issue, as with other methods for the quantification of WSS from 4D-flow MRI [182, 42, 214]. In Fig. 4.26, 2D histograms are used to compare the magnitude of WSS derived directly from the ground truth with that predicted using the PINN-GT configuration, across four phases of the cardiac cycle. From all four histograms, WSS magnitude appears to be uniformly under-predicted, with a slightly stronger bias at lower WSS values. Linear and parabolic extrapolation methods exhibit a strong bias at higher WSS magnitudes [182], with tail-off at under  $\sim 1\text{Pa}$  for the former and  $\sim 7\text{Pa}$  for the latter. The PINN model does not appear to be as adversely affected by such bias at the range of WSS magnitudes tested, with roughly uniform under-prediction shown. Whilst the work presented here suggests that PINNs could be a suitable method for quantification of WSS from 4D-flow MRI given more detailed knowledge of a boundary surface, the particular problem setting is highly complex, with moving boundaries and multi-scale flow. Full characterisation of errors arising due to limitations of the PINN model itself, rather than positional uncertainty, should be performed in simpler test cases, using rigid walls and exploring a larger range of WSS values. This has been briefly explored [191, 8], but a more detailed study across a wide range of flow states is required to full characterisation.



**Figure 4.26:** Case C: 2D histogram plots comparing WSS magnitude obtained using PINN-GT and the ground truth data, at mid-systole ( $T^{MS}$ ), early-diastole ( $T^{ED}$ ), mid-diastole ( $T^{MD}$ ) and late-diastole ( $T^{LD}$ ). The red line indicates a perfect match in magnitude between prediction and ground truth.

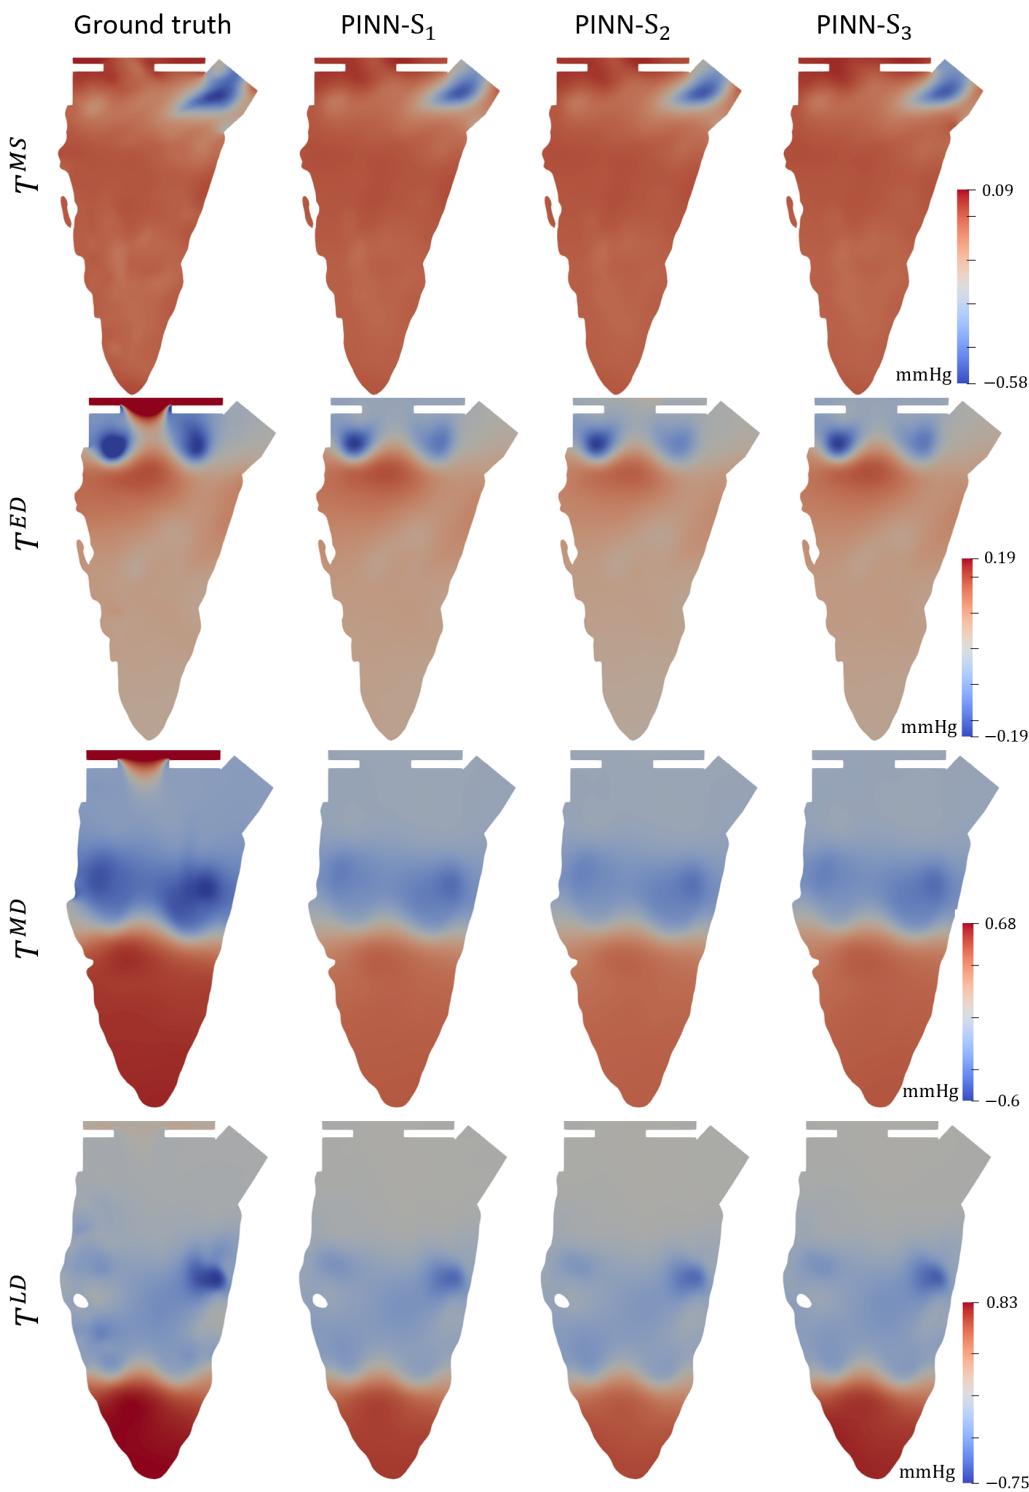
#### 4.3.2.5 Pressure

There is a consensus that intraventricular pressure and pressure gradients play an important role in the healthy function of LV flow [173, 243], however their usage in a clinical setting has been restricted by the difficulties faced in acquiring accurate measurements [173]. The invasive nature of cardiac catheterisation to measure absolute pressure [104, 46] has motivated the development of image-based pressure quantification techniques for use in both cardiac and vascular studies.

Although coupled to the velocity field through the Navier-Stokes equations, reconstruction of cardiovascular pressure fields from velocity measurement data is not a trivial task. Traditional methods utilised doppler echocardiography (ECG) to approximate the pressure drop,  $\Delta p$ , between two distinct points in the flow region, using the Bernoulli theorem, in either a simplified or modified form [46, 209, 70, 168, 95, 62, 77], or the Euler equations [237, 20]. Recent methods have exploited advancements in imaging technology to derive pressure from planar PC-MRI and 4D-flow MRI, with notable contributions including the use of a virtual work-energy function (vWERP) [59, 155, 74], pressure Poisson equations [242, 132, 69] and machine learning-based methods [204, 128].

Early approaches to quantify intraventricular pressure drop based on the simplified [209] and modified [77] Bernoulli equations have been shown to correlate poorly with catheter data [58], owing to the number of simplifications made about the intermediate flow field. The virtual Work-Energy Relative Pressure (vWERP) method has been successfully adapted for use in the cardiac chambers [155], utilising the arbitrary Lagrangian-Eulerian form of the Navier-Stokes equations across dynamic sub-domains of the chamber. However, although this model is significantly more accurate than previous approaches, it similarly only provides pressure drop values between two regions in the LV. As with velocity, the pressure field in the LV is complex and spatially heterogeneous [69], and cannot be fully characterised by relative differences. However, despite the potential usefulness of reconstructing localised pressure differences, there has been limited investigation into suitable methods. In [69, 179], the pressure Poisson equations are used to reconstruct pressure distributions from 4D-flow MRI data, however the resulting field is highly dependent on the accuracy of flow gradients, which can be limited [155].

In our PINN model, relative pressure fields are recovered naturally during training through the PDE loss constraint, without the use of any pressure data. We assess the point-wise accuracy of the derived pressure fields against the ground truth data, and use the PINN



**Figure 4.27:** Case B: 2D contours of the relative pressure field through the centre of the LV, at four distinct phases of the cardiac cycle, namely mid-systole ( $T^{MS}$ ), early-diastole ( $T^{ED}$ ), mid-diastole ( $T^{MD}$ ) and late-diastole ( $T^{LD}$ ). Fields are derived from the ground truth pressure field and the three PINN boundary configurations, where the mean pressure value is used to generate the relative pressure.

model to compute intraventricular pressure drop estimations by probing points in the pressure field to compare against the simplified Bernoulli (SB) approach [209].

In Fig. 4.27, relative pressure distributions at four distinct phases of the cardiac cycle are produced in case B using the ground truth pressure data and the three PINN boundary configurations. Generally, there is good qualitative agreement between PINN predictions and ground truth data, although there is some discrepancy in pressure magnitude in mid- and late-diastole ( $T^{MD}$  and  $T^{LD}$ ). The range of relative pressure values experienced within the ground truth data is slightly lower than would be expected in reality [155]. Therefore, given the under-estimation of extreme relative pressure values by the PINN model, validation with a larger range of pressure values would be useful in future work.

Visually, results across the three PINN boundary configurations are consistent, although there is a minor difference in magnitudes in late-diastole. This consistency across PINN configurations  $S_1$ ,  $S_2$  and  $S_3$  can also be observed quantitatively in Tab. 4.15, with each model obtaining max-normalised RMS values within 0.1% of each other.

In Fig. 5.6, pressure drop,  $\Delta p$ , is evaluated between the LV base and apex at four distinct cardiac phases using the three PINN boundary representations, ground truth pressure data and the SB approach. The SB method estimates pressure drop between two points,  $\mathbf{q}_1$  and  $\mathbf{q}_2$ , directly from the velocity field, as:

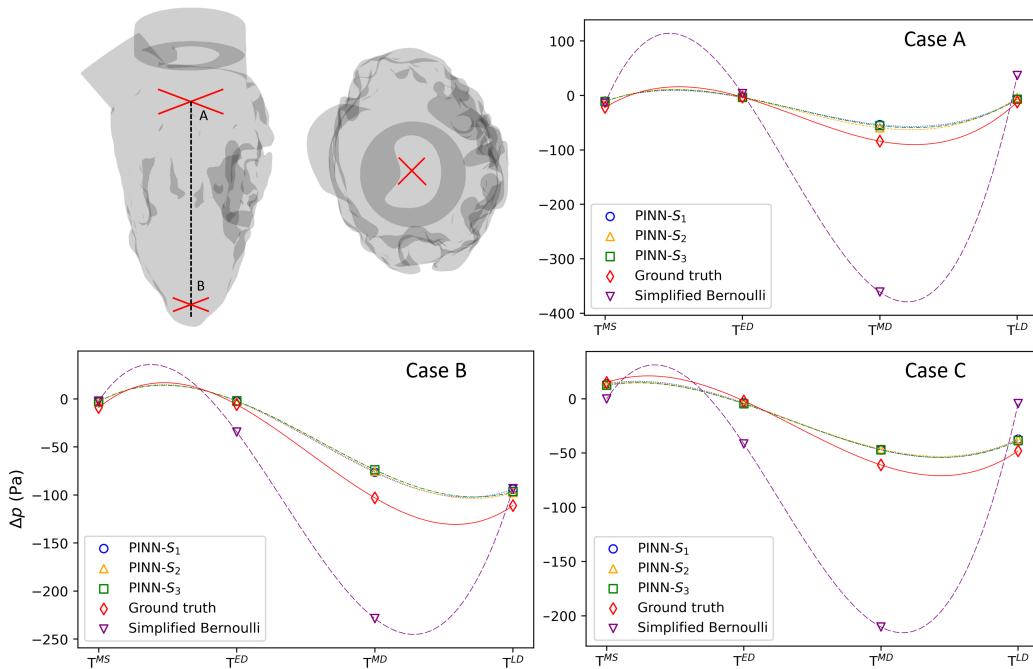
$$\Delta p = \frac{1}{2} \rho (u_2^2 - u_1^2), \quad (4.10)$$

where  $u_i = \mathbf{u}(\mathbf{q}_i) \cdot \hat{\mathbf{n}}$  with  $i = 1, 2$  and  $\hat{\mathbf{n}}$  the unit vector between points  $\mathbf{q}_1$  and  $\mathbf{q}_2$ . In Fig. 5.6, it is apparent that the PINN model offers a significant improvement over the SB method in quantifying  $\Delta p$ , whilst there is also minimal variation between PINN boundary representations in all cases. However,  $\Delta p$  is under-predicted by the PINN model, particularly at higher relative pressure magnitudes. This is also demonstrated in Fig. 4.27, where the greatest discrepancy between ground truth and PINN predictions is observed in mid-diastole, where the largest pressure gradients occur.

Overall, the ability of the PINN model to derive continuous relative pressure fields in the LV, without the use of any pressure data or additional constraints, provides a significant contribution. Most previous approaches are only able to characterise pressure drops between points or planes in the LV, and those that are able to derive the relative pressure distribution are computationally complex and heavily afflicted by corruption in the 4D-flow MRI data [155]. The PINN model also exhibits minimal variation in pressure

**Table 4.15:** Point-wise, max-normalised RMS error of the relative pressure field.

Case	$\text{RMS}_p$
A	PINN- $S_1$ 0.019
	PINN- $S_2$ 0.0199
	PINN- $S_3$ <b>0.018</b>
B	PINN- $S_1$ <b>0.014</b>
	PINN- $S_2$ <b>0.014</b>
	PINN- $S_3$ 0.015
C	PINN- $S_1$ <b>0.011</b>
	PINN- $S_2$ <b>0.011</b>
	PINN- $S_3$ <b>0.011</b>



**Figure 4.28:** Pressure drop,  $\Delta p$ , across all three cases at four distinct phases of the cardiac cycle, namely mid-systole ( $T^{MS}$ ), early-diastole ( $T^{ED}$ ), mid-diastole ( $T^{MD}$ ) and late-diastole ( $T^{LD}$ ).  $\Delta p$  is calculated between point A (distal basal region) and point B (distal apical region), as demonstrated in the upper left figure, using direct measurement of the relative pressure fields for ground truth and PINN quantification, and the simplified Bernoulli approach for direct quantification from synthetic 4D-flow MRI. The curves used in the plots are fit using cubic splines, and are intended to simply improve readability of results. Therefore, they are not necessarily representative of  $\Delta p$  at intermediate time points.

distributions derived using each of the boundary configurations, which is demonstrated quantitatively and qualitatively in the results presented in this section. As such, it can be concluded that positional uncertainty in the boundary condition plays a minimal role, demonstrating robustness across inter-operator variability in segmentation accuracy.

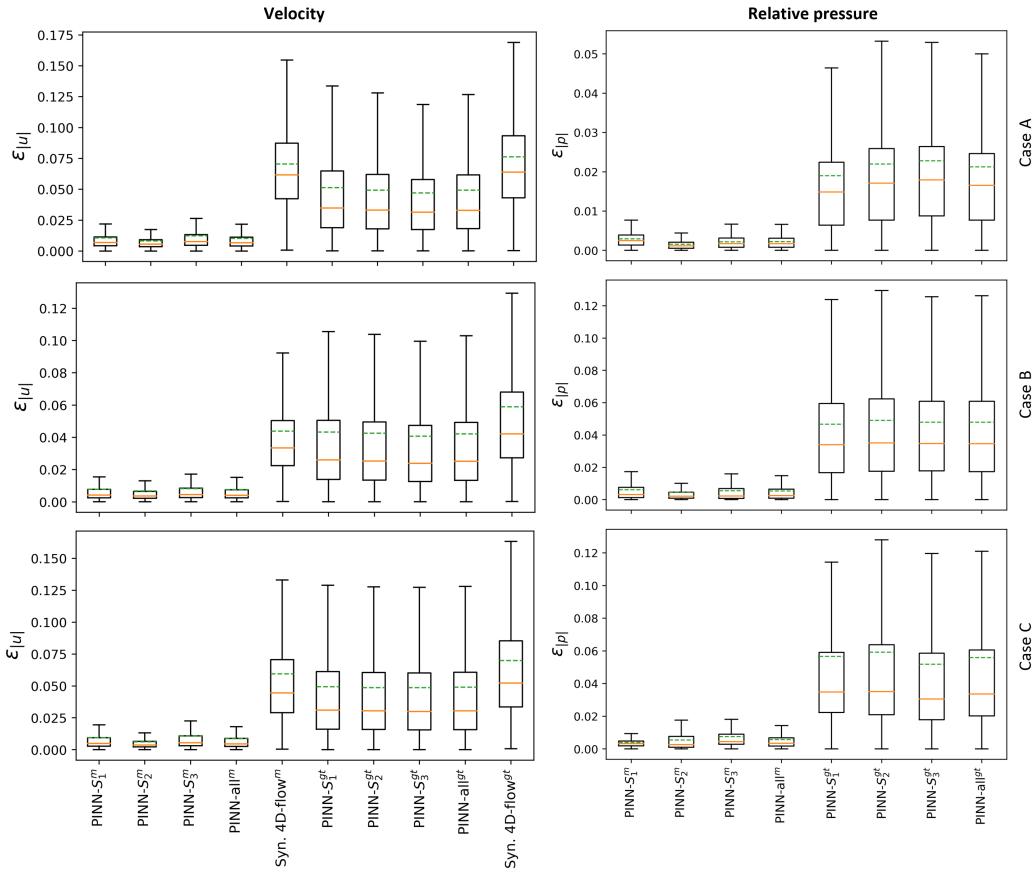
### 4.3.3 Positional Uncertainty

Positional uncertainty arises in MRI-based studies in the LV due to low spatial resolution of cine-MRI, resulting in inaccurate segmentation of the endocardium and other surfaces. For CFD-based studies, the consequences of such uncertainty can be extreme, as the simulated solution fields are highly dependent on the enforced boundary conditions [194, 222, 135]. While the PINN model introduced in this thesis is constrained by the no-slip boundary condition on the endocardial surface, it is also constrained by velocity data throughout the LV chamber, and as such, the impact of positional uncertainty on predicted fields may not be so detrimental.

To investigate the effects of positional uncertainty in PINN-based models in the LV, results were obtained and compared using distinct boundary representations in each case. In this section, these results are analysed to assess the variability across results obtained using each PINN model configuration. To do this, mean predicted fields are produced by averaging the results obtained using PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$  for each calculated variable. From this, error fields are generated for each boundary configuration (metrics denoted by  $(\cdot)^m$ ), allowing for quantification of the variability across PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$ . This is compared against the corresponding metrics against the ground truth (metrics denoted by  $(\cdot)^{gt}$ ), for comparison. Errors for the synthetic 4D-flow MRI data are also generated against the mean PINN configuration and the ground truth data. The ultimate purpose of this analysis is to assess variability across PINN boundary configurations, with the error against the ground truth data used as for comparison. Where errors against the mean PINN boundary configuration are significantly lower than errors against the ground truth data, it can be concluded that the variability across PINN boundary configurations is small, and therefore the impact of positional uncertainty is limited for the particular variable of interest.

Variability in the reconstructed velocity and relative pressure fields is assessed first. In Fig. 4.29, absolute error characteristics for velocity and relative pressure are compared across all boundary configurations. Max-normalised absolute errors against the mean

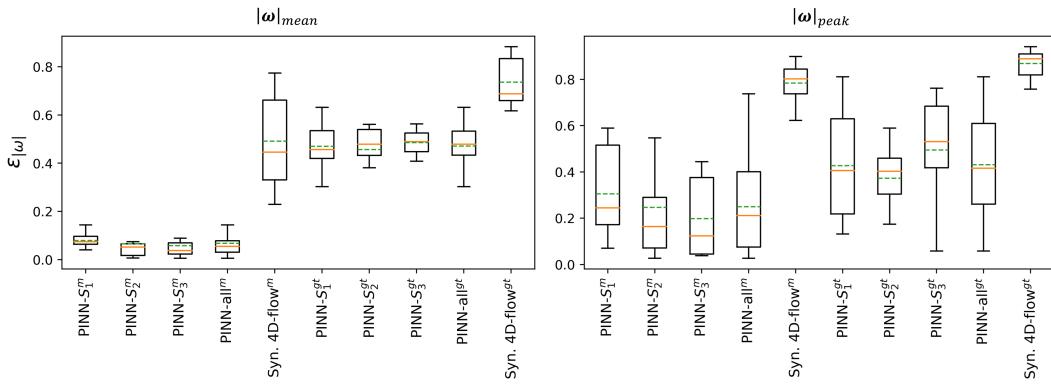
configuration (denoted  $(\cdot)^m$ ) and ground truth data (denoted  $(\cdot)^{gt}$ ) are produced, where it can be observed that there is minimal variability across PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$  for both predicted variables.



**Figure 4.29:** Box and whisker plots for point-wise, max-normalised error against the mean PINN configuration (denoted  $(\cdot)^m$ ) and ground truth (denoted  $(\cdot)^{gt}$ ) in velocity (left) and pressure (right) fields, for cases A (top row), B (centre row) and C (bottom row). The PINN-all configuration consists of all errors from PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$  concatenated.

A similar analysis is performed for mean and peak vorticity, denoted  $|\omega|_{peak}$  and  $|\omega|_{mean}$ , respectively, in Fig. 4.30. Here, relative errors against both the mean PINN configuration and ground truth data are calculated across all cardiac phases analysed, with results across all cases collated. As with velocity and relative pressure, the variability in mean vorticity calculation across PINN boundary configurations is low, although relative errors against the ground truth data are large. For peak vorticity, much larger variability across PINN boundary configurations is observed. As discussed earlier in this chapter, such discrepancies are expected, since the regions containing the highest-magnitude vorticity are those in the near-wall region. Given the differences in boundary topology, it is expected that near-wall flow will vary significantly across PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$ .

In Fig. 4.31, box and whisker plots are presented for max-normalised absolute error in mean and peak KE, measured individually at each cardiac phase. As with velocity

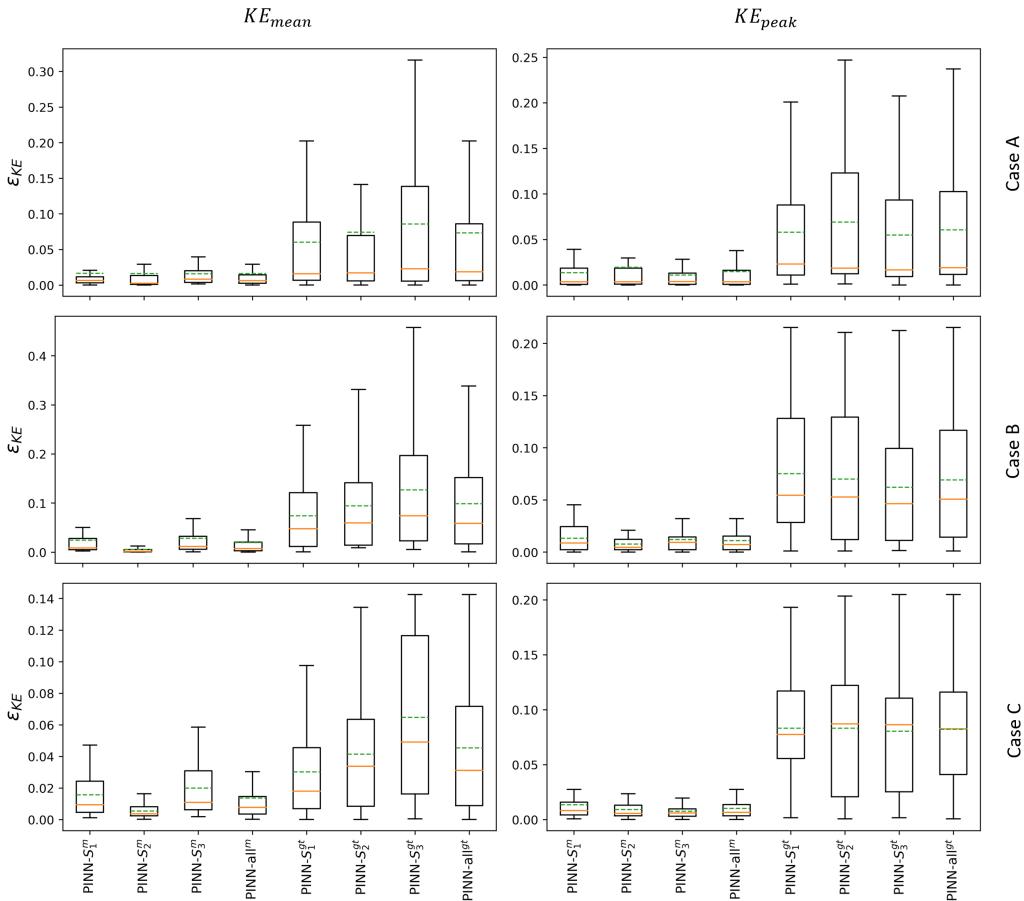


**Figure 4.30:** Box and whisker plots for relative error against the mean PINN configuration (denoted  $(\cdot)^m$ ) and ground truth (denoted  $(\cdot)^{gt}$ ) in mean vorticity (left) and peak vorticity (right) measurements, where results across all cases are collated. The PINN-all configuration consists of all errors from PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$  concatenated.

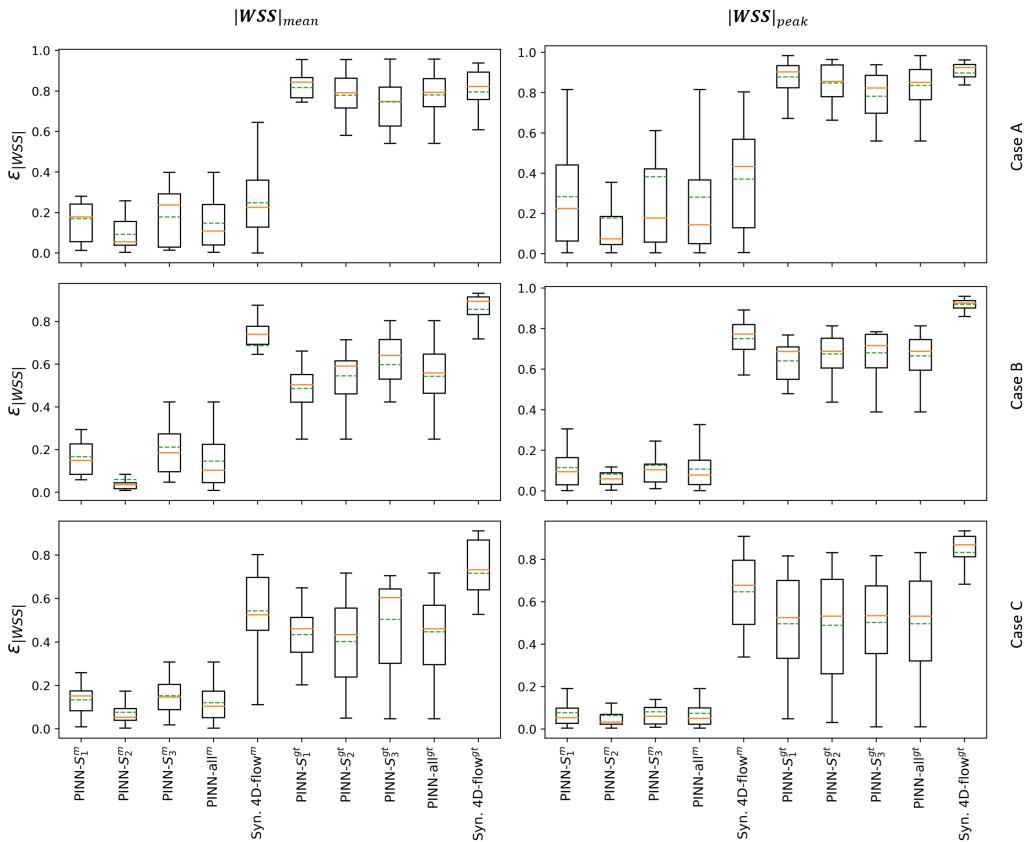
reconstructions, variability in both mean and peak KE errors across PINN boundary configurations is minimal, especially in relation to errors measured against the ground truth data. This result is generally consistent across all cases, although larger variation in mean KE error across boundary configurations can be observed in case C. However, in relation to errors computed against the ground truth data, such variations are relatively small.

Finally, the impact of positional uncertainty on WSS calculations is assessed in Fig. 4.32, where relative error in mean and peak WSS is analysed in mid-diastole for all three cases. Given the difference in surface topology, it would be expected that WSS varies greatly across PINN boundary configurations. However, although some variation is observed, it is not as pronounced as expected, at least in cases B and C. Lower errors computed against the mean configuration than ground truth are demonstrated for all PINN configurations, although in case A this feature alone may be misleading. Given the poor accuracy of WSS when calculated directly from 4D-flow MRI, we would like to see high errors when computed against both the mean PINN configuration and the ground truth data. However, the similarity in error characteristics between PINN configurations and the synthetic 4D-flow MRI data in case A indicates that the PINN model is not functioning well, even though variation across configurations is small. This is not true for cases B and C, where there is a pronounced difference between PINN results and those obtained through direct quantification. There is a greater variance mean and peak WSS across configurations in case B than case C, which is evidenced by larger quartiles in the PINN-all<sup>m</sup>.

Overall, results presented in this section demonstrate that positional uncertainty imparts a minimal effect on the quantification of flow variables away from the wall, which is char-



**Figure 4.31:** Box and whisker plots for max-normalised error against the mean PINN configuration (denoted  $(\cdot)^m$ ) and ground truth (denoted  $(\cdot)^{gt}$ ) in mean KE (left) and peak KE (right), measured at each cardiac phase for cases A (top row), B (centre row) and C (bottom row). The PINN-all configuration consists of all errors from PINN- $S_1$ , PINN- $S_2$  and PINN- $S_3$  concatenated. Results for synthetic 4D-flow MRI data are not included as errors were too large.



**Figure 4.32:** Box and whisker plots for relative error against the mean PINN configuration (denoted  $(\cdot)^m$ ) and ground truth (denoted  $(\cdot)^{gt}$ ) in mean WSS (left) and peak WSS (right), measured in mid-diastole for cases A (top row), B (centre row) and C (bottom row). Relative errors are calculated using the bull's eye segments outlined in Fig. 4.18.

acterised by low variance across PINN configurations for velocity, pressure, mean vorticity and KE predictions. Flow variables that are measured on or near the wall are more affected, which is evidenced by higher variance across PINN configurations for peak vorticity and mean and peak WSS errors. This suggests that errors arising from uncertainty in the no-slip boundary condition are generally confined to the near-wall region, and do not propagate into the LV chamber. This is likely due to the presence of the data loss term, which constrains the PINN model using velocity measurements throughout the entirety of the flow domain.

## 4.4 Conclusion

In this chapter, the PINN model introduced in chapter 3 is applied to a small cohort of synthetic LV cases in patient-specific geometries, where clinically-relevant variables are derived and predicted fields are compared with baseline methods utilised in the literature. The variables discussed include velocity, vorticity, kinetic energy, wall shear stress and relative pressure. To characterise the impact of positional uncertainty in the no-slip boundary condition, PINN results are obtained and compared in each case using three distinct endocardial representations. Variance in predicted fields across each boundary configuration is then quantified to assess the impact of positional uncertainty on the quantification of clinically-relevant variables.

Across all variables discussed, the PINN model is demonstrated to provide more accurate results than either direct measurement from the synthetic 4D-flow MRI data or competing methods. In particular, quantities that are calculated from spatial derivatives of the velocity field are reconstructed to a far higher quality. Further, it is demonstrated that positional uncertainty only affects the accuracy of variables measured in, or impacted by flow in, the near-wall region. Flow variables measured elsewhere in the domain remain largely unaffected by uncertainty in the endocardial position.

The application of PINNs in this domain presents a unique perspective for flow quantification in the LV. Although principally a super-resolution technique, the PINN model is able to accurately reconstruct a wide array of clinically-relevant variables that are not directly measurable from 4D-flow MRI, such as relative pressure and vorticity, in an inherently patient-specific manner. Competing methods are often computationally expensive and complex to apply, particularly in the case of pressure quantification, and as such have not been adopted in clinic. Further, work presented in this chapter suggests that the model

is robust to positional uncertainty in the no-slip boundary condition, at least for the majority of flow variables measured. Therefore, segmentation errors, which are inherent to MRI-based flow studies, impart less influence on results when compared to other modelling techniques, such as CFD.

# Chapter 5

## Super-Resolution of *in Vivo* Cardiac 4D-Flow MRI

### 5.1 Introduction

In this chapter we apply our PINN super-resolution model to real *in vivo* 4D-flow MRI data. The chapter is structured as so: firstly, the model is applied to a dual-resolution study, in which 4D-flow MRI acquisitions were made at two distinct spatial resolutions, where the dependence of the PINN model and predicted variables on the input spatial resolution is assessed; secondly, the model is applied to a small cohort of volunteer data with varying levels of LV remodelling, and clinically-relevant variables are assessed in relation to patient outcomes.

#### 5.1.1 Validation With Real 4D-Flow MRI

The validation of PC-MRI super-resolution models *in vivo* presents a number of challenges. 4D-flow MRI is the gold standard *in vivo* flow imaging modality, and thus the acquisition of reference data using a higher-quality imaging technique is infeasible. It is possible to acquire 4D-flow MRI data at two distinct resolutions, with the aim of recovering the high-resolution data by training the model with the low-resolution data. However, variability in heart rate and patient motion between acquisitions, for example, can result in unforeseen changes in the observed flow patterns, leading to poor repeatability. Further, the maximum spatial resolution in clinic for cardiac studies is limited to around  $2.5\text{mm}^3$  [23], which could

still be considered coarse for the quantification of flow derivatives and small-scale features. Another issue arising in *in vivo* studies is the presence of additional data corruption, in the form of *structured noise*, which is challenging to artificially replicate in synthetic studies and thus account for in the model. This poses an additional challenge for models that have been configured using synthetic data.

Previous publications addressing super-resolution of 4D-flow MRI, which are restricted to studies in the vasculature, have demonstrated the efficacy of proposed models *in vivo* and *in vitro*, with varying rigour. The use of *in vitro* phantoms allows for controlled flow conditions, leading to good repeatability between scans, and for this reason studies of this kind have been utilised in the literature. In [72], a single *in vitro* aneurysm phantom was used, for which both 4D-flow MRI (at spatial and temporal resolutions of  $0.57\text{mm}^3$  and 82.8ms respectively) and high-resolution particle image velocimetry (PIV) (at spatial and temporal resolutions of  $0.141\text{mm}^3$  and 20ms respectively) were acquired. Visually, velocity field reconstructions using their model appear accurate when compared with the PIV data, however quantitative analysis indicated no notable improvement in similarity indices (ASI and MSI) when compared to the raw 4D-flow MRI data. An *in vitro* study was also used in [75], where the trained model was applied to 4D-flow MRI data measured in a single bifurcation phantom, for which acquisitions at spatial resolutions of  $2\text{mm}^3$  and  $4\text{mm}^3$  were made. Qualitatively, effective de-noising was demonstrated, with the model able to reconstruct the  $2\text{mm}^3$  scan effectively. Relative errors in flow rate at three planes were used as a quantitative performance metric, where slight improvement over the competing method (sinc interpolation) was noted.

The model in [75] was also applied *in vivo* to flow in the aorta in a single volunteer. As noted in the publication, there was no course for rigorous validation in this example (aside from comparing flow rates across three planes), and thus it was used simply to demonstrate the visual improvement in the velocity reconstruction, with particular emphasis on the near-wall region. A larger *in vivo* study was performed in [203], producing super-resolved velocity reconstructions in 10 whole heart, 11 thoracic aorta and 3 intracranial aneurysm volunteers (24 volunteers in total). The intracranial aneurysm acquisitions were obtained at a spatial resolution of  $0.82\text{mm}^3$ , with the remaining 21 acquisitions obtained at a spatial resolution of  $2.1 - 2.5\text{mm}^3$ . The acquired 4D-flow MRI data were considered the high-resolution ground truth, from which low-resolution data were generated by synthetically downsampling in space by factors of  $2\times$ ,  $3\times$  and  $4\times$  to evaluate model performance. Whilst this approach allows for quantification of reconstruction errors with *in vivo* data, it should

still be considered a synthetic study, since the super-resolution task is performed only to reverse the artificial downsampling steps. In [195], a single *in vivo* study is performed in the thoracic aorta of a volunteer, at spatial and temporal resolutions of  $1.14\text{mm} \times 1.14\text{mm} \times 0.9\text{mm}$  and 30ms, respectively. The de-noising capabilities of their model are demonstrated qualitatively, although without high-resolution reference data, it is unclear whether small-scale flow features are also falsely removed. Quantitative performance is assessed by comparing mean flow rate, max flow rate and reversed flow index computed with both their model and the raw 4D-flow MRI.

### 5.1.2 This Chapter

#### 5.1.2.1 *In Vivo* Model Validation

The validation approach used in this chapter uses 4D-flow MRI data acquired at two distinct spatial resolutions, namely  $2.5\text{mm}^3$  and  $4\text{mm}^3$ , in a single volunteer. As discussed earlier in this section, variability between acquisitions means reconstructing the 4D-flow MRI velocity field at a higher resolution by training the model with low-resolution input data is challenging. As such, the validation approach is instead structured around obtaining PINN model results using both data sets, and assessing the variability in predicted quantities – in particular, quantities that are measured globally throughout the LV rather than point-wise. To this end, we would like to demonstrate that the prediction of important variables is not dependent on the input resolution of the training data.

It must be noted that this study, being limited in scope, is intended to serve a blueprint for a future *in vivo* validation study, rather than provide robust conclusions. Validation across more cases is required to do this.

#### 5.1.2.2 *In Vivo* Exploratory Study

In this part of the chapter, the PINN model is applied to a small cohort (4 cases) with varying levels of LV impairment. Here, the PINN model is used to investigate correlations between clinically-relevant variables and levels of impairment. As with the validation study, this section does not contain enough cases to rigorously conclude that certain variables are linked to cardiomyopathy. However, it does demonstrate the type of variables that are obtainable directly from the model.

## 5.2 Methodology

### 5.2.1 Model

The model, and associated components, used in this chapter is largely that established in chapter 3. As such, the PINN model architecture is composed of a Siren network [206] and utilises the self-adaptive PINN (SA-PINN) loss weighting scheme [158]. An additional loss term is included to account for temporal periodicity, and the type of loss function used is altered. The hyperparameters used were tuned using grid searches, which will be discussed later in this section.

#### 5.2.1.1 Loss Function

The loss function used in this chapter is similar in form to that established in chapters 3 and 4, however the type of loss used for each component is changed from mean-squared error (MSE) to Huber loss [103], and an additional loss component is included to account for temporal periodicity. The Huber loss is a piece-wise combination of MSE and mean absolute error (MAE), taking the form of MSE in the immediate neighbourhood of zero, and MAE outside of these bounds. As such, the Huber loss is more robust to significant outliers, as with MAE. This is a desirable property for training with real 4D-flow MRI data as structured noise, not accounted for in our synthetic studies, can introduce high-magnitude outliers which can significantly impact training convergence. In particular, when tissue regions intrude into the masked flow domain in phase images, they produce what is called 'phase offset error', where uniform noise drawn from  $\mathcal{U} \sim [-v_{enc}, v_{enc}]$  is generated in affected voxels. This can produce high-magnitude outliers in the training data which significantly hamper training convergence.

Due to temporal pseudo-periodicity in the acquired 4D-flow MRI data, we are able to enforce an additional constraint in the loss function by ensuring continuity of the velocity and pressure fields and their first derivatives in time. As such, the total loss function is given by:

$$\mathcal{L} = \mathcal{L}_{PDE}^{SA} + \mathcal{L}_{BC}^{SA} + \mathcal{L}_{data}^{SA} + \mathcal{L}_{PT}^{SA}, \quad (5.1)$$

for PDE loss,  $\mathcal{L}_{PDE}$ , boundary condition loss,  $\mathcal{L}_{BC}$ , data loss,  $\mathcal{L}_{data}$ , and the additional

temporal constraint,  $\mathcal{L}_{PT}$ . This additional constraint is given by:

$$\mathcal{L}_{PT}^{SA} = \mathcal{L}_{PT}^1 + \mathcal{L}_{PT}^2 + \mathcal{L}_{PT}^3 + \mathcal{L}_{PT}^4, \quad (5.2)$$

where

$$\mathcal{L}_{PT}^1 = \text{Huber}(\hat{\mathbf{u}}_{t_0} - \hat{\mathbf{u}}_{t_{N+1}}) \quad (5.3)$$

$$\mathcal{L}_{PT}^2 = \text{Huber}(\partial_t \hat{\mathbf{u}}_{t_0} - \partial_t \hat{\mathbf{u}}_{t_{N+1}}) \quad (5.4)$$

$$\mathcal{L}_{PT}^3 = \text{Huber}(\hat{p}_{t_0} - \hat{p}_{t_{N+1}}) \quad (5.5)$$

$$\mathcal{L}_{PT}^4 = \text{Huber}(\partial_t \hat{p}_{t_0} - \partial_t \hat{p}_{t_{N+1}}), \quad (5.6)$$

where  $t_i$  is the  $i^{th}$  time step,  $N$  is the number of time steps in the upsampled sequence of results,  $\hat{\mathbf{u}}_{t_i}$  is the velocity at time step  $t_i$  and  $\hat{p}_{t_i}$  is the pressure at time step  $t_i$ . Time step  $t_{N+1}$  is entirely artificial, and therefore there is no data present here. However, this additional constraint enforces periodicity by assuming the velocity and pressure fields at time  $t_{N+1}$  are equal to that at time  $t_0$ .

### 5.2.1.2 Hyperparameter Optimisation

In this chapter, hyperparameter optimisation was performed in each case individually, using grid searches instead of the Bayesian optimisation method used in the previous chapters. Extensive hyperparameter optimisation in this setting is ineffective, as the absence of ground truth data means there is no suitable metric for performance. To this end, the grid searches are performed to produce results that maintain as much of the peak velocity present in the 4D-flow MRI data without over-fitting. This is not intended to produce an ‘optimal’ model configuration, it is simply to eradicate obviously poor configurations.

### 5.2.2 Data Acquisition

The data utilised in this chapter were obtained from two sources, namely our own data set, referred to as the *dual-resolution* data set, and the DTI INMI data set (British Heart Foundation Intermediate Clinical Research Fellowship (FS/13/71/30378)) as used in [49]. In Tab. 5.1, information about the acquisitions and clinical outcome can be found for all cases. It should be noted that the scan date (acute or 3 months post-infarct) was selected

only based on the availability of data.

**Table 5.1:** Volunteer details (LV function).

Case	Data set	Acq. time	LV dysf.	EF	Infarct size	Infarct loc.
Case A	Dual-res.	N/A	Healthy	60 – 65%	0%	N/A
Case B	DTI INMI	3 months	Moderate	45%	23%	Inferior
Case C	DTI INMI	Acute	Normal	56%	0%	Inferior
Case D	DTI INMI	3 months	Severe	34%	36%	Anterior

### 5.2.2.1 Dual-Resolution Data

Two acquisitions were made in one subject in the dual-resolution 4D-flow MRI study, the details of which can be found in Tab. 5.2. Data were acquired at two spatial resolutions,  $2.5\text{mm}^3$  and  $4\text{mm}^3$ , using the SENSE parallel imaging acceleration technique. The temporal resolution was fixed at 30.7ms, with an encoding velocity of  $V_{enc} = 2\text{ms}^{-1}$ . The machine used was a 1.5 Tesla Philips Ingenia, with a flip angle of  $8^\circ$ , repetition time of 4.1ms, echo time of 2.24ms and number of signal averages of 1. For all acquisitions, the reconstructed in-plane spatial resolution was  $1.5625\text{mm} \times 1.5625\text{mm}$ , with through-plane resolution defined by the acquired spatial resolution. Case A in Tab. 5.1 corresponds to dual-resolution scan  $S_{2.5}$ .

**Table 5.2:** Dual-resolution 4D-flow MRI study: acquisition details for the two distinct scans, outlining acquired spatial and temporal resolutions (res.).

Scan name	Acceleration	Spatial res.	Scan duration
$S_{2.5}$	SENSE (factor 6)	$2.5\text{mm}^3$	16:05
$S_4$	SENSE (factor 6)	$4\text{mm}^3$	6:29

In addition to 4D-flow MRI, structural cine-MRI data were also acquired in long-axis (LA), short-axis (SA) and four-chamber (4CH) views. The SA stack was obtained with a slice thickness of 10mm with 10mm spacing between slices. Reconstructed in-plane resolution for all cine-MRI acquisitions was  $1.4063\text{mm} \times 1.4063\text{mm}$ .

### 5.2.2.2 DTI INMI Data Set

Alongside the dual-resolution acquisition, acquisitions in three additional subjects were used from the DTI INMI study. For the 4D-flow MRI data, the number of phases acquired per cardiac cycle was fixed at 30, with  $V_{enc} = 1.5\text{ms}^{-1}$ . The machine used was a 3 Tesla Philips Achieva, with a slice thickness of 3mm with 3mm spacing between slices. SENSE parallel acceleration was utilised, with a SENSE factor of 2, with a number of signal averages of 1. Further details for each subject can be found in Tab. 5.3.

**Table 5.3:** DTI INMI 4D-flow MRI data: acquisition details for the three distinct volunteers.

Subject	B (009)	C (019)	D (028)
Flip angle (°)	10	10	10
Repetition time (ms)	8.21	6.8	8.4
Echo time (ms)	3.51	3.67	3.46
Acquired spatial res. (mm)	$3.03 \times 3.08 \times 3$	$3.04 \times 3.09 \times 3$	$3.03 \times 3.03 \times 3$
Reconstructed spatial res. (mm)	$2.27 \times 2.27 \times 3$	$1.93 \times 1.93 \times 3$	$2.27 \times 2.27 \times 3$

For each subject, the SA stack was acquired with a slice thickness of 8mm, with 10mm spacing between slices, alongside LA and 4CH views. The reconstructed in-plane spatial resolution was  $1.18 \times 1.18$ mm for all cine-MRI acquisitions.

### 5.2.3 Data Processing

#### 5.2.3.1 Segmentation and Volume Interpolation

Two distinct procedures were used to segment the endocardial surface, dependent on the data set used. For the dual-resolution data set, segmentation and registration were performed by the author, under supervision from a trained clinician, using the open-source medical image processing software 3D Slicer (<http://www.slicer.org>) [73]. Contours were delineated in end-diastole using the LA, 4CH and SA stack using the grow from seeds algorithm, before the SlicerHeart module was used to interpolate the contours, generating a 3D volume. This volume was then registered in time to produce the beating LV and endocardium.

For the DTIINMI data set, segmentation and registration were performed by a trained clinician using the commercial software package Circle Cvi42 (<https://circlecvi.com>). Contours were delineated in all cardiac phases using the LA, 4CH and SA stack, producing a 3D volume. Then, the feature tracking module was used to register the endocardial motion in time. Using feature tracking, the extracted boundary motion included torsional effects, with the aim of improving the accuracy of near-wall velocity.

3D Slicer was used due to software availability issues, where Circle Cvi42 was only available for a brief period.

To enable temporal super-resolution using the PINN, the collocation and boundary condition point clouds must be sampled more densely in time. As such, interpolation of the LV volume and endocardium is required. For both data sets, this was performed using

node-wise cubic spline interpolation in time, sampling every second time point to reduce fluttering artefacts. The resulting volume was upsampled by a factor of 5 temporally in each case. Velocity values for the endocardial motion were calculated by taking the derivative of the cubic spline interpolant.

### 5.2.3.2 4D-flow MRI

Processing of 4D-flow MRI data was performed using a self-built Python script, based partially on that found in [75]. The process for generating the requisite PINN data was as follows: 1) 4D-flow MRI DICOM files, each structured on a single slice, were first imported; 2) DICOM headers were used to determine the heart rate (HR) and frames per cardiac cycle (FPC), allowing for the calculation of the temporal resolution as:

$$\text{Temp Res} = \frac{60}{\text{HR} \times \text{FPC}}; \quad (5.7)$$

3) pixel values ( $p$ ) in each DICOM image were converted to velocity using the DICOM headers for the rescale slope (RS) and rescale intercept (RI) using the formula:

$$\mathbf{u} = \frac{1}{100}(\text{RS} \times \mathbf{p} - \text{RI}); \quad (5.8)$$

4) the spatial locations of each voxel centre were obtained using the DICOM headers for pixel spacing (dx), patient position ( $S_x, S_y, S_z$ ) and patient orientation ( $X_x, X_y, X_z, Y_x, Y_y, Y_z$ ), which define the location of the upper left corner of the DICOM image and the orientation of the  $x$  and  $y$  axes. These values form the matrix:

$$M = \begin{bmatrix} X_x \text{dx} & Y_x \text{dx} & 0 & S_x \\ X_y \text{dx} & Y_y \text{dx} & 0 & S_y \\ X_z \text{dx} & Y_z \text{dx} & 0 & S_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5.9)$$

which generates an array of pixel locations using matrix multiplication with the voxel indices; 5) the corresponding LV volume at each phase is used to mask the LV flow domain in the reconstructed 4D-flow MRI data volume; 6) the velocity data are converted to comma-separated files for input to the PINN.

### 5.2.4 Variable Calculation

#### Kinetic Energy

To improve comparability with the wider literature, kinetic energy (KE) is quantified as KE per unit volume (KE/ml), using the units millijoules per millilitre (mJ/ml). As such, KE/ml (which may be interchangeably referred to as simply KE throughout this chapter), is given by:

$$KE = \frac{1}{2}\rho |\mathbf{u}|^2, \quad (5.10)$$

where a value of  $\rho = 1066\text{kgm}^{-3}$  is used.

In-plane KE (IPKE) is calculated using a plane with the normal,  $\hat{\mathbf{n}}$ , aligned in the apico-basal direction. Then, IPKE is calculated throughout the LV as:

$$IPKE = \frac{KE_{IP}^{sum}}{KE^{sum}}, \quad (5.11)$$

where

$$KE_{IP}^{sum} = \sum_{i=0}^{N_t} (KE^i - KE_{TP}^i) \quad (5.12)$$

$$KE^{sum} = \sum_{i=0}^{N_t} KE, \quad (5.13)$$

for  $N_t$  voxels at time  $t$ , and

$$KE_{TP} = \frac{1}{2}\rho |\mathbf{u} \cdot \hat{\mathbf{n}}|^2. \quad (5.14)$$

#### Relative Pressure

Without an additional pressure constraint, the PINN model can only predict pressure up to a constant, and as such, can only predict relative pressure. Relative pressure in this chapter is formed by subtracting the mean pressure, calculated across all time points, from the pressure field predicted by the PINN. Millimetres of mercury (mmHg) are used as the standard unit instead of Pascals (Pa) for ease of referencing across the literature. Direct quantification of pressure drop,  $\Delta p$ , from 4D-flow MRI data uses the simplified Bernoulli approach, as detailed in chapter 4.

#### Haemodynamic Forces

Haemodynamic forces (HDFs) are calculated in the LV by integrating the pressure gradient,  $\nabla p$ , over the flow domain. A transformation of coordinates is performed to align

the  $x$ -axis with the lateral-septal (LS) direction, the  $y$ -axis with the inferior-anterior (IA) direction and the  $z$ -axis with the apical-basal (AB) direction. Then, components of the HDF in the LS, IA and AB directions are taken to be the  $x$ ,  $y$  and  $z$  components of the integrated  $\nabla p$  field.

## 5.3 Results

The results in this chapter are constructed as follows: first, results from the *in vivo* validation study are presented; following this, the model is applied to a small cohort of patients with varying levels of LV impairment following MI.

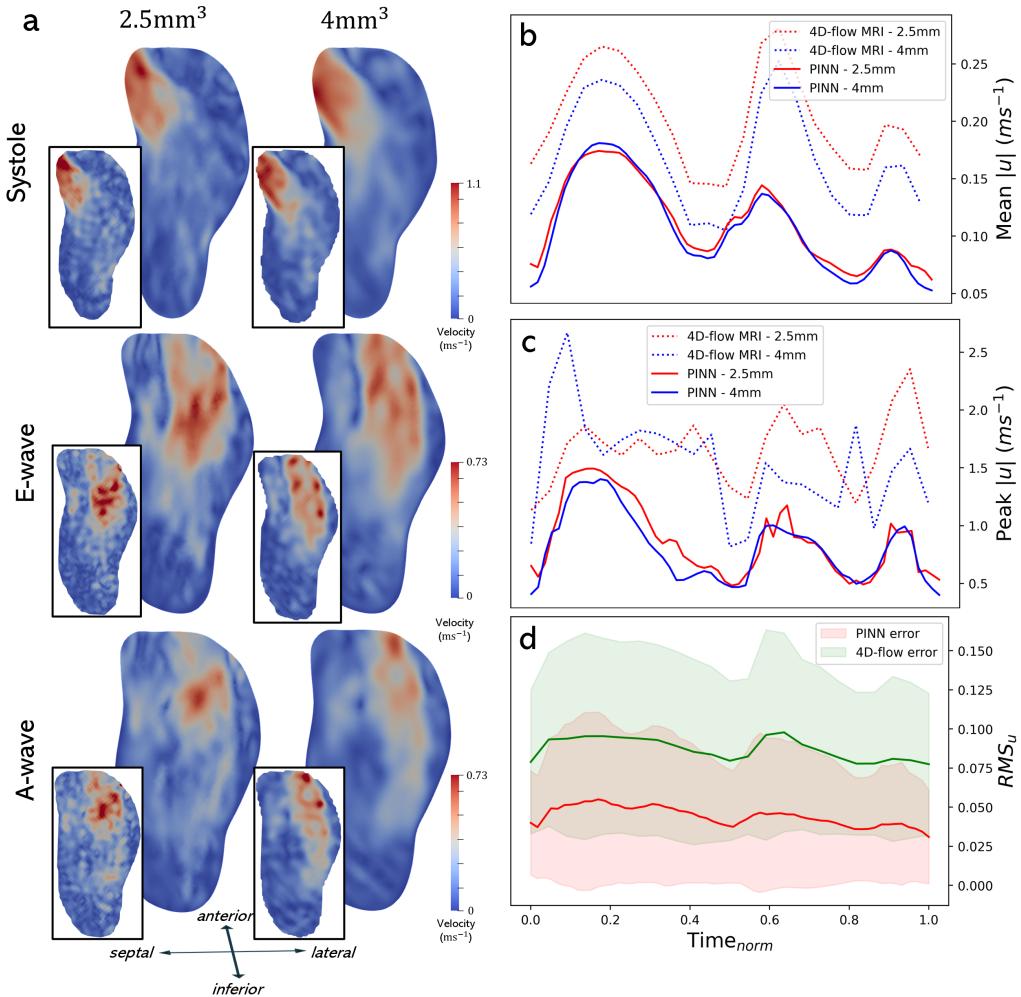
### 5.3.1 Model Validation: Dual-Resolution Study

In this section, results are presented using the dual-resolution 4D-flow MRI data set, for which low- ( $4\text{mm}^3$  voxel volume) and high-resolution ( $2.5\text{mm}^3$  voxel volume) 4D-flow MRI data were acquired in a single volunteer. A range of clinically-relevant variables are predicted using the PINN model with training data at the two different resolutions, which are compared against results obtained directly from the 4D-flow MRI data, where available.

#### 5.3.1.1 Velocity

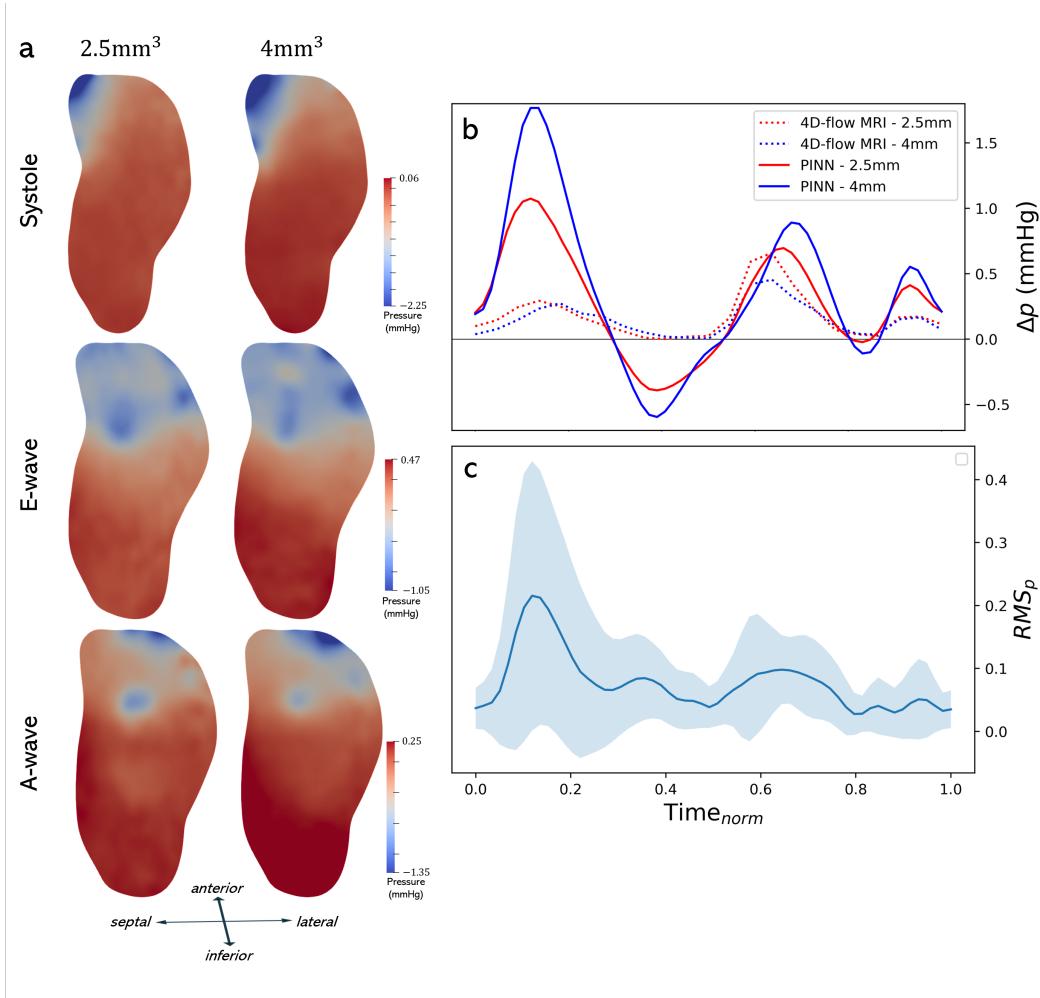
In Fig. 5.1(a), comparisons are made between PINN results obtained using the low- and high-resolution 4D-flow MRI data (namely spatial resolutions of  $4\text{mm}^3$  and  $2.5\text{mm}^3$ , respectively). As can be observed, in both cases the predicted velocity fields contain less noise than is observed in the corresponding 4D-flow MRI data, although it is not possible to conclusively discern noise from small-scale flow features without ground truth data. There are observable differences between the predicted velocity fields at each of the three cardiac phases. As seen in Figs. 5.1(b) and 5.1(c), there is a significantly greater similarity between PINN predicted fields than between the raw 4D-flow MRI data. Further, at both resolutions, the PINN predicted fields attain consistently lower mean and peak values than the corresponding 4D-flow MRI data throughout the cardiac cycle.

In Fig. 5.1(d), it can be seen that the magnitude of  $RMS_u$  (calculated point-wise) between



**Figure 5.1:** (a) 2D contours of the velocity magnitude taken at a long-axis slice in systole, peak E-wave and peak A-wave, comparing PINN results obtained using the high-resolution (left column) and low-resolution (right column) 4D-flow MRI data. Inserts display the raw 4D-flow data for the respective resolution; (b) Mean velocity magnitude throughout the cardiac cycle; (c) Peak velocity magnitude throughout the cardiac cycle; (d) Max-normalised root mean squared error in the velocity field ( $RMS_u$ )  $\pm 1$  standard deviation of the max-normalised relative error throughout the cardiac cycle, calculated directly between the low- and high-resolution 4D-flow MRI data (green) and between PINN fields predicted using the low- and high-resolution 4D-flow MRI data (red).

PINN predicted velocity fields correlates with the magnitude of  $RMS_u$  between the raw 4D-flow MRI data. This is similarly seen in Fig. 5.1, where the greatest qualitative differences can be observed during peak A-wave filling, where the 4D-flow MRI data appear least similar. Although the variance remains large, there is a marked reduction in error between the PINN predictions than between the raw 4D-flow MRI data, with a decrease in  $RMS_u$  and corresponding variance in diastole.



**Figure 5.2:** (a) 2D contours of the relative pressure field taken at a long-axis slice in systole, peak E-wave and peak A-wave, comparing PINN results obtained using the high-resolution (left column) and low-resolution (right column) 4D-flow MRI data; (b) Pressure drop ( $\Delta p$ ) between short-axis slices in the basal and apical region (where the simplified Bernoulli method [209] is used to calculate  $\Delta p$  directly from the 4D-flow MRI data); (c) Max-normalised root mean squared error in the pressure field ( $RMS_p$ )  $\pm 1$  standard deviation of the max-normalised relative error throughout the cardiac cycle, between PINN fields predicted using the low- and high-resolution 4D-flow MRI data.

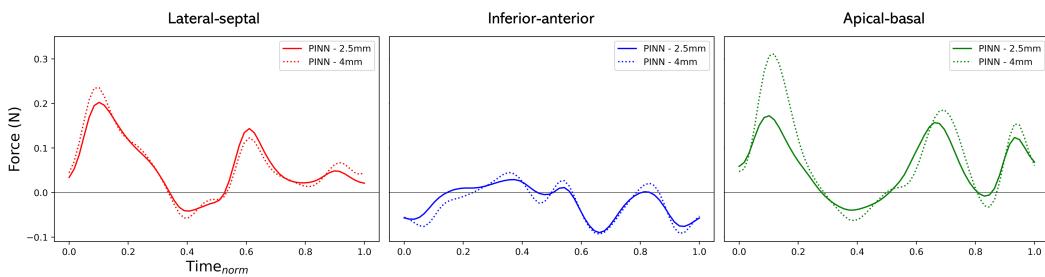
### 5.3.1.2 Relative Pressure

In Fig. 5.2(a), a qualitative analysis of the relative pressure field is made, in which the PINN predictions obtained using the low- and high-resolution data are compared. At the

three cardiac phases analysed, there are clear similarities between the predicted fields, with good agreement in the presence and position of localised regions of low pressure, particularly in diastole. There is, however, some discrepancy in the range of pressure values attained, which is most clear during systole. This is corroborated in Fig. 5.2(b), in which pressure drop ( $\Delta p$ ) between the basal and apical region is plotted throughout the cardiac cycle. Here, there is an obvious reduction in the maximum  $\Delta p$  when using the  $2.5\text{mm}^3$  4D-flow MRI data, particularly in the systolic phase, but also during E- and A-wave filling.

In Fig. 5.2(c),  $RMS_p$  ( $\pm 1$  standard deviation of the max-normalised relative error) between the PINN predicted fields is plotted throughout the cardiac cycle. Similar to  $RMS_u$ , there is a sharp reduction in the pressure error in diastole.

### 5.3.1.3 Haemodynamic Forces

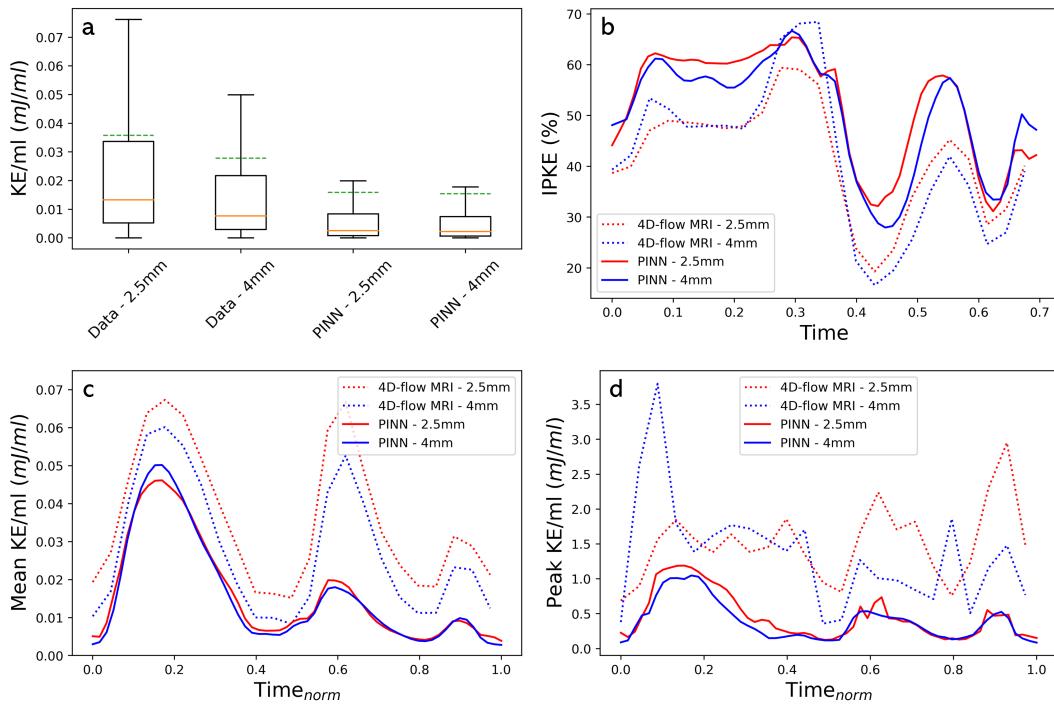


**Figure 5.3:** Haemodynamic forces (HDFs) in the LV in the lateral-septal, inferior-anterior and apical-basal direction throughout the cardiac cycle.

In Fig. 5.3, HDFs generated by the PINN using the low- and high-resolution 4D-flow MRI data are plotted throughout the cardiac cycle. In the lateral-septal and inferior-anterior directions, the predicted HDFs show good agreement. In the apical-basal direction, the agreement is acceptable in diastole but not in systole. This is in agreement with Fig. 5.2(b), where a large difference in peak  $\Delta p$  (which is intrinsically linked to HDFs) is observed in systole.

### 5.3.1.4 Kinetic Energy

In Fig. 5.4, KE is compared across low- and high-resolution 4D-flow MRI data and the corresponding PINN results obtained. In Fig. 5.4(a), box and whisker plots highlight the difference in KE/ml characteristics between PINN model predictions and raw 4D-flow MRI data, with PINN predictions exhibiting a significant reduction in peak and average



**Figure 5.4:** (a) Box and whisker plots for point-wise kinetic energy per unit volume (KE/ml) across the whole cardiac cycle, where the boxes represent the first and third quartiles, whiskers display maximum and minimum values (with outliers removed), the solid orange line represents the median value and the green dashed line represents the mean value; (b) In-plane KE (IPKE) throughout the cardiac cycle; (c) Mean KE/ml throughout the cardiac cycle; (d) Peak KE/ml throughout the cardiac cycle.

KE/ml when compared with the 4D-flow MRI data. In Fig. 5.4(b), IPKE is compared across the cardiac cycle. At both resolutions, there is a consistent increase in IPKE throughout the cycle. In Figs. 5.4(c) and (d), mean and peak KE/ml are compared across the cardiac cycle. In both cases, consistently good agreement between PINN model predictions is exhibited, with a marked reduction in mean and peak KE/ml from data to model predictions.

### 5.3.1.5 Limitations

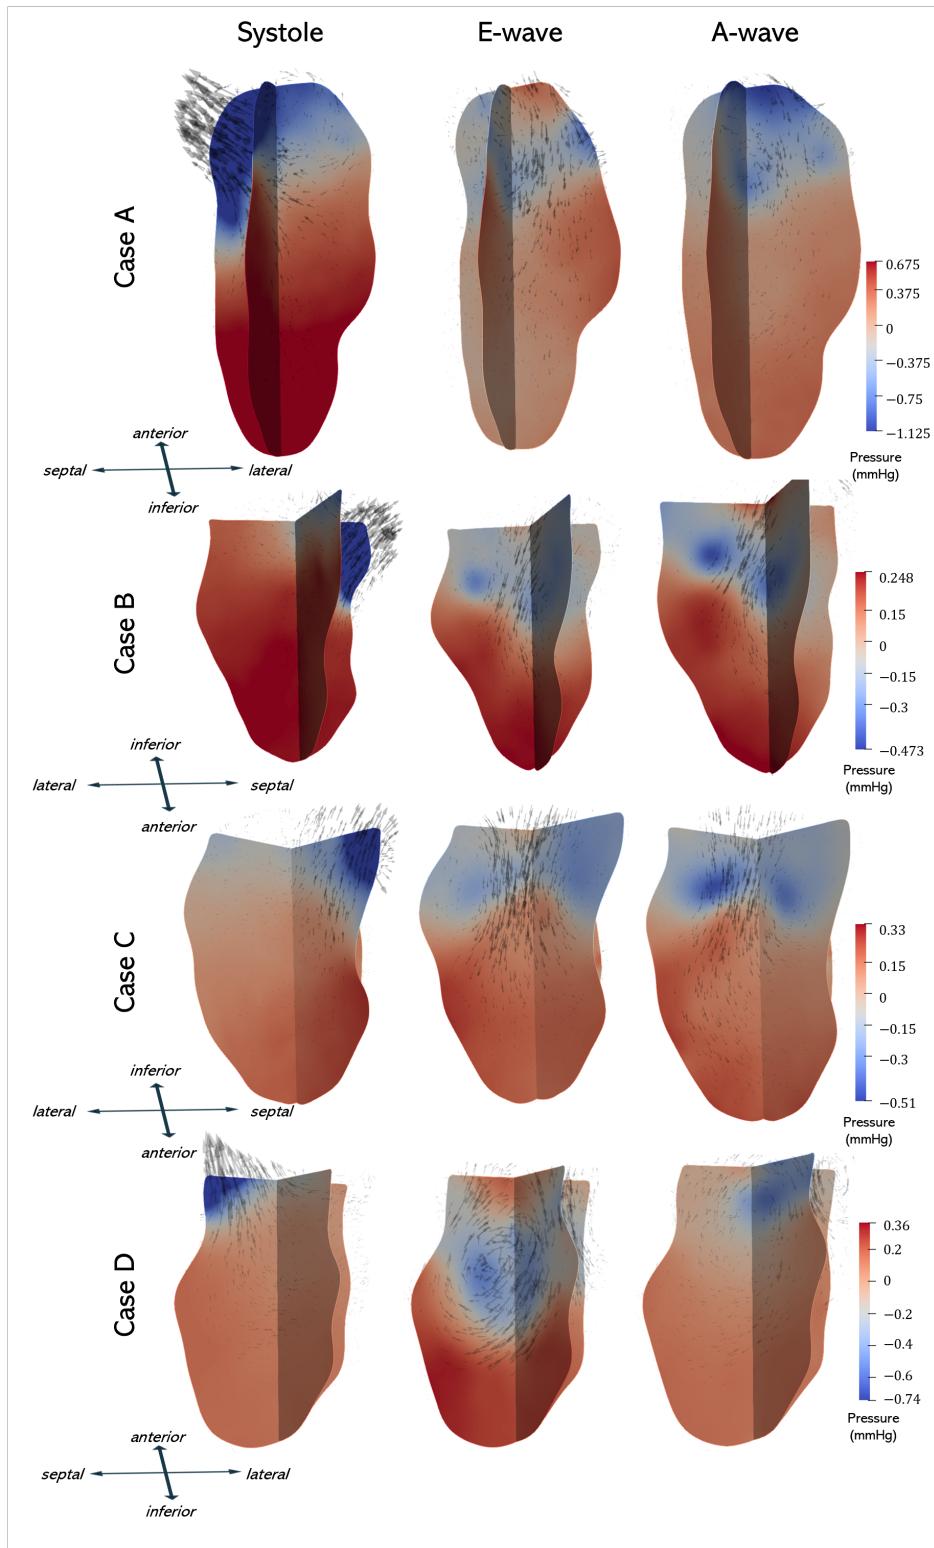
Both vorticity and WSS have not been considered in this section. Both variables were challenging to reconstruct without the effects of noise corrupting results. Therefore these have been omitted from this study.

### 5.3.2 Exploratory Study: Myocardial Infarction

In this section, the PINN model is used to produce results in a small cohort of patients with varying levels of LV impairment following MI. Clinically-relevant variables are generated

and assessed in connection with LV remodelling.

### 5.3.2.1 Relative Pressure

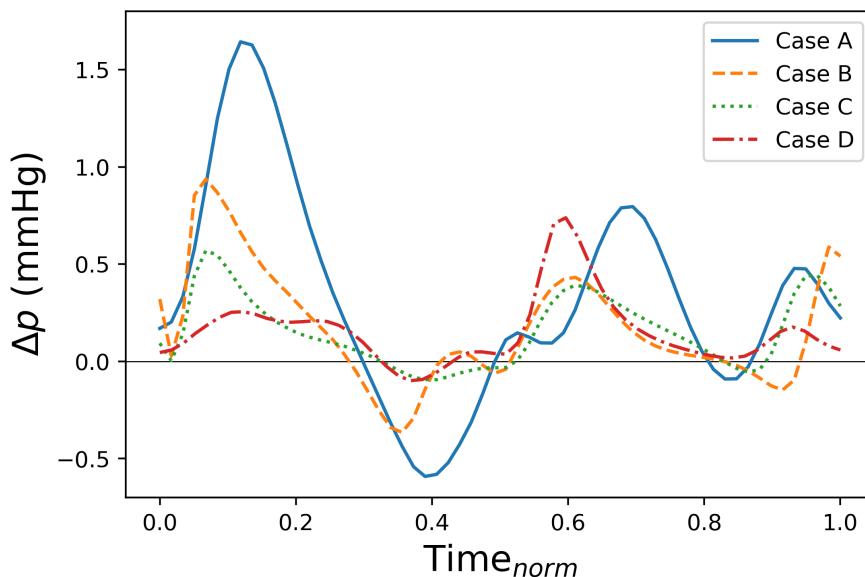


**Figure 5.5:** 2D contours of relative pressure on two long-axis planes in peak systole, peak E-wave filling and peak A-wave filling, with velocity vectors overlaid.

In Fig. 5.5, 2D contours of the relative pressure distribution are presented for all cases

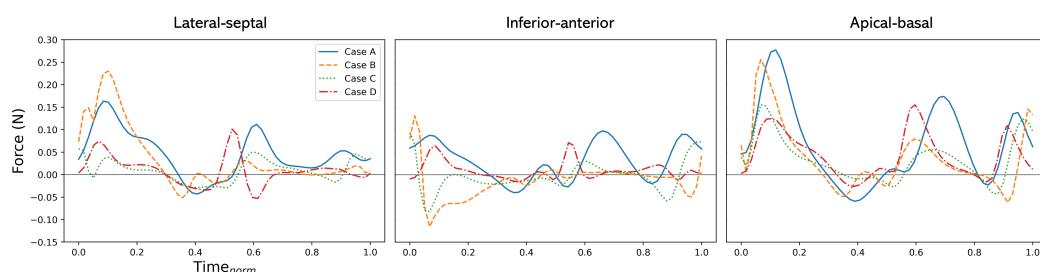
in systole, E-wave filling and A-wave filling. Pressure fields in both E-wave and A-wave filling are characterised by regions of relatively low pressure inside the mitral vortex ring, whereas the systolic pressure distribution is more uniform from apex to aortic valve, with low relative pressure in the outflow tract.

Curves of pressure drop,  $\Delta p$ , throughout the cardiac cycle, calculated between basal and apical slices, are shown in Fig. 5.6 for all cases. As expected, the greatest apico-basal pressure drop values are attained in the healthy subject (case A), with the smallest systolic  $\Delta p$  attained in the subject with severe LV dysfunction (case D), which is to be expected. Case D also experiences extreme pressure load in E-wave filling. This trend is visible in the relative pressure contours for cases A and D in Fig. 5.5.



**Figure 5.6:** Pressure drop,  $\Delta p$ , calculated between planes in the apical and basal region throughout the cardiac cycle for all four cases.

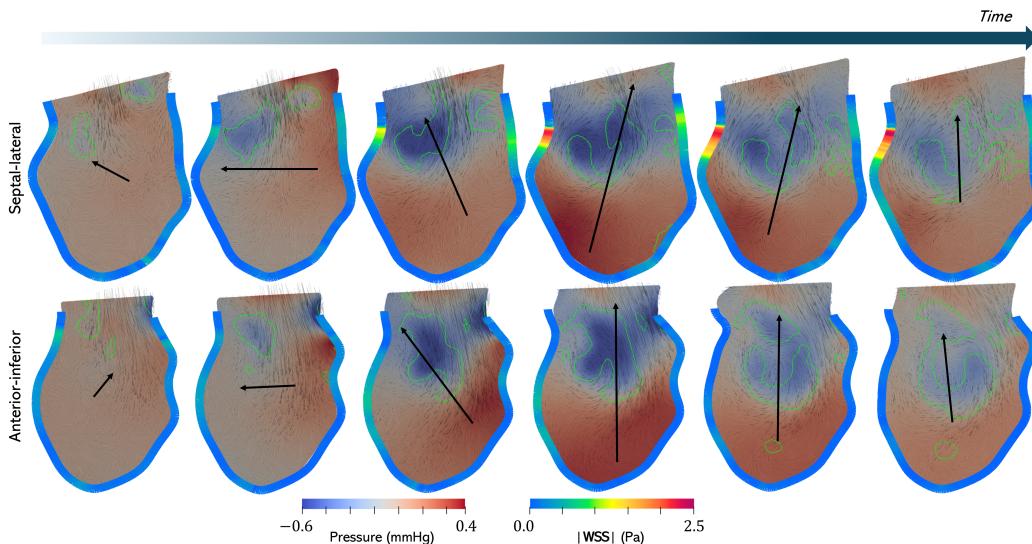
### 5.3.2.2 Haemodynamic Forces



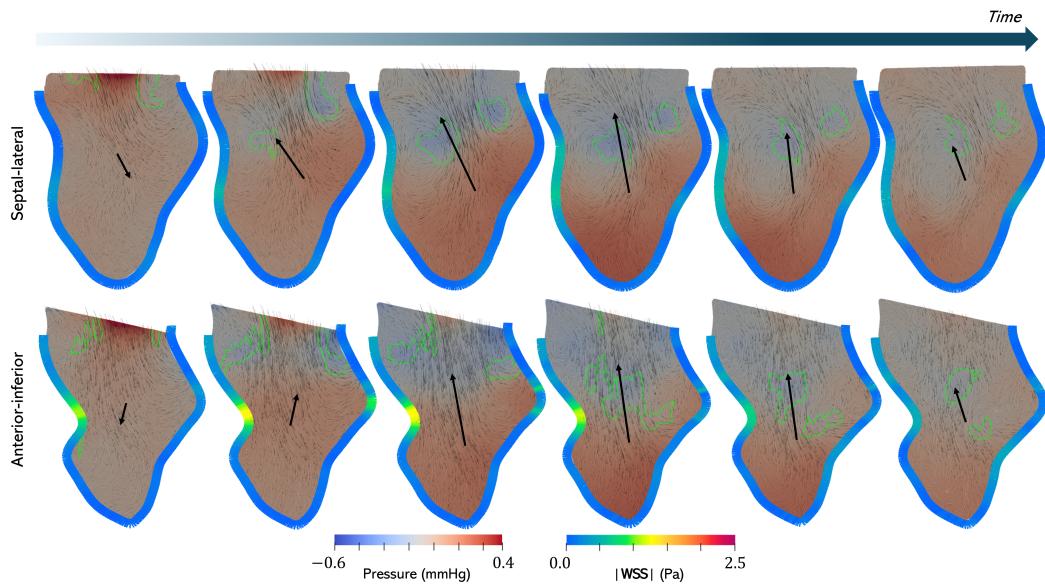
**Figure 5.7:** Lateral-septal, inferior-anterior and apical-basal HDFs throughout the cardiac cycle in all cases.

In Fig. 5.7, HDFs in the LS, IA and AB directions are presented for all cases. Curves for HDFs in the apico-basal direction resemble those of apico-basal pressure drop, although in systole, cases B and D produce more pronounced curves. This is likely a consequence of the LV in these cases having a greater volume, where the HDF is calculated by integrating the pressure gradient over the domain. As with pressure drop, case D experiences elevated HDFs in the apico-basal direction in E-wave filling, which is accompanied by an increase and reversal of HDFs in the latero-septal direction.

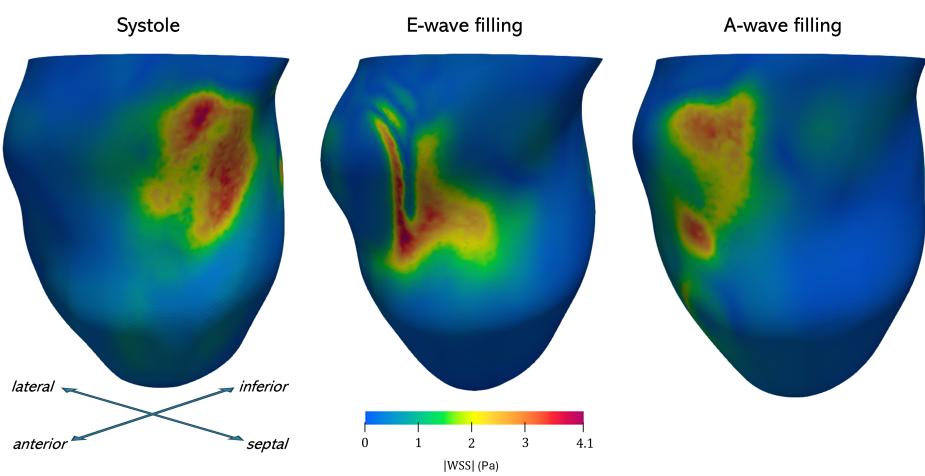
In Figs. 5.8 (case D) and 5.9 (case B), a selection of variables, namely relative pressure, velocity, vorticity, HDFs and WSS, are collated into one figure. These are plotted at 6 time points during E-wave filling, with the particular variables selected to demonstrate any potential interactions between them. In-plane coherent structures are delineated by taking isosurfaces of the Q-criterion ( $Q_{crit}$ ) at  $Q_{crit} = 200$ . In case D (severe LV remodelling, Fig. 5.8), the mitral vortex ring is unbalanced, with one dominant side producing transversal pressure gradients and thus transversal HDFs. This imbalance leads to elevated WSS on the septal wall where the larger vortex re-circulates upwards towards the base. Conversely, the mitral vortex ring in case B (moderate LV remodelling, Fig. 5.9) is more balanced, and as such, HDFs are more closely aligned with the apico-basal direction.



**Figure 5.8:** Case D: 2D contour plots of relative pressure on septal-lateral and anterior-inferior slices throughout E-wave filling, with the velocity vector field superimposed and WSS magnitude on the endocardium included around the boundary. Using in-plane isosurfaces of the Q-criterion, the green contours delineate the coherent structures of the mitral vortex ring. Large black arrows dictate the direction and magnitude (denoted by arrow length) of the HDFs in the LV.



**Figure 5.9:** Case B: See Fig. 5.8 for details.



**Figure 5.10:** Case C: WSS magnitude contours on the endocardium (bottom row) in systole, E-wave filling and A-wave filling, from left to right.

### 5.3.2.3 Wall Shear Stress

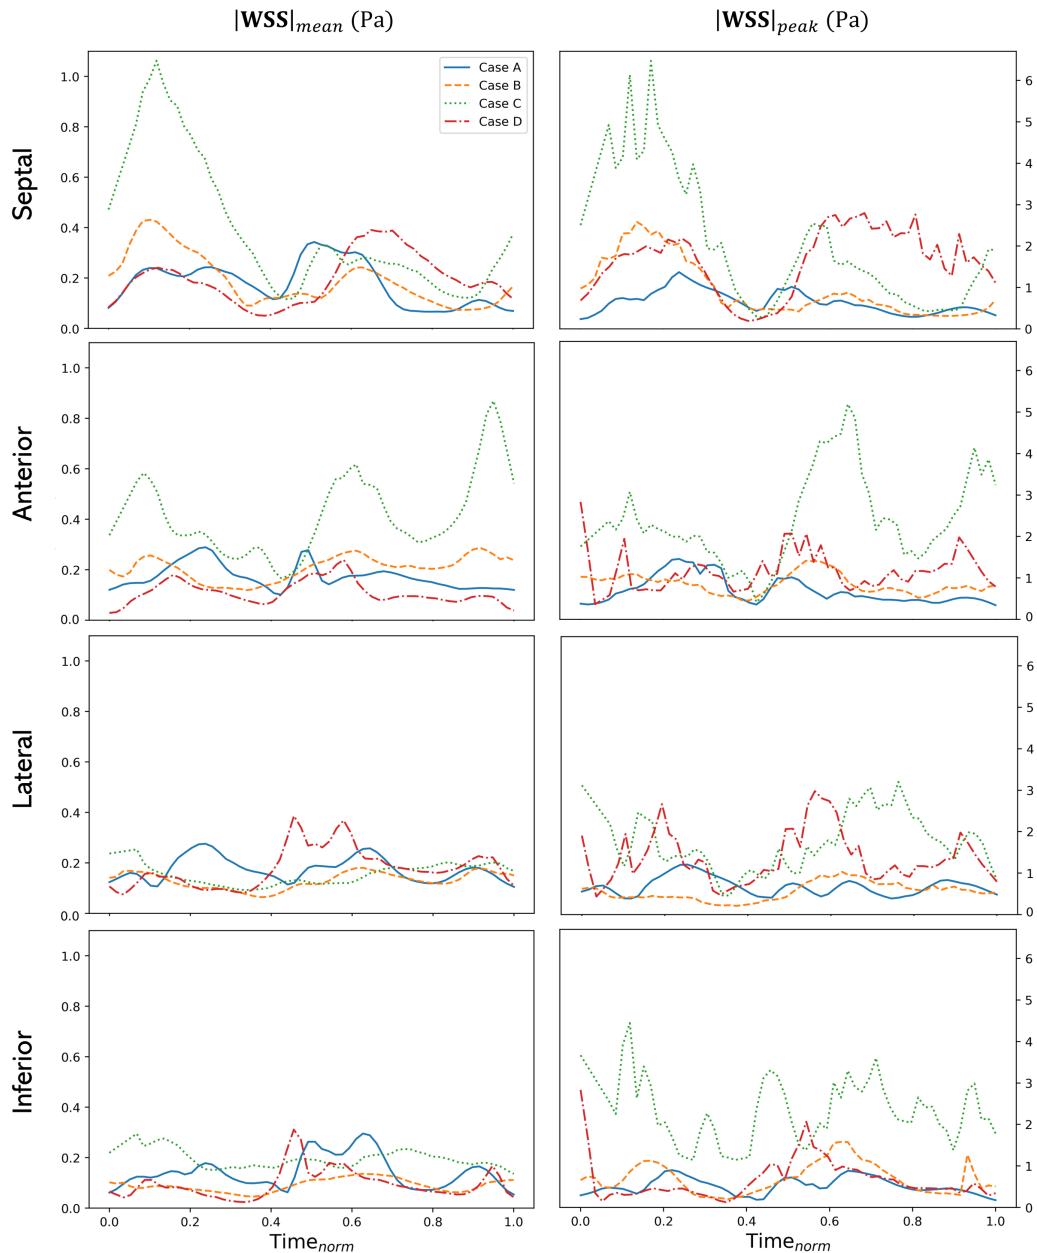
In Fig. 5.10, WSS contours are presented in systole, E-wave filling and A-wave filling for Case C. Elevated WSS magnitude can be observed in the septal region in systole, and the anterior region in both diastolic phases. In Fig. 5.11, mean and peak WSS is plotted throughout the cardiac cycle in the four major longitudinal regions of the endocardium, for all cases. In Fig. 5.11, mean and peak WSS are plotted through the cardiac cycle across the four major segments of the endocardium, in all cases. As depicted in Fig. 5.10, case C experiences relatively extreme WSS values in the septal and anterior regions. Case D displays high WSS in the septal region across diastole and in the lateral region in early E-wave filling.

### 5.3.2.4 Kinetic Energy

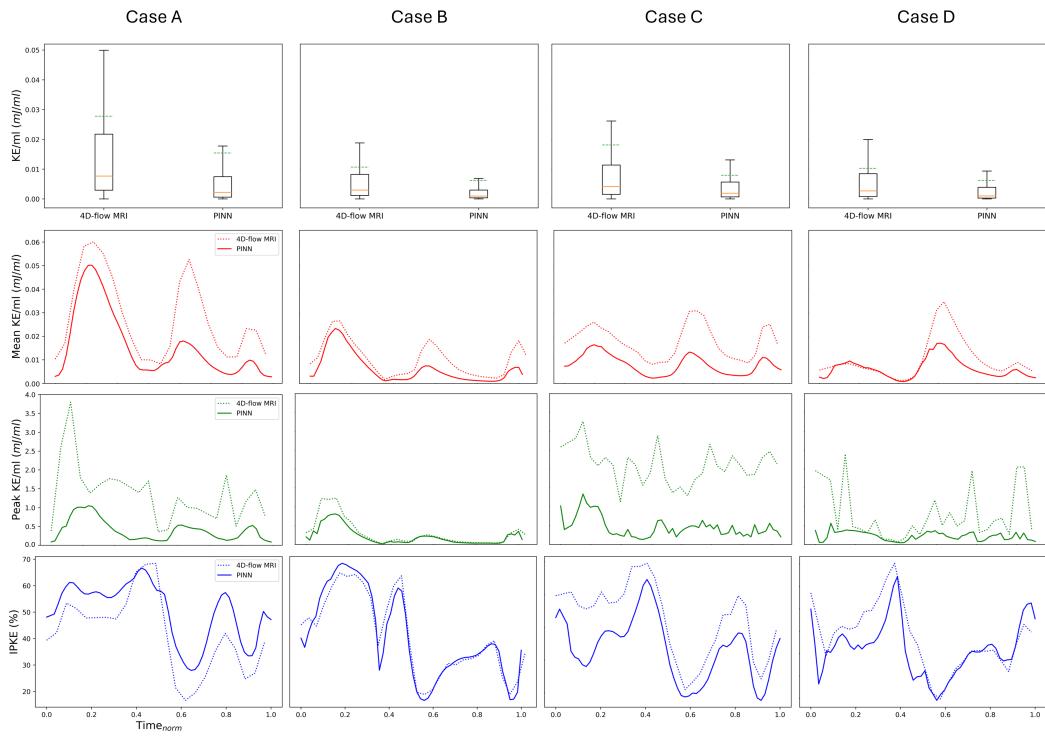
In Fig. 5.12, PINN-derived KE characteristics are presented across the cardiac cycle, with results compared against direct quantification from 4D-flow MRI data. KE per unit volume (KE/ml) is used instead of KE per unit mass to allow for comparison with other publications.

Across all cases, systole, E-wave filling and A-wave filling are discernible from mean KE and peak KE plots. IPKE is slightly elevated in systole, where flow ejected out of the aortic valve is generally not aligned in the vertical direction. Conversely, the mitral jet is typically directed from the base to the apex, and as such generates a higher proportion of through-plane KE. Case A (healthy) generates the largest KE values, with a particular increase in systolic KE. Case D experiences the lowest systolic KE, which correlates with the reduced pressure drop from apex to base in Fig. 5.6.

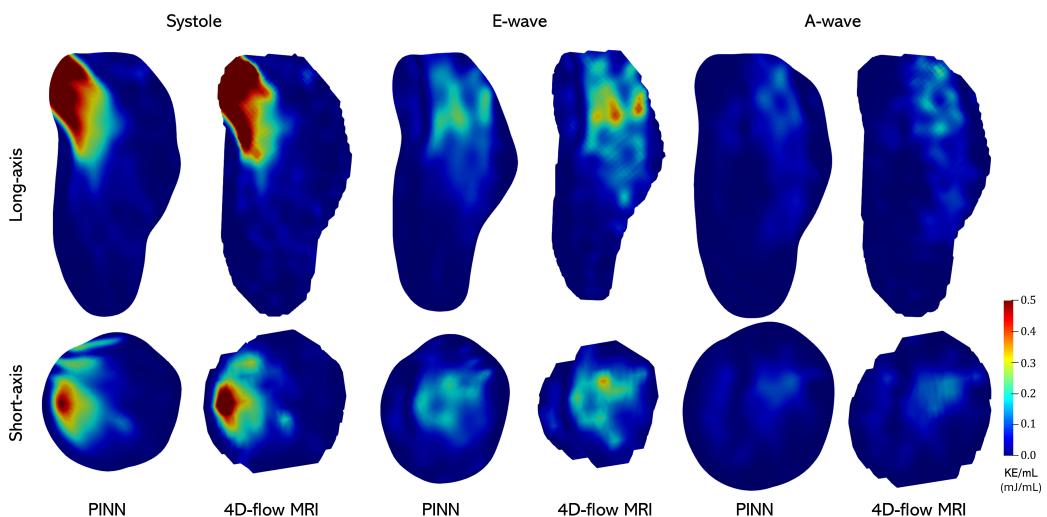
2D contour plots of KE in case A can be found in Fig. 5.13, with results obtained from the PINN model compared with direct quantification from 4D-flow MRI.



**Figure 5.11:** Mean WSS (left) and peak WSS (right) across the cardiac cycle in the septal, anterior, lateral and inferior regions of the endocardium for all cases.



**Figure 5.12:** Comparison of KE (per unit volume) characteristics of 4D-flow MRI data and PINN cases A,B,C and D, from left to right. (a - c) box and whisker plots for point-wise KE/mL across whole cardiac cycle, where the boxes represent the first and third quartiles, whiskers display maximum and minimum values (with outliers removed), the solid orange line represents the median value and the green dashed line represents the mean value. (d - f) mean KE/mL throughout the cardiac cycle. (g - i) peak KE/mL throughout the cardiac cycle. (j - l) in-plane KE throughout the cardiac cycle.



**Figure 5.13:** Case A: 2D contour plots of KE per unit volume (KE/mL) in peak systole (left), peak E-wave (centre) and peak A-wave (right), derived from the PINN and 4D-flow MRI at long-axis and short-axis planes.

## 5.4 Discussion

### 5.4.1 Model Validation

In this chapter, model validation was performed by producing results using the PINN model trained with 4D-flow MRI data of different spatial resolutions. Where possible, these results are compared against those directly quantified from the 4D-flow MRI data.

Across all variables analysed, there is generally good agreement between the PINN results obtained using the low- and high-resolution data. This contrasts with the results obtained directly from the 4D-flow MRI data, where results are inconsistent between the two spatial resolutions.

The largest discrepancies between PINN predictions occur in systole, which is particularly notable in the apico-basal direction as demonstrated in the  $\Delta p$  and corresponding HDFs curves. The largest pressure gradients experienced in the LV occur in this direction (and in phase of the cardiac cycle), and therefore it should be expected that this is where the greatest difference in pressure and HDFs should occur. However, this trend is not as clearly observed in the velocity predictions, where  $RMS_u$  is only slightly elevated in systole, and mean and peak  $|\mathbf{u}|$  achieve good agreement. Given that the enforced boundary motion is identical between the low- and high-resolution PINN results, and that the reconstructed velocity field is similar (globally, at least), it is unclear what drives the discrepancy in relative pressure and HDFs in systole. Additionally, it is unclear why this discrepancy is only exhibited in the apico-basal direction. Analysing the lateral-septal HDFs in Fig. 5.3, the forces experienced are of a similar magnitude to those seen in the apico-basal direction, although agreement is far better. This suggests that the discrepancy is not simply linked to the magnitude of the pressure difference. Overall, the qualitative similarities between PINN model predictions using low- and high-resolution data are displayed in Fig. 5.2(a). Given the qualitative differences exhibited in the velocity field, as shown in Fig. 5.1(a), the visual agreement shown in the pressure field is perhaps surprising.

As shown in Fig. 5.4, similar trends as observed with synthetic data in section 4.3.2.3 of chapter 4 emerge, where the overall KE is significantly higher when quantified directly from 4D-flow MRI data. In chapter 4, having access to ground truth data made it possible to assert that KE was significantly over-estimated when measured directly from 4D-flow MRI, where it was deduced that the presence of noise in the data was responsible for this. In the present study, it is not possible to conclusively state this fact in the absence

of reference data, however it appears likely that *in vivo* 4D-flow MRI is affected by noise in the same manner. This is corroborated by the clear differences between KE measured from the 4D-flow MRI data at low- and high-resolutions, with the former experiencing lower peak and average KE values. With 4D-flow MRI, SNR scales with voxel size, since the total signal received within a voxel is greater. Therefore, with the low-resolution 4D-flow MRI data exhibiting lower overall KE characteristics, it appears likely that higher levels of noise in the data contributes to the over-estimation of KE. Overall, there is very good agreement in KE between the PINN predictions at low- and high-resolutions. It is, of course, unclear in this study whether this is due to model bias or that the predicted values are representative of the ‘true’ KE field.

#### 5.4.2 LV Function Trends

As seen in Fig. 5.6, there is a clear discrepancy between healthy (case A) and severe dysfunction (case D) in systolic pressure drop, where the subjects with moderate (case B) and mild (case C) dysfunction attain systolic pressure drop values that lie between cases A and D. It would be expected that case C would display larger systolic  $\Delta p$  values given increased LV function, however, case C was obtained directly after infarct, whereas case B was obtained 3 months after infarct, and thus LV function was partially recovered. A pronounced negative pressure drop (high basal pressure and relatively lower apical pressure) can be observed for cases A and B during isovolumic relaxation, inducing closure of the aortic valve. Whilst present in cases C and D, this feature is less pronounced. Cases B (moderate LV dysfunction) and D (severe LV dysfunction) exhibit pressure reversal between systole and diastole, which has previously been linked with the onset of dilated cardiomyopathy [224]. Relative to systolic  $\Delta p$ , elevated  $\Delta p$  is observed during early E-wave filling in case D. This is accompanied by a sharp increase in HDF in the apical-basal direction, as seen in Fig. 5.7, and elevated mean KE relative to systolic KE, as seen in Fig. 5.12.

In Fig. 5.5, qualitative differences in the pressure distribution in E-wave filling are apparent between case D and the others. The re-circulation region around the mitral jet, characterised by low relative pressure, is unbalanced in case D, with one side of the vortex ring far larger than the other. This is studied in detail in Fig. 5.8, where the interplay between velocity, relative pressure, vorticity, HDFs and WSS is examined. The core of the mitral vortex ring in E-wave filling (denoted by the green contours) has a significantly larger area on the anteroseptal side when compared with the inferolateral side, which

dominates the pattern of flow throughout the LV. This is most apparent in the latter two time points in the anterior-inferior slice, where the large coherent structure in the centre of the domain corresponds almost entirely to the anteroseptal side of the vortex ring. Such imbalances are potentially induced by altered patterns of myocardial relaxation, caused by inhomogeneous tissue properties post-infarct. Imbalance in the mitral vortex ring produces an imbalance in pressure, resulting in strong transversal HDFs, particularly in early E-wave filling. This is discernible in Fig. 5.7, with a sharp increase in HDFs in both the lateral-septal and inferior-anterior directions in early E-wave filling. This finding is consistent with observations made in [172], where the development of transversal HDFs was hypothesised to anticipate LV remodelling. However, its onset is likely not driven through direct load on the myocardium, but through endothelial mechanotransduction via shear stresses on the endocardium that are disturbed by maladaptive HDF distributions. This interaction is investigated here, where significantly elevated WSS in the basal septal region can be observed in case D, with minor elevation in the basal lateral region. This is similarly observed in Fig. 5.11, with mean and peak WSS elevated throughout E-wave filling in the septal region and in early E-wave filling in the lateral region.

A similar study is demonstrated in a less impaired subject, namely case B, in Fig. 5.9. In this case, the mitral vortex ring is more balanced in size, particularly in the septal-lateral plane, with two low-pressure regions of approximately equivalent volume. As such, the transversal HDFs are comparatively smaller in magnitude than in case D, with stronger alignment in the apico-basal direction. There is a slight increase in WSS observed in the mid anterior region in early E-wave filling in Fig. 5.9, which appears to persist throughout diastole in the mean WSS in Fig. 5.11. However, from Fig. 5.11, it appears that mean and peak WSS values attained in the anterior region are only slightly elevated (mean values) or lower (peak values), compared with case D.

Given the difference in clinical outcome between cases B (moderate LV dysfunction) and D (severe LV dysfunction), the interplay between velocity, pressure, HDFs and WSS appears influential. It has been previously noted that large transversal HDFs may anticipate adverse LV remodelling [172], but the intermediate influence of WSS has not been assessed. Results presented here suggest that there is an interaction between elevated transversal HDFs and WSS, and WSS in general, which in turn may influence clinical outcome, but more evidence is required to state this conclusively.

Another interesting finding is the pronounced increase in WSS magnitude observed in Fig. 5.11 for case C, particularly in the septal and anterior regions. This can also be seen

qualitatively in Fig. 5.10, with areas of high local WSS in the septal region (in systole) and the anterior region (diastole). Although the outcome after 12 months was only mild LV dysfunction, case C was acquired acutely post-infarct, as opposed to 3 months post-infarct for cases B and D. As such, any structural adaptation of the myocardium (either positive or negative) as a consequence of disturbed haemodynamic flow would not have yet occurred. This could indicate that WSS, and subsequent mechanotransduction on the endocardium, does indeed drive remodelling in the LV, although analysis in a larger cohort of subjects is required to state this conclusively.

### 5.4.3 Variable Prediction Using PINNs

#### Kinetic Energy

In chapter 4, it was observed that KE calculated directly from synthetic 4D-flow MRI was over-predicted, which was deduced to be a consequence of background noise artefacts. A similar trend emerges in the present chapter, in this section and in section 5.4.1, as evidenced in the top row of Fig. 5.12, where a discrepancy in KE magnitude between PINN results and those derived directly from 4D-flow MRI is observed. Comparing against results presented with synthetic data in chapter 4, it could be hypothesised that noise corruption, both background and structured, is again responsible for such a pronounced discrepancy. However, unlike in section 5.4.1, some abnormal patterns here suggest there may be problems with the PINN model reconstruction. In diastole, particularly in E-wave filling, the discrepancy in mean KE between the PINN and direct quantification is larger than in systole. If background noise corruption was the sole reason for the difference in KE magnitude, a similar drop-off would be expected in systole. This effect can be observed qualitatively in Fig. 5.13, where regions of peak KE are not fully resolved by the PINN in diastole.

Additionally, a distinction must be made between background and structured noise, in particular uniform noise that is present when tissue regions intrude into the phase images (phase offset error). Case C encountered a large amount of tissue intrusions, and as such, introduced uniform noise drawn from  $\mathcal{U} \sim [-v_{enc}, v_{enc}]$ . As seen for case C in Fig. 5.12, this induces an additive increase in peak and mean KE when comparing direct quantification with PINN results, which is independent of KE itself. Cases A, B and D were less impacted by tissue intrusions, which is corroborated in the mean KE curves. It appears the noise intensity in the 4D-flow MRI data is dependent on the underlying velocity magnitude (and thus KE), as is typical of background noise artefacts, matching the behaviour observed in

chapter 4.

Without ground truth data reference, it is difficult to state conclusively whether the differences in KE characteristics arise as a consequence of de-noising by the PINN model, or incorrect prediction of the velocity field. However, based on the results in this section and those in section 5.4.1, it is likely that both contribute.

### Wall Shear Stress

As was discussed at length in chapter 4, positional uncertainty can hamper the accuracy of WSS calculations, particularly in attaining the correct magnitude. However, it was shown that the large-scale WSS distribution remained approximately the same across distinct boundary representations when using the PINN, with regions of highest and lowest WSS correctly predicted. As such, whilst the true WSS magnitude and distribution on the trabeculated endocardium may be impossible to recover using an approximated surface, there is at least consistency in the predicted distribution and magnitude using the PINN. Under the assumption that under-prediction of the magnitude is consistent between cases, the use of an ‘effective’ WSS may be suitable for the purpose of marker identification.

In the model validation study, WSS predictions were severely corrupted, and hence were omitted from this chapter. It is unclear why the results were so poor and whether this was a result of the data or the model itself, but factors such as phase offset error, segmentation errors and training issues likely contribute. Predictions with the DTI INMI data set were more consistent, although the validity of these results is unclear.

## 5.5 Conclusion

In this chapter, we have presented results from a limited *in vivo* validation study and produced some initial results in a small cohort of LV remodelling patients with varying levels of impairment. It was found that global measures of variables such as velocity, pressure and KE were largely independent of the input resolution of the training data, where significant differences were found when comparing direct quantification from the 4D-flow MRI data at different resolutions. Whilst the validation study used too few cases to rigorously validate the PINN model *in vivo*, it introduced a methodology for achieving this which may be expanded on in future work.

The PINN model was also used to produce clinically-relevant variables in four subjects across healthy, mild, moderate and severe outcomes with regards to LV dysfunction. The interplay between vorticity, relative pressure, HDFs and WSS was investigated in two cases, with severe and moderate adverse outcomes. It was demonstrated that imbalanced mitral vortex formation in E-wave filling contributed to an increase in transversal HDFs, which subsequently produced elevated WSS. Increased mean and peak WSS was also observed in the case acquired acutely post-infarct, indicating that WSS may be a trigger for remodelling.

The variables of interest presented in this chapter have been previously assessed *in vivo* in the literature, however, few approaches are able to derive all variables discussed in a unified manner. The PINN model is able to quantify a wide selection of variables that are not directly measured by 4D-flow MRI, for which multiple models would have been previously required to obtain. As such, it is possible to investigate the interplay between particular variables and their role in driving clinical outcomes.

Given apparent discrepancies between parameters calculated directly from 4D-flow MRI and those obtained with the PINN model, particularly those associate with KE, rigorous validation *in vivo* or *in vitro* is required to bridge the gap between results obtained in synthetic data and those obtained *in vivo*.

# Chapter 6

## Conclusions and Outlook

In this thesis, a physics-informed neural network (PINN) model is introduced for the super-resolution of cardiac 4D-flow MRI data. The PINN not only provides the velocity field at a higher resolution, but is also capable of directly exposing a wide array of solutions fields such as relative pressure and flow derivatives, without the need for numerical approximation. The model is validated across three distinct studies, comprised of two synthetic and one *in vivo* setups. Following this, the model is applied to a small cohort of volunteer and patient data to demonstrate the applicability of the model to real data. It is concluded that the use of this type of model is feasible in this setting, although certain limitations should be addressed and the model should be further validated using *in vivo* or *in vitro* data.

### 6.1 Summary

Phase-contrast magnetic resonance imaging (PC-MRI) has emerged as a powerful tool for the non-invasive quantification of haemodynamic markers in the cardiovascular system, improving understanding of the links between maladaptive blood flow and cardiovascular disease (CVD). By measuring changes in the spin angle of protons under two equal but opposite magnetic gradient pulses, PC-MRI is able to reconstruct blood velocity either on a 2D plane (planar PC-MRI) or within a 3D volume (4D-flow MRI). From the measured velocity field, it is possible to extract variables such as haemodynamic forces (HDFs), wall shear stress (WSS) and relative pressure drop using a variety of techniques.

For analyses in the heart chambers, 4D-flow MRI is the dominant PC-MRI technique

used today, owing to its ability to reconstruct time-resolved velocity fields in three spatial dimensions. This allows for comprehensive assessment of haemodynamic variables, which are often 3D in nature, throughout the whole cardiac chamber at once. However, the accuracy of both the velocity field and any quantities derived from 4D-flow MRI can be limited, owing to low spatio-temporal resolution and noise artefacts which are inherent to the modality. This has led to a recent interest in *super-resolution* methods that are designed to address the aforementioned limitations, although this field is emerging. As such, only a limited number of approaches have been developed at the time of writing, all of which are applied to the simpler task of vascular flow.

In this thesis, PINN model is introduced to provide super-resolution and de-noising of cardiac 4D-flow MRI. PINNs have previously been demonstrated to be an effective technique for 4D-flow MRI super-resolution [72]. Through weak enforcement via the loss function, the underlying physics governing the haemodynamic flow acts to regularise the network during training. This ensures that model outputs conform to the correct physics and boundary conditions, and allows the model to perform super-resolution without paired low- and high-resolution training data, which are not readily available for *in vivo* 4D-flow MRI studies. Further, the physics-based regularisation allows for the quantification of variables that are not directly measured, such as relative pressure, without the application of additional techniques which can be complex to use and may introduce numerical error.

In chapter 3, the PINN model was introduced and applied to an idealised synthetic data set. The overall aims for this chapter were two-fold: Firstly, it was demonstrated that a PINN model can provide super-resolution of 4D-flow MRI data and is robust to the effects of low spatio-temporal resolution and signal-to-noise ratio (SNR); Secondly, a model architecture was established that was able to operate effectively across data at different spatio-temporal resolutions and SNRs. To this end, synthetic 4D-flow MRI data were generated across a range of spatio-temporal resolutions and SNRs, from computational fluid dynamics (CFD) results generated in a 2D idealised left ventricle model. Then the PINN model was used to reconstruct the corresponding solution fields, where it was demonstrated that the optimised model configuration was suitably robust across the range of spatio-temporal resolutions and SNR levels examined.

In chapter 4, the PINN model was applied to a small cohort of synthetic cases generated using CFD in patient-specific LV geometries, for which three distinct boundary representations were produced in each case. The primary aim of the chapter was to assess the impact

of positional uncertainty (i.e. uncertainty that arises due to errors in the segmentation and registration of the deforming endocardium), but it was also used to demonstrate the effectiveness of the model in calculating clinically-relevant quantities that are not directly measured by 4D-flow MRI in a realistic setting. The accuracy of such quantities was compared against baseline methods from the literature, with the PINN model shown to outperform most competing approaches. It was also demonstrated that positional uncertainty played a minimal role in the quantification of variables measured away from the endocardium, although variables measured in the near-wall region differed significantly across boundary representations.

In chapter 5, a limited *in vivo* validation study was presented, for which PINN model results were obtained using 4D-flow MRI data at two distinct spatial resolutions. Then, the PINN model was used to obtain a variety of clinically-relevant quantities in a small cohort of LV remodelling patients, at varying levels of impairment. It was shown that PINN model predictions were largely independent of the resolution of the input data, with the validation study providing a blueprint for a larger study in the future. Further, whilst the number of cases was limited, there were a number of interesting trends observed across the diseased population. However, it is noted that such studies were limited and are required to be extended in future work to make any conclusive statements.

## 6.2 Limitations

### 6.2.1 Hyperparameter Sensitivity

Hyperparameter tuning is a process that is ubiquitous throughout ML, however, from experience, PINN models appear particularly susceptible to sub-optimal hyperparameter choices. This is particularly true for hyperparameters related to the weighting of loss terms, where gradient imbalances impart a significant influence on the stability of model training. In this thesis, over 150 training runs were required in each study to obtain acceptable results. Many of the model configurations obtained were unusable, due to either over-fitting to data noise or under-fitting.

Whilst dynamic weighting methods are an active area of research within the field of physics-informed ML, many approaches still introduce sensitive parameters that impart a strong influence on the quality of results. Additionally, much of the related literature

introduces such methods using a small number of simplified cases, and does not demonstrate that a single parameter configuration is able to operate across differing domains and conditions.

To progress the current work, more effort should be placed into identifying or formulating novel dynamic loss weighting strategies that require less case-specific tuning. This is critical for this type of model to be adopted within research and clinic, where extensive hyperparameter tuning will not be possible.

### 6.2.2 Inter-Case Generalisation

As discussed in chapter 3, PINNs (in the generalised form) have no mechanism for inter-case generalisation, with case-specific training required for each new geometry, set of boundary conditions or set of physical parameters. In the present work, this is a hindrance when attempting to scale up the number of LV cases used within a study, particularly in relation to the previous section on hyperparameter sensitivity.

To address this, a future avenue of work could be to adopt a transfer learning approach, in which a shared set of model weights are used across distinct, but similar, domains. This would significantly reduce the training time required to ‘fine-tune’ the model in each case, where typically only weights in the last layer of the network are trained.

### 6.2.3 Wall Shear Stress

Of the variables that the PINN model is able to generate, WSS is the most problematic to accurately reconstruct. In chapter 5, it was shown that WSS was consistently under-predicted, particularly when positional uncertainty was a factor, but also somewhat when the true LV boundary was used for model training and evaluation. Although partially investigated in a handful of small studies [191, 8], the suitability of PINNs for the quantification of WSS has not been fully characterised.

Positional uncertainty exerts a significant influence on WSS calculations due to the sensitivity of the variable to inaccuracies in the boundary layer. This is problematic in the present application, where significant positional uncertainty is present in endocardial surfaces segmented from MRI. The validity of WSS when calculated on a surface that is affected by such large amounts of uncertainty is questionable.

To address the aforementioned issues, the capability of PINN models to accurately reconstruct WSS should be assessed across a range of simplified fluid flow scenarios at various Reynolds numbers, with synthetic 4D-flow MRI data at varying spatial resolutions. Since WSS is typically under-predicted when calculated from 4D-flow MRI, either directly or using curve fitting, it should be assessed whether PINNs are able to successfully overcome this.

#### 6.2.4 Synthetic Data Limitations

The synthetic data used within this thesis include significantly idealised, 2D CFD models and CFD models within patient-specific geometries, albeit with simplifications. These simplifications include contraction and relaxation of the myocardium being driven by a sinusoidal function in time instead of realistic volume curves with distinct E-wave and A-wave phases, which results in a deep penetration of the mitral jet into the ventricle. This is not typically observed in real LV flow, where deceleration of the mitral jet occurs in the basal region in both E-wave and A-wave filling. However, it is unlikely that the PINN model would face any additional challenge modelling LV flow with these characteristics, which is the primary purpose of these studies.

An additional modelling assumption made within the synthetic data is the use of a fixed mitral valve, defined on the mitral valve plane, without the inclusion of flexible valve leaflets. Whilst this construction is not representative of the true LV physiology, it produces flow patterns that are characteristic of LV flow, primarily in the formation of the mitral vortex ring. While it is certain that the inclusion of flexible valve leaflets would provide a more accurate representation, they would introduce significant computational expense, given that a two-way fluid-structure interaction would be required to model them. It was determined that, for the purposes of the study, the fixed mitral valve presented an accurate enough representation of LV flow.

Finally, the synthetic data used within this thesis included artefacts introduced by background MRI noise only, and did not consider structured noise. Structured noise, arising due to patient motion and ghosting, for example, is more challenging to synthesise but may impart a strong influence on results. The addition of these artefacts should be considered in future work.

## 6.3 Outlook

The key obstacle in the adoption of computational methods in cardiovascular research is appropriate validation of models. This thesis has introduced a technique for super-resolution of cardiac 4D-flow MRI data and demonstrated its applicability in a small number of synthetic and real cases. However, extensive validation is still required for such a model to be used in practice.

### 6.3.1 Model Validation

It is the authors opinion that model validation should be undertaken in four stages: Validation of variable predictions should first be demonstrated using synthetic, *in vitro* and *in vivo* data; This should be followed by a larger *in vivo* study utilising a patient population with diverse clinical outcomes, in which the capability of the model to produce haemodynamic metrics that can be used to correctly classify patient outcomes is assessed.

#### 6.3.1.1 Validation With Synthetic Data

Synthetic data sets are useful for validation as they provide ground truth data that are of an arbitrarily high resolution and free of noise artefacts. They also allow for the introduction of data corruption in a controlled manner, where the level of degradation can be rigorously controlled. However, they should not be used in isolation, as not all data degradation types can be readily simulated. Examples include structured noise artefacts, such as those introduced by patient motion, and the complex averaging of the velocity field experienced throughout the MRI acquisition. With these effects missing, synthetically generated data presents a simplified picture of 4D-flow MRI limitations.

In chapters 3 and 4, synthetic data have been used to configure the PINN model and demonstrate model efficacy. Whilst the number of cases used was limited, synthetic 4D-flow MRI data were generated using different spatio-temporal resolutions and SNR values in each case, which helped to replicate variability between acquisitions. Still, this type of study should be extended to explore the variability in flow patterns and values, such as Reynolds number, experienced with real LV flow. The generation of CFD data can be a significant bottleneck, where simulating fluid flow in a 3D, patient-specific heart can take weeks to complete. However, whilst it is important to include some complex synthetic

cases, replicating similar flow characteristics in simpler models may suffice, allowing greater exploration of flow patterns and characteristics.

### 6.3.1.2 Validation With *In Vivo* Data

As discussed in chapter 5, validation of a 4D-flow MRI super-resolution method *in vivo* can be challenging. Firstly, cardiac 4D-flow MRI is limited by the spatio-temporal resolution that is possible to achieve within a reasonable scan time, with an upper limit on the spatial resolution of around  $2.5\text{mm}^3$ . There is no alternative modality that can provide flow data at a higher resolution *in vivo*, and therefore a voxel size of  $2.5\text{mm}^3$ , which is still relatively coarse, is the best resolution available. Secondly, variability in heart rate and patient motion between acquisitions leads to poor repeatability and can produce artefacts in the data, particularly the case of cardiac 4D-flow MRI.

Due to the above, using higher-resolution 4D-flow MRI scans as reference data is not a plausible approach. Where in the synthetic studies in this thesis it was possible to attempt to reconstruct the reference CFD results (which are of a very high resolution and free from noise) from corrupted flow data, such an approach is not recommended for *in vivo* data. Instead, a more suitable approach is to produce PINN model results using both the low- and high-resolution 4D-flow MRI data and assess the consistency of predictions, focused more on globally measured variables instead of point-wise velocity values. Whilst this has been done at two resolutions in this thesis, this could be extended to include three or four distinct scans at varying spatial resolutions. Additionally, more cases are required in a study of this type.

### 6.3.1.3 Validation With *In Vitro* Data

In addition to *in vivo* validation with 4D-flow MRI, it is possible to provide model validation *in vitro* using a physical phantom. This type of study has been undertaken to validate 4D-flow MRI methods previously [72, 75], where silicone phantoms were used to replicate flow in a section of the vasculature. In [75], low- and high-resolution 4D-flow MRI data were acquired at spatial resolutions of  $4\text{mm}^3$  and  $2\text{mm}^3$ , respectively. In [72], PIV imaging was used to obtain the high-resolution reference data, at a spatial resolution of  $0.141\text{mm}^3$ , with low-resolution training data obtained using 4D-flow MRI at a spatial resolution of  $0.57\text{mm}^3$ .

With a phantom setup, issues around repeatability are alleviated, with no limitation on scan duration and no variation in heart rate or patient position. This allows for acquisitions at a higher resolution and SNR than is possible *in vivo*, with limited artefacts arising in the form of structured noise. Further, the use of more accurate flow imaging techniques, such as PIV, is possible *in vitro* by using transparent (or translucent) silicone to construct the flow domain walls. It would also be possible to measure additional variables, such as pressure, directly, which would allow for validation of such variables when approximated with the PINN model.

Constructing a beating LV phantom would likely be a challenging task. Instead, replicating the key flow conditions in a rigid-wall setup would likely be sufficient alongside synthetic and *in vivo* validation studies in deforming domains.

#### 6.3.1.4 Patient Classification

Given the PINN model introduced in this thesis is intended to be used in stratifying LV remodelling patients, it is important to assess the ability of the model to correctly perform this task. To do this, a cohort of patients with known clinical outcomes may be used, with 4D-flow MRI data acquired acutely post-infarct. The validation task is to attempt to correctly classify the diseased subsets of the population (using a classification model that is distinct to the PINN), using the PINN results as input. This may be performed by analysing variables that have been previously linked with ALVR, but may also allow for new haemodynamic marker discovery.

# Bibliography

- [1] Martín Abadi et al. “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems”. In: (2016). URL: [www.tensorflow.org..](http://www.tensorflow.org..)
- [2] Haruhiko Abe et al. “Contrast echocardiography for assessing left ventricular vortex strength in heart failure: A prospective cohort study”. In: *European Heart Journal Cardiovascular Imaging* 14.11 (Nov. 2013), pp. 1049–1060. ISSN: 20472404. DOI: [10.1093/ehjci/jet049](https://doi.org/10.1093/ehjci/jet049).
- [3] Mahbubul Alam et al. “Assessment of left ventricular function using mitral annular velocities in patients with congestive heart failure with or without the presence of significant mitral regurgitation”. In: *Journal of the American Society of Echocardiography* 16.3 (Mar. 2003), pp. 240–245. ISSN: 08947317. DOI: [10.1067/mje.2003.52](https://doi.org/10.1067/mje.2003.52).
- [4] Sokratis J. Anagnostopoulos et al. “Residual-based attention and connection to information bottleneck theory in PINNs”. In: (July 2023). URL: <http://arxiv.org/abs/2307.00379>.
- [5] Sanjeev Arora et al. “Fine-Grained Analysis of Optimization and Generalization for Overparameterized Two-Layer Neural Networks”. In: (Jan. 2019). URL: <http://arxiv.org/abs/1901.08584>.
- [6] Per M. Arvidsson et al. “Kinetic energy of left ventricular blood flow across heart failure phenotypes and in subclinical diastolic dysfunction”. In: *Journal of Applied Physiology* 133.3 (Sept. 2022), pp. 697–709. ISSN: 15221601. DOI: [10.1152/JAPPLPHYSIOL.00257.2022](https://doi.org/10.1152/JAPPLPHYSIOL.00257.2022).
- [7] Per M. Arvidsson et al. *The shape of the healthy heart is optimized for vortex ring formation*. 2016. DOI: [10.1186/1532-429X-18-S1-023](https://doi.org/10.1186/1532-429X-18-S1-023).
- [8] Amirhossein Arzani, Jian-Xun Wang, and Roshan M. D’Souza. “Uncovering near-wall blood flow from sparse data with physics-informed neural networks”. In: (Apr. 2021). URL: <http://arxiv.org/abs/2104.08249>.

[9] Ali Bakhshinejad et al. “Merging computational fluid dynamics and 4D Flow MRI using proper orthogonal decomposition and ridge regression”. In: *Journal of Biomechanics* 58 (June 2017), pp. 162–173. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2017.05.004.

[10] Yalin Baştanlar and Mustafa Özysal. “Introduction to machine learning”. In: *Methods in Molecular Biology* 1107 (2014), pp. 105–128. ISSN: 10643745. DOI: 10.1007/978-1-62703-748-8\\_\\_7.

[11] R. J.M. Bastiaans, G. A.J. Van der Plas, and R. N. Kieft. “The performance of a new PTV algorithm applied in super-resolution PIV”. In: *Experiments in Fluids* 32.3 (Mar. 2002), pp. 346–356. ISSN: 07234864. DOI: 10.1007/s003480100363.

[12] Rene Bastkowski et al. “Self-gated golden-angle spiral 4D flow MRI”. In: *Magnetic Resonance in Medicine* 80.3 (Sept. 2018), pp. 904–913. ISSN: 15222594. DOI: 10.1002/mrm.27085.

[13] Atilim Güneş Baydin, Barak A Pearlmutter, and Jeffrey Mark Siskind. *Automatic Differentiation in Machine Learning: a Survey*. Tech. rep. 2018, pp. 1–43.

[14] Filipe de Avila Belbute-Peres, Yi-fan Chen, and Fei Sha. “HyperPINN: Learning parameterized differential equations with physics-informed hypernetworks”. In: (Oct. 2021). URL: <http://arxiv.org/abs/2111.01008>.

[15] Hadar Ben-Arzi et al. “Longitudinal Changes in Left Ventricular Blood Flow Kinetic Energy After Myocardial Infarction: Predictive Relevance for Cardiac Remodeling”. In: *Journal of Magnetic Resonance Imaging* 56.3 (Sept. 2022), pp. 768–778. ISSN: 15222586. DOI: 10.1002/jmri.28015.

[16] Emelia J. Benjamin et al. *Heart Disease and Stroke Statistics’2017 Update: A Report from the American Heart Association*. Mar. 2017. DOI: 10.1161/CIR.000000000000485.

[17] Albert S Berahas, Jorge Nocedal, and Martin Takáč. *A Multi-Batch L-BFGS Method for Machine Learning*. Tech. rep.

[18] M J Berger and Lawrence Livermore. *Local Adaptive Mesh Refinement for Shock Hydrodynamics P. COELLA*. Tech. rep. 1989, p. 6484.

[19] J Bergstra, D Yamins, and D D Cox. *Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures*. Tech. rep. 2013.

[20] Javier Bermejo et al. “Spatio-temporal mapping of intracardiac pressure gradients. A solution to Euler’s equation from digital postprocessing of color Doppler M-mode echocardiograms”. In: *World Federation for Ultrasound in Medicine & Biology* (2001).

[21] Christian H. Bischof et al. “Automatic differentiation of the general-purpose computational fluid dynamics package FLUENT”. In: *Journal of Fluids Engineering, Transactions of the ASME* 129.5 (May 2007), pp. 652–658. ISSN: 00982202. DOI: 10.1115/1.2720475.

[22] Rafael Bischof and Michael A Kraus. “Multi-Objective Loss Balancing for Physics-Informed Deep Learning”. In: *Pre-print* (2022). URL: [https://github.com/rbischof/relative\\_balancing..](https://github.com/rbischof/relative_balancing..)

[23] Malenka M. Bissell et al. *4D Flow cardiovascular magnetic resonance consensus statement: 2023 update*. Dec. 2023. DOI: 10.1186/s12968-023-00942-z.

[24] Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. “Machine learning for fluid mechanics”. In: *Annu. Rev. Fluid Mech* 52 (2020), pp. 477–508. DOI: 10.1146/annurev-fluid-010719. URL: <https://doi.org/10.1146/annurev-fluid-010719->.

[25] Burch G. “The Trabeculae Carneae—A Thought”. In: *Am Heart J.* (1975).

[26] Fanny Buyens et al. “Calculation of left ventricle relative pressure distribution in MRI using acceleration data”. In: *Magnetic Resonance in Medicine* 53.4 (2005), pp. 877–884. ISSN: 07403194. DOI: 10.1002/mrm.20415.

[27] Anil Can and Rose Du. “Association of hemodynamic factors with intracranial aneurysm formation and rupture: Systematic review and meta-analysis”. In: *Neurosurgery* 78.4 (Apr. 2016), pp. 510–519. ISSN: 15244040. DOI: 10.1227/NEU.0000000000001083.

[28] Federico Canè et al. “A CFD study on the interplay of torsion and vortex guidance by the mitral valve on the left ventricular wash-out making use of overset meshes (Chimera technique)”. In: *Frontiers in Medical Technology* 4 (Dec. 2022). DOI: 10.3389/fmedt.2022.1018058.

[29] Arturo Cárdenas-Blanco et al. “Noise in magnitude magnetic resonance images”. In: *Concepts in Magnetic Resonance Part A: Bridging Education and Research* 32.6 (Nov. 2008), pp. 409–416. ISSN: 15466086. DOI: 10.1002/cmr.a.20124.

[30] *Cardiovascular Diseases (CVDs) - World Health Organization (WHO)*. 2021. URL: [https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)).

[31] Bee Ting Chan et al. “Impact of myocardial infarction on intraventricular vortex and flow energetics assessed using computational simulations”. In: *International Journal for Numerical Methods in Biomedical Engineering* 35.6 (June 2019). ISSN: 20407947. DOI: 10.1002/cnm.3204.

[32] Bee Ting Chan et al. “Quantitative analysis of intraventricular flow-energetics and vortex in ischaemic hearts”. In: *Coronary Artery Disease* 29.4 (2018), pp. 316–324. ISSN: 14735830. DOI: 10.1097/MCA.0000000000000596.

[33] Jeffrey Shi Kai Chan et al. “Fragmented Vortex in Heart Failure With Reduced Ejection Fraction: A Prospective Vector Flow Mapping Study”. In: *Ultrasound in Medicine and Biology* 49.4 (Apr. 2023), pp. 982–988. ISSN: 1879291X. DOI: 10.1016/j.ultrasmedbio.2022.12.001.

[34] John J. Charonko et al. “Vortices formed on the mitral valve tips aid normal left ventricular filling”. In: *Annals of Biomedical Engineering* 41.5 (May 2013), pp. 1049–1061. ISSN: 00906964. DOI: 10.1007/s10439-013-0755-0.

[35] Haotian Chen et al. “A Critical Evaluation of Using Physics-Informed Neural Networks for Simulating Voltammetry: Strengths, Weaknesses and Best Practices”. In: *Journal of Electroanalytical Chemistry* 925 (Nov. 2022). ISSN: 15726657. DOI: 10.1016/j.jelechem.2022.116918.

[36] W. W. Chen et al. “Study of cardiovascular function using a coupled left ventricle and systemic circulation model”. In: *Journal of Biomechanics* 49.12 (Aug. 2016), pp. 2445–2454. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2016.03.009.

[37] Zhao Chen et al. “GradNorm: Gradient Normalization for Adaptive Loss Balancing in Deep Multitask Networks”. In: (Nov. 2017). URL: <http://arxiv.org/abs/1711.02257>.

[38] Christopher P. Cheng, David Parker, and Charles A. Taylor. “Quantification of wall shear stress in large blood vessels using lagrangian interpolation functions with cine phase-contrast magnetic resonance imaging”. In: *Annals of Biomedical Engineering* 30.8 (2002), pp. 1020–1032. ISSN: 00906964. DOI: 10.1114/1.1511239.

[39] C. Chnafa, S. Mendez, and F. Nicoud. “Image-based large-eddy simulation in a realistic left heart”. In: *Computers and Fluids* 94 (May 2014), pp. 173–187. ISSN: 00457930. DOI: 10.1016/j.compfluid.2014.01.030.

[40] C. Chnafa, S. Mendez, and F. Nicoud. “Image-Based Simulations Show Important Flow Fluctuations in a Normal Left Ventricle: What Could be the Implications?” In: *Annals of Biomedical Engineering* 44.11 (Nov. 2016), pp. 3346–3358. ISSN: 15739686. DOI: 10.1007/s10439-016-1614-6.

[41] Merih Cibis et al. “Relation between wall shear stress and carotid artery wall thickening MRI versus CFD”. In: *Journal of Biomechanics* 49.5 (Mar. 2016), pp. 735–741. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2016.02.004.

[42] Merih Cibis et al. “The effect of resolution on viscous dissipation measured with 4D flow MRI in patients with Fontan circulation: Evaluation using computational fluid dynamics”. In: *Journal of Biomechanics* 48.12 (Sept. 2015), pp. 2984–2989. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2015.07.039.

[43] Dario Collia et al. “Simplified mitral valve modeling for prospective clinical application of left ventricular fluid dynamics”. In: *Journal of Computational Physics* 398 (Dec. 2019). ISSN: 10902716. DOI: 10.1016/j.jcp.2019.108895.

[44] *Computer Vision and Pattern Recognition, 2009, CVPR 2009, IEEE Conference on : dates: 20-25 June 2009*. IEEE, 2009. ISBN: 9781424439911.

[45] Thomas E Conturo and Bruce H Robinsons. *Analysis of Encoding Efficiency in MR Imaging of Velocity Magnitude and Direction*. Tech. rep. 1992, pp. 233–247.

[46] Michael Courtois, Sandor J Kovacs Jr, and Philip A Ludbrook. *Transmitral Pressure-Flow Velocity Relation Importance of Regional Pressure Gradients in the Left Ventricle During Diastole*. Tech. rep. 1987. URL: <http://ahajournals.org>.

[47] Salvatore Cuomo et al. “Scientific Machine Learning through Physics-Informed Neural Networks: Where we are and What’s next”. In: (Jan. 2022). URL: <http://arxiv.org/abs/2201.05624>.

[48] Arka Das et al. “Acute intra-cavity 4D flow cardiovascular magnetic resonance predicts long-term adverse remodelling following ST-elevation myocardial infarction”. In: *Journal of Cardiovascular Magnetic Resonance* 24.1 (Dec. 2022). ISSN: 1532429X. DOI: 10.1186/s12968-022-00889-7.

[49] Arka Das et al. “Pathophysiology of LV Remodeling Following STEMI: A Longitudinal Diffusion Tensor CMR Study”. In: *JACC: Cardiovascular Imaging* 16.2 (Feb. 2023), pp. 159–171. ISSN: 18767591. DOI: 10.1016/j.jcmg.2022.04.002.

[50] Arka Daw et al. “Rethinking the Importance of Sampling in Physics-informed Neural Networks”. In: (). DOI: 10.48550/arXiv.2207.02338. URL: <https://www.researchgate.net/publication/361808020>.

[51] Kalyanmoy Deb. “Multi-Objective Optimization Using Evolutionary Algorithms: An Introduction”. In: *Wang, L., Ng, A., Deb, K. (eds) Multi-objective Evolutionary Optimisation for Product Design and Manufacturing. Springer, London* (2011). DOI: [https://doi.org/10.1007/978-0-85729-652-8\\\_\\\_1](https://doi.org/10.1007/978-0-85729-652-8\_\_1). URL: <http://www.iitk.ac.in/kangal/deb.htm>.

[52] Ahmet Demirkiran et al. “Association of left ventricular flow energetics with remodeling after myocardial infarction: New hemodynamic insights for left ventricular remodeling”. In: *International Journal of Cardiology* 367 (Nov. 2022), pp. 105–114. ISSN: 18741754. DOI: 10.1016/j.ijcard.2022.08.040.

[53] Ahmet Demirkiran et al. “Association of left ventricular flow energetics with remodeling after myocardial infarction: New hemodynamic insights for left ventricular remodeling”. In: *International Journal of Cardiology* 367 (Nov. 2022), pp. 105–114. ISSN: 18741754. DOI: 10.1016/j.ijcard.2022.08.040.

[54] Ahmet Demirkiran et al. “Left ventricular four-dimensional blood flow distribution, energetics, and vorticity in chronic myocardial infarction patients with/without left ventricular thrombus”. In: *European Journal of Radiology* 150 (May 2022). ISSN: 18727727. DOI: 10.1016/j.ejrad.2022.110233.

[55] Zhiwen Deng et al. “Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework”. In: *Physics of Fluids* 31.12 (Dec. 2019). ISSN: 10897666. DOI: 10.1063/1.5127031.

[56] Shaan Desai et al. “One-Shot Transfer Learning of Physics-Informed Neural Networks”. In: (Oct. 2021). URL: <http://arxiv.org/abs/2110.11286>.

[57] Anagha Deshmane et al. “Parallel MR imaging”. In: *Journal of Magnetic Resonance Imaging* 36.1 (July 2012), pp. 55–72. ISSN: 10531807. DOI: 10.1002/jmri.23639.

[58] Fabrizio Donati et al. “Beyond Bernoulli: Improving the Accuracy and Precision of Noninvasive Estimation of Peak Pressure Drops”. In: *Circulation: Cardiovascular Imaging* 10.1 (Jan. 2017). ISSN: 19420080. DOI: 10.1161/CIRCIMAGING.116.005207.

[59] Fabrizio Donati et al. “Non-invasive pressure difference estimation from PC-MRI using the work-energy equation”. In: *Medical Image Analysis* 26.1 (Dec. 2015), pp. 159–172. ISSN: 13618423. DOI: 10.1016/j.media.2015.08.012.

[60] Siamak N. Doost et al. “Heart blood flow simulation: A perspective review”. In: *BioMedical Engineering Online* 15.1 (Aug. 2016). ISSN: 1475925X. DOI: 10.1186/s12938-016-0224-8.

[61] Petter Dyverfeldt et al. “4D flow cardiovascular magnetic resonance consensus statement”. In: *Journal of Cardiovascular Magnetic Resonance* 17.1 (Aug. 2015). ISSN: 1532429X. DOI: 10.1186/s12968-015-0174-5.

[62] Tino Ebbers et al. “Estimation of relative cardiovascular pressures using time-resolved three-dimensional phase contrast MRI”. In: *Magnetic Resonance in Medicine* 45.5 (2001), pp. 872–879. ISSN: 07403194. DOI: 10.1002/mrm.1116.

[63] Mohammed S.M. Elbaz et al. “Assessment of viscous energy loss and the association with three-dimensional vortex ring formation in left ventricular inflow: In vivo evaluation using four-dimensional flow MRI”. In: *Magnetic Resonance in Medicine* 77.2 (Feb. 2017), pp. 794–805. ISSN: 15222594. DOI: 10.1002/mrm.26129.

[64] Mohammed S.M. Elbaz et al. “Vortex flow during early and late left ventricular filling in normal subjects: Quantitative characterization using retrospectively-gated 4D flow cardiovascular magnetic resonance and three-dimensional vortex core analysis”. In: *Journal of Cardiovascular Magnetic Resonance* 16.1 (Sept. 2014). ISSN: 1532429X. DOI: 10.1186/s12968-014-0078-9.

[65] Alan D. Elster. *How to Measure SNR - Questions and Answers in MRI*. 2024. URL: <https://mriquestions.com/signal-to-noise.html>.

[66] Alan D. Elster. “Magnetism - Questions and Answers in MRI”. In: (2024). URL: <https://mriquestions.com/choice-of-trteetc.html>.

[67] Alan D. Elster. “PROPELLER/BLADE - Questions and Answers in MRI”. In: (2024). URL: <https://mriquestions.com/propellerblade.html>.

[68] Alan D. Elster. “Radial Sampling - Questions and Answers in MRI”. In: () .

[69] Jonatan Eriksson et al. “Spatial heterogeneity of four-dimensional relative pressure fields in the human left ventricle”. In: *Magnetic Resonance in Medicine* 74.6 (Dec. 2015), pp. 1716–1725. ISSN: 15222594. DOI: 10.1002/mrm.25539.

[70] Fulvio Faccenda et al. “Doppler Measurement of the Pressure Drop Caused by Arterial Stenosis: An Experimental Study: A Case Report”. In: *Angiology* (1985). DOI: 10.1177/000331978503601211.

[71] Mojtaba F. Fathi et al. “Denoising and spatial resolution enhancement of 4D flow MRI using proper orthogonal decomposition and lasso regularization”. In: *Computerized Medical Imaging and Graphics* 70 (Dec. 2018), pp. 165–172. ISSN: 18790771. DOI: 10.1016/j.compmedimag.2018.07.003.

[72] Mojtaba F. Fathi et al. “Super-resolution and denoising of 4D-Flow MRI using physics-Informed deep neural nets”. In: *Computer Methods and Programs in Biomedicine* 197 (Dec. 2020). ISSN: 18727565. DOI: 10.1016/j.cmpb.2020.105729.

[73] Andriy Fedorov et al. “3D Slicer as an image computing platform for the Quantitative Imaging Network”. In: *Magnetic Resonance Imaging* 30.9 (Nov. 2012), pp. 1323–1341. ISSN: 0730725X. DOI: 10.1016/j.mri.2012.05.001.

[74] E. Ferdian et al. “Cerebrovascular super-resolution 4D Flow MRI – Sequential combination of resolution enhancement by deep learning and physics-informed image processing to non-invasively quantify intracranial velocity, flow, and relative pressure”. In: *Medical Image Analysis* 88 (Aug. 2023). ISSN: 13618423. DOI: 10.1016/j.media.2023.102831.

[75] Edward Ferdian et al. “4DFlowNet: Super-Resolution 4D Flow MRI using Deep Learning and Computational Fluid Dynamics”. In: *Frontiers in Physics* 8 (2020). DOI: 10.3389/fphy.2020.00138.

[76] Edward Ferdian et al. “WSSNet: Aortic Wall Shear Stress Estimation Using Deep Learning on 4D Flow MRI”. In: *Frontiers in Cardiovascular Medicine* 8 (2021). ISSN: 2297055X. DOI: 10.3389/fcvm.2021.769927.

[77] Michael S. Firstenberg et al. “Noninvasive estimation of transmitral pressure drop across the normal mitral valve in humans: Importance of convective and inertial forces during left ventricular filling”. In: *Journal of the American College of Cardiology* 36.6 (Nov. 2000), pp. 1942–1949. ISSN: 07351097. DOI: 10.1016/S0735-1097(00)00963-3.

[78] A E Fischer, J Balcom, and E J Fordham. “Spin warp NMR and applications to human whole-body imaging”. In: *Phys. Med. Biol.* 25 (1980), p. 751.

[79] Kai Fukami, Koji Fukagata, and Kunihiko Taira. “Assessment of supervised machine learning methods for fluid flows”. In: *Theoretical and Computational Fluid Dynamics* 34.4 (Aug. 2020), pp. 497–519. ISSN: 14322250. DOI: 10.1007/s00162-020-00518-y.

[80] Kai Fukami, Koji Fukagata, and Kunihiko Taira. “Super-resolution reconstruction of turbulent flows with machine learning”. In: (Nov. 2018). DOI: 10.1017/jfm.2019.238. URL: <http://arxiv.org/abs/1811.11328> <http://dx.doi.org/10.1017/jfm.2019.238>.

[81] Kunihiko Fukushima. “Cognitron: A Self-organizing Multilayered Neural Network”. In: *Biol. Cybernetics* 20 (1975), pp. 121–136.

[82] Maurizio Galderisi. “Diastolic dysfunction and diastolic heart failure: Diagnostic, prognostic and therapeutic aspects”. In: *Cardiovascular Ultrasound* 3 (Apr. 2005). ISSN: 14767120. DOI: 10.1186/1476-7120-3-9.

[83] Han Gao, Luning Sun, and Jian-Xun Wang. “Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels”. In: (Nov. 2020). URL: <http://arxiv.org/abs/2011.02364>.

[84] Pankaj Garg et al. “Left ventricular blood flow kinetic energy after myocardial infarction - Insights from 4D flow cardiovascular magnetic resonance”. In: *Journal of Cardiovascular Magnetic Resonance* 20.1 (Aug. 2018). ISSN: 1532429X. DOI: 10.1186/s12968-018-0483-6.

[85] Morgane Garreau et al. “Accelerated sequences of 4D flow MRI using GRAPPA and compressed sensing: A comparison against conventional MRI and computational fluid dynamics”. In: *Magnetic Resonance in Medicine* 88.6 (Dec. 2022), pp. 2432–2446. ISSN: 15222594. DOI: 10.1002/mrm.29404.

[86] MB Giles, D Ghate, and MC Duta. “Using Automatic Differentiation for Adjoint CFD Code Development”. In: *Computational Fluid Dynamics Journal* 16 (2008).

[87] Xavier Glorot and Yoshua Bengio. *Understanding the difficulty of training deep feedforward neural networks*. Tech. rep. URL: <http://www.iro.umontreal..>

[88] Ian J. Goodfellow et al. “Generative Adversarial Networks”. In: (June 2014). URL: <http://arxiv.org/abs/1406.2661>.

[89] Miroslawa Gorecka et al. *Rationale and clinical applications of 4D flow cardiovascular magnetic resonance in assessment of valvular heart disease: a comprehensive review*. Dec. 2022. DOI: 10.1186/s12968-022-00882-0.

[90] Leonid Goubergrits et al. “CT-Based Analysis of Left Ventricular Hemodynamics Using Statistical Shape Modeling and Computational Fluid Dynamics”. In: *Frontiers in Cardiovascular Medicine* 9 (July 2022). ISSN: 2297055X. DOI: 10.3389/fcvm.2022.901902.

[91] Mark A. Griswold et al. “Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA)”. In: *Magnetic Resonance in Medicine* 47.6 (2002), pp. 1202–1210. ISSN: 07403194. DOI: 10.1002/mrm.10171.

[92] Hákon Gudbjartsson and Samuel Patz. *The Rician Distribution of Noisy MRI Data*. Tech. rep.

[93] Aakash N. Gupta, Michael Markl, and Mohammed S.M. Elbaz. *Intracardiac and Vascular Hemodynamics with Cardiovascular Magnetic Resonance in Heart Failure*. Jan. 2021. DOI: 10.1016/j.hfc.2020.08.010.

[94] David Ha, Andrew Dai, and Quoc V. Le. “HyperNetworks”. In: (Sept. 2016). URL: <http://arxiv.org/abs/1609.09106>.

[95] L Hatle et al. “Noninvasive assessment of pressure drop in mitral stenosis by Doppler ultrasound”. In: *British Heart Journal* 40 (1978), pp. 131–140.

[96] Fan He, Lu Hua, and Li jian Gao. “Computational analysis of blood flow and wall mechanics in a model of early atherosclerotic artery”. In: *Journal of Mechanical Science and Technology* 31.2 (Feb. 2017), pp. 1015–1020. ISSN: 1738494X. DOI: 10.1007/s12206-017-0154-9.

[97] *Health matters: preventing cardiovascular disease - Public Health England (PHE)*. 2019. URL: <https://www.gov.uk/government/publications/health-matters-preventing-cardiovascular-disease/health-matters-preventing-cardiovascular-disease>.

[98] Paul A. Heidenreich et al. “Forecasting the future of cardiovascular disease in the United States: A policy statement from the American Heart Association”. In: *Circulation* 123.8 (Mar. 2011), pp. 933–944. ISSN: 00097322. DOI: 10.1161/CIR.0b013e31820a55f5.

[99] A. Ali Heydari, Craig A. Thompson, and Asif Mehmood. “SoftAdapt: Techniques for Adaptive Loss Weighting of Neural Networks with Multi-Part Loss Functions”. In: (Dec. 2019). URL: <http://arxiv.org/abs/1912.12355>.

[100] Geu Ru Hong et al. “Characterization and Quantification of Vortex Flow in the Human Left Ventricle by Contrast Echocardiography Using Vector Particle Image Velocimetry”. In: *JACC: Cardiovascular Imaging* 1.6 (Nov. 2008), pp. 705–717. ISSN: 1936878X. DOI: 10.1016/j.jcmg.2008.06.008.

[101] Albert Hsiao et al. “Rapid pediatric cardiac assessment of flow and ventricular volume with compressed sensing parallel imaging volumetric cine phase-contrast MRI”. In: *American Journal of Roentgenology* 198.3 (Mar. 2012). ISSN: 0361803X. DOI: 10.2214/AJR.11.6969.

[102] Feng Huang et al. “K-t GRAPPA: A k-space implementation for dynamic MRI with high reduction factor”. In: *Magnetic Resonance in Medicine* 54.5 (2005), pp. 1172–1184. ISSN: 07403194. DOI: 10.1002/mrm.20641.

[103] Peter. J. Huber. “Robust Estimation of a Location Parameter”. In: *The Annals of Mathematical Statistics* (1964).

[104] Hussein M El-Husseiny et al. “Intraventricular pressure gradient: A novel tool to assess the post-infarction chronic congestive heart failure”. In: *Frontiers in Cardiovascular Medicine* 9 (2022). DOI: [10.3389/fcvm.2022.944171](https://doi.org/10.3389/fcvm.2022.944171).

[105] IEEE Engineering in Medicine and Biology Society, IEEE Signal Processing Society, and IEEE International Conference on Service-Oriented Computing and Applications. *IEEE 10th International Symposium on Biomedical Imaging (ISBI), 2013*. IEEE. ISBN: 9781467364553.

[106] Pablo Irarrazaval et al. “Noise estimation for the velocity in MRI phase-contrast”. In: *Magnetic Resonance Imaging* 63 (Nov. 2019), pp. 250–257. ISSN: 18735894. DOI: [10.1016/j.mri.2019.08.028](https://doi.org/10.1016/j.mri.2019.08.028).

[107] Arthur Jacotécole et al. “Neural Tangent Kernel: Convergence and Generalization in Neural Networks”. In: *Advances in neural information processing systems* (2018).

[108] Ameya D. Jagtap and George Em Karniadakis. “Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations”. In: *Communications in Computational Physics* 28.5 (Nov. 2020), pp. 2002–2041. ISSN: 19917120. DOI: [10.4208/CICP.OA-2020-0164](https://doi.org/10.4208/CICP.OA-2020-0164).

[109] Ameya D. Jagtap, Kenji Kawaguchi, and George Em Karniadakis. “Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 476.2239 (July 2020). ISSN: 14712946. DOI: [10.1098/rspa.2020.0334](https://doi.org/10.1098/rspa.2020.0334).

[110] Ameya D. Jagtap, Kenji Kawaguchi, and George Em Karniadakis. “Adaptive activation functions accelerate convergence in deep and physics-informed neural networks”. In: *Journal of Computational Physics* 404 (Mar. 2020). ISSN: 10902716. DOI: [10.1016/j.jcp.2019.109136](https://doi.org/10.1016/j.jcp.2019.109136).

[111] Ameya D. Jagtap et al. “Deep Kronecker neural networks: A general framework for neural networks with adaptive activation functions”. In: *Neurocomputing* 468 (Jan. 2022), pp. 165–180. ISSN: 18728286. DOI: [10.1016/j.neucom.2021.10.036](https://doi.org/10.1016/j.neucom.2021.10.036).

[112] Oren N. Jaspan, Roman Fleysher, and Michael L. Lipton. *Compressed sensing MRI: A review of the clinical literature*. 2015. DOI: [10.1259/bjr.20150487](https://doi.org/10.1259/bjr.20150487).

[113] Jinhee JEong and Fazle Hussain. “On the identification of a vortex”. In: *Journal of Fluid Mechanics* 285 (1995), pp. 69–94. ISSN: 14697645. DOI: 10.1017/S0022112095000462.

[114] Max E Jerrell. *Automatic Differentiation and Interval Arithmetic for Estimation of Disequilibrium Models*. Tech. rep. 1997, pp. 295–316.

[115] Xiaowei Jin et al. “NSFnets (Navier-Stokes Flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations”. In: (Mar. 2020). URL: <http://arxiv.org/abs/2003.06496>.

[116] Xiaowei Jin et al. “NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations”. In: *Journal of Computational Physics* 426 (Feb. 2021). ISSN: 10902716. DOI: 10.1016/j.jcp.2020.109951.

[117] Zhi-Qin John Xu et al. “Frequency principle: Fourier analysis sheds light on deep neural networks”. In: *Communications in Computational Physics* 28 (2020). DOI: <https://doi.org/10.4208/cicp.OA-2020-0085>. URL: <https://ins.sjtu.edu.cn/people/xuzhiqin/>.

[118] Kevin M. Johnson and Michael Markl. “Improved SNR in phase contrast velocimetry with five-point balanced flow encoding”. In: *Magnetic Resonance in Medicine* 63.2 (2010), pp. 349–355. ISSN: 15222594. DOI: 10.1002/mrm.22202.

[119] Mikael Kanski et al. “Left ventricular fluid kinetic energy time curves in heart failure from cardiovascular magnetic resonance 4D flow data”. In: *Journal of Cardiovascular Magnetic Resonance* 17.1 (Dec. 2015). ISSN: 1532429X. DOI: 10.1186/s12968-015-0211-4.

[120] George Em Karniadakis et al. “Physics-informed machine learning”. In: *Nature Reviews Physics* 3.6 (June 2021), pp. 422–440. ISSN: 25225820. DOI: 10.1038/s42254-021-00314-5.

[121] S P Karwatowski et al. *A comparison of left ventricular myocardial velocity in diastole measured by magnetic resonance and left ventricular filling measured by Doppler echocardiography*. Tech. rep. 1996, pp. 795–802.

[122] Ali Kashefi and Tapan Mukerji. “Physics-informed PointNet: A deep learning solver for steady-state incompressible flows and thermal fields on multiple sets of irregular geometries”. In: *Journal of Computational Physics* 468 (Nov. 2022). ISSN: 10902716. DOI: 10.1016/j.jcp.2022.111510.

[123] R D Keane and R J Adrian. “Super-resolution particle imaging velocimetry”. In: *Meas. Sci. Technol* 6 (1995), p. 754.

[124] Alex Kendall, Yarin Gal, and Roberto Cipolla. “Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics”. In: (May 2017). URL: <http://arxiv.org/abs/1705.07115>.

[125] Arash Kheradvar et al. *Diagnostic and prognostic significance of cardiovascular vortex formation*. Nov. 2019. DOI: 10.1016/j.jjcc.2019.05.005.

[126] Hyojin Kim et al. “Unsupervised deep learning for super-resolution reconstruction of turbulence”. In: (July 2020). URL: <http://arxiv.org/abs/2007.15324>.

[127] Diederik P. Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In: (Dec. 2014). URL: <http://arxiv.org/abs/1412.6980>.

[128] Georgios Kissas et al. “Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks”. In: *Computer Methods in Applied Mechanics and Engineering* 358 (Jan. 2020). ISSN: 00457825. DOI: 10.1016/j.cma.2019.112623.

[129] Taha Sabri Koltukluoglu. “Fourier Spectral Dynamic Data Assimilation: Interlacing CFD with 4D Flow MRI”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 11765 LNCS. Springer, 2019, pp. 741–749. ISBN: 9783030322441. DOI: 10.1007/978-3-030-32245-8{\\_}82.

[130] Christopher M. Kramer et al. “Standardized cardiovascular magnetic resonance imaging (CMR) protocols: 2020 update”. In: *Journal of Cardiovascular Magnetic Resonance* 22.1 (Feb. 2020). ISSN: 1532429X. DOI: 10.1186/s12968-020-00607-1.

[131] Corina Kräuter et al. “Automated mitral valve vortex ring extraction from 4D-flow MRI”. In: *Magnetic Resonance in Medicine* 84.6 (Dec. 2020), pp. 3396–3408. ISSN: 15222594. DOI: 10.1002/mrm.28361.

[132] Sebastian B.S. Krittian et al. “A finite-element approach to the direct computation of relative cardiovascular pressure from time-resolved MR velocity data”. In: *Medical Image Analysis* 16.5 (July 2012), pp. 1029–1037. ISSN: 13618415. DOI: 10.1016/j.media.2012.04.003.

[133] Scott Kulp et al. “Using high resolution cardiac CT data to model and visualize patient-specific interactions between trabeculae and blood flow”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 6891 LNCS. PART 1. 2011, pp. 468–475. ISBN: 9783642236228. DOI: 10.1007/978-3-642-23623-5{\\_}59.

[134] Chang Quan Lai et al. “Fluid mechanics of blood flow in human fetal left ventricles based on patient-specific 4D ultrasound scans”. In: *Biomechanics and Modeling in Mechanobiology* 15.5 (Oct. 2016), pp. 1159–1172. ISSN: 16177940. DOI: 10.1007/s10237-015-0750-5.

[135] Jonas Lantz et al. “Patient-Specific Simulation of Cardiac Blood Flow from High-Resolution Computed Tomography”. In: *Journal of Biomechanical Engineering* 138.12 (Dec. 2016). ISSN: 15288951. DOI: 10.1115/1.4034652.

[136] Y.A. LeCun et al. *Neural Networks: Tricks of the Trade*. Ed. by Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller. Vol. 7700. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012. ISBN: 978-3-642-35288-1. DOI: 10.1007/978-3-642-35289-8. URL: <http://link.springer.com/10.1007/978-3-642-35289-8>.

[137] Christian Ledig et al. “Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network”. In: (Sept. 2016). URL: <http://arxiv.org/abs/1609.04802>.

[138] Adrian T. Lee, G. Bruce Pike, and Norbert J. Pelc. “Three-Point Phase-Contrast Velocity Measurements with Increased Velocity-to-Noise Ratio”. In: *Magnetic Resonance in Medicine* 33.1 (1995), pp. 122–126. ISSN: 15222594. DOI: 10.1002/mrm.1910330119.

[139] Sébastien Levilly, Said Moussaoui, and Jean Michel Serfaty. “Navier-Stokes-Based Regularization for 4d Flow MRI Super-Resolution”. In: *Proceedings - International Symposium on Biomedical Imaging*. Vol. 2022-March. IEEE Computer Society, 2022. ISBN: 9781665429238. DOI: 10.1109/ISBI52829.2022.9761510.

[140] Y. Li, B. Sixou, and F. Peyrin. *A Review of the Deep Learning Methods for Medical Images Super Resolution Problems*. Apr. 2021. DOI: 10.1016/j.irbm.2020.08.004.

[141] Lukas Liebel and Marco Körner. “Auxiliary Tasks in Multi-task Learning”. In: (May 2018). URL: <http://arxiv.org/abs/1805.06334>.

[142] Bo Liu et al. “Deep learning methods for super-resolution reconstruction of turbulent flows”. In: *Physics of Fluids* 32.2 (Feb. 2020). ISSN: 10897666. DOI: 10.1063/1.5140772.

[143] Dan Liu et al. “Longitudinal strain bull’s eye plot patterns in patients with cardiomyopathy and concentric left ventricular hypertrophy”. In: *European Journal of*

*Medical Research* 21.1 (May 2016). ISSN: 2047783X. DOI: 10.1186/s40001-016-0216-y.

- [144] Dong C Liu and Jorge Nocedal. “On the limited memory BFGS method for large scale optimization”. In: *Mathematical Programming* 45 (1989), pp. 503–528.
- [145] Shikun Liu, Edward Johns, and Andrew J. Davison. “End-to-End Multi-Task Learning with Attention”. In: (Mar. 2018). URL: <http://arxiv.org/abs/1803.10704>.
- [146] Ziqi Liu, Wei Cai, and Zhi-Qin John Xu. “Multi-scale Deep Neural Network (MscaleDNN) for Solving Poisson-Boltzmann Equation in Complex Domains”. In: (July 2020). URL: <http://arxiv.org/abs/2007.11207>.
- [147] Lu Lu et al. “DeepXDE: A deep learning library for solving differential equations”. In: *SIAM Review* 63.1 (2021), pp. 208–228. ISSN: 00361445. DOI: 10.1137/19M1274067.
- [148] Lu Lu et al. “Physics-informed neural networks with hard constraints for inverse design”. In: (Feb. 2021). URL: <http://arxiv.org/abs/2102.04626>.
- [149] Liliana E. Ma et al. “Aortic 4D flow MRI in 2 minutes using compressed sensing, respiratory controlled adaptive k-space reordering, and inline reconstruction”. In: *Magnetic Resonance in Medicine* 81.6 (June 2019), pp. 3675–3690. ISSN: 15222594. DOI: 10.1002/mrm.27684.
- [150] Muazzam Maqsood et al. “Transfer learning assisted classification and detection of alzheimer’s disease stages using 3D MRI scans”. In: *Sensors (Switzerland)* 19.11 (June 2019). ISSN: 14248220. DOI: 10.3390/s19112645.
- [151] Stefano Markidis. “The Old and the New: Can Physics-Informed Deep-Learning Replace Traditional Linear Solvers?” In: *Frontiers in Big Data* 4 (Nov. 2021). ISSN: 2624909X. DOI: 10.3389/fdata.2021.669097.
- [152] Michael Markl, Philip J. Kilner, and Tino Ebbers. *Comprehensive 4D velocity mapping of the heart and great vessels by cardiovascular magnetic resonance*. 2011. DOI: 10.1186/1532-429X-13-7.
- [153] Michael Markl et al. “4D flow MRI”. In: *Journal of Magnetic Resonance Imaging* 36.5 (Nov. 2012), pp. 1015–1036. ISSN: 10531807. DOI: 10.1002/jmri.23632.
- [154] R. Timothy Marler and Jasbir S. Arora. “The weighted sum method for multi-objective optimization: New insights”. In: *Structural and Multidisciplinary Optimization* 41.6 (June 2010), pp. 853–862. ISSN: 1615147X. DOI: 10.1007/s00158-009-0460-7.

[155] David Marlevi et al. “Non-invasive estimation of relative pressure for intracardiac flows using virtual work-energy”. In: *Medical Image Analysis* 68 (Feb. 2021). ISSN: 13618423. DOI: 10.1016/j.media.2020.101948.

[156] Pablo Martínez-Legazpi et al. “Contribution of the Diastolic Vortex Ring to Left Ventricular Filling”. In: *J Am Coll Cardiol.* (2014). DOI: 10.1016/j.jacc.2014.06.1205.

[157] Kevin Mc Namara, Hamzah Alzubaidi, and John Keith Jackson. “Cardiovascular disease as a leading cause of death: how are pharmacists getting involved?” In: *Integrated Pharmacy Research and Practice* Volume 8 (Feb. 2019), pp. 1–11. DOI: 10.2147/iprp.s133088.

[158] Levi McClenny and Ulisses Braga-Neto. “Self-Adaptive Physics-Informed Neural Networks using a Soft Attention Mechanism”. In: (Sept. 2020). URL: <http://arxiv.org/abs/2009.04544>.

[159] Siddhartha Mishra and Roberto Molinaro. “Estimates on the generalization error of physics-informed neural networks for approximating PDEs”. In: *IMA Journal of Numerical Analysis* 43.1 (Feb. 2023), pp. 1–43. ISSN: 0272-4979. DOI: 10.1093/imanum/drab093.

[160] Rajat Mittal et al. “Computational modeling of cardiac hemodynamics: Current status and future outlook”. In: *Journal of Computational Physics* 305 (Jan. 2016), pp. 1065–1082. ISSN: 10902716. DOI: 10.1016/j.jcp.2015.11.022.

[161] Suguru Miyauchi et al. “Numerical analysis of the blood flow in the left ventricle with internal structures: Effect of trabeculae carneae models and atrial fibrillation”. In: *AIP Advances* 9.10 (Oct. 2019). ISSN: 21583226. DOI: 10.1063/1.5124730.

[162] K L Moore. *Clinically Oriented Anatomy*. Williams & Wilkins, 1992. ISBN: 9780683094855. URL: <https://books.google.co.uk/books?id=XN9qAAAAMAAJ>.

[163] Paul R Moran et al. *Verification and Evaluation of Internal Flow and Motion True Magnetic Resonance Imaging by the Phase Gradient Modulation Method1*. Tech. rep. 1985.

[164] John C. Morud et al. “The effect of chordae tendineae on systolic flow”. In: *Computers in Biology and Medicine* 109 (June 2019), pp. 91–100. ISSN: 18790534. DOI: 10.1016/j.combiomed.2019.04.020.

[165] Mohammad Amin Nabian, Rini Jasmine Gladstone, and Hadi Meidani. “Efficient training of physics-informed neural networks via importance sampling”. In: *Computer-*

*Aided Civil and Infrastructure Engineering* 36.8 (Aug. 2021), pp. 962–977. ISSN: 14678667. DOI: 10.1111/mice.12685.

[166] David Nolte and Cristóbal Bertoglio. “Reducing the impact of geometric errors in flow computations using velocity measurements”. In: *International Journal for Numerical Methods in Biomedical Engineering* 35.6 (June 2019). ISSN: 20407947. DOI: 10.1002/cnm.3203.

[167] Lukas Obermeier et al. “CT-Based Simulation of Left Ventricular Hemodynamics: A Pilot Study in Mitral Regurgitation and Left Ventricle Aneurysm Patients”. In: *Frontiers in Cardiovascular Medicine* 9 (Mar. 2022). DOI: 10.3389/fcvm.2022.828556.

[168] John N. Oshinski et al. “Improved measurement of pressure gradients in aortic coarctation by magnetic resonance imaging”. In: *Journal of the American College of Cardiology* 28.7 (Dec. 1997), pp. 1818–1826. ISSN: 07351097. DOI: 10.1016/S0735-1097(96)00395-6.

[169] Lauren Partin, Daniele E. Schiavazzi, and Carlos A. Sing Long. “An analysis of reconstruction noise from undersampled 4D flow MRI”. In: *Biomedical Signal Processing and Control* 84 (July 2023). ISSN: 17468108. DOI: 10.1016/j.bspc.2023.104800.

[170] A Pasipoularides. *Heart’s Vortex: Intracardiac Blood Flow Phenomena*. Pmph USA Ltd Series. People’s Medical Publishing House, 2009. ISBN: 9781607950332. URL: <https://books.google.co.uk/books?id=eMKqdIvxEmQC>.

[171] Ashitha Pathrose et al. “Highly accelerated aortic 4D flow MRI using compressed sensing: Performance at different acceleration factors in patients with aortic disease”. In: *Magnetic Resonance in Medicine* 85.4 (Apr. 2021), pp. 2174–2187. ISSN: 15222594. DOI: 10.1002/mrm.28561.

[172] Gianni Pedrizzetti et al. “Cardiac fluid dynamics anticipates heart adaptation”. In: *Journal of Biomechanics* 48.2 (Jan. 2015), pp. 388–391. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2014.11.049.

[173] Gianni Pedrizzetti et al. “On estimating intraventricular hemodynamic forces from endocardial dynamics: A comparative study with 4D flow MRI”. In: *Journal of Biomechanics* 60 (July 2017), pp. 203–210. ISSN: 18732380. DOI: 10.1016/j.jbiomech.2017.06.046.

[174] Gianni Pedrizzetti et al. *The vortex - An early predictor of cardiovascular outcome?* 2014. DOI: 10.1038/nrcardio.2014.75.

[175] Norbert J Pelc et al. “Encoding Strategies for Three-Direction Phase-Contrast MR Imaging of Flow”. In: *J Magn Reson Imaging*. (1991). DOI: 10.1002/jmri.1880010404.

[176] Raphaël Pellegrin et al. “Transfer Learning with Physics-Informed Neural Networks for Efficient Simulation of Branched Flows”. In: (Oct. 2022). URL: <http://arxiv.org/abs/2211.00214>.

[177] Wei Peng et al. “RANG: A Residual-based Adaptive Node Generation Method for Physics-Informed Neural Networks”. In: (May 2022). URL: <http://arxiv.org/abs/2205.01051>.

[178] Eva S. Peper et al. “Highly accelerated 4D flow cardiovascular magnetic resonance using a pseudo-spiral Cartesian acquisition and compressed sensing reconstruction for carotid flow and wall shear stress”. In: *Journal of Cardiovascular Magnetic Resonance* 22.1 (Jan. 2020). ISSN: 1532429X. DOI: 10.1186/s12968-019-0582-z.

[179] X Per et al. “Hemodynamic forces using four-dimensional flow MRI: an independent biomarker of cardiac function in heart failure with left ventricular dyssynchrony?” In: *Am J Physiol Heart Circ Physiol* 315 (2018), pp. 1627–1639. DOI: 10.1152/ajpheart.00112.2018. –Patients. URL: [www.ajpheart.org](http://www.ajpheart.org).

[180] X Per et al. “Left and right ventricular hemodynamic forces in healthy volunteers and elite athletes assessed with 4D flow magnetic resonance imaging”. In: *Am J Physiol Heart Circ Physiol* 312 (2017), pp. 314–328. DOI: 10.1152/ajpheart.00583.2016. –Intracardiac. URL: <http://ajpheart.physiology.org/>.

[181] Isaac Perez-Raya et al. “Towards multi-modal data fusion for super-resolution and denoising of 4D-Flow MRI”. In: *International Journal for Numerical Methods in Biomedical Engineering* 36.9 (Sept. 2020). ISSN: 20407947. DOI: 10.1002/cnm.3381.

[182] Sven Petersson, Petter Dyverfeldt, and Tino Ebbers. “Assessment of the accuracy of MRI wall shear stress estimation using numerical simulations”. In: *Journal of Magnetic Resonance Imaging* 36.1 (July 2012), pp. 128–138. ISSN: 10531807. DOI: 10.1002/jmri.23610.

[183] James G Pipe. “Motion Correction With PROPELLER MRI: Application to Head Motion and Free-Breathing Cardiac Imaging”. In: *Magnetic Resonance in Medicine* 42 (1999), p. 963. DOI: 10.1002/(sici)1522-2594(199911)42:5<963::aid-mrm17>3.0.co;2-1.

[184] Konstantinos Prantikos et al. “Physics-informed neural network with transfer learning (TL-PINN) based on domain similarity measure for prediction of nuclear reactor transients”. In: *Scientific Reports* 13.1 (Dec. 2023). ISSN: 20452322. DOI: 10.1038/s41598-023-43325-1.

[185] Klaas P. Pruessmann et al. “SENSE: Sensitivity encoding for fast MRI”. In: *Magnetic Resonance in Medicine* 42.5 (1999), pp. 952–962. ISSN: 07403194. DOI: 10.1002/(SICI)1522-2594(199911)42:5;952::AID-MRM163.0.CO;2-S.

[186] Charles R. Qi et al. “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation”. In: (Dec. 2016). URL: <http://arxiv.org/abs/1612.00593>.

[187] Nasim Rahaman et al. *On the Spectral Bias of Neural Networks*. Tech. rep. URL: <https://github.com/nasimrahaman/SpectralBias>.

[188] Ali Rahimi and Ben Recht. “Random Features for Large-Scale Kernel Machines”. In: *Advances in Neural Information Processing Systems 20 (NIPS)* (2007).

[189] 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”. In: *Journal of Computational Physics* 378 (Feb. 2019), pp. 686–707. ISSN: 10902716. DOI: 10.1016/j.jcp.2018.10.045.

[190] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. “Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations”. In: (Nov. 2017). URL: <http://arxiv.org/abs/1711.10561>.

[191] Maziar Raissi, Alireza Yazdani, and George Em Karniadakis. “Hidden Fluid Mechanics: A Navier-Stokes Informed Deep Learning Framework for Assimilating Flow Visualization Data”. In: (Aug. 2018). URL: <http://arxiv.org/abs/1808.04327>.

[192] Prajit Ramachandran, Barret Zoph, and Quoc V Le Google Brain. “Searching for activation functions”. In: (2017).

[193] Nikoo R. Saber et al. “Computational flow modeling of the left ventricle based on in vivo MRI data: Initial experience”. In: *Annals of Biomedical Engineering* 29.4 (2001), pp. 275–283. ISSN: 00906964. DOI: 10.1114/1.1359452.

[194] Federica Sacco et al. “Left ventricular trabeculations decrease the wall shear stress and increase the intra-ventricular pressure drop in CFD simulations”. In: *Frontiers in Physiology* 9 (Apr. 2018). ISSN: 1664042X. DOI: 10.3389/fphys.2018.00458.

[195] Simone Saitta et al. “Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI”. In: (Feb. 2023). URL: <http://arxiv.org/abs/2302.12835>.

[196] Ali Sarrami-Foroushani et al. “In-silico trial of intracranial flow diverters replicates and expands insights from conventional clinical trials”. In: *Nature Communications* 12.1 (Dec. 2021). ISSN: 20411723. DOI: 10.1038/s41467-021-23998-w.

[197] F. Scarano. “A super-resolution particle image velocimetry interrogation approach by means of velocity second derivatives correlation”. In: *Measurement Science and Technology* 15.2 (2004), pp. 475–486. ISSN: 09570233. DOI: 10.1088/0957-0233/15/2/023.

[198] Michal Schäfer et al. “4D-flow cardiac magnetic resonance-derived vorticity is sensitive marker of left ventricular diastolic dysfunction in patients with mild-to-moderate chronic obstructive pulmonary disease”. In: *European heart journal cardiovascular Imaging* 19.4 (Apr. 2018), pp. 415–424. ISSN: 20472412. DOI: 10.1093/ehjci/jex069.

[199] Torsten Schenkel et al. “MRI-Based CFD analysis of flow in a human left ventricle: Methodology and application to a healthy heart”. In: *Annals of Biomedical Engineering* 37.3 (Mar. 2009), pp. 503–515. ISSN: 00906964. DOI: 10.1007/s10439-008-9627-4.

[200] Vincent F.M. Segers, Dirk L. Brutsaert, and Gilles W. De Keulenaer. *Cardiac remodeling: Endothelial cells have more to say than just NO*. Apr. 2018. DOI: 10.3389/fphys.2018.00382.

[201] Partho P. Sengupta. “Reclassifying Heart Failure: Predominantly Subendocardial, Subepicardial, and Transmural”. In: *Heart Failure Clinics* 4.3 (July 2008), pp. 379–382. ISSN: 1551-7136. DOI: 10.1016/J.HFC.2008.03.013. URL: <https://www.sciencedirect.com/science/article/abs/pii/S1551713608000585?via%3Dihub>.

[202] Jung Hee Seo and Rajat Mittal. “Effect of diastolic flow patterns on the function of the left ventricle”. In: *Physics of Fluids* 25.11 (Aug. 2013). ISSN: 10897666. DOI: 10.1063/1.4819067.

[203] Suprosanna Shit et al. “SRflow: Deep learning based super-resolution of 4D-flow MRI data”. In: *Frontiers in Artificial Intelligence* 5 (2022). DOI: 10.3389/frai.2022.928181.

[204] Suprosanna Shit et al. “Velocity-To-Pressure (V2P) - Net: Inferring Relative Pressures from Time-Varying 3D Fluid Flow Velocities”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 12729 LNCS. Springer Science and Business Media

Deutschland GmbH, 2021, pp. 545–558. ISBN: 9783030781903. DOI: 10.1007/978-3-030-78191-0\}42.

[205] Fergus Shone et al. “Deep Physics-Informed Super-Resolution of Cardiac 4D-Flow MRI”. In: *Information Processing in Medical Imaging*. Ed. by Alejandro Frangi et al. Cham: Springer Nature Switzerland, 2023, pp. 511–522. ISBN: 978-3-031-34048-2.

[206] Vincent Sitzmann et al. “Implicit Neural Representations with Periodic Activation Functions”. In: (June 2020). URL: <http://arxiv.org/abs/2006.09661>.

[207] Aparna Sodhi et al. “Highly accelerated compressed sensing 4D flow MRI in congenital and acquired heart disease: comparison of aorta and main pulmonary artery flow parameters with conventional 4D flow MRI in children and young adults”. In: *Pediatric Radiology* (Dec. 2023). ISSN: 14321998. DOI: 10.1007/s00247-023-05788-2.

[208] Aurélien F. Stalder et al. “Quantitative 2D and 3D phase contrast MRI: Optimized analysis of blood flow and vessel wall parameters”. In: *Magnetic Resonance in Medicine* 60.5 (2008), pp. 1218–1231. ISSN: 15222594. DOI: 10.1002/mrm.21778.

[209] R. Brad Stamm and Randolph P. Martin. “Quantification of pressure gradients across stenotic valves by Doppler ultrasound”. In: *Journal of the American College of Cardiology* 2.4 (1983), pp. 707–718. ISSN: 07351097. DOI: 10.1016/S0735-1097(83)80311-8.

[210] Zoran Stankovic et al. “4D flow imaging with MRI”. In: *Cardiovascular Diagnosis and Therapy* (2014).

[211] Adel Stitou and M L Riethmuller. “Extension of PIV to super resolution using PTV”. In: *Meas. Sci. Technol* 12 (2001), pp. 1398–1403.

[212] Shashank Subramanian et al. “Adaptive Self-supervision Algorithms for Physics-informed Neural Networks”. In: (July 2022). URL: <http://arxiv.org/abs/2207.04084>.

[213] Jeremy Szajer and Kevin Ho-Shon. *A comparison of 4D flow MRI-derived wall shear stress with computational fluid dynamics methods for intracranial aneurysms and carotid bifurcations — A review*. May 2018. DOI: 10.1016/j.mri.2017.12.005.

[214] Jeremy Szajer and Kevin Ho-Shon. *A comparison of 4D flow MRI-derived wall shear stress with computational fluid dynamics methods for intracranial aneurysms and carotid bifurcations — A review*. May 2018. DOI: 10.1016/j.mri.2017.12.005.

[215] K Takehara et al. “A Kalman tracker for super-resolution PIV”. In: *Experiments in Fluids* 29 (2000). DOI: <https://doi.org/10.1007/s003480070005>.

[216] Matthew Tancik et al. “Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains”. In: (June 2020). URL: <http://arxiv.org/abs/2006.10739>.

[217] Jeffrey Tsao, Peter Boesiger, and Klaas P. Pruessmann. “k-t BLAST and k-t SENSE: Dynamic MRI With High Frame Rate Exploiting Spatiotemporal Correlations”. In: *Magnetic Resonance in Medicine* 50.5 (2003), pp. 1031–1042. ISSN: 07403194. DOI: 10.1002/mrm.10611.

[218] Sergio Uribe et al. “Four-dimensional (4D) flow of the whole heart and great vessels using real-time respiratory self-gating”. In: *Magnetic Resonance in Medicine* 62.4 (2009), pp. 984–992. ISSN: 15222594. DOI: 10.1002/mrm.22090.

[219] Muthiah Vaduganathan et al. *The Global Burden of Cardiovascular Diseases and Risk: A Compass for Future Health*. Dec. 2022. DOI: 10.1016/j.jacc.2022.11.005.

[220] Pim Van Ooij et al. “Aortic Valve Stenosis Alters Expression of Regional Aortic Wall Shear Stress: New Insights From a 4-Dimensional Flow Magnetic Resonance Imaging Study of 571 Subjects”. In: (). DOI: 10.1161/JAHA.117. URL: <http://ahajournals.org>.

[221] Ashish Vaswani et al. “Attention Is All You Need”. In: *Advances in Neural Information Processing Systems 30 (NIPS 2017)* 30 (2017).

[222] Vijay Vedula et al. “Effect of trabeculae and papillary muscles on the hemodynamics of the left ventricle”. In: *Theoretical and Computational Fluid Dynamics* 30.1-2 (Apr. 2016), pp. 3–21. ISSN: 14322250. DOI: 10.1007/s00162-015-0349-6.

[223] Arun Verma. “An introduction to automatic differentiation”. In: *Current Science* 78.7 (Apr. 2000), pp. 804–807. ISSN: 00113891. DOI: 10.1002/pamm.200310012.

[224] Jacqueline L. Vos et al. “CMR-derived left ventricular intraventricular pressure gradients identify different patterns associated with prognosis in dilated cardiomyopathy”. In: *European Heart Journal Cardiovascular Imaging* 24.9 (Sept. 2023), pp. 1231–1240. ISSN: 20472412. DOI: 10.1093/ehjci/jead083.

[225] Sifan Wang, Yujun Teng, and Paris Perdikaris. “Understanding and mitigating gradient flow pathologies in physics-informed neural networks”. In: *SIAM Journal on Scientific Computing* 43.5 (2021), pp. 3055–3081. ISSN: 10957197. DOI: 10.1137/20M1318043.

[226] Sifan Wang, Yujun Teng, and Paris Perdikaris. “Understanding and mitigating gradient pathologies in physics-informed neural networks”. In: (Jan. 2020). URL: <http://arxiv.org/abs/2001.04536>.

[227] Sifan Wang, Xinling Yu, and Paris Perdikaris. “When and why PINNs fail to train: A neural tangent kernel perspective”. In: (July 2020). URL: <http://arxiv.org/abs/2007.14527>.

[228] Sifan Wang, Xinling Yu, and Paris Perdikaris. “When and why PINNs fail to train: A neural tangent kernel perspective”. In: *Journal of Computational Physics* 449 (Jan. 2022). ISSN: 10902716. DOI: [10.1016/j.jcp.2021.110768](https://doi.org/10.1016/j.jcp.2021.110768).

[229] Zhihao Wang, Jian Chen, and Steven C. H. Hoi. “Deep Learning for Image Super-resolution: A Survey”. In: (Feb. 2019). URL: <http://arxiv.org/abs/1902.06068>.

[230] A. C.G. Wenink and A. C. Gittenberger-De Groot. “Left and right ventricular trabecular patterns. Consequence of ventricular septation and valve development”. In: *British Heart Journal* 48.5 (1982), pp. 462–468. ISSN: 00070769. DOI: [10.1136/hrt.48.5.462](https://doi.org/10.1136/hrt.48.5.462).

[231] Chenxi Wu et al. “A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks”. In: *Computer Methods in Applied Mechanics and Engineering* 403 (Jan. 2023). ISSN: 00457825. DOI: [10.1016/j.cma.2022.115671](https://doi.org/10.1016/j.cma.2022.115671).

[232] Xiaolin Wu et al. “4D Flow Patterns and Relative Pressure Distribution in a Left Ventricle Model by Shake-the-Box and Proper Orthogonal Decomposition Analysis”. In: *Cardiovascular Engineering and Technology* (2023). ISSN: 18694098. DOI: [10.1007/s13239-023-00684-0](https://doi.org/10.1007/s13239-023-00684-0).

[233] David T. Wymer et al. “Phase-contrast MRI: Physics, techniques, and clinical applications”. In: *Radiographics* 40.1 (Jan. 2020), pp. 122–140. ISSN: 15271323. DOI: [10.1148/rg.2020190039](https://doi.org/10.1148/rg.2020190039).

[234] You Xie et al. “TempoGAN: A temporally coherent, volumetric GAN for super-resolution fluid flow”. In: *ACM Transactions on Graphics* 37.4 (2018). ISSN: 15577368. DOI: [10.1145/3197517.3201304](https://doi.org/10.1145/3197517.3201304).

[235] Chen Xu et al. “Transfer learning based physics-informed neural networks for solving inverse problems in engineering structures under different loading scenarios”. In: *Computer Methods in Applied Mechanics and Engineering* 405 (Feb. 2023). ISSN: 00457825. DOI: [10.1016/j.cma.2022.115852](https://doi.org/10.1016/j.cma.2022.115852).

[236] Y Yamamoto and T Uemura. *A New Super-Resolution PIV Accelerated by Characteristic Pixel Selection*. Tech. rep. 3. 2003, pp. 263–272.

[237] Raquel Yotti et al. “Noninvasive assessment of ejection intraventricular pressure gradients”. In: *Journal of the American College of Cardiology* 43.9 (May 2004), pp. 1654–1662. ISSN: 07351097. DOI: 10.1016/j.jacc.2003.09.066.

[238] Tianhe Yu et al. “Gradient Surgery for Multi-Task Learning”. In: (Jan. 2020). URL: <http://arxiv.org/abs/2001.06782>.

[239] Gizem Yüce et al. “A Structured Dictionary Perspective on Implicit Neural Representations”. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2022* (2022).

[240] Bastian Zapf et al. “Investigating molecular transport in the human brain from MRI with physics-informed neural networks”. In: *Scientific Reports* 12.1 (Dec. 2022). ISSN: 20452322. DOI: 10.1038/s41598-022-19157-w.

[241] Guangtao Zhang et al. *DASA-PINNs: Differentiable adversarial self-adaptive pointwise weighting scheme for physics-informed neural networks*. Tech. rep. URL: <https://ssrn.com/abstract=4376049>.

[242] Jiacheng Zhang et al. “4D Flow MRI Pressure Estimation Using Velocity Measurement-Error-Based Weighted Least-Squares”. In: *IEEE Transactions on Medical Imaging* 39.5 (May 2020), pp. 1668–1680. ISSN: 1558254X. DOI: 10.1109/TMI.2019.2954697.

[243] An-Sheng Zheng and Hong-Xia Yu. “Value of clinical applications of differential pressure and relative pressure imaging in the left ventricle”. In: *World Journal of Clinical Cases* 11.17 (June 2023), pp. 3967–3975. ISSN: 2307-8960. DOI: 10.12998/wjcc.v11.i17.3967.

[244] Xiahai Zhuang. “Multivariate Mixture Model for Myocardial Segmentation Combining Multi-Source Images”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41.12 (Dec. 2019), pp. 2933–2946. ISSN: 19393539. DOI: 10.1109/TPAMI.2018.2869576.