Translating computational modelling tools for clinical practice in congenital heart disease

Endrit Pajaziti

Student Number: 13035736

Primary Supervisor: Prof. Silvia Schievano

Secondary Supervisor: Dr. Claudio Capelli

Dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Institute of Cardiovascular Science

University College London

July 5, 2023

I, Endrit Pajaziti, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Acknowledgements

This thesis would not have been possible without the guidance and support I have received from a large number of people. First and foremost, I am forever grateful to my PhD supervisors Prof. Silvia Schievano and Dr. Claudio Capelli, and my undergraduate supervisor Dr. Emilie Sauvage. For the past 6 years, they have played a pivotal part in my personal development and growth. I am deeply thankful for all the mentorship, knowledge, feedback and opportunities they have given me. However, most of all, I am grateful for their kindness, and how I was made to feel like a valued and important member of their team at every step of the way.

I would also like to acknowledge and thank Prof. Andrew Cook and Prof. Vivek Muthurangu for their constant support and availability in helping me formulate and conduct my research. Special thanks to Dr. Raphaël Sivera for his crucially important help with statistical shape modelling, Dr. Javier Montalt-Tordera for his machine learning wisdom, Liam Swanson for his advice on CFD modelling (and for paying rent as my flatmate!) and Nikos Zavaropoulos for helping with the online functionality of VheaRts. From the clinical side, I would like to thank Dr. Elena Milano, Mr. Georgios Belitsis, Dr. Beatrice Bonello, Mr. Owase Jeelani and more, there are too many to list.

Many thanks to my friends and colleagues for the memories forged during this journey: Liam, Claudia, Raphaël, Ebba, Steven, Benni, Chloe, Garazi, María, Claire and many others! Lastly, I must thank my family, for my achievements mean the most to them of all, and without their unwavering support, this piece of work could never have been possible.

Abstract

Increasingly large numbers of medical centres worldwide are equipped with the means to acquire 3D images of patients by utilising magnetic resonance (MR) or computed tomography (CT) scanners. The interpretation of patient 3D image data has significant implications on clinical decision-making and treatment planning. In their raw form, MR and CT images have become critical in routine practice. However, in congenital heart disease (CHD), lesions are often anatomically and physiologically complex. In many cases, 3D imaging alone can fail to provide conclusive information for the clinical team. In the past 20-30 years, several image-derived modelling applications have shown major advancements. Tools such as computational fluid dynamics (CFD) and virtual reality (VR) have successfully demonstrated valuable uses in the management of CHD. However, due to current software limitations, these applications have remained largely isolated to research settings, and have yet to become part of clinical practice. The overall aim of this project was to explore new routes for making conventional computational modelling software more accessible for CHD clinics. The first objective was to create an automatic and fast pipeline for performing vascular CFD simulations. By leveraging machine learning, a solution was built using synthetically generated aortic anatomies, and was seen to be able to predict 3D aortic pressure and velocity flow fields with comparable accuracy to conventional CFD. The second objective was to design a virtual reality (VR) application tailored for supporting the surgical planning and teaching of CHD. The solution was a Unity-based application which included numerous specialised tools, such as mesh-editing features and online networking for group learning. Overall, the outcomes of this ongoing project showed strong indications that the integration of VR and CFD into clinical settings is possible, and has potential for extending 3D imaging and supporting the diagnosis, management and teaching of CHD.

Impact statement

Medical imaging is critical for the diagnosis, treatment planning and monitoring of patients with congenital heart disease (CHD). However, more advanced patient-specific image-based computer modelling tools have not yet made a full transition into the clinical field. The work in this thesis proposes suitable methods of integrating two main tools into CHD clinical/educational settings: computational fluid dynamics (CFD) and virtual reality (VR).

CFD is capable of non-invasively modelling detailed hemodynamic flows. For CHD, this has numerous powerful applications, such as for better understanding of patient patho-physiological conditions, evaluating theoretical post-operative outcomes and simulating exercise conditions. However, software is typically incompatible with clinical practice due to time/computational requirements and manual labour. This limits the feasibility of CFD for routine clinical practice. I applied machine learning (ML) techniques to build a data-driven pipeline for fast simulation of 3D pressure and velocity flows in post-coarctation aortas. This automated pipeline will open up new possibilities for 'precision medicine' ML tools constructed using physically-constrained simulation data.

VR is a growing phenomenon and has seen a rapid increase in applications within healthcare. This has notable ramifications for the field of CHD, where patients present with complex and often unique anatomy. As part of my research, I developed an application for visualising cardiac structures in VR. This was then integrated into the clinical and teaching practices at Great Ormond Street Hospital for Children and UCL, resulting in over 50 surgical cases planned pre-operatively with the aid of VR and more than 240 students taught CHD anatomy/morphology via VR. This experience has shown that VR is received well, and could significantly improve CHD clinical practice/education if implemented on a global scale.

1. For a research manuscript that has already been published (if not yet published, please skip to section 2)**:**		
a) Where was the work published? (e.g. journal name)	Journal of Cardiovascular Magnetic Resonance	
b) Who published the work? (e.g. Elsevier/Oxford University Press):	BioMed Central	
c) When was the work published?	November 2022	
d) Was the work subject to academic peer review?	Yes	
e) Have you retained the copyright for the work?	Yes	
[If no, please seek permission from the relevant publisher and check the box next to the below statement]: *I acknowledge permission of the publisher named under 1b to include in this thesis portions of the publication named as included in 1a.*		
2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3)**:**		
a) Has the manuscript been uploaded to a preprint server? (e.g. medRxiv):	N/A	**If yes, which server?** N/A
b) Where is the work intended to be published? (e.g. names of journals that you are planning to submit to)	N/A	
c) List the manuscript's authors in the intended authorship order:	N/A	
d) Stage of publication	N/A	
3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4**):**		
Javier Montalt-Tordera (1): original manuscript, formal analysis, software **Endrit Pajaziti (1):** original manuscript, formal analysis, software **Rod Jones:** data curation **Emilie Sauvage:** review, methods **Rajesh Puranik:** data curation **Aakansha Ajay Vir Singh:** data curation **Claudio Capelli:** review, methods **Jennifer Steeden:** review, methods **Silvia Schievano:** review, methods **Vivek Muthurangu:** review, conceptualisation, supervision		
4. In which chapter(s) of your thesis can this material be found?		
Chapter 4		
5. e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):		

Candidate:	Endrit Pajaziti	**Date:**	01/02/2023
Supervisor/ Senior Author	Silvia Schievano	**Date:**	01/02/2023

<table>
<tr><td colspan="2">1. For a research manuscript that has already been published (if not yet published, please skip to section 2):</td></tr>
<tr><td>a) Where was the work published? (e.g. journal name)</td><td>PLOS Computational Biology</td></tr>
<tr><td>b) Who published the work? (e.g. Elsevier/Oxford University Press):</td><td>PLOS</td></tr>
<tr><td>c) When was the work published?</td><td>24/04/2023</td></tr>
<tr><td>d) Was the work subject to academic peer review?</td><td>Yes</td></tr>
<tr><td>e) Have you retained the copyright for the work?</td><td>Creative Commons 4.0 License</td></tr>
<tr><td colspan="2">[If no, please seek permission from the relevant publisher and check the box next to the below statement]:
I acknowledge permission of the publisher named under 1b to include in this thesis portions of the publication named as included in 1a.</td></tr>
</table>

<table>
<tr><td colspan="3">2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):</td></tr>
<tr><td>a) Has the manuscript been uploaded to a preprint server? (e.g. medRxiv):</td><td>N/A</td><td>If yes, which server?
N/A</td></tr>
<tr><td>b) Where is the work intended to be published? (e.g. names of journals that you are planning to submit to)</td><td colspan="2">N//A</td></tr>
<tr><td>c) List the manuscript's authors in the intended authorship order:</td><td colspan="2">N/A</td></tr>
<tr><td>d) Stage of publication</td><td colspan="2">N/A</td></tr>
<tr><td colspan="3">3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):</td></tr>
<tr><td colspan="3">Endrit Pajaziti: Writing, formal analysis, software Javier Montalt-Tordera: methodology, review Claudio Capelli: review Raphael Sivera: methodology, review, software Emilie Sauvage: methodology, review Michael Quail: Data collection, review, Silvia Schievano: conceptualisation, supervision, review Vivek Muthurangu: conceptualisation, supervision, review</td></tr>
<tr><td colspan="3">4. In which chapter(s) of your thesis can this material be found?</td></tr>
<tr><td colspan="3">Chapter 5</td></tr>
</table>

<table>
<tr><td colspan="4">5. e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):</td></tr>
<tr><td>Candidate:</td><td>Endrit Pajaziti</td><td>Date:</td><td>01/05/2023</td></tr>
<tr><td>Supervisor/ Senior Author</td><td>Silvia Schievano</td><td>Date:</td><td>01/05/2023</td></tr>
</table>

<table>
<tr><td colspan="4">1. For a research manuscript that has already been published (if not yet published, please skip to section 2):</td></tr>
<tr><td colspan="2">a) Where was the work published? (e.g. journal name)</td><td colspan="2">Circulation: Cardiovascular Imaging</td></tr>
<tr><td colspan="2">b) Who published the work? (e.g. Elsevier/Oxford University Press):</td><td colspan="2">Lippincott Williams & Wilkins</td></tr>
<tr><td colspan="2">c) When was the work published?</td><td colspan="2">July 2019</td></tr>
<tr><td colspan="2">d) Was the work subject to academic peer review?</td><td colspan="2">Yes</td></tr>
<tr><td colspan="2">e) Have you retained the copyright for the work?</td><td colspan="2">No</td></tr>
<tr><td colspan="4">[If no, please seek permission from the relevant publisher and check the box next to the below statement]:
✓ I acknowledge permission of the publisher named under 1b to include in this thesis portions of the publication named as included in 1a.</td></tr>
<tr><td colspan="4">2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):</td></tr>
<tr><td colspan="2">a) Has the manuscript been uploaded to a preprint server? (e.g. medRxiv):</td><td>N/A</td><td>If yes, which server?
N/A</td></tr>
<tr><td colspan="2">b) Where is the work intended to be published? (e.g. names of journals that you are planning to submit to)</td><td colspan="2">N/A</td></tr>
<tr><td colspan="2">c) List the manuscript's authors in the intended authorship order:</td><td colspan="2">N/A</td></tr>
<tr><td colspan="2">d) Stage of publication</td><td colspan="2">N/A</td></tr>
<tr><td colspan="4">3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):</td></tr>
<tr><td colspan="4">Elena Giulia Milano: Writing, revision, formal analysis Endrit Pajaziti: writing, revision, software Emilie Sauvage: revision Andrew Cook: revision Silvia Schievano: conceptualisation, revision Kristian H. Mortensen: revision Andrew M Taylor: revision Jan Marek: revision Martin Kostolny: revision, investigation Claudio Capelli: conceptualisation, revision</td></tr>
<tr><td colspan="4">4. In which chapter(s) of your thesis can this material be found?</td></tr>
<tr><td colspan="4">Chapter 7</td></tr>
<tr><td colspan="4">5. e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):</td></tr>
<tr><td>Candidate:</td><td>Endrit Pajaziti</td><td>Date:</td><td>01/02/2023</td></tr>
<tr><td>Supervisor/ Senior Author</td><td>Silvia Schievano</td><td>Date:</td><td>01/02/2023</td></tr>
</table>

1. For a research manuscript that has already been published (if not yet published, please skip to section 2)**:**	
a) Where was the work published? (e.g. journal name)	European Heart Journal: Digital Health
b) Who published the work? (e.g. Elsevier/Oxford University Press):	Oxford University Press
c) When was the work published?	December 2021
d) Was the work subject to academic peer review?	Yes
e) Have you retained the copyright for the work?	No

[If no, please seek permission from the relevant publisher and check the box next to the below statement]:

✓ *I acknowledge permission of the publisher named under 1b to include in this thesis portions of the publication named as included in 1a.*

2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3)**:**		
a) Has the manuscript been uploaded to a preprint server? (e.g. medRxiv):	N/A	**If yes, which server?** N/A
b) Where is the work intended to be published? (e.g. names of journals that you are planning to submit to)	N/A	
c) List the manuscript's authors in the intended authorship order:	N/A	
d) Stage of publication	N/A	

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4**):**

Elena Giulia Milano: writing, formal analysis, investigation **Martin Kostolny:** methdology, review **Endrit Pajaziti:** methodology, software, writing **Jan Marek:** review **William Regan:** review **Massimo Caputo:** investigation, review **Giovanni Battista Luciani:** investigation, review **Kristian H Mortensen:** review **Andrew C Cook:** review **Silvia Schievano:** conceptualisation, review **Claudio Capelli:** conceptualisation, review

4. In which chapter(s) of your thesis can this material be found?

Chapter 7

5. e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

Candidate:	Endrit Pajaziti	**Date:**	01/02/2023
Supervisor/ Senior Author	Silvia Schievano	**Date:**	01/02/2023

<table>
<tr><td colspan="4">1. For a research manuscript that has already been published (if not yet published, please skip to section 2):</td></tr>
<tr><td colspan="2">a) Where was the work published? (e.g. journal name)</td><td colspan="2">Electronics</td></tr>
<tr><td colspan="2">b) Who published the work? (e.g. Elsevier/Oxford University Press):</td><td colspan="2">Multidisciplinary Digital Publishing Institute (MDPI)</td></tr>
<tr><td colspan="2">c) When was the work published?</td><td colspan="2">August 2021</td></tr>
<tr><td colspan="2">d) Was the work subject to academic peer review?</td><td colspan="2">Yes</td></tr>
<tr><td colspan="2">e) Have you retained the copyright for the work?</td><td colspan="2">Yes</td></tr>
<tr><td colspan="4">[If no, please seek permission from the relevant publisher and check the box next to the below statement]:
I acknowledge permission of the publisher named under 1b to include in this thesis portions of the publication named as included in 1a.</td></tr>
<tr><td colspan="4">2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):</td></tr>
<tr><td colspan="2">a) Has the manuscript been uploaded to a preprint server? (e.g. medRxiv):</td><td>N/A</td><td>If yes, which server?
N/A</td></tr>
<tr><td colspan="2">b) Where is the work intended to be published? (e.g. names of journals that you are planning to submit to)</td><td colspan="2">N/A</td></tr>
<tr><td colspan="2">c) List the manuscript's authors in the intended authorship order:</td><td colspan="2">N/A</td></tr>
<tr><td colspan="2">d) Stage of publication</td><td colspan="2">N/A</td></tr>
<tr><td colspan="4">3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):</td></tr>
<tr><td colspan="4">Endrit Pajaziti: writing, formal analysis, software Silvia Schievano: conceptualisation, review Emilie Sauvage: conceptualisation, review Andrew Cook: investigation, review Claudio Capelli: conceptualisation, supervision, review</td></tr>
<tr><td colspan="4">4. In which chapter(s) of your thesis can this material be found?</td></tr>
<tr><td colspan="4">Chapter 8</td></tr>
<tr><td colspan="4">5. e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):</td></tr>
<tr><td>Candidate:</td><td>Endrit Pajaziti</td><td>Date:</td><td>01/02/2023</td></tr>
<tr><td>Supervisor/ Senior Author</td><td>Silvia Schievano</td><td>Date:</td><td>01/02/2023</td></tr>
</table>

Contents

List of Figures

List of Tables

Abbreviations

ANN	artifical neural network
Ao	aorta
AOCA	anomalous origin of the coronary arteries
ASD	atrial septal defect
ASO	arterial switch operation
CCTGA	congenitally corrected transposition of the great arteries
CFD	computational fluid dynamics
CHD	congenital heart disease
CMR	cardiac magnetic resonance
CNN	convolutional neural network
CoA	coarctation of the aorta
CT	computed tomography
DNN	deep neural network
DORV	double outlet right ventricle
GT	ground truth
HD	Hausdorff distance
HLHS	hypoplastic left heart syndrome
HMD	head-mounted display
IoU	Intersection over Union
LV	left ventricle
LVOT	left ventricular outflow tract

MAPCAs	major aortopulmonary collateral arteries
MAPE	mean absolute percentage error
ML	machine learning
MSE	mean squared error
PA	pulmonary artery
PCA	principal component analysis
PDA	patent ductus arteriosus
PTA	persistent truncus arteriosus
ReLU	rectified linear unit
RV	right ventricle
RVOT	right ventricular outflow tract
SSM	statistical shape model
SVD	singular value decomposition
TCPC	total cavopulmonary connections
TGA	transposition of the great arteries
TOE	transoesophageal two-dimensional echocardiography
TOF	tetralogy of Fallot
TTE	transthoracic two-dimensional echocardiography
VMTK	vascular modelling toolkit
VR	virtual reality
VSD	ventricular septal defect
XR	extended reality

Chapter 1

Introduction

1.1 Motivation for this research

Congenital heart disease (CHD) affects almost 1% of newborns every year. Lesions can vary widely both in terms of severity and morphological features; however, patients often require intervention such as open-heart surgery, which carries significant risks. The advent of cardiac medical imaging modalities has revolutionised the way patients are diagnosed, treated and monitored. The application of image-based, patient-specific modelling tools for supporting clinical practice have also shown to have proven benefits in clinical practice. Despite this, many modelling approaches have failed so far to become integrated into routine practice due to reasons that include: (i) high computational resource requirements, (i) long processing times, (iii) the need for a specialist to operate the software, (iv) high barrier of accessibility for clinicians, and (v) tools which are not specialised for CHD.

1.2 Objectives of the work presented

The overall aim of this work is to take patient-specific modelling tools which have shown potential benefits for routine use in healthcare, and to develop and test more readily-deployable applications for possible clinical integration, with the ultimate goal of improving outcomes in CHD. The focus regarding modelling tools in this thesis is primarily on computational fluid dynamics (CFD) and virtual reality (VR). Computer methods including machine learning (ML), automatic segmentation, statistical shape modelling, data dimensionality reduction and 3D game engines are used to facilitate the translation of these tools. In this thesis, the main objectives are:

1. **To develop a pipeline that enables accelerated and clinically suitable solutions of CFD, capable of inferring 3D flow in the great arteries**

   The end-goal is to create models capable of rapidly and automatically computing flow fields in the great arteries using only routinely acquired 3D medical image data. This relies on two main components: (i) an automatic segmentation model which can produce surface meshes suitable for CFD, and (ii) an ML model which can infer CFD flow fields from any given shape.

2. **To develop and deploy a VR platform for clinical and educational applications of CHD**

   The second objective is to demonstrate the feasibility of developing and integrating a VR application into clinical and educational settings related to CHD. This requires: (i) the design and development of a novel VR platform for visualisation of CHD, (ii) application of the tool in supporting the pre-operative planning of real patient clinical cases, and (iii) the demonstration of VR in enhancing courses, workshops and training sessions related to CHD.

1.3 Thesis outline and structure

Chapter 2 provides a general overview of CHD, including major clinical challenges and diagnostic medical imaging. The challenges are grouped into two main categories around clinical decision-making, and education and training, with an emphasis on unmet needs which computer modelling may address. This Chapter includes information related to CHD relevant to the context of this thesis.

Chapter 3 summarises some important theoretical concepts behind the main methodologies used in this thesis (i.e. CFD, ML, statistical shape modelling and VR). A general overview of each method is provided, including relevant applications reported in the literature for CHD.

Chapter 4 investigates the use of ML for automatically segmenting CFD-suitable surface meshes of the great arteries. The methods used to create and evaluate the automatic segmentation model are described in detail. CFD simulations on the ground-truth (manual segmentation) and prediction (ML segmentation) are compared, to assess the suitability of the models for accelerating clinical CFD.

Chapter 5 describes an ML approach for rapidly predicting 3D pressure and velocity flow fields in aortas with repaired coarctation. Novel aspects to the study include generation of a synthetic training cohort and low-dimensional representations of large unstructured 3D flow fields. Results from the ML models are compared against conventional CFD, and differences between the two approaches are discussed.

Chapter 6 describes the development of a novel VR platform for supporting teaching and surgical planning in CHD. Functionalities of the platform are detailed, including the methods used to achieve them.

Chapter 7 is an overview of the clinical applications in which VR was used. A review of CHD cases for whom treatment was pre-operatively planned with the help of VR is provided. In particular, I present three cases of double outlet right ventricle, detailed as case studies. Additionally, a retrospective evaluation of the potential added advantage of VR compared to other conventional image modalities is reported.

Chapter 8 documents how VR was used for teaching cardiac anatomy. An overview of the library of 3D heart models developed for education is given. Applications in undergraduate/postgraduate teaching are demonstrated, including a multi-user VR workshop for CHD, and in clinical professional courses, including a cardiac morphology course and an echocardiography course. A study assessing the potential benefits of VR for improving MSc student knowledge is presented.

Chapter 2

Background

2.1 Congenital heart disease

2.1.1 Introduction

Congenital anomalies are described as structural or functional abnormalities that develop prenatally and are present upon birth. Of these defects, Congenital heart disease (CHD) is the most common type with an incidence of $\sim$8 per every 1000 live births [9], therefore constituting $\sim$28% of all major congenital anomalies [10]. Congenital heart disease (CHD) is defined as a gross structural abnormality of the heart or intrathoracic vessels which is actually, or potentially, of functional significance [11]. Genetic factors are commonly thought to be associated with the development of CHD during embryonic gestation. However, the mechanics through which genetic deficiencies are translated into structural lesions is not fully understood and remains a widely studied field [12].

CHD anomalies are highly variable and can even present as a combination of multiple lesions. In approximately 25% of infants afflicted with CHD, invasive surgical intervention is required in the first year of life [13]. During the early era of cardiovascular surgery (1950s - 1970s), the mortality for patients undergoing intervention was high (25% - 50%) [14]. However, in recent years, the prognosis for patients diagnosed with CHD has greatly improved. This is due to advances in diagnostic methods, equipment, surgical techniques and cardiothoracic anaesthesia. With mortality rates decreasing, much of modern surgery is now focused on improving patient surgical outcomes and minimising functional morbidity. However, despite low average surgical mortality and morbidity rates, there are still certain lesions and high-risk CHD types which are challenging to treat. Current mortality and morbidity statistics vary (especially between regions). Berger et al. analysed 1,550 patient ICU records from 7 different centres (between 2011 and 2013), and reported a 3.2% mean morality rate and a 4.8% mean incidence of functional morbidity post-surgery [15]. A meta-analysis by Best et al. indicated that the chance of 1-year survival for the lowest-risk CHD lesion was 95.5%, while for the highest-risk lesion it was 17.4% [16].

2.1.2 Main classifications

Most CHD arises due to defective cardiovascular development during the 3rd to 8th gestational week [13]. Abnormalities vary widely, ranging from small pinholes between heart chambers to significant aberrations in the configuration or morphology of the great arteries. Classifications and subgroups of CHD are not unanimously agreed upon. A widely-accepted segmentation of CHD divides types into two main groups, cyanotic (right to left shunts) and acyanotic (left to right shunts) [17], with further subgroups under cyanotic CHD. In this thesis, the classification of CHD types from Leiner et al. [18] is adopted. This favours dividing CHD based on specific clinical indications that form four main subgroups:

- **Shunt lesions:** abnormal holes or passages, commonly in the septum, that enable blood to pass between pulmonary and systemic circulations
- **Valve lesions:** abnormal valvular anatomy with functional implications on the heart and hemodynamics
- **Conotruncal lesions:** abnormalities of the cardiac outflow tracts and great arteries
- **Arterial lesions:** vascular abnormalities of the aorta (Ao), pulmonary artery or coronary arteries

Some examples for the main CHD types (excluding complex disease) are shown in Table 2.1. In this thesis, the focus is mostly around two CHD categories, conotruncal lesions and arterial lesions, further described in the following sections.

Shunt lesions	**Valve lesions**	**Conotruncal lesions**	**Arterial lesions**
Atrial septal defects	Mitral valve lesions	Double outlet right ventricle	Coarctation of the aorta
Ventricular septal defects	Aortic valve disease	Transposition of the great arteries	Double aortic arch
Sinus venosus defects	Tricuspid valve disease (including Ebstein)	Tetralogy of Fallot	Coronary fistula
Patent ductus arteriosus	Pulmonary valve disease	Persistent truncus arteriosus	Pulmonary artery stenosis

Table 2.1: Examples of CHD types for the four primary categories.

2.1.3 Conotruncal lesions

Conotruncal lesions concern primarily the outflow tracts of the heart, and to be defined as such, there must be the presence of anatomic abnormalities of the great arteries and the conus (infundibulum) [19]. This section will briefly describe the anatomy of some relevant types of conontruncal lesions, including the general landscape of current diagnostic and treatment options.

2.1.3.1 Overview

Anomalies related to the cardiac outflow tracts are some of the most frequently diagnosed CHD types, with prevalence observed to be $\sim$30% [20]. During embryonic life, the cardiac outflow tract comprises a single arterial structure known as the bulbus cordis, connected to the embryonic ventricle. The proximal segment of this structure later forms the infundibulum (conus), while the distal sections septate and spiral to form the aortic and pulmonary arterial trunks. The formation of fully septated, properly arranged great arteries and spiralling outflow tracts requires several complex morphogenetic events to take place during gestation. It is during this process that defects may occur, which can result in abnormalities such as an incorrect configuration of the great arteries or incomplete septation.

Many conotruncal lesions present with atrial or ventricular septal defects that enable blood to flow between chambers. A defect in the septum which enables blood to shunt between atria is known as an atrial septal defect (ASD), and similarly for ventricles is referred to as a ventricular septal defect (VSD). Narrowing of the great vessels and valves (stenosis) is also commonly observed in conotruncal lesions. Classic examples of conotruncal CHD include:

- Double outlet right ventricle (DORV) is characterised by a ventriculo-arterial connection, whereby both great arteries arise entirely or predominantly from the right ventricle (RV). Position/size of the great arteries and VSDs can be highly variable (Fig. 2.1). Prevalence is $\sim$0.9% of all CHDs [13].

- Transposition of the great arteries (TGA) involves an abnormal arrangement of the great arteries. The most common type (d-loop TGA, 4% of all CHD cases) results in an aorta arising from the right side of the heart and a pulmonary artery (PA) arising from the left side [13]. Multiple variations of TGA exist, such as congenitally corrected transposition of the great arteries (CCTGA) where the ventricles are also swapped.

- Tetralogy of Fallot (TOF) is characterised by four features: (i) a VSD, (ii) aorta positioned over VSD (overriding aorta) allowing blood flow from both ventricles, (iii) stenosis of the pulmonary outflow tract and (iv) right ventricular hypertrophy. TOF is one of the most common conotruncal lesions, comprising $\sim$5% of all CHD cases [13].

- Persistent truncus arteriosus (PTA) is defined as an incomplete division of the great arteries, resulting in a single arterial trunk. VSDs are almost always present. Multiple variations of the disease exist, resulting in differing levels of separation and spatial arrangements. PTA represents $\sim$1% of all CHD cases [13].

2.1.3.2 Diagnosis and treatment

Diagnosis for conotruncal lesions can occur prenatally during ultrasound imaging [21]. In some CHD pathologies, this can be of critical importance. For example, in TGA or TOF, babies after birth develop cyanosis if there is insufficient or no left to right shunting of blood, therefore requiring immediate treatment and possibly also perinatal surgical intervention [14]. TOF and TGA account for $\sim$10% of all

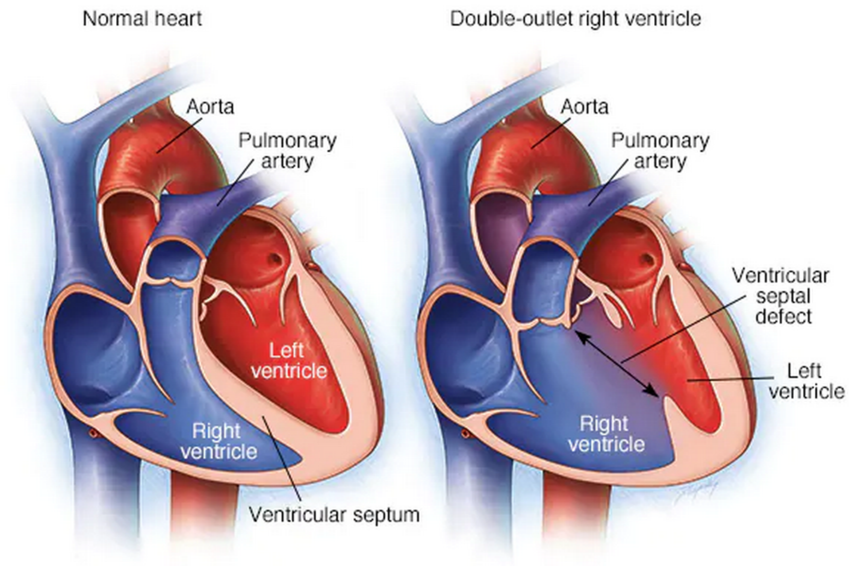

Figure 2.1: The anatomy of double outlet right ventricle [1].

CHD cases; however, only 17% of TGA cases and 50% of TOF cases are diagnosed in-utero [21].

Due to the complex spatial arrangements in many conotruncal lesions, treatment, often within the first year, is challenging and requires appropriate planning and surgical preparation. Lowered life expectancy and quality of life is often still not restored to normal levels even post-surgical intervention [13]. Patients regularly require follow-up and functional assessment of intracardiac implants and ventricular function. In some cases, re-intervention may be required in order to replace implants such as conduits when beginning to exhibit poor function or regurgitation. Examples of some surgical approaches for the previously aforementioned lesions are detailed in Table 2.2.

In the repair of conotruncal disease, the main aim is to restore the normal circulatory outflow from the heart, as much as possible. Since cases present with widely different anatomical arrangements, surgical approaches often need to be tweaked and personalised to the patient. This is compounded by the fact that CHD class divisions are not always clear (e.g. the controversy in differentiating between TOF and DORV) [22], which can result in the feasibility of multiple surgical approaches across a wide spectrum of cases. In general CHD surgery, 18% of cases necessitate

Diagnosis	Possible treatment	Description
DORV	Intraventricular LV-Ao baffle	Creation of an intraventricular tunnel (baffle) to divert and isolate the outflow of the left ventricle (LV) to the right-sided aorta through the pre-existing VSD.
TGA	Arterial switch	Transection of the great arteries above the sinuses (including with the coronary 'buttons') followed by repositioning in the correct ventricular-arterial arrangements.
TOF	Total repair of TOF	VSD patch closure and right ventricular outflow tract (RVOT) obstruction relief by means of pulmonary valvotomy and resection of hypertrophied muscle bundles.
PTA	Rastelli procedure	Closure of the VSD (using a patch) and insertion of an extracardiac conduit connecting the right ventricle to the pulmonary artery.

Table 2.2: Examples of surgical interventions for correcting congenital conotruncal lesions.

an approach that combines two or more procedures. The proportion of cases requiring a combined surgical approach increases to nearly 30% when considering the types of interventions commonly performed for conotruncal lesions. These include intracardiac baffling, arterial switch operation (ASO) and more [23].

In conotruncal CHD, the spatial arrangement of the great arteries and the intracardiac anatomy are of critical importance. Restoring normal cardiac outflow and preserving the function of both ventricles (biventricular repair) is considered a high priority and to produce better patient outcomes than single ventricle repair [24]. However, in DORV for example, the approach for biventricular repair is strongly dependent on the location of the VSD, great arteries, valves and their spatial relationship to one another. In some cases, biventricular repair may be impossible to perform, whereas in others it may be highly challenging due to the presence of comorbidities. For example, DORV may present with a stenotic RVOT (TOF-type) and/or transposed great arteries (TGA-type) which further complicates the route for performing successful biventricular repair. Despite the availability of surgical approaches for dealing with these issues (e.g. ASO), complex intracardiac baffles can be prone to complications such as left ventricular outflow tract (LVOT) obstruction, valvular disease and ventricular dysfunction [25]. Therefore, this might not be a suitable approach for all patients. Bradley et al. reported that biventricular repair (especially Rastelli-type) was associated with higher late mortality than single ventricle repair [26]. Indeed, the procedure with the highest risk of early mortality in DORV is intracardiac tunnelling with ASO [27]. This finding was particularly evident in borderline candidates that were nominated for biventricular repair. Nevertheless, biventricular repairs remain to be the primary choice for DORV and other conotruncal lesions, especially with favourable anatomy that does not necessitate the creation of convoluted or complex intracardiac outflow tracts.

2.1.4 Arterial lesions

Arterial lesions in CHD can be defined as structural abnormalities of the aorta, PA or coronary arteries. This section briefly describes some common types including modern diagnostic and treatment options, with particular emphasis on coarctation of the aorta.

2.1.4.1 Overview

Congenital abnormalities of the aorta, pulmonary artery, and coronary arteries may occur in isolation or in association with other CHD lesions. Typical arterial abnormalities include vascular narrowing (stenosis), dilation, improper anatomical connections and hypoplasia, with implications ranging from asymptomatic incidental findings to sudden cardiac deaths. These defects originate during embryonic development. At the 29th day of gestation, six pairs of aortic arches are present. After a series of complex processes, these eventually go on to form the aorta, pulmonary arteries and ductus arteriosus (which seals after birth). During this phase, abnormalities of the great arteries may develop, although further changes can occur postnatally, such as increased vascular elasticity due to a pressure imbalance [28]. Similarly, coronary artery abnormalities typically occur during the embryonic phase when the coronary buds (later sinuses) do not originate on either side of the undivided proximal bulbus cordis. Pathways of the vessels may also be abnormal if the routes do not follow the correct pattern through the epicardial atrioventricular and interventricular grooves [29]. Some examples of arterial lesions can be seen in Table 2.3.

Coarctation of the aorta (CoA) is an example of a fairly common arterial lesion

Aortic anomalies	Pulmonary artery anomalies	Coronary artery anomalies
Coarctation of the aorta	Patent ductus arteriosus	Coronary fistula
Aortic stenosis	Pulmonary artery stenosis	Duplication of the left coronary artery
Interrupted aortic arch	Interruption of the pulmonary artery	Anomalous origin of the coronary arteries
Hypoplastic aortic arch	Pulmonary sling	Single coronary artery

Table 2.3: Examples of congenital arterial lesions.

which occurs in 5-8% of all CHD cases [30]. CoA is characterised by a narrowing in the descending aortic lumen as a result of wall thickening and/or infolding (Fig. 2.2). The underlying factors leading to the development of CoA are still not fully understood. Common symptoms are systemic arterial hypertension, with more rare severe risk factors including aortic dissection and aneurysms. As with all arterial CHD lesions, CoA is commonly associated with other disease, such as patent ductus arteriosus (PDA), bicuspid aortic valve and hypoplastic aortic arch [31].

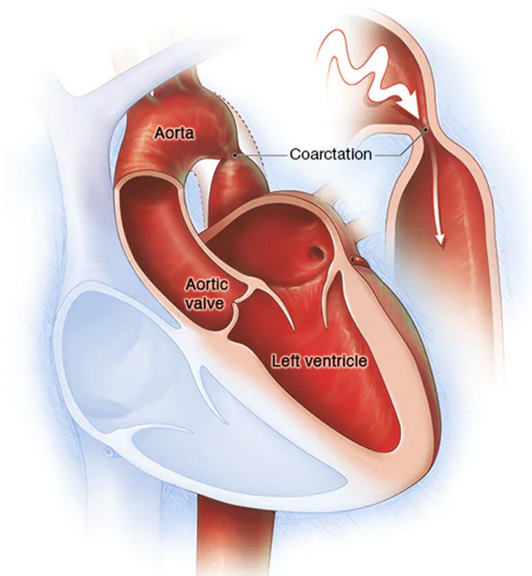

Figure 2.2: The anatomy of coarctation of the aorta [1].

2.1.4.2 Diagnosis and treatment

Diagnosis of arterial lesions are commonly missed, especially when isolated and not associated to a more easily diagnosed disease such as TGA or TOF. Undetected coronary abnormalities are one of the leading causes of sudden death in young adults [32]. The challenge of diagnosis is exacerbated by the fact that only 18-30% of children with coronary artery abnormalities experience symptoms [33]. In most cases, these children are diagnosed using transthoracic echocardiography. This technique is also commonly used to diagnose other arterial lesions. However, when used prenatally for diagnosis of CoA, it was found that two thirds of cases were not identified with TTE [30], resulting in CoA being the most commonly missed fetal CHD diagnosis. Once diagnosed, catheter angiography is considered the gold standard in assessment of CoA and for deciding intervention [30]. This minimally invasive catheter-based procedure allows for detailed pressure data to be recorded. The pressure gradient across the coarctation is typically used as an indication for intervention [34]. Other imaging techniques such as cardiac magnetic resonance (CMR) and computed tomography (CT) are also frequently utilised for diagnosis of arterial lesions. In general, congenital abnormalities of the great arteries result in clearer symptoms (e.g. hypertension, low arterial blood saturation, cyanosis) compared to congenital coronary anomalies and are thus easier to diagnose, at birth and later in life [35, 36].

The main aim of intervention in arterial lesions is to improve vascular function and reduce the risk of adverse events. Unlike conotruncal disease, arterial lesions can be repaired in certain instances with a catheter-based interventional approach that avoids open-heart surgery. This is done by expanding a balloon at a site of vascular narrowing in order to relieve stenosis/coarctation. In some cases, this procedure will be accompanied by the deployment of a balloon-expandable stent for preventing re-stenosis or re-coarctation. Some examples of surgical procedures are reported in Table 2.4.

In CoA for example, surgical repair by end-to-end anastomosis is usually the preferred choice for neonates [37]. Balloon angioplasty is usually performed after 3–6 months, due to the high risk of re-coarctation from highly elastic aortic tissue that is present in younger patients. Stent placement is generally performed on older children (>25 kg) as a more effective alternative to balloon angioplasty [30]. All

Diagnosis	Possible Treatment	Description
Coarctation of the aorta	Balloon angioplasty	Minimally invasive. A catheter-fed balloon is inflated at the site of coarctation to relieve narrowing. A stent may also be deployed.
Interrupted aortic arch	End-to-end anastomosis	The PDA is doubly-ligated. Both ends of the aortic arch are anastomosed to form an uninterrupted aorta.
Pulmonary sling	Resection of the PA with possible tracheoplasty	The PDA is doubly-ligated. The left PA branch is resected and sutured in the correct position. The trachea is repaired if necessary.
Anomalous origin of the RCA	Translocation of the RCA	The aorta is resected. The anomalous RCA is detached (including the coronary button) and reimplanted in the correct anatomical location.

Table 2.4: Examples of treatments for congenital arterial lesions.

of these treatments aim to eliminate the aortic pressure gradient and reduce hypertension in the left ventricle. The choice of intervention is highly dependent on the anatomy, presence of comorbidities and severity. For CoA, long term complications may occur after all types of treatment. Recoarctation can occur even after ”successful” interventions, often requiring reoperation. A review of 11 major studies by Rothman et al. found that rates of recoarctation were on average 16% [38], with a more recent publication by Dias et al. having placed this number closer to 12% [39]. Other serious complications include aortic aneurysms, which develop in 12% of patients with recurrent CoA, requiring long-term follow-up [34]. Systemic hypertension is another post-op complication which requires management, although this usually presents in patients operated after 20 years of age.

Other arterial lesions may also necessitate an aggressive interventional approach, followed by life-long monitoring. Davis et al. found that 0.17% of asymptomatic patients evaluated with TTE presented with anomalous origin of a coronary artery [33]. While the prevalence is low, if undiagnosed and untreated the consequences can be severe, with clinical symptoms including cardiomyopathy, arrhythmia, myocardial infarction and sudden death. Anatomical features such as acute take-off angles and slit-like ostia have been proposed as risk factors for severe adverse events [29] [32]. Surgery such as ‘unroofing’ of the coronary artery (exposing tunnelled sections in the aorta) may be used to reduce the risk of complications. Other options include re-implantation of the anomalous coronary artery into the correct location. However, indications and routes for treatment are often unclear, due to the difficulty in establishing the existence and severity of such a congenital defect. For arterial lesions, this extends even to post-surgical management, where monitoring and further functional assessment is equally important. For example, in patients diagnosed with pulmonary stenosis (present in 8% of all CHD), surgery may involve resection of the trunk and insertion of a transannular patch [40]. Pulmonary regurgitation is a common post-operative complication; yet, because it can be well tolerated, might not require reintervention for numerous years. However, this can lead to irreversible RV damage in the long-term. Overall, whilst the surgical management of arterial lesions has greatly improved, proper diagnosis, planning and follow-up are all equally important for ensuring the long-term health of patients, even of those presenting with common or 'easily treatable' arterial lesions.

2.2 Clinical challenges in CHD

The burden of CHD is rising on a global scale. One estimation placed $\sim$12 million people to be living with CHD in 2017, displaying an increase of 18.7% since 1990 [41]. Much of this can be attributed to a growing global population, along with improved treatments for sufferers of CHD that grant better long-term survival. Indeed, in the same time-frame, CHD mortality rates are understood to have decreased by 50% in countries with access to treatment. In some countries, adult patients with CHD are now more prevalent than pediatric patients. For example, in the United States alone, adults with CHD represent two-thirds of the entire national CHD population [13]. Another reason for the rising prevalence may be the inclusion of new data from regions with low socio-demographic indices that were previously underrepresented. Indeed, only 25% of the world has access to cardiac surgery, with the distribution being very skewed and favoured towards countries with high socio-demographic indices [42]. However, the increasing spread of diagnostic tools has led to previously unreported cases of CHD being identified.

Despite great advancements in the diagnosis and management of CHD, significant clinical challenges remain. Crucially, the education and training of clinical professionals requires more optimal and accessible resources in order to fill the growing need for CHD practitioners. Furthermore, improvements in the clinical approach are also required, even in centres which already treat CHD patients with low mortality/morbidity rates. Notably, greater clinical support should be made accessible for decision-making activities related to the timing of intervention, surgical planning and prediction of the patient outcome.

2.2.1 Decision-making

The landscape of clinical diagnosis, treatment, planning and long-term management in CHD is continually changing. Knowing **what** decisions to make and **how** to make them is a problem frequently encountered. This even extends past pediatric settings, since the rising prevalence of adult CHD has now created an almost entirely new set of clinical challenges. For example, the development of hypertrophy, arrhythmia and other morbidities can develop later on in life and take different forms in adults versus children [13]. Complications can also develop as a consequence of improper treatment/palliation, and are in some cases not fully understood or predictable due to current gaps in knowledge. This adds to the already challenging environment of clinical decision-making. In practice, clinicians typically need to follow medically defined guidelines when assessing patients. However, in CHD this is not always simple, as cases are often: (i) complex anatomically and sometimes unique, (ii) present with multiple comorbidities that need to be considered, and (iii) lie on the border between classifications, resulting in multiple possible treatment options. Due to these reasons, intuition and experience play an important role in the clinical management of CHD patients. In two clinical studies (not related to CHD), it was observed that intuition played an important role in clinical risk prediction, and in some cases led to incorrect or erroneous decision-making [43, 44]. In a field of surgery as complex as CHD, minimising the risk to the patient is crucial, not only for reducing mortality but also for producing better outcomes and quality of life. For this reason, there is still a need for further refinement of clinical decision-making related to CHD. Regarding interventional decision-making, the main problems that a cardiologist or surgeon might face have been distilled into three simplified questions:

- *Should* there be a surgical intervention?
- If so, *what interventional approach* should be taken?
- With the proposed approach, *what are the expected outcomes* of the intervention?

In CHD, answering these questions often requires a comprehensive analysis of the patient using medical imaging, physical examination, clinical measurements (age, weight, vital signs), medical history and more. In highly complex cases, the choice of intervention is sometimes not completely decided until the surgery is underway, due to the lack of strong clinical evidence in favouring one approach over another. For this reason, useful and informative data is needed at all levels to help guide the clinical decision-making process. Even with modern technological advances, treatment planning in CHD remains a major challenge.

2.2.1.1 Diagnosis and decision to intervene

At the very first step, diagnosis is required to establish the condition of the patient. This can be particularly difficult in arterial lesions, such as CoA, where diagnosis is missed in 85% of older children referred to a hospital with murmurs/hypertension [45]. Some of these patients are asymptomatic and will remain so, although $\sim$50% die under 30 due to sudden events such as aortic dissection/rupture.

Currently, the accepted medical indication for intervention in CoA is the existence of a transcoarctation pressure gradient of $>$20 mmHg [30, 31]. Catheterisation is considered the ‘gold-standard’ method to extract this value, although TTE is favoured in some centres since it prevents unnecessary hospitalisation. The binary threshold of 20 mmHg has been shown to not always give a clear indication for intervention. For example, it has been observed that some patients do not experience an LV strain reduction post-intervention [46], while other patients which are observed to be normotensive remain at risk of adverse events [47].

In post-op CoA patients which possibly require re-intervention, the approach on determining an intervention remains the same, regardless of the introduction of new factors. For example, Ou et al. showed that post-repair CoA arch shape is correlated to adverse outcomes [48], and Bruse et al. correlated post-repair CoA shape features with hemodynamic risk factors [49]. Despite this, there is an absence of shape-based indices to be used as part of the decision-making process. Further non-invasive metrics are needed to give clinicians a clearer indication of the patient physiological condition for diagnosis and intervention decision.

2.2.1.2 Pre-operative planning

Following diagnosis and a decision to intervene, the next step to consider is **what** procedure should be performed and **how** it should be conducted. In complex conotruncal lesions such as DORV, anatomies are often unique and require individualised surgical repair. In this case, the decision between a biventricular repair (preserving both ventricles) and a univentricular repair (preserving only one ventricle) has significant implications. For example, the perceived superiority of biventricular physiology against univentricular physiology has sometimes led to the inappropriate pursuit of biventricular repair in borderline candidates, resulting in elevated mortality rates [50]. Indeed, in DORV, the non-committed VSD types are associated with early and late morality [27]. The presence of non-commited/remote VSDs can also have other consequences; in some cases driving the choice unnecessarily towards a univentricular repair even with an adequately sized LV and RV [51].

For complex surgical repairs, each patient is often treated as a 'bespoke' case that necessitates its own tailored approach. For example, Corno et al. detailed a narrative-based approach to be used by surgeons for constructing a bespoke interventional plan in DORV repair [52]. With such an approach, the burden of problem-solving during surgery can be minimised through extensive planning, thus allowing the surgeon to focus more on the delivery of the procedure. However, since the lesions that require this kind of approach are complex, guidelines tend to be highly generalisable, which may limit their effectiveness for constructing individualised repair. For example, when trimming the patch during intraventricular baffling in DORV, Corno et al. recommends sizing it using the "most posterior point of the intraventricular communication and the most anterior point of the aortic valve" [52].

However, in reality, multiple patch sizes can fit this description, and this guideline fails to give strong constraints for the baffle design. This is important since improper baffle design may have significant clinical ramifications, such as LVOT obstruction. Indeed, surgical outcomes for CHD have been observed to differ strongly, even between centres within the same country [53]. In modern medicine, guidelines are needed in conjunction with medical imaging, which is the primary tool for planning surgery. However, for conotruncal lesions in particular, it is often challenging to relate 2D cross-sectional images to the actual 3D anatomy. The burden of needing to mentally visualise the patient's 3D cardiac anatomical configuration also greatly complicates the pre-operative phase. Therefore, to better support surgeons during pre-operative planning, improved accessibility to advanced 3D visualisation/simulation tools is essential.

2.2.1.3 Prognosis

The development of a prognosis which includes an expectation of surgical outcomes is an ongoing challenge for pediatric cardiologists. In CoA for example, it is mostly assumed that the elimination/minimisation of the transcoartation pressure gradient will improve the patient's health and reduce the risk of adverse events. However, it is not uncommon for post-op patients to experience complications, such as elevated LV strain or hypertension [30] [46] [54]. Additionally, if multiple interventions are required (e.g. due to recoarctation), further complications may develop. For example, aneurysms are $3\times$ more likely to form in patients with persistent coarctation than those with native coarctation [34]. Dijkema et al. reported that 4-14% of CoA patients show re-coarctation, 13% form aortic aneurysms and 35%-68% have persistent chronic hypertension [30]. Even though surgical repair is generally the favoured option over balloon angioplasty and stenting (especially in infants), it is unclear exactly how the surgical approach relates to post-op complications.

Being able to accurately predict the results of surgical intervention is useful not only for properly informing the patient/parents, but also for feeding back into the previous two clinical questions. For example, an infant diagnosed with post-op CoA (pressure gradient $>$20mmHg) would likely be nominated for surgical re-intervention. However, multiple compounding factors could contribute to a poor outcome, including an improper diagnosis, poor surgical approach choice or even an unclear understanding of the effect of the pre-existing arch geometry on the patient condition. Therefore, being able to accurately predict the post-intervention physiology on a patient-specific basis would be crucial for optimising the treatment and management plan. This fits into the vision of 'precision cardiology', which is of growing interest in CHD and relies on the use of computational modelling to improve diagnosis, procedure planning and prognosis on a patient-specific basis [55].

2.2.2 Education and training

Finding more effective methods to teach anatomy and clinical practice of CHD is of high interest. Historically, cadaveric dissection was the primary medium through which anatomical education was delivered. The traditional model of education was primarily an apprenticeship-like one, with instruction typically given on a one-to-one basis [56]. Changes in legislation during the early 19$^{\text{th}}$ century resulted in stricter regulations around the procurement of cadavers, resulting in a shift that favoured teaching anatomy to groups of students instead of individuals. The use of cadaveric specimens remains a crucial aspect of anatomical education, especially in CHD, where post-mortem hearts are still considered the 'gold standard' for teaching cardiac morphology. However, the wide breadth of CHD types makes it difficult for educators to consolidate a sufficiently diverse collection of samples for teaching. This is additionally hampered by the difficulty in acquiring and preserving rare lesions, resulting in a skewed distribution where the majority of resources remain isolated within specific centres and countries. Additionally, the global increased demand for cadaveric resources is in line with the growing numbers of medical students and trainees [57]. While samples are exceptionally useful in education, they are also not without their faults. For example, it is argued how much cadaveric specimens can be related to living anatomy [58]. This is especially the case in CHD, where collapsed heart chambers/vessels and the process of fixation may produce important alterations in the overall appearance/texture. Additionally, the diminishing numbers of trained anatomists is making it increasingly challenging to deliver and sustain this teaching format [58]. Finally, all anatomical specimens are prone to degradation over time. The dependency on anatomical specimens (despite their scarcity and limitations) has made it increasingly difficult to scale up CHD teaching in-line with growing global demand.

Alternative pedagogical methods of teaching which include diagrams, images and videos have become crucial for counteracting the scarcity of anatomical specimens. In modern educational settings, these resources have become the primary tools for delivering morphological teaching. With the advent of the World Wide Web, the dissemination and exchange of this content has been widespread, enabling easier access to educational resources for CHD than ever before. The rise of medical imaging has also played a role in generating visual teaching material which relates directly to living anatomy, as opposed to cadaveric specimens. However, all of these tools still have their limitations. Notably, the three-dimensional appreciation of anatomy is diminished when perceiving diagrams, images or videos on a screen. This is of particular importance in a field such as CHD, where depth perception is key to be able to properly delineate complex spatial arrangements of structures. Secondly, many resources do not contain any patient-specific relevance, for example, idealistic diagrams/videos which fail to provide a complete representation of true individual cardiac anatomy.

Recently, technological advances such as 3D-printed models [59, 60] have allowed for the creation of patient-specific, medical image-derived 3D models to be produced. A number of studies have recorded highly positive feedback related to the use of 3D printing for teaching CHD [60]. However, despite previously observed benefits, some limitations hamper the widespread use of 3D printing as an alternative to conventional CHD teaching methods. Primarily, the biggest challenge is related to time-requirements and high cost, which have made it difficult to scale up the production of 3D printing for CHD. Printers that create high-fidelity, deformable models are very expensive, with servicing/material costs also high. In a world which is becoming increasingly dependent on easily-distributable digital resources for independent learning, 3D printing does little to facilitate this transition. While it has been seen that lower-cost, rigid 3D printed options are useful [59], Yuen et al. found that educators greatly prefer anatomical specimens and, in some cases, do not consider 3D printed models as sufficiently accurate representations for anatomical teaching [61]. In general, whilst the usage of 3D printing is indeed useful as an accessory to conventional teaching methods, it remains to be seen if it can make a significant impact in revolutionising CHD education. A solution which solves many of the limitations of conventional CHD teaching and 3D printed models should: **(i)** be low cost and digital, **(ii)** enable independent/collective learning and **(iii)** enable 3D interaction with patient-specific anatomical models.

2.3 Medical imaging and post-processing

Medical imaging is crucial for the management of CHD from fetal development all the way into adulthood. It allows for the retrieval of both anatomical and physiological information related to the heart and circulation of patients. In addition, images are also important precursors for image-based computational modelling, discussed in later sections. In this section, the four main types of medical imaging used in CHD are briefly described, with advantages and disadvantages of each modality highlighted, and the specific use-cases where they are most effective (Table 2.5).

Imaging Modality	**Vascular Anatomy**	**Intracardiac Anatomy**	**Blood Flow**	**Ventricular Function**	**Interventions**
Echocardiography	×	✓	✓	✓	×
Catheter Angiography	✓	✓	✓	×	✓
Computed Tomography	✓	✓	×	✓	×
Magnetic Resonance	✓	✓	✓	✓	✓

Table 2.5: The capacity of the main imaging modalities used in CHD [8].

2.3.1 Echocardiography

Echocardiography is the primary imaging modality for initial diagnosis of CHD [62], most commonly used to diagnose and inspect the anatomy and function of structures such as valves, large arteries, chambers, myocardium and more [8], ranging from fetal screening all the way into adulthood. Echocardiography principles are based upon high frequency sound waves from a probe being emitted and directed into heart tissue. The ultrasound pulses reflect off tissues with different material properties and return to the probe. This data is then used to project an image on a monitor (Fig. 2.3).

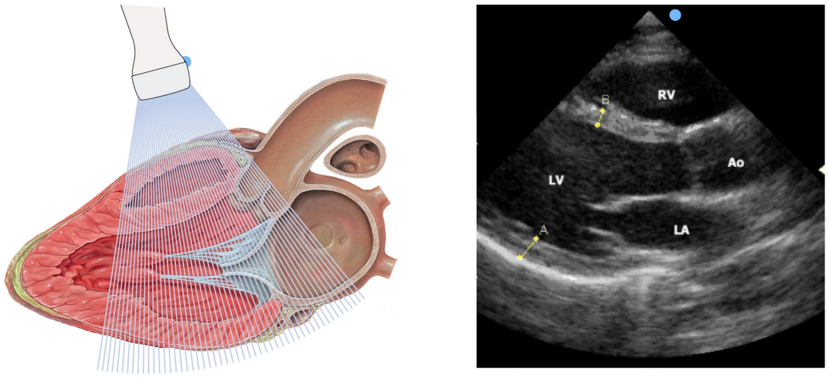

Figure 2.3: A parasternal long-axis view using two-dimensional (2D) echocardiography [2].

There are two main types of echocardiography commonly deployed: transthoracic two-dimensional echocardiography (TTE) and transoesophageal two-dimensional echocardiography (TOE). TTE is conducted by placing a probe over the patient's chest, whereas in TOE the probe is passed down the oesophagus and positioned behind the heart. TTE is the primary method for diagnosing cardiac conditions, whilst TOE is more commonly used during surgical interventions to monitor specific functions of the heart and guide catheterisation [63]. Benefits to TOE include a clearer acoustic window, due to the positioning of the probe which avoids the lungs and ribs. Information regarding the hemodynamics can also be retrieved using TTE or TOE, in a process known as Doppler echocardiography. This is common for identifying the presence of valvular regurgitation, shunting between chambers and calculate measures of cardiac function. In CoA for example, 2D Doppler echo enables velocity information to be recorded for estimating the pressure gradient across the coarctation [46]. However, the accuracy of this method is contested, with some reports observing Doppler echocardiography to overestimate the transvalvular gradient in patients with aortic stenosis when compared to CMR [64]. Errors due to inter-user variability also exist, as the image acquisition is highly dependent on the probe position/angle. This has been observed to be particularly evident in prenatal screening [65]. Lastly, the development of 3D echocardiography enables volumetric image acquisition, which gives better understanding of three-dimensional anatomy. However, limitations related to spatial and temporal resolution make 2D ultrasound the preferred choice, especially since the advantages of 3D echo for initial diagnosis and non-invasive measurements are not well established.

2.3.2 Catheter angiography

Catheter angiography is conducted in a catheterisation laboratory (cath-lab) and is typically used to visualise arteries, veins and heart chambers. A catheter is inserted into the patient via a blood vessel before being guided towards the heart or area of interest (such as a coarctation). X-ray fluoroscopy is used to provide a real-time indication of the catheter location. Once the catheter tip arrives at the destination of choice (e.g. a coarctation or stenosis), contrast agents can be injected and further images taken, usually at a frame rate of 2-3 per second. This provides a high resolution capture of vascular anatomy and dynamics. Additionally, the catheter enables direct measurement of pressures.

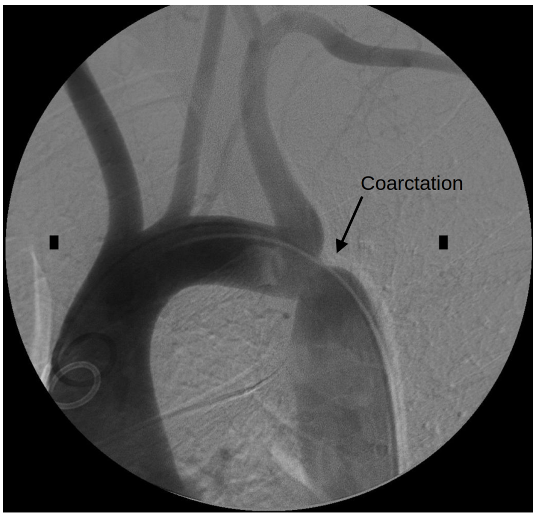

Figure 2.4: Coarctation of the aorta as viewed during a cardiac angiography procedure [3].

Catheter angiography is generally performed in situations where critical hemodynamic measurements are necessary, for example if there is a high degree of suspicion of coronary artery abnormalities [8] or for the quantification of transcoarctation pressures in CoA [30]. Since the approach is invasive, it is usually avoided in modern CHD as a method of diagnosis, not least due to the harmful ionising radiation which the patient is subjected to. Additionally, catheter angiography can only provide 2D projections of image data, which is suboptimal for the visualisation of 3D anatomy [66]. For these reasons, there has been a decline in the use of catheter angiography in favour of three-dimensional and non-invasive modalities, such as cardiac computed tomography and cardiac magnetic resonance.

2.3.3 Cardiac computed tomography

Cardiac computed tomography (CT) is a well-established imaging modality, frequently used for the management of patients with CHD. In CT, an X-ray source rotates around a patient while the attenuation is measured by a series of detectors on the other side. In modern scanners (multidetector CT), scans are performed in a spiralling fashion which has enabled acquisition during a single breath-hold and during the first pass of contrast, resulting in the ability to quickly reconstruct cardiac images in any 2D plane or in 3D [8]. Data is acquired in an axial spiral volume, and reconstructed into sagittal, coronal and oblique planes. Iodinated contrast is typically administered during the procedure in order to better visualise the heart, and is also referred to as computed tomography angiography (CTA). Capturing certain cardiac structures often requires the acquisition to occur at a specific time point in the cardiac cycle. In these cases, an electrocardiogram (ECG) is used for calibration, and is known as 'ECG-gating'. In modern scanners, the capture of multiple frames over a cardiac cycle is also possible and is known as four-dimensional CT (4DCT). Some indications for CT in CHD are reported in a single-centre study [8] and include: aortic vascular rings, PA anatomy (if no functional information needed), pulmonary venous anatomy, post-operative systemic to pulmonary shunts and more.

A major advantage of CT is the rapid acquisition time. This reduces the need for general anaesthesia, which is especially of importance in critically ill patients, as the risk of prolonged sedation can be higher than the radiation risk [8]. Another advantage is the high spatial resolution, which is particularly suitable for assessing

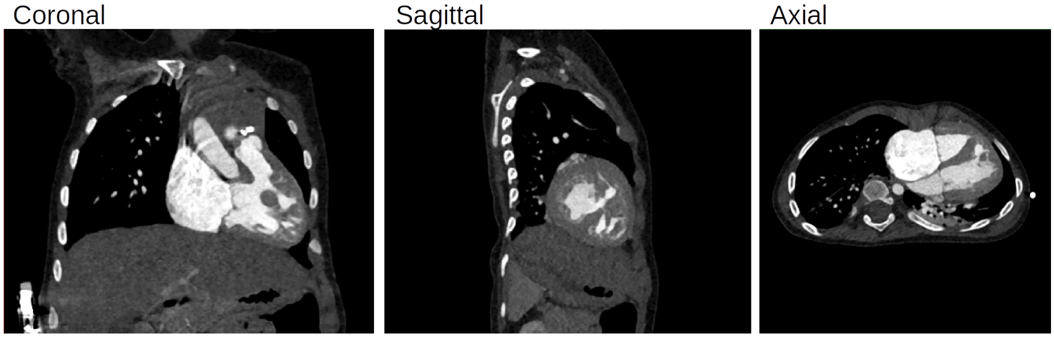

Figure 2.5: CTA dataset of a patient diagnosed with VSD. Coronal, axial and sagittal planes are shown.

intracardiac anatomy and in neonates. However, drawbacks to CT include radiation dosing, artefacts arising from medical implants and limitations when functional/hemodynamic information is necessary, especially at high heart rates.

2.3.4 Cardiac magnetic resonance

Cardiac magnetic resonance (CMR) imaging is a highly powerful imaging modality for CHD management. CMR is based upon the principles of 'nuclear spin', a fundamental property of atomic nuclei and elementary particles that can be affected by the presence of a strong magnetic field. CMR is a versatile tool which allows for detailed assessment of both cardiovascular anatomy and function. Other advantages include a large field of view, unlimited choice of imaging planes and low operator-dependency [8]. Different MR sequences allow for acquisition of different parameters and information. Steady state free precession (SSFP) is especially popular since it can create high levels of contrast between blood and nearby cardiac tissue using rapid pulse frequencies, thus massively decreasing acquisition time. Another application of SSFP includes bright-blood cine CMR that enables acquisition of dynamic images of the heart [67].

Using CMR, clinical metrics such as ventricular mass, end-diastolic volume and ejection fraction can be reliably assessed. In some cases, magnetic resonance angiography (MRA) may be performed by injecting a gadolinium-based contrast agent to improve the delineation of arterial lesions. Furthermore, phase-encoded imaging is commonly performed to measure flow volumes, peak velocities and regurgitation fractions. Velocity can be encoded in 2D, which is known as 2D phase-contrast (PC) CMR. Importantly, this has shown to be accurate (even more so than TEE) when compared to catheter angiography measurements [68]. Recent developments have also enabled 3D encoding of the flow field over time (4D flow CMR). However, this tool currently suffers from limitations preventing regular clinical usage; notably long acquisition times (10–20 minutes), large amounts of post-processing and low spatial resolution [69].

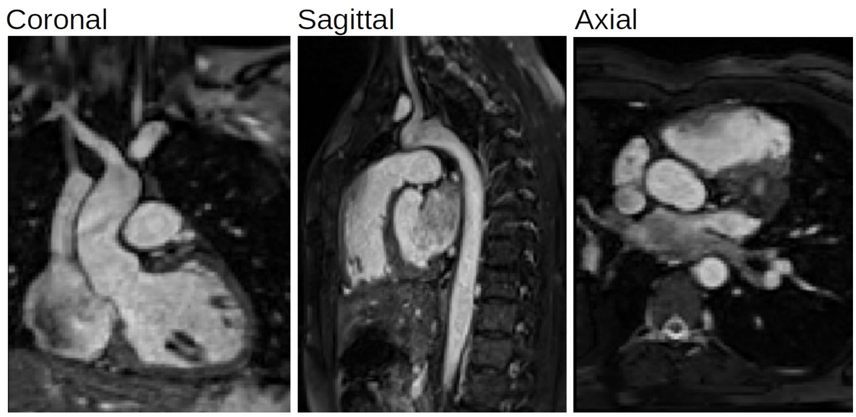

Figure 2.6: Whole-heart SSFP cine CMR in a patient with surgically corrected CoA.

CMR is appropriate for use in multiple scenarios for patients with CHD. For example, patients with surgically corrected CoA are recommended for lifelong follow-ups in order to screen for recoarctation, aneurysm formations and more. With CMR, the 3D anatomy including the location, length and diameter of the coarctation can be accurately established. CMR is also an excellent imaging choice for the assessment of patients with conotruncal lesions such as TOF, due to the ability to evaluate ventricular function. However, motion artefacts (due to respiratory and cardiac artefacts) often pose challenges, potentially requiring measures such as alternative pulse sequences, breath-holding and ECG/respiratory gating. Furthermore, CMR scanning of a patient with complex CHD can take up to an hour. In small children (especially those with high heart rates) this often proves difficult, thus necessitating sedation [66]. Additionally, patients with ferromagnetic metallic implants are unable to undergo CMR due to the strong magnetic fields. Despite these limitations, CMR remains a crucial form of assessment for CHD patients, and is constantly evolving due to developments in sequencing and machine learning.

2.3.5 3D image segmentation

3D images from CMR or CT are typically stored and transmitted in a universal format known as Digital imaging and Communications in Medicine (DICOM). This format contains image data and attributes related to the patient (e.g. age, weight, date of scan). Three-dimensional image information is stored as voxels with 8 or 16 bit greyscale (pixel intensity) values. As seen previously in Fig. 2.6 and 2.5, data can be viewed in three planes (coronal, sagittal, axial), with spacing between slices depending on the type of scan and commonly varying between 1-10 mm. Typically, the planar images are displayed on a monitor for clinical and diagnostic assessment. However, if needed, specific regions of interest may be extracted from the 3D volume via *image segmentation*, one of the fundamental processes for image analysis. The first stage of image segmentation is the generation of a 'mask', a binary 3D image which retains only the voxels in the structure of interest (e.g. the aorta). The creation of the mask can be done manually, by selecting individual voxels, or using semi-automatic tools present in most modern software, such as seeded region-growing and level-set contouring. Segmented regions of the 3D image can be converted into a 3D surface mesh with algorithmic approaches (Fig. 2.7). This transformation results in a three-dimensional grid of nodes with connectivity information (for defining triangular elements) that enables computer graphics visualisation.

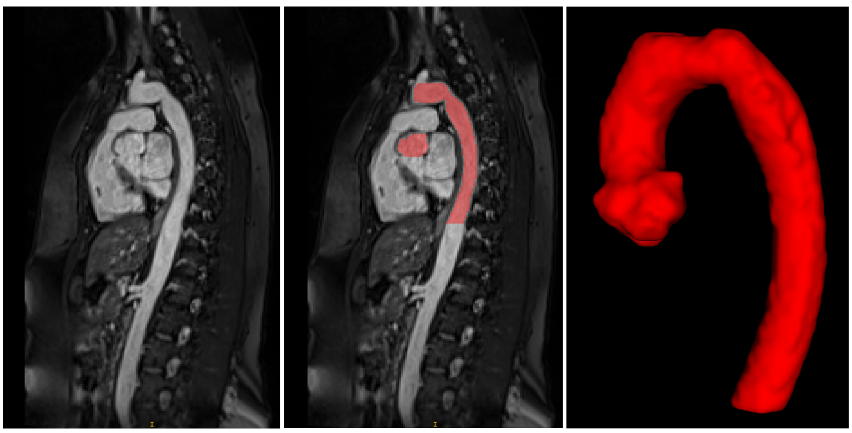

Figure 2.7: Image segmentation of an aorta from CMR in a patient with surgically repaired CoA. The generation of the mask is followed by surface mesh reconstruction.

An algorithmic approach is typically deployed (e.g. marching cubes [70]) to facilitate the conversion of a binary image volume into a 3D triangular surface mesh. Following surface reconstruction, it is common for grids to be further refined using algorithms for remeshing and smoothing. Several modern implementations of re-triangulating surfaces exist, many of which are based upon the principle of Delaunay triangulation and Laplacian smoothing [71]. Other mesh post-processing functions may be necessary in order to remove floating erroneous elements or holes in the surface. In cardiovascular applications, segmentations are used for object classification tasks for disease, denoising and organ measurement. In general, the use of image segmentation including subsequent surface mesh representation is also crucial in many cardiac computational modelling applications, some of which include 3D printing, virtual reality and finite-element simulations.

2.4 Summary

In this chapter, a general background to CHD was provided, starting with a classification of the primary lesion types. Special emphasis was given to conotruncal and arterial abnormalities, which were explored in greater depth due to their relevance in this thesis. Notable lesions within these classifications included double outlet right ventricle and coarctation of the aorta. Clinical challenges related to decision-making, education and training within the field of CHD were highlighted. It was shown that there is a lack of accurate non-invasive tools for diagnosis and prognosis of CHD lesions, especially for predicting post-operative outcomes. Similarly, it was seen that the repair of complex CHD is often hampered by challenging 3D anatomical configurations, which are not currently fully evaluated in 3D prior to intervention. In educational settings, the need for new methods to support the growing rise of CHD was highlighted, in particular with digital approaches that support 3D visualisation of patient-specific abnormalities and are widely accessible. Following this, a brief overview into the primary imaging modalities used in CHD was provided, along with image segmentation processes, critical for many computational modelling pipelines. In this context, the general concepts relevant to the methodologies adopted in this thesis will be described in the following chapter.

Chapter 3

Methodology

3.1 Introduction

Within the scope of this thesis, multiple methodologies were used (sometimes in conjunction with one another) to facilitate the translation of computational modelling tools into clinics. This chapter is used to describe the general theory of each modelling tool, including some recorded applications within the field of CHD from the literature. The four main methodologies discussed are: (i) computational fluid dynamics, (ii) machine learning, (iii) statistical shape modelling, and (iv) virtual reality.

3.2 Computational fluid dynamics

CFD is a subset of fluid mechanics which combines mathematical formulations of fluid structures with computational methods, for the purpose of quantifying flow properties in a given domain. During its history, the application of this technique has seen uses in several fields, such as aeronautical and automotive engineering. In CHD, the use of CFD is well documented for assessing the hemodynamics of patient cardiac structures, improving understanding of CHD, optimising medical device design and more. At its foundation, CFD relies upon fundamental principles of physical fluid behaviour.

3.2.1 Governing equations

The governing equations of fluid flow represent mathematical formulations of the conservation laws of physics:

- The mass of a fluid remains constant over time in any closed system
- The rate of change of momentum equals the sum of the forces on a fluid particle (Newton's second law)
- The rate of change of energy equals the sum of the rate of heat addition and the rate of work done on a fluid particle (first law of thermodynamics)

Fluid behaviour is described in terms of macroscopic properties, such as velocity, pressure and density. A fluid particle can be considered the smallest possible element of fluid, where the influence of microscopic properties is negligible. A small element of fluid can be defined with sides δx, δy and δz (Fig. 3.1). The approach of defining a control volume which contains fluid elements inside a fixed region is known as the Eulerian approach. Flow properties are expressed as fields, as opposed to tracking individual particles (Lagrangian approach). In CFD, the Eulerian framework is the most common approach, and requires forms of the conservation laws which are concerned with flow properties for fluid elements that are stationary in 3D space.

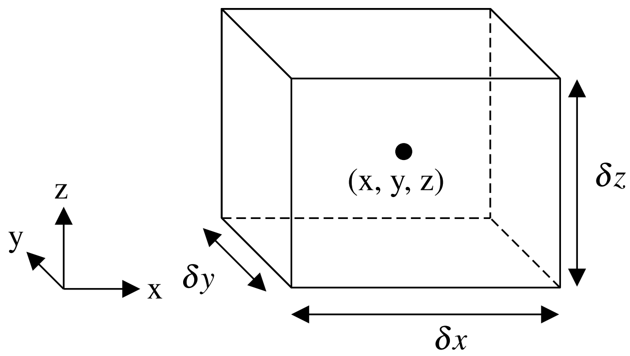

Figure 3.1: Fluid element for conservation laws.

In a fluid element, the centre is located at position (x, y, z). Fluid flow equations are derived from any systematic changes in mass, momentum and energy in the fluid element due to flow across its boundaries, and where appropriate, due to the action of sources inside the element. All fluid properties are functions of space and time. Since fluid elements are so small, properties at the faces can be sufficiently expressed with the first two Taylor series expansion terms. For example, the pressure at the left and right YZ faces (both at a distance of $\frac{1}{2}\delta x$ from the centre), can be expressed respectively as [72]:

$$\mathrm{p} - \frac{\partial \mathrm{p}}{\partial x}\frac{1}{2}\delta x \quad \text{and} \quad \mathrm{p} + \frac{\partial \mathrm{p}}{\partial x}\frac{1}{2}\delta x \tag{3.1}$$

3.2.1.1 Conservation of mass

In a fluid element, the rate of increase of mass is equivalent to the net rate of flow of mass into the element. The mass flow across each face also needs to be accounted for, and is given by the product of density, area and the velocity component normal to the face (Fig. 3.2). The sum of all these components equals the net rate of flow of mass into the element across its boundaries.

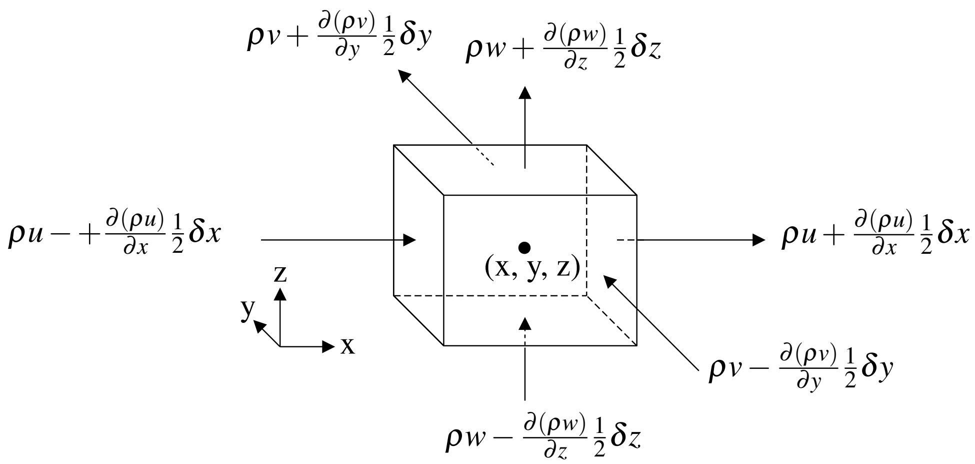

Figure 3.2: Conservation of mass in a fluid element.

The expression for the net rate of flow of mass into an element across its faces is then equated to the rate of increase of mass inside the element. Following this, the following equation can be derived (where **u** is the three-dimensional velocity vector):

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) \tag{3.2}$$

Eq. 3.2 is also known as the *continuity equation*. It describes the unsteady, three-dimensional conservation of mass at a point in a compressible fluid. The first term describes the rate of change of the density. The second term (convection term) describes the net flow of mass across the element boundaries. For an incompressible fluid (e.g. a liquid), the density p is constant, thus eliminating the first term.

3.2.1.2 Conservation of momentum

Newton's second law dictates that for a particle, the rate of change in momentum equals the sum of the forces on the particle. Rates of increase of x, y and z momentum per unit volume of a fluid particle are given as $\rho\frac{Du}{Dt}, \rho\frac{Dv}{Dt}, \rho\frac{Dw}{Dt}$, respectively. The stress acting on a fluid particle can be defined in terms of the pressure p and nine viscous stress components τ_{ij} that act in a direction j on a boundary ij. The product of stress and area yields the forces. For example, Fig. 3.3) shows the stress components in the x direction [72].

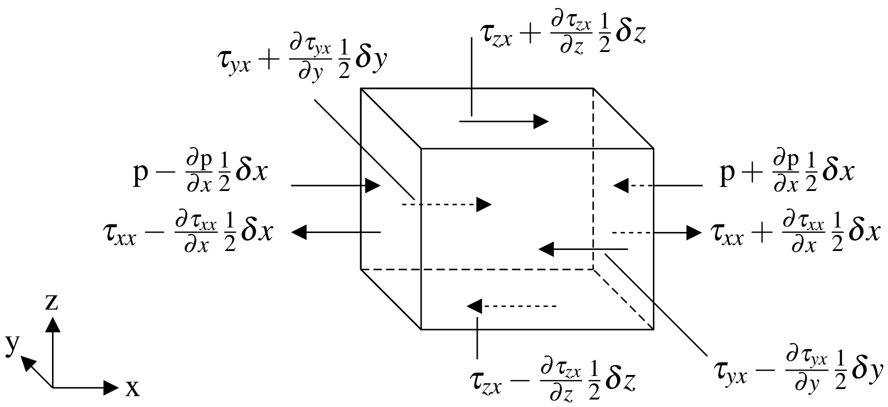

Figure 3.3: Stress components acting in the x direction on a fluid element.

After resolving all the forces in the x direction, the total force per unit volume on the fluid due to surface stresses is equal to the sum divided by the volume $\delta x \delta y \delta z$:

$$\frac{\partial(-\mathrm{p} + \tau_{xx})}{\partial x} + \frac{\partial \tau_{yx}}{\partial y} + \frac{\partial \tau_{zx}}{\partial z} \tag{3.3}$$

In order to avoid considering body forces in detail (Coriolis force, electromagnetic force etc.), their overall effect can be represented as a source S_{Mx} of x

momentum. By setting the rate of change of x momentum equal to the total force in the x direction due to surface stresses (Eq. 3.3) plus sources, the x component of the *momentum equation* can be defined:

$$\rho \frac{Du}{Dt} = \frac{\partial(-\mathrm{p} + \tau_{xx})}{\partial x} + \frac{\partial \tau_{yx}}{\partial y} + \frac{\partial \tau_{zx}}{\partial z} + S_{Mx} \tag{3.4}$$

The same approach can be used to derive the y and z components of the momentum equation:

$$\rho \frac{Dv}{Dt} = \frac{\partial \tau_{xy}}{\partial x} + \frac{\partial(-\mathrm{p} + \tau_{yy})}{\partial y} + \frac{\partial \tau_{zy}}{\partial z} + S_{My} \tag{3.5}$$

$$\rho \frac{Dw}{Dt} = \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial(-\mathrm{p} + \tau_{zz})}{\partial z} + S_{Mz} \tag{3.6}$$

3.2.1.3 The energy equation

The energy equation is derived from the first law of thermodynamics. It states that the rate of increase of energy for a fluid particle is equivalent to the net rate of heat added to the particle and the net rate of work done on the particle. The rate of increase of energy is denoted by $\rho \frac{DE}{Dt}$. The derivation of the energy equation involves a similar approach as previously shown, by resolving components of heat flux across the boundaries. The final energy equation is given in Eq. 3.7, where parameters include internal energy i, Boltzmann's constant k, temperature T, dissipation Φ and an energy source term S_i.

$$\frac{\partial(\rho i)}{\partial t} + \nabla \cdot (\rho i \mathbf{u}) = -\mathrm{p}\nabla \cdot \mathbf{u} + \nabla \cdot (k \nabla T) + \Phi + S_i \tag{3.7}$$

Liquids and gases flowing at low speeds behave as *incompressible fluids*. Without density variation, there is no connection between the energy equation and mass conservation and momentum equations. The flow field in these cases can be solved using only mass conservation and momentum conservation formulations. If the problem does not involve heat transfer, then the energy equation can be discarded [72].

3.2.1.4 The Navier-Stokes formulation

In the conservation equations thus far, the viscous stress components τ_{ij} have been additional unknowns. By using models to describe these viscous stresses, more useful forms of the conservation equations can be formed. One of the most common methods is to describe viscous stresses as a function of the fluid's deformation (strain) rate. In three-dimensional flows, the deformation rate is composed of a linear deformation rate and a volumetric deformation rate. For incompressible liquids, the effect of the volumetric deformation rate can be neglected. In a *Newtonian fluid*, the linear deformation rate is directly proportional to the viscous stress. Therefore, a constant of proportionality (dynamic viscosity μ) can be introduced to express all nine viscous stress components as functions of the linear deformation rate.

Substituting these new expressions for the viscous stresses into the conservation of momentum equations (Eqs. 3.4, 3.5 and 3.6) yield the *Navier-Stokes equations*. These are more useful forms of the conservation of momentum and energy equations, and describe the motion of a viscous, heat-conducting fluid. Convenient expressions of the Navier-Stokes equations can be constructed by 'hiding' the lengthy viscous stress terms in the momentum source. Therefore, the new source is defined $S_M = S_M + [s_M]$, where $[s_M]$ is the viscous forces. Finally, the Navier-Stokes equations can be written as:

$$\rho \frac{Du}{Dt} = -\frac{\partial \mathrm{p}}{\partial x} + \nabla \cdot (\mu \nabla u) + S_{Mx} \tag{3.8}$$

$$\rho \frac{Dv}{Dt} = -\frac{\partial \mathrm{p}}{\partial y} + \nabla \cdot (\mu \nabla v) + S_{My} \tag{3.9}$$

$$\rho \frac{Dw}{Dt} = -\frac{\partial \mathrm{p}}{\partial z} + \nabla \cdot (\mu \nabla w) + S_{Mz} \tag{3.10}$$

3.2.2 The finite volume method

The governing equations (conservation of mass, momentum and energy) have all been shown to share a similar form. By introducing a generic property ϕ, the conservative form of all fluid equations follow this structure:

$$\boxed{\frac{\partial(\rho\phi)}{\partial t}} + \boxed{\nabla\cdot(\rho\phi\mathbf{u})} = \boxed{\nabla\cdot(\Gamma\nabla\phi)} + \boxed{S_\phi} \tag{3.11}$$

Equation 3.11 is also known as the *transport equation* for property ϕ. Each of the four terms describes a different transport process in systems of fluid flow. The term in green describes the rate of change of ϕ in the fluid element, in red is the convective term, in blue is the diffusive term (Γ=diffusive coefficient) and finally in yellow is the source term. This is the starting point for the finite volume method, which is an approach for representing and evaluating a system of partial differential equations. By setting ϕ equal to 1, u, v, w and i, special forms of the conservation equations (mass, momentum, energy) can be formed. The key step in the finite volume method is to integrate these partial differential equations over a three-dimensional control volume (CV):

$$\boxed{\int_{CV}\frac{\partial(\rho\phi)}{\partial t}\mathrm{d}V} + \boxed{\int_{CV}\nabla\cdot(\rho\phi\mathbf{u})\mathrm{d}V} = \boxed{\int_{CV}\nabla\cdot(\Gamma\nabla\phi)\mathrm{d}V} + \boxed{\int_{CV}S_\phi\mathrm{d}V} \tag{3.12}$$

A crucial step in this approach is the conversion of the convective/diffusive volume integrals into surface integrals with the use of Gauss' divergence theorem. For a vector **a**, this theorem states that:

$$\int_{CV}\nabla\cdot(\mathbf{a})\mathrm{d}V = \int_A a_n\mathrm{d}A \tag{3.13}$$

Where A is the bounding surface and a_n is the component of vector **a** normal to the surface element dA. By applying this rule to equation 3.12, improved versions of equation 3.12 are found. In steady-state problems (no time component), the rate of change term can be ignored, resulting in the integrated steady transport equation:

$$\int_A(\rho\phi\mathbf{u})_n\mathrm{d}A = \int_A(\Gamma\nabla\phi)_n\mathrm{d}A + \int_{CV}S_\phi\mathrm{d}V \tag{3.14}$$

These terms can now be evaluated as fluxes at the surface of each fluid element. In order to solve for most fluid problems, conditions of flow at the domain boundaries are important to specify. The following auxiliary boundary conditions are typically required for solving steady/unsteady incompressible viscous flow inside a control volume (with no heat transfer):

- For unsteady flows, $\mathbf{u}$ must be given at time $t = 0$ everywhere in the solution region
- On solid walls, $\mathbf{u} = \mathbf{u}_{wall}$ (no-slip condition)
- At the inlet, $\mathbf{u}$ must be known as a function of position
- At the outlet, $-\mathrm{p} + \mu \frac{\partial u_n}{dn} = F_n$ and $\mu \frac{\partial u_t}{dn} = F_t$ (stress continuity)

In fully-developed flows (e.g. through a duct Fig. 3.4), it can be assumed there is no change in any of the velocity components in the direction across the outlet boundary ($\frac{\partial u_n}{\partial n} = 0$). Therefore, the gradient of all variables ϕ can be assumed 0 in the flow direction at the outlet, except for pressure.

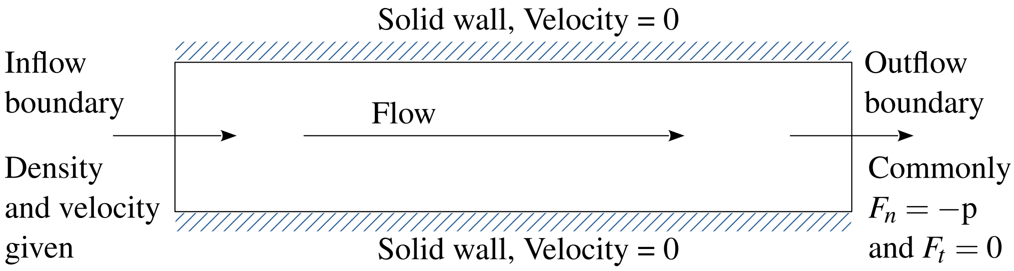

Figure 3.4: Boundary conditions for an internal flow problem.

In the finite volume method, the domain is typically divided into a grid of discrete cells, with properties conserved across element boundaries. Flow properties are computed at the centroid of each cell. In order to solve non-linear transport equations for thousands/millions of elements, a numerical approach is needed. Typically, this involves using a linear solver which iteratively solves and updates the flow variables until sufficient convergence is observed (when residual errors stop differing between iterations). Multiple algorithmic approaches have been demonstrated in CFD for performing this process, such as the semi-implicit method for pressure linked equations (SIMPLE), which is commonly used and is often available in commercial CFD solvers. Once convergence has been achieved, interpolation methods are used to describe the variation of properties between cell centroids in the volume. In the solution process, additional physical models for phenomena may be implemented, such as turbulence when modelling non-laminar flow.

3.2.3 The CFD procedure

In the previous sections, the general fundamental concepts behind CFD using a finite volume approach were described. However, in real applications, pre- and post-processing are also needed in the overall CFD software pipeline. The first step involves the creation of a geometry with boundaries that define the overall fluid domain. Common formats include surface meshes and computer aided design (CAD) models. Following this, the domain is typically discretised into numerous smaller volume mesh elements. The accuracy and convergence of the simulation is heavily dependent on the construction of this mesh. Typically, as the number of elements increases, the accuracy of the simulation also increases (at the expense of computational time). Multiple other aspects of a mesh also influence the solution, such as the mesh skewness, type of element, refinement level near walls and more. In real-world applications, it is not always possible to discretise a geometry with a structured mesh where nodes are uniformly distributed. The use of unstructured mesh solvers that build grids using tetrahedral elements (3D) is a robust approach commonly used to tackle this issue.

After mesh generation, the bounding surfaces of the control volume are defined, typically as inlets, outlets and walls. Inlet and outlet boundary conditions are applied and material properties (viscosity, density) of the fluid are also set. In the solver, additional modelling considerations (e.g. turbulence) are implemented, if necessary. The overall pipeline is shown in Fig. 3.5 [73]. Once the solution has converged, post-processing is essential in order to extract information of relevance, such as pressure/velocity distributions inside the flow field, from the 3D domain. Some popular commercial CFD solvers include Fluent and CFX (Ansys, Inc), Star-CCM+ (CD-adapco), OpenFOAM (openCFD .Ltd).

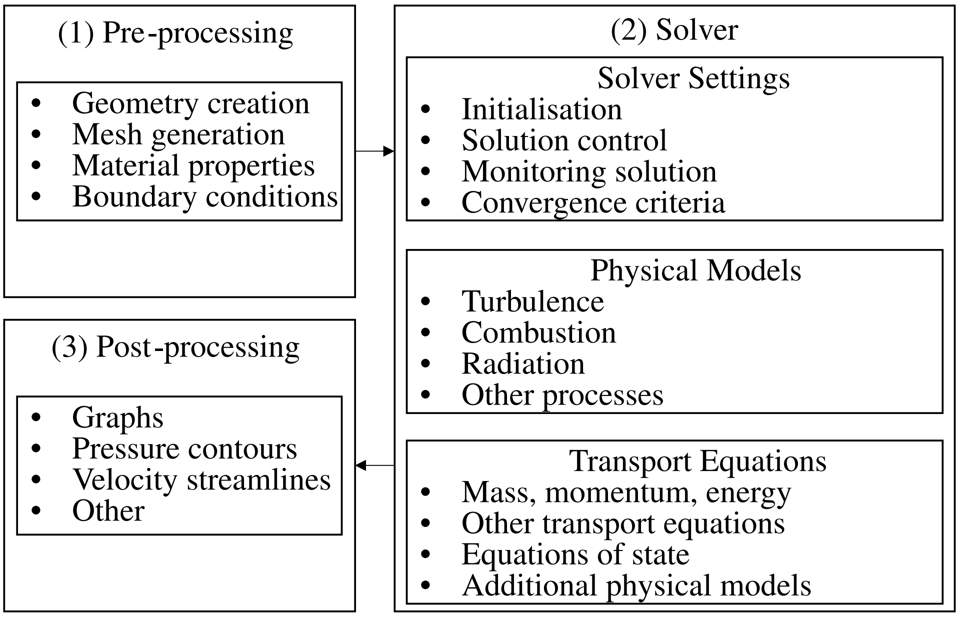

Figure 3.5: The overall CFD software pipeline from a user perspective.

3.2.4 Recent applications of CFD in CHD: a review

Over the past 20 years, there has been an increased interest in the use of CFD for cardiovascular applications. Improved CPU hardware, more accessible simulation software and a rising need for deeper insights into hemodynamic flows for research have all contributed to this growth. The ability to simulate in-vivo conditions has enabled low-cost, reproducible and non-invasive testing of stents, valve prostheses and other devices. For example, regulatory bodies such as the U.S. Food and Drug Administration have published guidelines on the use of modelling techniques for testing prosthetic valve designs [74]. The increased interest in simulation for device testing aligns with a growing focus on in-silico clinical trials. Simultaneously, the increased capabilities of image segmentation during this time period has enabled patient-specific CFD, allowing for non-invasive quantification of hemodynamics in humans. Importantly, several studies have validated CFD measurements against conventional clinical methods such as 4D cardiovascular magnetic resonance (CMR) imaging or catheter-based pressure measurements [75, 76].

Patient-specific CFD has become an essential component for the research and study of hemodynamic flow patterns in abnormal vascular structures. These insights can then be used to better understand lesions and their associated physio-

Application	CFD metrics	Findings	Author
Blood flow in shunts	Pressure, Velocity, WSS	Graft angulation presents a risk for WSS-induced thrombosis, which is more likely to occur in elongated central shunts	Ascuitto et al. [77]
HLHS	Pressure, energy loss, WSS	Creation of a smooth aortic arch angle reduced WSS and energy loss, and improves performance after Norwood	Itatani et al. [78]
RV pathologies (DORV, TOF)	Energy loss	Greater insight into the causes of RV-pulmonary circulation failure and new metrics for understanding dysfunction	Lee et al. [79]
Fontan circulation	Energy loss	Using CFD to optimise geometry and improve the efficiency therefore the clinical outcome	Rijnberg et al. [80]

Table 3.1: Examples of CFD applications for CHD.

logical features. For example, Youssefi et al. used patient-specific CFD to show that aortic stenosis is correlated with higher levels of wall shear stress and helicity in the ascending aorta [81]. A different CFD study by Osorio et al. showed that the placement location/angle of LV assist devices is correlated with risk of cerebral thromboembolism [82]. Ascuitto et al. showed that wall shear stress-induced thrombosis is more likely to occur in elongated central surgical shunts [77]. Other examples are shown in Table 3.1 [83, 84].

An emerging application of patient-specific CFD is to obtain physiological insight for guiding clinical diagnosis, decision-making and surgery. The predictive power of CFD in these cases may enable alteration of the conventional clinical procedure. For example, LaDisa et al. described the feasibility of CFD for estimating the pressure gradient non-invasively in coarctation of the aorta, potentially reducing the need for catheterisation [85]. This is especially significant since such an approach may be a suitable alternative to cardiac catheterisation with pharmacological stress testing for assessing conditions during exercise [86]. CFD has shown to be suitable for modelling flow even in highly clinically challenging 3D anatomy, such as coronary arteries with anomalous aortic origins [87] or post-repair structures including corrected coarctation of aorta (see Chapter 2, Section 2.2.1.1) and Fontan circulation [88,89]. Furthermore, 3D CFD can be used to derive metrics which have potential for aiding the risk stratification of patients vulnerable to adverse events.

For example, Qiu et al. associated a specific CFD helical hemodynamic flow pattern with a significantly increased risk of abdominal aortic aneurysm rupture [90].

In some cases, virtual surgery may also be performed alongside with CFD to assess the optimal surgical configuration and maximise the predictive capabilities of CFD. For example, in patients with hypoplastic left heart, surgical correction involving the creation of a total cavopulmonary connections (TCPC) may be necessary (Fontan procedure). CFD has shown to be suitable for better understanding how the hemodynamics is related to TCPC geometry, and in some cases has been used for pre-operative planning on a patient-specific basis [91–93]. Another example of CFD for virtual surgery was presented by Itu et al. who described the use of CFD for predicting the hemodynamic response to aortic stenting in patients with CoA [94]. Thus, integration of CFD into clinical settings may have important ramifications when used: (i) for detailed assessment of patient-specific hemodynamics in response to stress or interventions, (ii) to estimate CFD-derived indices for supporting clinical decision-making and risk-stratification, and (iii) alongside virtual surgery for optimising the interventional approach.

Sufficient evidence exists to indicate that CFD can potentially provide significant benefits in the routine clinical management of CHD [95, 96]. However, the lack of an automatic and fast CFD platform suitable for clinical implementation has prevented this hypothesis to be adequately explored. Namely, high computational resources and long simulation times make it difficult to facilitate integration in hospitals. Specialist knowledge is often needed in order to setup the simulation and post-process the data. Due to these challenges, CFD has remained primarily a research tool, and has yet to be properly implemented in routine clinical practice. In order to encourage this transition, fully automatic, fast, reproducible and robust CFD prediction tools would be necessary to lower the barriers for entry into the healthcare environment.

3.3 Machine learning

ML can be defined as the scientific field of creating computer models which can learn to perform a task from data, without being explicitly programmed. A computer program is said to learn from experience E with respect to a task T and some performance measure P, if its performance on T (measured by P) improves with experience E. Much of ML relies on the combination of both statistical theory and computational algorithm/optimisation principles. The modern landscape of ML is vast, with applications ranging from self-driving cars, to email spam filters and advertising. In CHD, ML has recently found applications in accelerated CMR acquisition, automatic image segmentation, diagnostic support and a few more.

ML can be subdivided into two categories: supervised ML and unsupervised ML. In supervised ML, the training data E used for model training is labelled. Therefore, inputs are paired with their corresponding output, which makes it possible to use specific algorithms to learn patterns. Prediction tasks for supervised models typically constitute of either classification (discrete output) or regression (continuous output) problems. For example, at its most simplistic form, linear regression can be considered a form a supervised ML. For an existing numerical dataset with two variables (x and y), a line of best fit with the form $y = B_0 + B_1 x$ can be found by using a numerical approach that looks to minimise some error between the expected and actual results. This principle forms the basic foundational backbone of how modern ML algorithms train models. Several different types of supervised ML algorithms exist, such as k-nearest neighbours, decision trees, logistic regression and more, each with their own advantages and disadvantages. Unsupervised ML models work by finding patterns within unpaired data, and are outside of the scope of this thesis.

3.3.1 Artificial neural networks

An artifical neural network (ANN) is a highly flexible type of ML modelling architecture which can be used in numerous applications. The principles are inspired by the networks of biological neurons found in the brain. One of the earliest and simplest forms of an ANN is the *perceptron*, which is based upon the concept of an artificial neuron called a linear threshold unit (LTU). Each input is a numerical value associated with a weight coefficient. The LTU computes a weighted sum of its inputs, then applies a step function to that value for the final output. This can be used for binary classification tasks. The output is given by the formula $H_\Theta(\mathbf{x}) = \text{step}(\mathbf{x}^\text{T}\Theta)$, where the inputs are $\mathbf{x}$ and the weights are Θ.

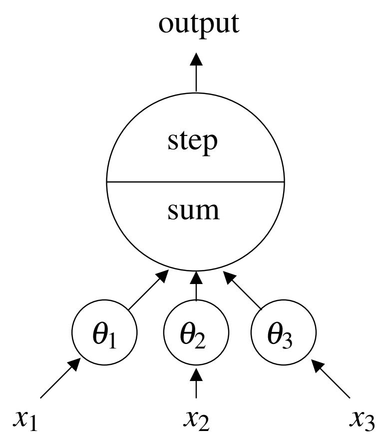

Figure 3.6: Architecture of an LTU.

Training an LTU in this case means finding the correct values for the weights so that the step function classifies input data as best as possible. A perceptron is composed of a single layer of LTUs, with each unit connected to all inputs. This forms a *dense* layer, while input neurons form an *input layer*. A trainable bias feature containing a constant value is generally added (e.g. $b = 1$). The role of the bias unit is to effectively act as the 'y-intercept' variable for the layer, enabling for a more flexible model. Perceptrons can be stacked and connected between layers to form more complex models for predicting multiple outputs, known as multi-layer perceptrons (MLPs) [97].

Equation 3.15 shows how the outputs of an MLP can be expressed as a function of the inputs $\mathbf{x}$, weights Θ, bias b and *activation function* ϕ. In modern ML, much more complex and effective architectures have been developed, however the perceptron remains a fundamental conceptual building block in ANNs.

$$h_{\Theta,b}(\mathbf{x}) = \phi(\mathbf{x}\Theta + b) \tag{3.15}$$

3.3.2 Deep neural networks

When an ANN contains more than one layer (excluding input and output layers), it is commonly referred to as a deep neural network (DNN). If the network parses data in a sequential manner with no looping connections, it can be considered a *feed-forward* network. DNNs are highly popular due to their extremely configurable structure, which can be tailored in order to solve complex non-linear problems.

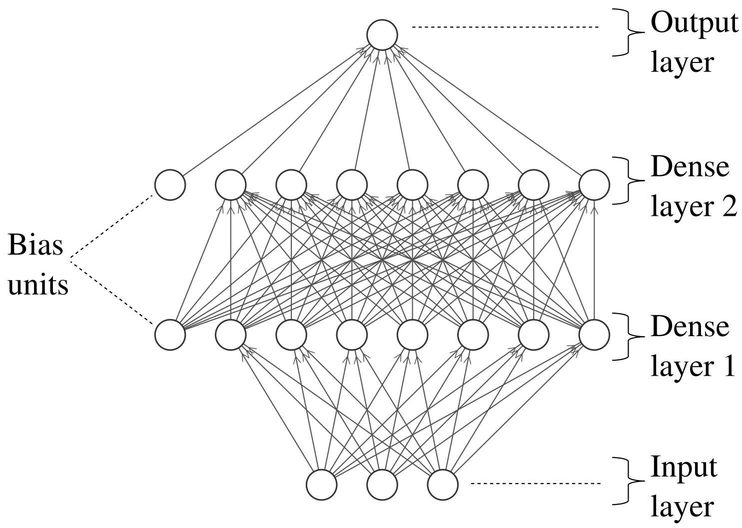

Figure 3.7: Example architecture of a single output regression DNN. The input layer contains 3 units. Two hidden layers contain 8 units each. The bias units are shown, although are often omitted in diagrams.

An essential algorithm for computing the errors of all individual network units was first described in 1986, and termed *backpropagation* [98]. This approach uses two passes: a forward pass which computes the prediction error of the model, and a second pass which goes through all network layers in reverse to measure the contribution from each connection to the overall error. After this, the weights are tweaked in each unit to reduce the error. This is done using an iterative process called gradient descent (GD), which measures the local gradient of a cost (error) function with respect to a weight vector Θ, and seeks to adjust Θ in a direction that reduces the error. It is based on the observation that a differentiable multivariate function F_Θ decreases the fastest if it goes in the direction of the gradient. Equation 3.16 expresses this relationship, with γ representing a constant learning rate [97].

$$\Theta_{i+1} = \Theta_i - \gamma \nabla F(\Theta_i) \tag{3.16}$$

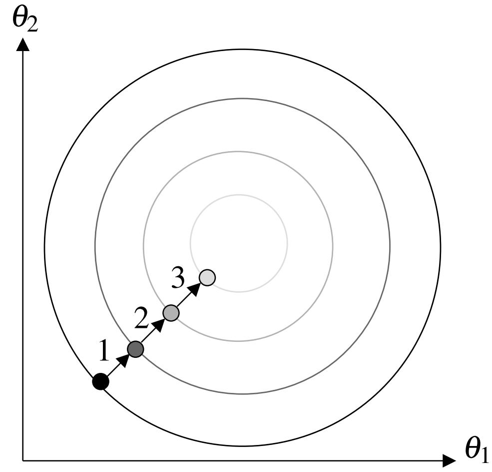

Figure 3.8: Gradient descent in a convex function. A random starting point is initialised. With each iteration, the local gradient is computed and used to adjust the weights θ_1 and θ_2 towards the minima. Three example iterations are shown.

For example, if using the commonly applied mean squared error (MSE) formula as the cost function in a simple linear regression model, F_Θ would be considered convex. In a convex function, the line connecting any two points on the curve cannot cross it. For two parameters, this can be visually imagined as a bowl, since the function is continuous, smooth and has a single global minimum at the centre (Fig. 3.8). In this case, gradient descent would eventually converge towards the global minimum as each iteration adjusts the weights in ways that minimise the cost function.

As the step function is non-differentiable, alternative activation functions are needed in MLPs. Some commonly used non-linear activations include the rectified linear unit (ReLU) and logistic function. ReLU has become widely popular due to its fast performance and good applicability to most problems. Also, since ReLU has a constant positive value when differentiated, it can solve a common issue in backpropagation where gradients at lower layers receive small values and remain unchanged (known as vanishing gradients). However, ReLU also has drawbacks, and can cause GD to bounce around due to its sharp slope change about the origin

$(0,0)$. In addition, units can sometimes stop learning (dying ReLUs) if the network is tweaked in a way that the weighted sum of inputs to a unit is negative. Because the derivative of the ReLU function is 0 at negative values, that neuron would effectively stop training. Despite these issues, ReLU remains highly useful in DNN modelling.

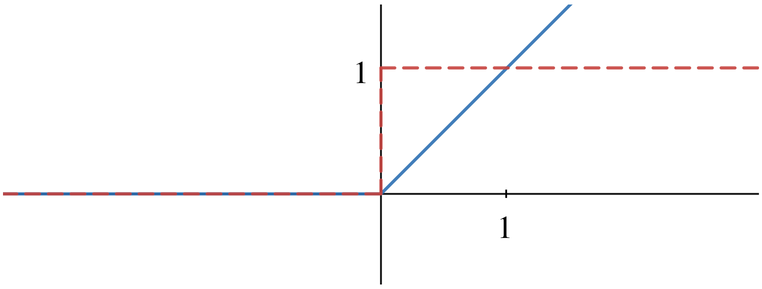

Figure 3.9: ReLU activation function (blue) and its derivative (red).

Unlike the previous linear regression example where the cost function formed a smooth bowl (Fig. 3.8), gradient descent in DNNs is usually a highly dimensional and non-linear problem. This only worsens the more parameters and weights there are to tweak. The multivariate landscape of a DNN solution can contain several local minima that the optimiser can get stuck in. Large plateaus and sharp variations in gradients can also make convergence unstable or time-consuming. A highly important parameter when training DNNs is the learning rate (LR) which influences the size of each iteration step. LR is an example of a *hyperparameter*, which needs to be tuned to improve the performance of the optimiser. If the selected LR is too large, the solution will bounce around without converging; too small and it will take a long time to converge.

A common practice is to use momentum-based optimisers such as adaptive moment estimation (Adam) and RMSProp to make convergence easier. Momentum optimisation works by subtracting the local gradient from a momentum vector **m**, and updating the weights by adding **m** (Eq 3.17).

$$\mathbf{m} = \beta \mathbf{m} - \gamma \nabla F(\Theta_i) \tag{3.17}$$

Referring back to the bowl analogy, this effectively means the gradient is used only for informing *acceleration*, not speed. To simulate some sort of friction in the system and prevent momentum from growing too large, a new parameter β is introduced. The expression for momentum is shown in Equation 3.17, with the gradient descent weight-update step effectively becoming Equation 3.18.

$$\Theta_{i+1} = \Theta_i + \mathbf{m} \tag{3.18}$$

The use of momentum in gradient descent effectively allows the optimiser to more quickly cross 'plateaus' and 'roll' over local minima. In optimisation schemes like Adam, there are also added benefits such as a reduced need for LR tuning, since Adam adaptively adjusts the LR while running. Another problem other than convergence is overfitting. If model weights are tweaked too stringently with respect to the training data, the DNN will be unable to perform inference on unseen data. Techniques such as stopping optimisation early, regularisation, batch-normalisation and dropout all have been designed to reduce overfitting and improve performance [97].

3.3.3 Convolutional neural networks

A convolutional neural network (CNN) is a form of DNN, widely used for computer vision applications. CNNs make it possible to perform complex image-based tasks through the use of *convolutional layers*. The role of a convolutional layer is to assemble small low-level features into larger higher-level features. This requires a kernel (filter) to parse through the image and produce a new matrix of features with reduced dimensionality (feature map). This is shown in Fig. 3.10, where a 6x6 image is converted into a 4x4 feature map using a 3x3 convolutional filter. In this example, the filter starts from the top-left corner and *strides* one pixel at a time, eventually ending in the bottom-right corner. The filter performs an element-wise dot product with each section. The resulting matrix values are summed up and form the new pixel value in the feature map [97].

After the convolutional layer, it is common to perform *pooling* on the resulting feature map. This is used to downsize the matrix and reduce computational load through dimensionality reduction. Additionally, it is useful for extracting dominant

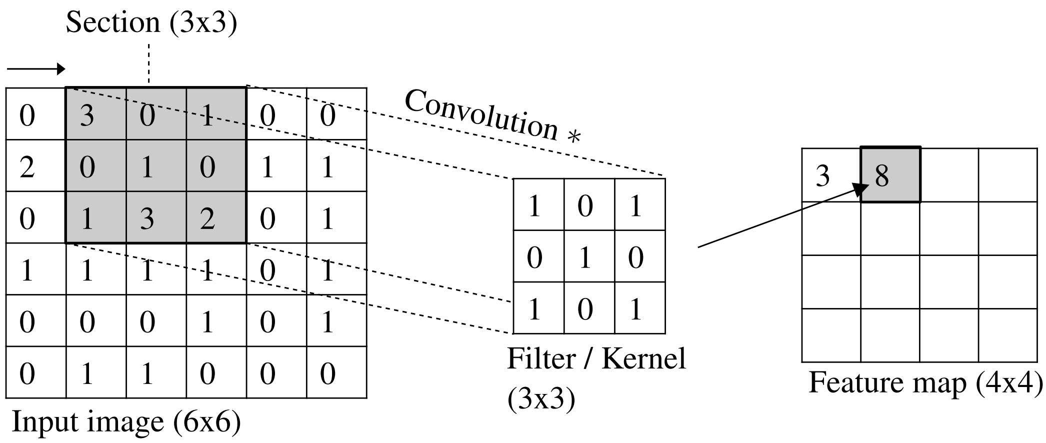

Figure 3.10: Convolutional filter operation.

features which are positionally and rotationally invariant. This can help the model become more robust and perform better. Max pooling is the most popular method and involves retaining only the highest element in a section. This has a secondary effect of acting as a de-noiser, as it discards all pixels in each section apart from one.

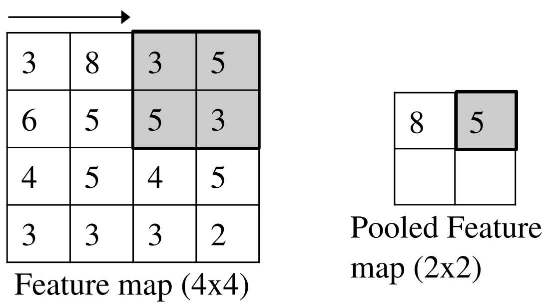

Figure 3.11: Max pooling with a filter size (2x2) and a stride of 2.

CNN architectures may involve several layers of convolutional and pooling layers. The convolved and pooled features are then passed into a fully connected network, which is responsible for either classification or regression, depending on its design. In an image-recognition classifier with four possible classes, the network may follow a general architecture as seen in Fig. 3.12. During training, each weight in the hidden layer is adjusted using backpropagation as normal. In a convolutional layer, each filter (array of numbers) can be considered a weight, and is also optimised using gradient descent. This allows the network to prioritise specific features which are of relevance to the problem solution.

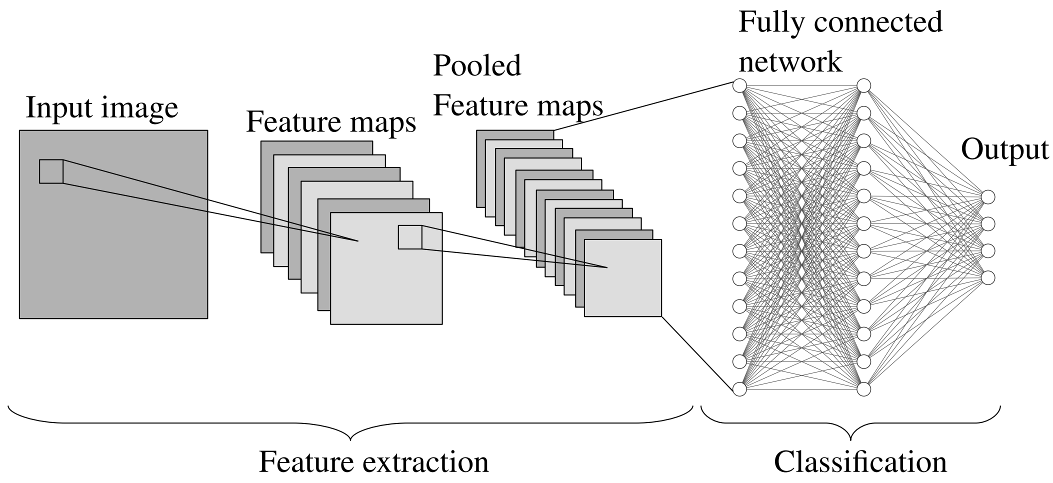

Figure 3.12: Example architecture for a CNN classifier with 4 output classes.

Progress in the field of CNNs in recent decades has been rapid, resulting in numerous types of highly specialised network architectures for different applications. In the field of automatic medical image segmentation, the U-Net architecture has become widely popular. This multiscale network is composed of an encoder (consisting of convolutional and downsampling layers), and a decoder (which relies on convolutional and upsampling layers). Each convolutional layer is followed by batch normalisation and ReLU activation. The number of convolutional filters doubles with each downsampling layer, and halves with each upscaling layer. After every pooling step, the spatial resolution is halved, and vice versa. The final convolutional layer contains a softmax activation function and performs the classification of the features into an image mask. In this step, each pixel is assigned a value corresponding to a class (e.g. 0 for background, 1 for organ).

Importantly, this network makes use of skip connections by appending data from earlier layers onto later layers. This effectively creates an uninterrupted gradient flow from the first layer to the last layer, thus diminishing the 'vanishing gradients' problem. Additionally, this improves the models ability to retain information that may have been lost during downsampling, enabling more stable weight optimisation and improving model prediction granularity [99].

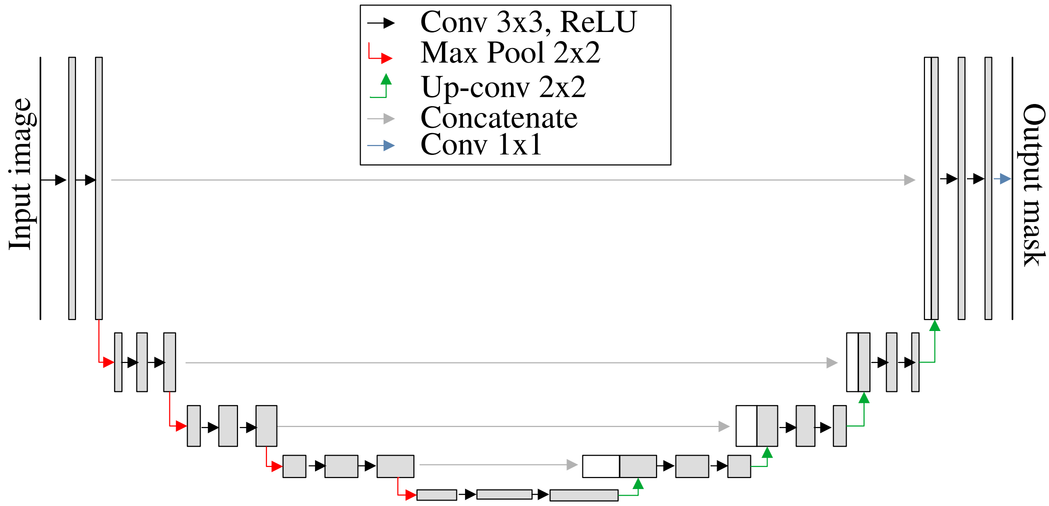

Figure 3.13: Overview of standard U-Net architecture.

3.3.4 Recent applications of ML in CHD: a review

Due to growing evidence of the usefulness and accuracy of ML models, a rising number of ML algorithms are being approved for clinical use by regulatory bodies worldwide [100]. A common application of ML is for image-based tasks. For example, Hauptmann et al. described a CNN-based method which accelerates whole-heart CMR image acquisition and removes motion artefacts [101]. Steeden et al. showed that images produced using an accelerated CMR acquisition method were not statistically different to the ground truth data, and could be acquired 3x faster than conventional scanning [102]. The use of image-based models for aiding diagnosis/prognosis has also been documented [103]. Arnaout et al. presented models which were able to identify abnormal fetal ultrasound cases with comparable accuracy to manual clinical evaluation [104]. In a separate study, Arnaout et al. showed that CNN-based classifiers could potentially support echocardiography through accurate identification of the viewing plane in a TTE scan [105]. Other documented cardiovascular imaging applications include detection of hypertrophic cardiomyopathy, quantification of aortic valve area from ultrasound and assessment of calcium deposits [100].

A recent application of ML in medical imaging is for automatic image segmentation. Automatic or semi-automatic segmentation methods based on shape models or level-set algorithms have been extensively used in the past [106–108].

However, they often require some type of user input or manual post-correction, rely on priors which are not readily available, or may struggle to adapt to abnormal anatomies (e.g. CHD). Automatic CMR image segmentation of cardiovascular structures such as ventricles, atria, great arteries has been possible with the use of DNN architectures such as the U-net [109–111]. This has been achievable even with very limited datasets, due to image-augmentation and synthetic data generation facilitated via deep learning techniques [112]. Quantitative metrics derived from ML segmentations (e.g. ventricular volumes) compare well with manual segmentations [113, 114], and these techniques are now entering clinical practice.

The acceleration of image segmentation has other, indirect consequences on the clinical management of CHD, by helping to facilitate the translation of modelling tools. For example, image segmentation remains one of the most user-intensive and time-consuming parts of the CFD modelling workflow, and automating this step would lower the barriers for CFD clinical adoption [115]. ML can be used to accelerate the CFD pipeline in even further ways, such as by building models which entirely replace the need for a conventional CFD solver. For example, Yevtushenko et al. presented a centreline-based DNN model capable of inferring pressure gradients in aortic geometries, with results equivalent to those from a conventional CFD solver [116]. Feiger et al. described a ML-based approach for developing models that infer accurate pressure drop and WSS values in patients with CoA [117]. Liang et al. detailed a shape-driven approach in which a synthetic cohort of aortas was generated in order to produce a sufficiently large dataset for ML model training. In the same study, ML models were shown to be capable of reproducing aortic pressure and velocity flow fields with high accuracy, when compared to conventional CFD [118, 119].

ML has also been used in non-imaging based applications of CHD healthcare, such as for predictive and diagnostic modelling. For example, Qu et al. built an ensemble ML model that uses over 1000 clinical indicators (glucose, coagulation levels, etc.) to predict the chance of CHD development during early pregnancy. Preliminary results were promising and also highlighted some specific clinical parameters that displayed strong correlations with CHD development [120]. The use of ECG-based DNNs for detecting arrhythmia and left ventricular dysfunction have also been demonstrated [100]. Diller et al. presented a single-centre study showing that a combination of inputs (including from ECG, clinical metrics and demographic information) could be used to train models capable of prognosticating patient disease and suggesting therapy options [121]. By augmenting electronic patient records using these models, in the future it may be possible to better manage patients and provide more personalised care.

Overall, it has been seen that ML and DNNs are being used more and more frequently in clinical CHD environments in order to perform image-based tasks, accelerate the CFD workflow, provide diagnostic support and help the clinician to better understand patient anatomy/physiology. Of notable relevance is the proven ability of ML for facilitating the translation of existing computational tools into clinical settings (automatic image segmentation, CFD). Challenges for developing new ML solutions include large requirements of data, time-consuming data labelling and opaque 'black box' models which can limit how interpretable the results are for a clinician. Despite these limitations, ML applications in CHD are increasingly applied due to the vast potential benefits.

3.4 Statistical shape modelling

A statistical shape model (SSM) describes the distribution of a population of shapes. One of the most simple applications of statistical shape modelling is to represent a group of objects (e.g. skulls) as a mean shape. SSMs are also used widely in medical imaging for tasks such as clustering and exploring shape variability within a cohort. Multiple frameworks for constructing SSMs exist, each with differing approaches. Shapeworks [122] is a popular solution which uses a particle-based method for the parameterisation of shapes. Another popular framework is Deformetrica [123] which is based upon the large deformation diffeomorphic metric mapping (LDDMM) method.

3.4.1 Deformetrica

With Deformetrica, shapes can be expressed as deformations of ambient space (the space surrounding and including a mathematical object). D'Arcy Thompson (1917) was the first to propose using ambient space deformations as a way of comparing biological shapes. In this context, a *diffeomorphic* transformation of ambient space is a deformation that is both differentiable and invertible.

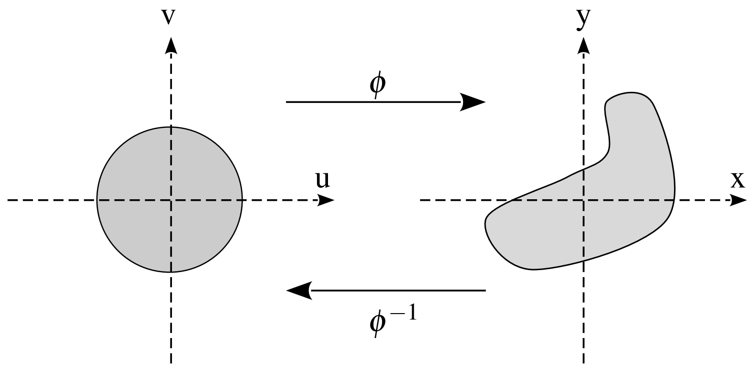

Figure 3.14: Example diffeomorphic transformation which maps a smooth manifold into another. The function is invertible and fully differentiable at all points (no sharp corners).

In Deformetrica, a space deformation is described as a time-varying diffeomorphism Φ_t. Practically, if $x \in \mathbb{R}^d$ represents a point in ambient space, the deformed point is expressed as $\Phi_{t=1}(x)$. A time-varying diffeomorphism Φ_t can be given as the integral of an 'energy-conserving' set of differential equations (Hamiltonian equations), with an initial condition that $\dot{\Phi}_{t=0} = v_{t=0}$. In this case, the possible values of $v_{t=0}$ are set to belong to a linear space built using a convolution kernel K, and a paired set of control points $(q_i)_{i=1,\dots,n}$ and vectors $(\mu_i)_{i=1,\dots,n}$ (known as momenta) distributed in ambient space. Using the kernel K and the paired sets $(q_i)_{i=1,\dots,n}$ and $(\mu_i)_{i=1,\dots,n}$, a 'velocity' vector field $v(x)$ is defined through:

$$v(x) = \sum_{i=1}^{n} K(x, q_i) \cdot \mu_i \tag{3.19}$$

Where K is typically a Gaussian kernel $K(x,y) = \exp((-\|x-y\|)^2/\sigma^2)$ with width $\sigma > 0$ [123]. Therefore, with this approach, diffeomorphisms can be fully characterised using an initial set of control points and momenta. This transformation can therefore be denoted as $\Phi_{q,\mu}$. Consequently, each model will be formulated as an optimisation problem on the variables $(q_i)_{i=1,\dots,n}$ and $(\mu_i)_{i=1,\dots,n}$. For the optimisation process, non-parametric 'current' or 'manifold' distances can be used to evaluate the error between the deformation and the target (do not require point correspondence).

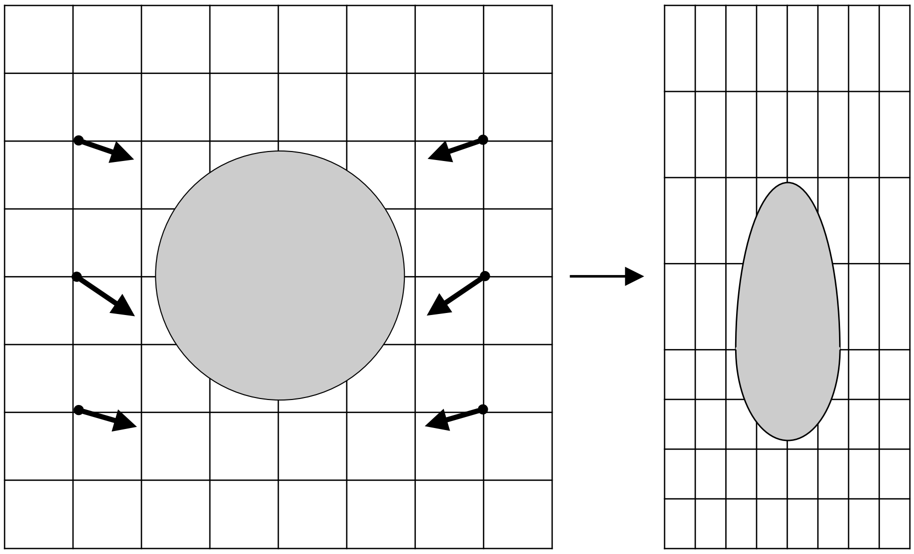

Figure 3.15: Simplified diagram of a control point based diffeomorphism of ambient space. The transformation that maps the two shapes to one another can be fully parameterised by a control point grid q (9x9) and corresponding momentum vectors μ (only six out of 81 momenta vectors are drawn).

Registration is a common application of Deformetrica and shape modelling in general. For a given target S_j, an appropriate diffeomorphism Φ_j can be estimated by minimising a given cost function. When approaching a collection of shapes S_i, an atlas approach can be used to compute a set of diffeomorphisms S_j and a mean diffeomorphism T, or mean template. In this case, the cost function seeks to find optimal values of T, q and u_i. In an atlas model, all diffeomorphisms computed are respective to the mean T.

3.4.2 Principal component analysis

Principal component analysis (PCA) is a statistical technique that is used to identify patterns in a dataset. It is often used to reduce the dimensionality of large data sets, by projecting the data into a smaller set of uncorrelated axes called principal components. The first principal component accounts for as much of the variability in the data as possible, and each succeeding component has the highest variance possible under the constraint that it is orthogonal (uncorrelated) to the preceding components. In Fig. 3.16, it can be seen how a two-dimensional dataset can be expressed using a single new axis (C1) of maximal variance capture. The usefulness of this technique becomes apparent when data is much more highly dimensional (thousands or millions of parameters). PCA can be useful for data visualisation, feature extraction, and noise reduction. PCA may also be used to improve the interpretability of large datasets or for extracting useful features to make DNN model training easier. In statistical shape modelling, a common application of PCA is to represent deformations with lower-dimensional approximations, allowing for shapes to be clustered by their PCA component scores.

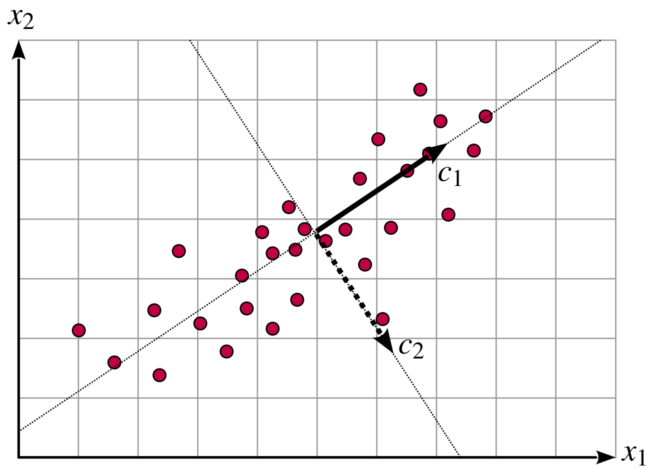

Figure 3.16: Selection of new axes to project data on based on variance. C1 is the first principal component (maximal variance), C2 is the second.

For any given square matrix $\mathbf{A}$, there exist a number of unique *eigenvectors* $\mathbf{x}$ that only undergo scaling (no rotation) when multiplied by $\mathbf{A}$ (i.e. $\mathbf{Ax} = \lambda \mathbf{x}$). The scaling factor λ is often referred to as an *eigenvalue*. Properties of eigenvectors are highly related to principal components. For example, a dataset with the mean subtracted from all its elements (standardised) can be given as a matrix $\mathbf{M}$ (m samples $\times$ n features). The $n \times n$ covariance matrix $\mathbf{C}$ can then be given as $\mathbf{C} = \mathbf{M}^{\mathrm{T}}\mathbf{M}/(n-1)$.

The *eigendecomposition* of $\mathbf{C}$ yields its eigenvectors, which also happen to be the principal components of $\mathbf{M}$. However, in practice, the numerical formulation of $\mathbf{M}^{\mathrm{T}}\mathbf{M}$ can cause instability and precision loss for some matrices, therefore an alternative matrix factorisation method called singular value decomposition (SVD) is often applied.

In SVD, the matrix $\mathbf{M}$ can be expressed as a factorisation so that $\mathbf{M} = \mathbf{U}_{m\times m}\mathbf{S}_{m\times n}\mathbf{V}^{\mathrm{T}}_{n\times n}$. By using this new definition of $\mathbf{M}$, computing the covariance matrix directly is avoided, and the eigendecomposition of $\mathbf{C}$ is instead equivalent to $\mathbf{V}\frac{\mathbf{S}^2}{(n-1)}\mathbf{V}^{\mathrm{T}}$. The principal components are given as the rows in the matrix $\mathbf{V}^{\mathbf{T}}$. The projection of the data $\mathbf{M}$ along the new coordinate axes is given in the rows $\mathbf{US}$. Numerous algorithmic approaches exist for efficient computation of SVD. Most modern implementations in linear algebra packages are developments of the original Golub-Reinsch algorithm, which uses an iterative approach to find a stable solution for the SVD of a given matrix [124].

PCA is essential in numerous applications such as statistical shape modelling. Previously, it was described how momenta could be used to describe a population as deformations relative to a mean shape. When applying PCA to the grouped momenta, the principal components of shape deformation can be found. These are also known as the *modes* of variation. The projection of each subject along the PCA modes allows each shape to be represented as a *low-dimensional* deformation from the mean. Often, these 'modes' take on a unique form (warping, bending, etc.) which gives further indication into the shape variance within the population. This can be coupled with methods such as cluster analysis to investigate how shape features correlate with other metrics. It also enables a way of parameterising unseen subjects, by expressing them in terms of PCA mode scores.

3.4.3 Recent applications of SSMs in CHD: a review

Due to the highly morphological nature of CHD abnormalities, statistical shape modelling has been widely used in conjunction with PCA for research into the relationship between CHD and shape. For example, Bruse et al. explored the possibility of using an SSM to automatically extract shape-based biomarkers for diagnosis and risk-stratification in CoA. Correlations between specific shape PCA modes and clinical metrics such as ejection fraction were found, indicating potential suitability for this technique in routine clinical settings [125]. Sophocleous et al. leveraged longitudinal SSMs to assess the change in shape during growth in diseased aortas [126]. Similarly, studies have been conducted on other pathologies such as TOF. For example, Kollar et al. identified ventricular shape modes that correlated with intracardiac biomechanics and physiological dysfunction in the RV [127].

Another alternative use for SSMs is for supporting other modelling workflows, such as automatic image segmentation. Contrary to a CNN, this approach works by identifying image landmarks and computing an SSM deformation (from the mean) to help guide the segmentation process. An early implementation of this approach was detailed by Weese et al. [107]. Alba et al. found that this approach enabled accurate automatic ventricular segmentation of severely diseased hearts even with a normal morphological heart SSM model [128]. In recent years, SSMs for automatic segmentation have been superseded by deep-learning models, which have shown much faster inference times and work directly on the pixels without the need for a prior [129]. SSMs and PCA have also been used for encoding shape features, an important application in the development of shape-driven models for replicating CFD [118]. This approach relies on the fact that in the past it has been shown hemodynamic flow structures can be regressed from shape features [130]. Using this approach, shapes can be represented as low-dimensional deformations, making them suitable for use with DNNs. The ability to generate synthetic shapes using an SSM also enables the expansion of a pre-existing dataset of real patient geometries. Hoeijmakers et al. reported findings in which their computed atlas model was used to randomly sample over 2,000 synthetic aortic valve geometries, which were then simulated with CFD and used to build a DNN for inferring pressure drop [119]. The ability to generate realistic patient geometries based on a pre-existing population also has ramifications in the growing field of in-silico modelling, e.g. the optimisation of left atrial appendage closure devices [131].

3.5 Virtual reality

Extended reality (XR) is an umbrella term which encompasses all forms of computer simulated experiences that rely on visual, auditory and sensory stimulation. Virtual realty (VR) is a subset of XR, and unique in that it creates an enclosed synthetic environment with no outer stimulus from the real world, hence facilitating complete immersion. With modern technology, this is typically accomplished with the use of a head-mounted display (HMD), tracked hand controllers and audio headphones. Recent and rapid developments in the design of commercial devices have opened up the possibility to create applications targeted for the medical field. In this section, an overview of modern VR will be given, including hardware/software aspects and some general applications of VR for treating and teaching CHD.

3.5.1 Depth perception

Depth perception is one of the fundamental factors for allowing humans to perceive the world in three dimensions. In order to perceive depth, there are two main types of visual signals that the brain processes. The first are monocular cues, which can be perceived with a singular eye. The second type of stimuli are binocular cues, which require two eyes. Some important monocular cues for depth perception include:

- motion parallax: the relative motion of an object with respect to its background gives some indication of its distance to the observer
- perspective: when parallel lines converge towards a horizon, this makes it easier for an observer to estimate the relative distance between objects
- curvilinear perspective: parallel lines in a visual field tend to become curved as they near the outer extremities. This effect produces a slight natural distortion in the visuals which enhances the feeling of being present in a 3D scene
- lighting and shading: the way that light falls on objects can provide a visual cue to help the brain determine the size, shape and distance of that object in the scene

- texture gradient: nearby objects have higher levels of visible detail, whereas further objects appear denser and less detailed

These monocular cues (and others not mentioned) play an important role in three-dimensional vision. However, in species with two-front facing eyes, binocular cues enable far superior depth perception capabilities. Stereopsis is the process whereby two images of the same scene from slightly different perspectives are combined to perceive distance and visualise depth. The separation of the eyes is known as the interpupillary distance (IPD), and is crucial in this process. The difference between the left and right eye images (binocular disparity) is used to estimate distance [4]. The binocular disparity, δ, between images is approximated by computing the relative phase shift between objects in the scene, as shown in Fig. 3.17 and Eq. 3.20. The sensation of depth Z has been suggested to be approximately related to binocular disparity through a psychophysical formulation (Equation 3.20), where D is the viewing distance to the object. Note, if D changes (e.g. observer moves away from object), maintaining the same depth perception will require δ to change. Since δ is inversely proportional to D^2, at large distances this the brain would need to be able to delineate extremely small binocular disparities, due to the similarity in images. When this distinction cannot be detected in the brain, the stereoscopic depth of the object is not perceived.

$$\delta = \phi_L - \phi_R \qquad \delta \approx \frac{IPD \cdot Z}{D^2} \tag{3.20}$$

Stereopsis occurs specifically in the visual cortex following a number of complex neurological processes which are not fully understood. The difficulty in understanding this mechanism is also complicated by the lack of a clear distinction between the actual 3D scene (geometric model), and the perceived 3D scene in the observer's mind (perceptual model) [132]. Apart from stereopsis, another binocular cue that the brain processes is *convergence*, which is the result of both eyes having to fixate on an object in order to maintain stereoscopic vision. Since the eyeballs converge as the object appears closer, nearer objects require more extraocular muscle exertion to focus. The brain interprets this muscular exertion as an additional

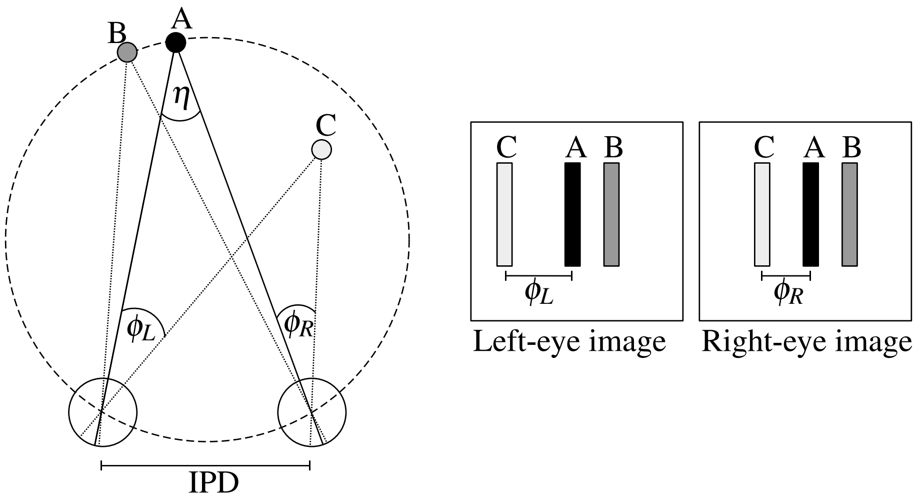

Figure 3.17: The geometry of binocular vision [4]. Both eyes fixate on a pillar (A) in the scene, with two other pillars (B and C) also visualised. Vergence η is the angle of fixation. The angle ϕ between A and a relative object (e.g. C) varies in each eye. Binocular disparity is computed as the difference between these two angles.

indicator for perceiving distance.

3.5.2 Head-mounted displays

The earliest binocular devices for simulating 3D environments began with the development of the stereoscope, designed by Sir Charles Wheatstone in 1838. The principle of a stereoscope relies on two images simulating left and right eye views of an object. If both eyes view the images intended for them at once, the brain is 'tricked' into perceiving a 3D projection. This principle is the foundation of VR and is relied upon even today. Up until recent years, HMD devices struggled to enter the wider commercial market due to multiple limitations. These included issues related to low image contrast, narrow field of view (FOV), poor comfort and low latency, which often caused motion sickness and discomfort during head movement. In 2010, the first prototype of the Oculus Rift HMD (Meta Platforms Inc.) was revealed, which featured a previously unseen 90 degree FOV and eliminated lens-based distortion artefacts that other products at the time suffered from. This device is often accredited for producing a revival in the interest of VR from commercial sectors.

The Oculus Rift and similar devices around the time typically relied on a ca-

Figure 3.18: Quest 2 head-mounted VR headset, first released in 2020 by Meta Platforms Inc. (formerly Facebook).

bled connection to a desktop with a high-end graphics processing unit (GPU) in order to offload the low-latency rendering tasks typically required in VR. Additionally, external ‘base stations’ were required for HMD and controller tracking tasks. Over time, advancements have enabled the manufacture of ‘standalone’ devices which contain onboard computers to perform computing tasks without the aid of cabled components [5]. Additionally, many modern solutions now incorporate inbuilt tracking systems instead of the more conventional external base stations. One of the most popular current standalone headsets is the Meta Quest 2 (released 2020, Meta Platforms Inc.), which was the primary HMD of choice for the majority of activities described within this dissertation (Fig. 3.18).

3.5.2.1 Head-mounted display input

Hardware components in VR devices typically fulfil certain roles that place them in one of two categories: input devices and output devices. In modern HMDs, user input is primarily conducted via handheld controllers. Controllers are typically tracked with 6 degrees of freedom (DOF), recording both translation and rotation. The combined positional and rotational information of an object in Euclidean space is also referred to as a ‘pose’. Controller tracking is typically achieved with the use of inertial sensors such as accelerometers, gyroscopes and magnetometers to measure orientation. Due to their fast refresh rate, inertial sensors are a popular

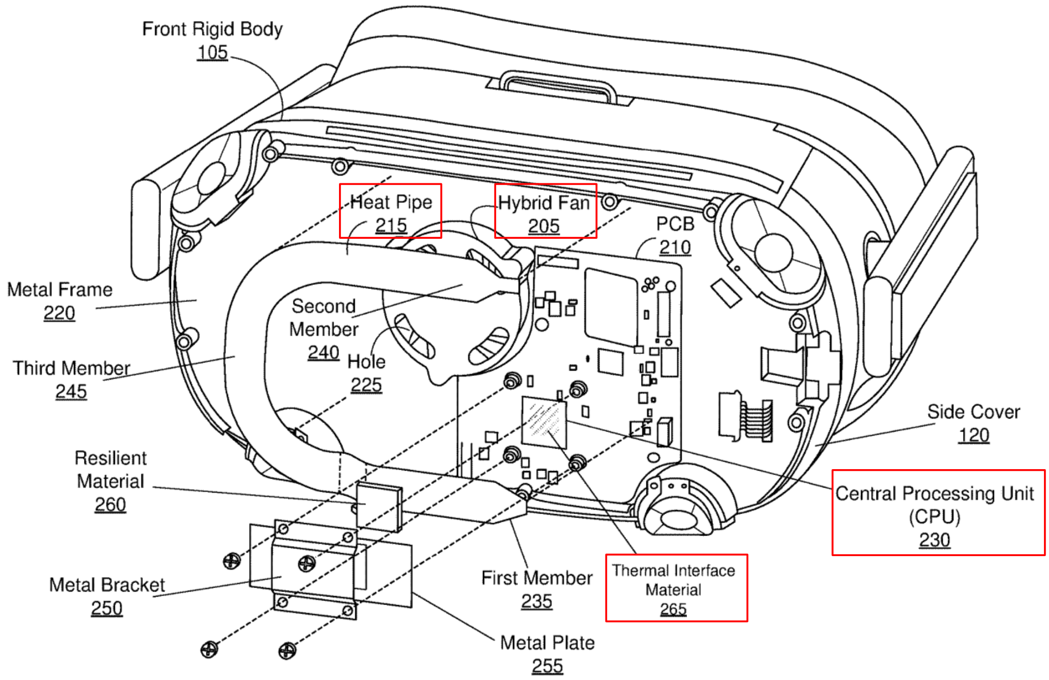

Figure 3.19: Exploded view of Quest 2 headset revealing the onboard Qualcomm Snapdragon XR2 central processing unit and heat removal system lying under the front cover [5]. The Quest 2 runs on an Android-based operating system.

choice in modern devices such as the Quest 2. However, positional errors typically accumulate over time (drift), requiring an additional tracking system to compensate.

Optical tracking (or camera-based tracking) has become increasingly popular because of its high accuracy and flexibility [133]. In the Quest 2, this is implemented using an “inside-out” approach, which utilises four head-mounted cameras with wide-angle lenses to detect the pose of targets situated in space. In this context, each target is a controller fitted with a ring (Fig. 3.20) that contains infrared light-emitting diodes (LEDs) arranged in a specific pattern. Flashes from the IR LEDs are synchronised with the cameras, which transform the measured radiation into greyscale images. Pattern recognition algorithms detect the 2D positions of the LED markers in each image. These are then passed onto a central tracking controller, which uses triangulation to compute the 3D pose (with an accuracy of < 1 mm and $0.1\ ^{\circ}$) [133]. The latency between the tracked and rendered controller may also be reduced with kinematic prediction algorithms, which use motion to estimate the future pose of the handsets for improving smoothness [134].

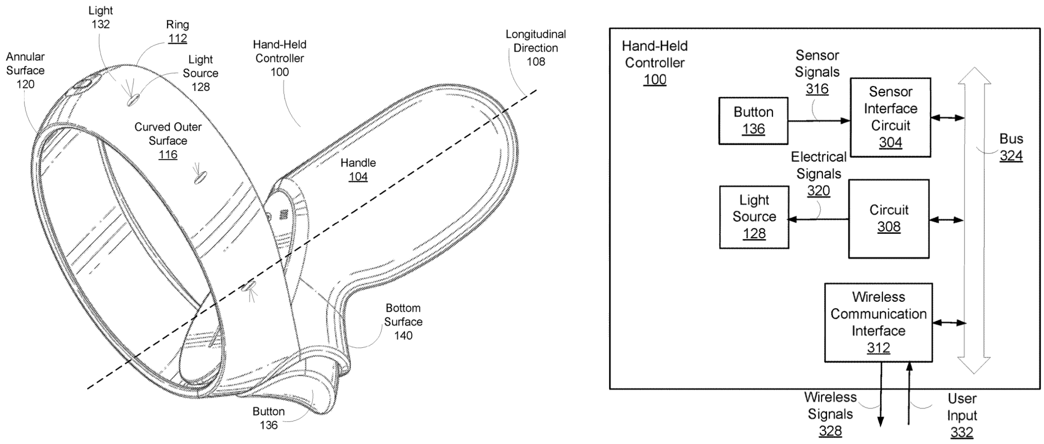

Figure 3.20: Schematic block diagram and technical drawing of Quest 2 controller. Buttons and LEDs can communicate via a bus. Wireless communication allows signals to be received to and from the HMD [6].

In the Quest 2, the cameras run at 60 Hz and perform both controller and headset pose tracking on alternating frames (effectively running at 30 Hz each). Low exposure images are taken to highlight the positions of the LEDs for optical tracking. On alternate frames, high exposure images are taken of the surroundings for the HMD pose estimation. These images are needed for constructing a 3D map of the surroundings while simultaneously estimating the observer's respective location; a well-established computational problem known as simultaneous localisation and mapping (SLAM). For the solution, a continual update of point cloud landmarks (such as corners of tables) is processed by extracting features from the image data. Additionally, a constant stream of inertial sensor measurements from the HMD is recorded. Using a history of landmark observations and headset movements, a recursive method can be used to estimate the most probable pose and 3D environment map. Thus, given a set of control input vectors $U_{0:t}$ and sensor observations $Z_{0:t}$ up to and including time t, the user's current state (pose) x_t and environmental map of time-invariant landmarks m can be found by solving the following probabilistic distribution (joint posterior) [135]:

$$P(x_t m | Z_{0:t}, U_{0:t}, x_0) \tag{3.21}$$

This problem can be expressed as a two-step sequential process involving a prediction (time-update) and a correction (measurement-update). For the solution, expectation-maximisation algorithms are used to iteratively search for a local maximum likelihood of the joint posterior. Numerous algorithmic approaches (Kinect, Intel RealSense, Google Tango) have been applied to solve this problem [133]. Importantly, this method assumes that the location of the user is deterministic at all times and treats computed landmark locations as 3D points known with absolute certainty.

3.5.2.2 Head-mounted display output

In HMDs, optical and display components are of critical importance for delivering an immersive visual experience to the user. High resolution, colour reproduction and contrast ratios in displays are important for creating a sense of immersion in VR. High frame rates (refresh rates) in screens are desirable for preventing motion sickness, with the Quest 2 displays running natively at 90Hz.

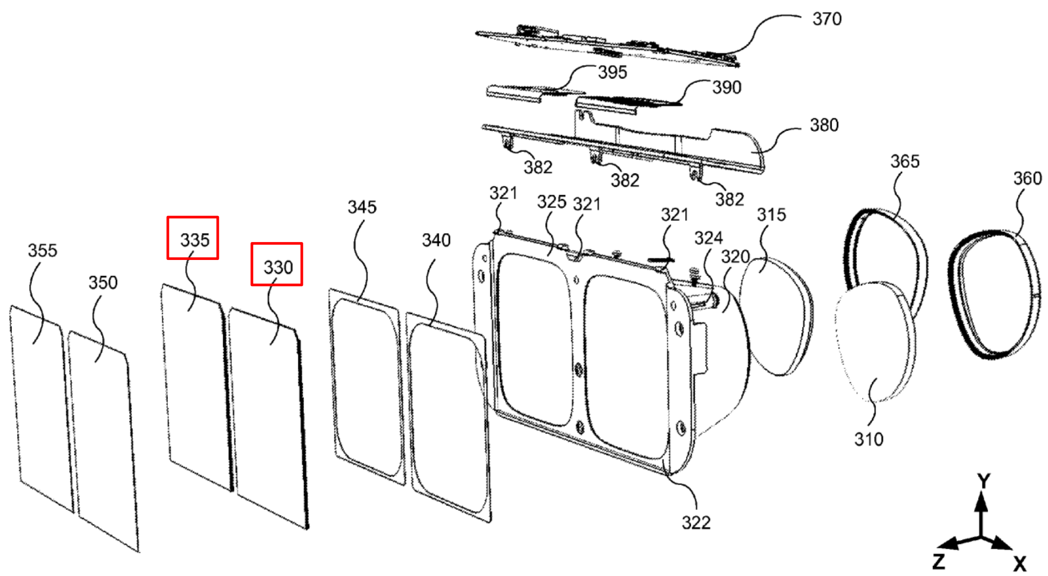

Figure 3.21: Exploded diagram of a Meta Quest 1 display system consisting of two organic LED screens (330 & 335). The lenses (310 & 315) can be seen, along with other parts of the housing such as frames, mounts, brackets and more [7].

In the human eye, objects at varying distances can be brought into focus (made sharper) in a process known as *accommodation*. This mechanism is managed by ciliary muscles in the eye, which contract and alter the shape of the lens, thus making incoming light from the object refract and converge on the retina. At optical infinity (object distance $> \sim 6$ m), light rays entering the eye can be assumed as parallel, and the cornea and lens (in a relaxed state) refract light to produce focus. As objects approach the observer, accommodation is required and the lens changes shape to maintain image clarity. At very close distances, the lens cannot accommodate any further and the focal point does not coincide with the retina, causing a blurred image. This becomes a problem in stereoscopic VR HMDs, where the displays need to be positioned close to the eyes in order to produce a sufficiently large FOV. For this reason, lenses are commonly built into HMDs in order to restore image focus (Fig. 3.22).

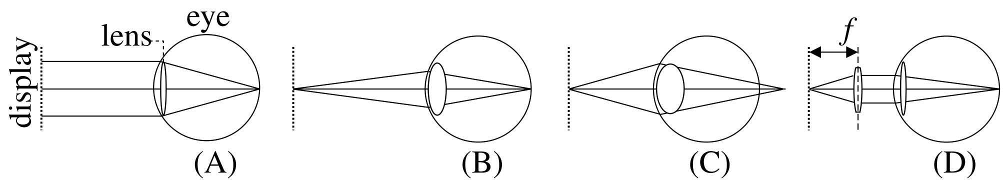

Figure 3.22: Principles of accommodation in the eye. A) Display is at a distance of optical infinity from the eye (no accommodation), (B) display is nearby requiring accommodation and is in focus, (C) display is very near and out of focus even at maximum accommodation, (D) lens is positioned at distance f from the display, thus refracting rays as parallel from display and enabling focusing.

The measure of a lens' optical power is the reciprocal of its focal length f, and is an important consideration in lens design. Focal length can be computed using inner/outer lens curvature R_1 and R_2, lens width d, and the refractive index of the lens material n with the Lensmaker's equation:

$$\frac{1}{f} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2} + \frac{(n-1)d}{nR_1R_2}\right) \tag{3.22}$$

A lens design which refracts light from a display to produce parallel rays would theoretically produce an accommodation distance of optical infinity (Fig. 3.22D). However, such a configuration would practically have issues due to the discrepancy between the accommodation distance and the vergence angle (Fig. 3.17). Essentially, fixation on an object that is perceived to be nearby would not correlate with the accommodation distance (infinity), resulting in 'visual fatigue/strain' over time. This effect is known as the vergence-accommodation conflict (VAC). In the Quest

2, lenses have a focal distance of $\sim$1.3 m, as a compromise to reduce eye strain for objects held at an arm's length and objects viewed in the distance. Current research aims to solve this with the use of eye-tracking technology and variable focal length lenses [136].

In HMDs, minimisation of the size and weight while maintaining satisfactory lens magnification is a significant design challenge. For example, reduction of the lens-display distance can be achieved if the lens focal length f is decreased (increased optical power) with thicker and more curved lenses (Eq. 3.22). Fresnel lenses approximate a curved surface with stepwise elements and can be used instead of regular biconvex lenses to counteract the increased thickness/curvature, such as in the Quest 2. Other methods include using a material with a higher refractive index, although this can pose manufacturing difficulties. More recent pancake optics use polarisation techniques to bounce light between optical elements, thus reducing the effective focal length (Meta Quest Pro).

The FOV is another important consideration of lens and HMD design, with higher FOVs contributing to improved virtual immersion. The FOV of a HMD increases with larger displays S and shorter focal lengths f, as shown in Equation 3.23. However, maximising these design parameters has the undesirable effect of making the optics larger, and introduces issues such as lens distortion. In order to counteract this, software post-processing (barrel distortion) is applied to VR images so that the projection appears more rectilinear after passing through the lens.

$$FOV = 2arctan\left(\frac{S}{2f}\right) \tag{3.23}$$

This technique becomes more difficult to calibrate as the focal length increases, making users more likely to experience discomfort due to distortive effects. In addition to this, chromatic aberration may occur in higher focal length optics, which is the effect of the lens refracting colours differently according to their wavelength. This can result in blurring and rainbow-coloured artefacts appearing at area of high contrast and at the fringes of the image. Software-based correction is often required to minimise this effect. Developments in eye-tracking technology, pancake lenses

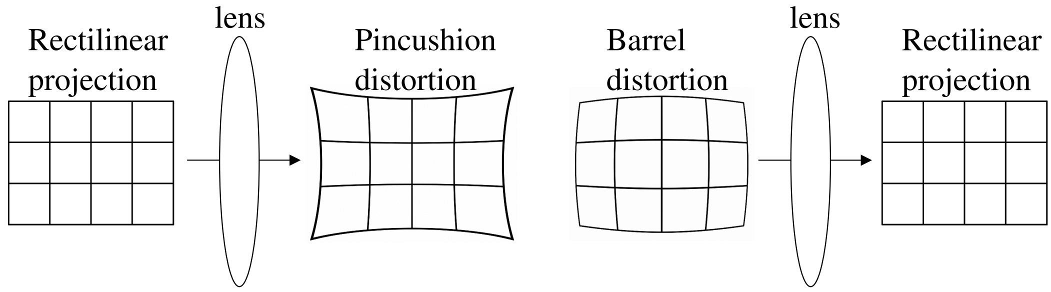

Figure 3.23: Pin-cushion distortion resulting from a lens (left). Application of barrel pre-distortion to the images to undo the warping effect (right).

and software-based artefact correction show great promise for furthering the experience of VR visual output [133].

3.5.3 Virtual environments

Virtual environments can be described as arrangements of 3D objects with programmable behaviours. In VR, 2D projections of the 3D scene are displayed in the headset through graphical rendering. Input from the user is also processed in real-time to enable interactions and allow the user to affect the surrounding world. Overall, these two processes are essential and must run concurrently to maintain an immersive and reactive virtual environment.

3.5.3.1 Computer graphics

The visualisation of a 3D virtual environment relies on foundational principles from computer graphics. Cartesian surface meshes are widely used for rendering 3D topological surfaces. In game engines such as Unity (Unity Technologies), surface meshes contain vertex positions, triangle connectivity, face normals and more. Triangles are defined by indexing three points from the array of vertices. The direction of the triangular normal vector of each triangle is defined by the order the vertices are indexed (winding order). This information can be used to determine whether or not a polygon should be rendered to the display (back-face culling). For a set winding order convention, the computation of the dot product between the camera-triangle vector O and the normal N can describe which way the triangle is facing, and whether it needs rendering. Specific instructions for how the GPU should render an object (e.g. back-face culling) are given through a type of scripting known as *shaders*. These are an essential part of the rendering pipeline in game engines.

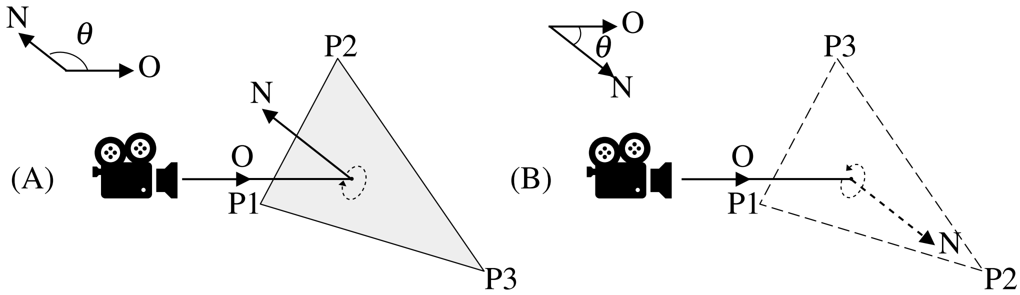

Figure 3.24: Back-face culling using clockwise winding order. Front-facing triangles are rendered (A), while back-facing triangles are discarded (B). N = triangle normal, O = observer-triangle vector, P1 = first vertex, P2 = second vertex, P3 = third vertex, θ = O N angle.

Algorithm 1 Backface culling

% Compute dot product between O and N
ON = dot(O, N)
% if ON is less than or equal to 0, render, else cull
if ON $<=$ 0 **then**
    render triangle
else
    cull triangle
end if

Shaders can be defined as computer programs which describe how a scene should be coloured, shaded, lit and ultimately, rendered. The two primary types of shaders are vertex and fragment (pixel) shaders. Vertex shaders are used in the first step and transform a 3D scene of objects into a 2D camera projection composed of pixels in a process known as rasterisation. At this stage, various properties can be assigned to vertices or modified, such as colour or texture values. Mathematical functions can also be implemented in order to create specific effects, such as custom illumination models. For example, a popular shader technique is *Phong shading*, which uses normal interpolation to approximate smooth shading on a faceted surface (Fig. 3.25). Once the 2D projection is finalised, the vertex shader passes on the rasterised image onto the fragment shader. The fragment shader then computes the colour of each pixel for the display, with considerations such as object transparency, depth and lighting also taken into account.

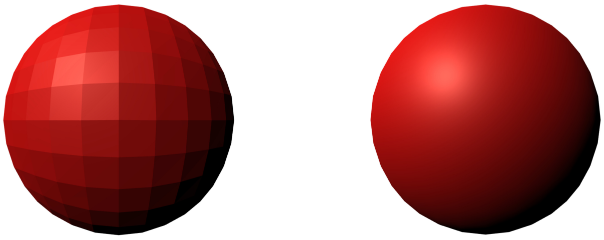

Figure 3.25: Standard flat polygonal shading (left) compared to Phong shading (right).

3.5.3.2 Interaction

Human-computer interaction is an essential component of VR. The ability to build in environmental responses to user input has allowed for the possibility of creating applications which are not purely perceptual. Typically, interaction in virtual environments is conducted through the combination of a user interface (UI) and a collision based event system. User interfaces are especially important for creating menus (buttons, toggles) that control various functions. Selection of UI elements in 3D is performed using *raycasting*. This is typically managed by firing a ray from the position of the user's handheld controller towards a flat panel in 3D space. In order to determine if a UI element was triggered, geometric checks between rays and polygonal entities are used. For example, a ray-plane intersection is widely applied in such situations. A plane P can be described by the dot product of a position on the plane X and the planar normal N. A ray (or segment) is given by the parametric equation $S(t) = A + t(B - A)$, where A and B are the start and end points respectively. Expressions for computing the parameter t and the intersection point Q can thus be derived [137]:

$$t = \frac{X \cdot N - N \cdot A}{N \cdot (B - A)} \qquad Q = A + t(B - A) \tag{3.24}$$

In addition to UI interaction, 3D objects in space may interact with one another in VR. This is typically managed using collision detection, an approach in which each object has an associated 'collider' that is triggered when contact is made with another collider. Colliders are typically primitive shapes such as cubes or spheres that are not rendered by the GPU. For example, if two objects containing sphere colliders make contact, a sphere-sphere algorithm can be used in order to flag this

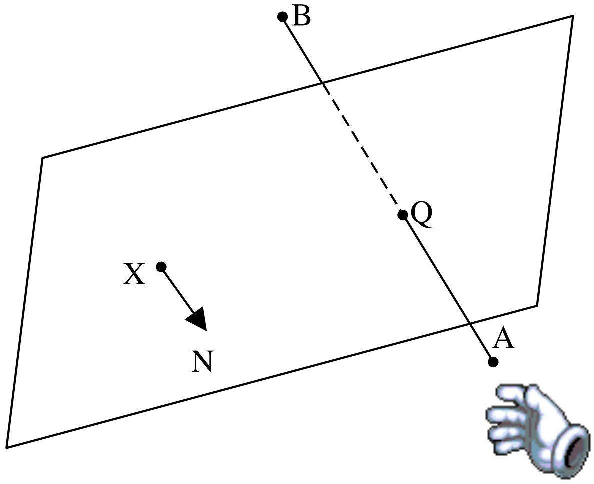

Figure 3.26: Intersection of a ray (or segment) with a plane during raycasting of a UI element in VR. The hand (A) represents the ray source (controller position in 3D space).

Algorithm 2 Ray-plane intersection

% Compute t value in parametric equation for segment
AB = B - A
t = (dot(N, X) - dot(N, A)) / dot(N, AB)
% Intersection exists if t between 0 and 1, compute and return value
if t >= 0 and t <= 1 **then**
  Q = A + t * AB
else
  Q = null
end if

interaction [137]. Detecting collision is important for creating interactions such as 'grabbing' or 'handling' of 3D objects with the handheld controllers (Fig. 3.27).

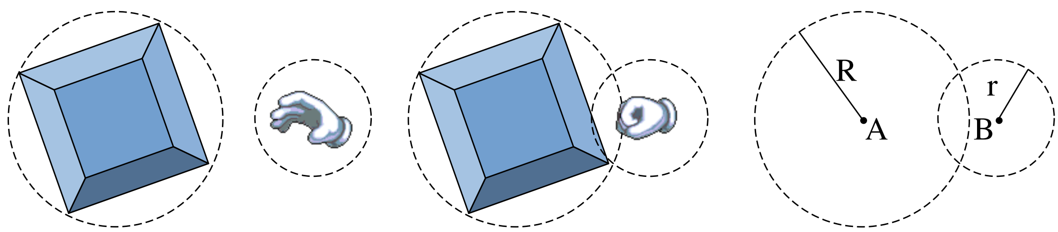

Figure 3.27: Detecting the intersection of two spheres. This can be used to allow for model handling and spatial manipulation.

Similarly to ray-plane intersection, a mathematical formulation exists for determining the occurrence of a sphere-sphere collision. This overlap test involves a simple comparison of the Euclidean distance between the sphere centres A and B against the sum of the sphere radii R and r.

Algorithm 3 Sphere-sphere intersection

```
% Compute squared distance between centres
D = A - B
D_squared = dot(D, D)
radii_sum = r + R
% Spheres intersect if squared distance is less than sum of radii
if D_squared <= radii_sum then
    Spheres intersect
else
    Spheres do not intersect
end if
```

3.5.4 Recent applications of VR in CHD: a review

Recent improvements to HMD hardware have enabled breakthroughs in medical applications. Bulky headsets with poor latency and visual clarity have been replaced by commercial HMD solutions which mitigate those issues. This has lowered the barrier for VR to be used as a means of visualising patient-specific 3D anatomical models. Understanding the spatial relationships between anatomical structures is of particular importance in the realm of complex CHD.

Prior to VR, 3D printed models have been successfully used in the past to aid 3D visualisation of cardiovascular anatomies [60, 138]. Numerous applications of 3D printing have been shown for interventional planning [139–143] and anatomical evaluation [144–146]. However, VR has begun to show several advantages over 3D printing, such as lowered costs and no waiting times. Additionally, VR has opened up the opportunity for integrating conventional 3D-assisted tools in an interactive environment, such as measuring, dynamically clipping meshes and more [147]. The use of shared online virtual rooms for interacting with 3D models also offers the possibility to perform joint 3D assessment of patient cases remotely. Other advantages of VR include the ability to digitally preserve 3D models, easily distribute resources and generate simulated training environments.

VR solutions for planning cardiac surgery are becoming increasingly common. Numerous applications for planning cardiothoracic interventions have been demonstrated, including coronary bypass grafting, and mitral valve annuloplasty ring sizing amongst others [148, 149]. Interest in VR for CHD is growing due to the pronounced morphological challenges frequently faced. Reported findings related to the use of VR for visualisation and analysis of CHD are promising, with or without preliminary image segmentation [150–153]. Lesions requiring intracardiac repairs (e.g. VSD, DORV, etc.) have been frequently reported defects in literature, related to surgical planning with VR. For example, Mendez et al. describes the use of VR for planning repair of a large inlet-type VSD [150]. Similarly, Ghosh et al. described a patient case of multiple muscular VSDs where VR was used pre-operatively [154]. Using VR, surgeons were able to (i) locate the VSDs in 3D, (ii) determine the RV incision site, (iii) identify the muscle bundle requiring division to access the anterior defect, and (iv) ultimately decide on a hybrid transcatheter and surgical approach. Ayerbe et al. presented a complex DORV case-study in which the patient was assessed with VR in order to decide the surgical approach to repair an obstructed RV-PA conduit [155]. Navigation of the 3D model confirmed that the obstruction was related to the subaortic conal prominence. Cen et al. reported a further five surgical cases of patients with diagnosed pulmonary atresia, VSDs and major aortopulmonary collateral arteries (MAPCAs), all planned pre-operatively using mixed reality headsets [156].

In Section 2.2.2, it was mentioned how pathological specimens are regarded as the ‘gold standard’ resource for learning cardiac morphology [157]. Unfortunately, the number of available specimens is low compared to needs, the samples degrade over time and thus remain an inaccessible resource to many. VR offers the opportunity to digitally ”preserve” specimens, view models from non-conventional views and interact with other users in online virtual 3D environments. The integration of VR in courses for teaching CHD is, however, still in its infancy [158–160]. Increasing numbers of single and multi-user VR platforms for anatomical exploration of cardiac models are being reported [147]. This is in part due to a decreased barrier to game engines such as Unity, which have enabled researchers to more easily develop tailored 3D solutions for VR. Axelrod et al. demonstrated the early-stage feasibility of a VR platform for CHD (Stanford virtual heart project) [158]. Lau et al. described a Unity3D-based project in which four patient-specific models of the heart were explored by medical professionals (n=35), with 72% of participants reporting positive feedback [161]. Kim et al. detailed the use of a multi-user Unity-based application for enabling collaborative discussion of patient cases in VR [160]. Most of the participants (68%) reported that VR would be their preferred method of choice for examining 3D models of CHD and for determining a diagnosis.

In recent years, rapid advancements and improved accessibility to VR have led to a rising need for research-backed evidence that supports its implementation in clinical and teaching environments. When used for pre-operative planning, the difficulty in finding tangible clinical metrics for evaluating VR has made qualitative methods highly popular [160, 161]. For example, Pushparajah et al. reported the following findings: surgeons (n=5) were presented with retrospective cases in VR (n=15) requiring repair of the atrioventricular valves. Following this, 67% reported 'more' or 'much more' confidence in anatomical understanding following VR assessment [162]. Despite the challenge in finding clinical parameters to correlate against VR usage, some studies have made efforts in the past. For example, Ye et al. recently conducted a prospective study assessing the use of mixed-reality devices for planning the repair of complex DORV, with patients split into a group assessed with VR (n=17) and a control group (n=17) [163]. Despite all patients undergoing the same repair, the time for surgery was 21.4% lower in the intervention group ($p < 0.05$). In order to build greater support for integration of VR in clinical settings, more studies with similar approaches are needed. In the literature, the assessment of VR for CHD education is also primarily questionnaire-based, although recent studies have shown greater efforts to assess knowledge uptake with quantitative metrics. Notably, Lim et al. assessed the usefulness of the Stanford virtual heart project [158] for improving the knowledge acquisition of medical trainees [164]. Participants were divided into an intervention group (n=80) and a control group (n=51), with knowledge being tested using a series of assessments relating to anatomical and visuospatial concepts of CHD. Findings showed higher scores for the intervention group, with 100% of participants recommending integration of the tool into the residency programme. However, a contrasting study by Patel et al. found that VR provided no improvement in CHD assessment scores for students when compared to 3D models on a 2D screen [165]. The sparsity of studies in the field means that it is still not clear as to what extent VR is beneficial in CHD pre-operative planning and education, especially when compared to other 3D modalities.

In research, the range of software used for VR visualisation is highly variable, with many authors presenting competing applications. While commercial applications for viewing 3D models exist, they are typically focused on the teleconferencing market and do not contain the sufficient tools or ease-of-use for integration in CHD clinical/teaching situations. In order to have access to the full capabilities of VR, the development of a comprehensive platform with a sufficient array of functionalities is necessary. Features such as online virtual rooms, visualisation tools and a curated selection of 3D models for education are important in order to be able to study various applications of VR. Furthermore, integration into clinical settings necessitates the platform to have intuitive interactions and to run independently, without the support of an engineer. Compatibility with current and future commercial headsets is necessary.

3.6 Summary and Conclusion

In this section, four main computational modelling techniques of importance for this thesis were described, along with their application in CHD and remaining challenges: (i) CFD, (ii) ML, (iii) statistical shape modelling and (iv) VR. A brief overview of CFD theory and methods were presented. Applications of CFD in CHD showed high suitability for routine clinical practice, although CFD is hampered by long simulation times, high computational resources and the need for an engineer to set up simulations. Automating and accelerating the CFD pipeline were identified as critical tasks for facilitating the translation of CFD to CHD clinical settings.

Following this, a description of ML and DNNs was given. Clinical applications were shown to be vast and rapidly growing, with possible uses for automatic segmentation and creation of data-driven CFD models. Statistical shape modelling and PCA were then described, including some applications which involved representing shapes as low-dimensional deformation modes for clustering and analysis. These methodologies will be adopted in the next two Chapters to automate CMR image segmentation (Chapter 4) and immediately predict pressure and velocity information (Chapter 5), in order to develop a CFD pipeline for full clinical implementation and improved management of CHD. Finally, the fundamental and technical principles behind modern VR were introduced in this Chapter. This included some background to human depth perception, HMD design and the creation of virtual environments. Applications of VR in the CHD field were shown, displaying great promise for pre-operative and educational settings. In Chapter 6 of this thesis, these VR principles will be leveraged to develop an in-house VR application, with a set of tools designed to respond to CHD specific clinical and educational demands and unmet needs. Chapters 7 and 8 will be centred around clinical and educational applications respectively, to demonstrate the capabilities of the technology when implemented into those settings.

Chapter 4

Automatic segmentation of the great arteries for computational hemodynamic assessment

4.1 Introduction

Interest in the use of CFD for the assessment of cardiovascular disease has been increasing over the past two decades. CFD models have previously been shown to be suitable for predicting hemodynamic response to interventions, thereby aiding therapeutic planning (Section 3.2.4). Patient-specific reconstructions are often derived from 3D CMR, particularly cardiac and respiratory gated whole-heart sequences. The use of ML for accurately segmenting ventricles and great vessels from CMR images was discussed in Section 3.3.4. Implementation of automatic segmentation for clinical CFD would significantly reduce the time requirements for simulation. However, the suitability of ML segmentation for CFD has not previously been investigated.

4.2 Aims and objectives

The aims of this study were to: (i) develop an ML method for simultaneous segmentation of the aorta and PAs from whole heart CMR images in patients with paediatric or adult CHD, (ii) compare conventional and ML segmentations using traditional image-based scores, (iii) compare CFD metrics derived from both conventional and ML segmentations, and (iv) investigate the association between image-based scores and CFD errors. In this Chapter, I was responsible for the mesh pre-processing, CFD simulation and flow analysis, while Dr. Javier Montalt-Tordera contributed with the development and image-based evaluation of the ML segmentation model.

4.3 Methods

4.3.1 Subjects

90 cardiac triggered, respiratory navigated, 3D whole heart, balanced, steady state, free precession (WH-bSSFP) data from previously scanned children and adults with paediatric CHD (excluding patients with single ventricles) were used for this study. All patients were scanned on a 1.5T scanner (Avanto, Siemens Healthineers AG, Erlangen, Germany) using a standard whole-heart balanced steady state free precession sequence [166]. The imaging protocol was as follows: orientation: sagittal, matrix size: 256 x 144 x 96 (head-foot, anterior-posterior, left-right), acquired voxel size: 1.6 mm (isotropic), flip angle: 90°. Image acquisition was accelerated using GRAPPA (factor of 2 along phase encoding dimension) and partial Fourier (factor of 6/8 along both phase and slice encoding dimensions).

Additionally, 10 external examples were retrospectively collected from a different centre. These were scanned on a 1.5 T scanner (Ingenia, Philips Healthcare, Amsterdam, Netherlands) with the following imaging protocol: orientation: axial, matrix size: 240 x 240 x 110 (left-right, anterior-posterior, head-foot), acquired voxel size: 1.44 mm (isotropic), flip angle: 90°. Image acquisition was accelerated using SENSE (reduction factor of 2) and partial Fourier (6/8).

4.3.2 Ground truth segmentation

Reference standard conventional segmentation (Section 2.3) of the aorta and PAs was performed using a semi-automatic technique with manual correction (Plug-ins built in Horos v4.0, Horosproject.org sponsored by Nimble Co LLC d/b/a Purview, Maryland, USA). Initial segmentation was done using the fast level-set method [106]. This requires the user to: (i) set a threshold, (ii) place seeds in the vessel of interest, and (iii) add blocking regions to prevent segmentation of unwanted structures. The quality of this initial segmentation is dependent on both the underlying anatomy and the image quality, but manual correction is always required to remove unwanted structures and clip vessels. The proximal limit of both the aortic and PA segmentations was the semi-lunar valve. The distal limit of the segmentations were the diaphragmatic level of the aorta, and hilar branches of the PAs. Head and neck arteries were manually removed at their origin.

All 90 datasets were segmented by a primary observer (10 years' experience in CMR post-processing). The primary observer's segmentations are referred to as the ground truth (GT). In addition, a secondary observer (19 years' experience in CMR post-processing) segmented 10 of these images (test set, see below) to investigate inter-observer variability. These are referred to as second observer (SO) data. The 10 external examples were segmented by the secondary observer only, following the same procedure.

4.3.3 Data preparation

Prior to ML training, the pixel intensities of the WH-bSSFP data were normalised (range [0, 1]). The aortic and pulmonary binary segmentation masks were concatenated in the channel dimension and combined with a third channel containing a binary "background" mask (one-hot encoding). The images and 3-channel segmentation masks were either centrally cropped or symmetrically zero-padded to a fixed matrix size of 160×96×64 (superior-inferior, anterior-posterior, left-right). Finally, the image-label pairs were randomly split into a training set (70 examples, 78%), a validation set (10 examples, 11%) and a test set (10 examples, 11%). This split was used to maximize the size of the training set, while providing sufficient data for validation and statistical analysis. The examples from the external test set were reoriented, interpolated and cropped to match the orientation, matrix size and voxel size of the in-house data.

4.3.4 Network architecture

A U-Net [167] CNN was used to simultaneously segment the aorta and PAs from WH-bSSFP data. The network architecture is shown in Section 3.3.3. Each convolutional layer was followed by a batch normalization layer and a rectified linear unit (ReLU) activation. Downscaling was performed using max-pooling layers, and upscaling was performed using transpose convolution layers. The number of convolutional filters after the first layer was set to double after each downscaling layer and halve after each upscaling step. The final convolutional layer has three filters (equalling the number of possible classes – aorta, PA and background), followed by a softmax activation. Final predicted labels were obtained by assigning each pixel to the class with the highest probability.

4.3.5 Model training and evaluation

The network implementation and training scheme was formed to allow the investigation of multiple hyperparameter values, with the full search space shown in Table 4.1. A Hyperband algorithm [168] was used to perform efficient hyperparameter optimization. This method samples the search space randomly and adaptively allocates more computational resources to the most promising hyperparameters combinations. Dice score was used by the Hyperband to assess performance and choose the final model.

Table 4.1: Hyperparameter search space.

Parameter	Type	Domain
Scales	Architecture	{2, 3, 4}
Layers per block	Architecture	{2, 3, 4}
Initial filters	Architecture	{32, 64}
Learning rate	Training	[0.0001, 0.01][1]
Batch size	Training	{2, 4}
Loss function	Training	{CCE[2], Dice, IoU, Tversky, focal Tversky}

1 2 3

The neural network and related functionality were implemented and trained

[1] $\{\cdot\}$ represents a discrete set of values; $[\cdot]$ represents a continuous interval of real values.

[2] The learning rate was sampled using a log-uniform distribution.

[3] CCE: categorical cross-entropy.

using TensorFlow [169]. In particular, the model implementation, losses and metrics are available in TensorFlow MRI [170], an open-source framework developed in-house. Weights were initialised using He's method [171] and optimised using the Adam algorithm [172]. Training, including hyperparameter optimization, took $\sim$24 hr on an Nvidia Titan RTX GPU with 24 GB of onboard RAM (Nvidia Corporation, Santa Clara, CA, USA).

The optimised ML model was evaluated on the test dataset against the ground truth segmentations (ML vs GT). The accuracy of segmentation was quantified using several image-based segmentation metrics: Dice score, Intersection over Union (IoU), Hausdorff distance (HD) and average surface distance (ASD). Each metric was computed independently for each vessel. Additionally, the same metrics were calculated for the secondary observer's segmentation against the GT (SO vs GT), and between the ML model and the secondary observer (ML vs SO).

To assess generalization ability to data from other sources, the ML model was also evaluated on the external test set. The same set of metrics were computed (Dice, IoU, HD and ASD) between the ML predictions (Ext-ML) and the ground truth segmentations (referred to as Ext-SO, since the manual segmentation was performed by the same person as the SO data). Prior to evaluation, ML masks were filtered to remove all but the largest connected component, as identified using 3D connected component labelling with 26-connectivity [173]. This post-processing step was used to eliminate small background regions which had been misclassified as vessels.

4.3.6 Surface and volume meshing pipeline

The resultant segmentation masks were converted into finite element volume meshes, using the processes shown in Fig. 4.1. The masks (GT, ML and SO) were first transformed into surface meshes, by applying the marching cubes algorithm using an implementation from vascular modelling toolkit (VMTK) [174]. Following this, re-meshing and smoothing was performed with consistent parameters. The surface meshes were clipped manually at inlets and outlets to create planar surfaces; 30 mm extensions were added to encourage the flow profile to develop past an initial flat plug condition. Finally, these were meshed with tetrahedral elements to build the final unstructured grid for CFD analysis. The grid resolution was determined through a sensitivity analysis (see Appendix A.1). To assess the effect of the manual clipping of the anatomies on the CFD, the simulations from the ML segmentations were also re-run using the same inlet and outlet locations as the GT data, defined by overlaying the GT clipped geometry to the corresponding ML geometry.

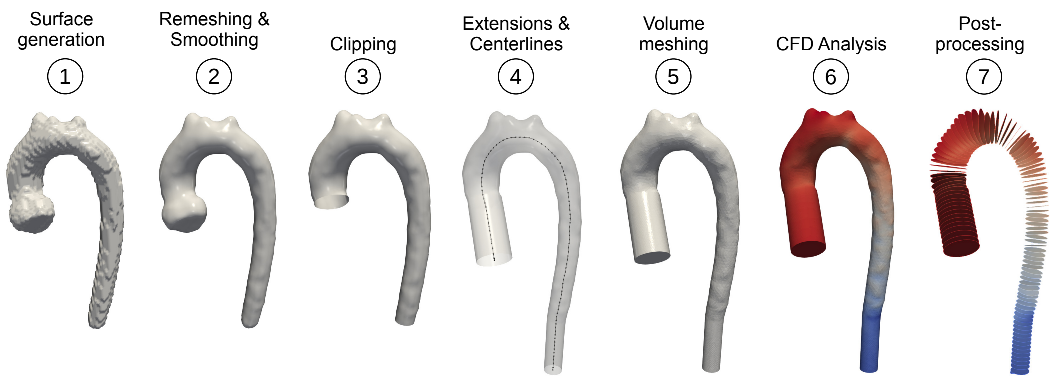

Figure 4.1: Automatic mesh processing pipeline from segmentation to CFD analysis, followed by post-processing to reshape the data in a consistent format between subjects (99 planes from inlet to outlet containing average pressure and velocity).

4.3.7 CFD boundary conditions

CFD simulations were carried out using the solver Fluent 19.0 (Ansys, Canonsburg, PA, USA). Blood was modelled as an incompressible Newtonian fluid with density 1060 kg/m^3 and 0.004 Pa·s dynamic viscosity [175]. Vessel walls were considered rigid, and no-slip conditions were imposed. A laminar, steady-state model was selected to simulate blood flow at peak systole [176, 177]. A generalisable inlet condition for the aorta and PA was applied to all subjects, with a uniform (plug) inlet velocity profile of 0.66 m/s for the aorta and 0.57 m/s for the PA [178]. The outlets for all cases were assumed to be at zero pressure and the convergence criteria was set at 10^{-4} for the residual errors. Simulations were run on a Dell workstation, with a Xeon CPU E5-2630 (24 processors at 2.3 GHz), 32 GB RAM and an Nvidia GeForce GTX 1080 Ti.

4.3.8 CFD post-processing

To compare the flow field between the pairs of different unstructured meshes (ML vs GT, SO vs GT and ML vs SO), correspondence was created by subdividing each vessel with 99 planes orthogonal to the centrelines, calculated in VMTK, and equally distanced. Static pressure and velocity magnitude were averaged in each plane (see Fig. 4.1) and a percentage error was calculated for each plane pair, using ML as reference (GT in the comparison between SO and GT). The mean absolute percentage error (MAPE) for pressure and velocity were computed for each vessel pair.

$$MAPE = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{True_i - Pred_i}{True_i}\right| \times 100\% \tag{4.1}$$

4.3.9 Statistical analysis

Shapiro-Wilk tests were used to test the normality of the different segmentation metrics and CFD errors, grouped by vessel (aorta and PA), and segmentation pair (ML vs GT, ML vs SO and SO vs GT). Wilcoxon signed rank tests were used to compare the pressure and velocity errors for the ML vs GT group. Mann-Whitney U-tests were used to compare segmentation metrics and flow field errors between the aorta and the PA, for the ML vs GT group. Friedman tests for repeated measurements were performed to compare segmentation metrics and flow field errors between the ML vs GT, ML vs SO and SO vs GT groups, for both aorta and PA segmentations. Significant Friedman test results were followed up by pairwise Wilcoxon post-hoc tests. Additionally, Wilcoxon signed rank tests were used to compare ML vs GT and SO vs GT metrics for both aorta and PA segmentations. Mann-Whitney U-tests were used to compare Ext-ML vs Ext-SO segmentation metrics against ML vs SO metrics. Wilcoxon signed rank tests were used to compare the pressure and velocity errors for the manually clipped ML vs GT data and the equally clipped ML vs GT data. Pearson's correlation coefficient was used to measure the linear relationship between each pair of a segmentation metric (i.e. Dice, IoU, HD or ASD) and a flow field error (pressure or velocity MAPEs), for both aorta and PA segmentations.

Pearson correlation coefficients were also used for estimating the linear relationship between the flow errors and the radius/angulation error in the mesh outlet flow extensions (ML vs GT). This was performed to assess the influence of outlet geometry error on the flow field error. The p-value was calculated for each comparison to test non-correlation. Throughout this work, a p-value <0.05 was considered statistically significant.

4.4 Results

4.4.1 Hyperparameter optimisation

A total of 124 hyperparameter configurations were sampled during the neural network optimization procedure (see Appendix A.2). The best performing configuration was as follows: scales = 3, layers per block = 2, initial filters = 64, learning rate = 3.46x10-4, batch size = 2, and loss function = focal Tversky. This model was selected and used in all further experiments.

4.4.2 ML segmentation

The ML segmentation was successful in all 10 test datasets. The specific diagnoses for these patients were: repaired tetralogy of Fallot ($n = 1$), repaired Tetralogy of Fallot with mild right PA stenosis ($n = 1$), Marfan syndrome with dilated aorta ($n = 1$), Marfan syndrome with pectus excavatum ($n = 1$), dilated PA ($n = 1$), bicuspid aortic valve with dilated aorta and unrepaired VSD ($n = 1$), repaired double outlet right ventricle with right sided arch ($n = 1$), unrepaired atrial septal defect ($n = 1$), aortic regurgitation with dilated aorta ($n = 1$), post-Ross procedure with mechanical aortic valve ($n = 1$). Inference time for the ML model was approximately 160 ms for simultaneous segmentation of aorta and PAs, compared to approximately 30 minutes for manual segmentation of aorta and PAs. There was good agreement between the ML and GT segmentation, with a median Dice score of 0.945 (interquartile range: 0.929–0.955) for the aorta and 0.885 (0.851–0.899) for the PAs. The Dice score was significantly higher for the aorta than the PAs ($p = 0.002$) with similar findings observed for IoU, HD and ASD (Fig. 4.2A–D).

The best, median and worst segmented images in terms of Dice score are shown

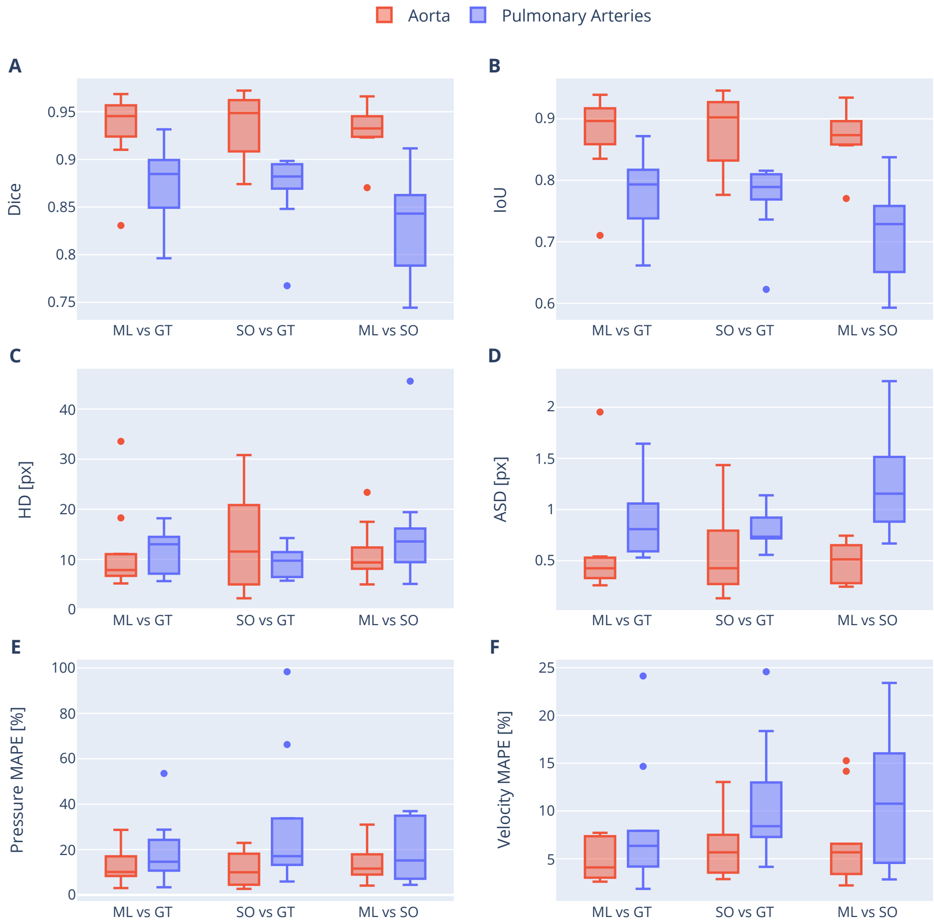

Figure 4.2: Segmentation metrics and flow field errors. Three segmentations are compared in a pairwise fashion: machine learning (ML), ground truth (GT) and secondary observer (SO). (A–B): Confusion-based similarity metrics: Dice score and IoU. (C–D) Distance-based similarity metrics: Hausdorff distance (HD) and average surface distance (ASD), measured in pixels. (E–F) CFD-derived pressure and velocity mean average percentage errors (MAPE).

in Fig. 4.3. The three main differences were: (i) the length of the vessel segmented, (ii) differences in pixel labelling that resulted in small deviations of the vessel border, and (iii) small protrusions at the origin of the carotid and subclavian arteries in the ML segmentations of the aorta.

The aortic inter-observer Dice score (SO vs GT) was 0.949 (0.916–0.960) and was not significantly different from ML vs GT ($p = 0.575$). The pulmonary Dice score for the SO vs GT was 0.882 (0.870–0.894) and was also not significantly different from ML vs GT ($p = 0.721$). The ML vs SO Dice score was 0.933 (0.924–0.944) for the aorta, which was not significantly different from ML vs GT

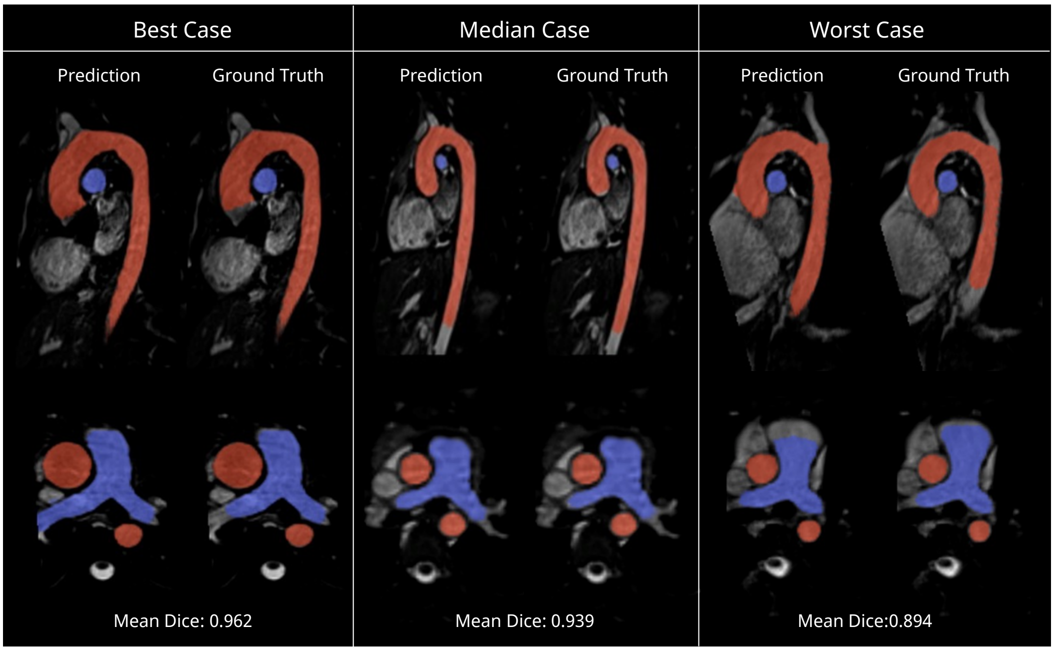

Figure 4.3: Test set segmentation overlays. Predicted and ground truth masks are overlaid on the original images for the best, median and worst test cases. Aorta and PA masks are shown in red and blue, respectively. Multiplanar reformats of the original 3D volume were manually selected on a case-by-case basis to be most informative. The best case had Ross procedure and mechanical aortic valve, the median case had an atrial septal defect and the worst case had a dilated PA.

and SO vs GT ($p = 0.741$), and 0.843 (0.791–0.860) for the PAs, which trended towards being lower than ML vs GT and SO vs GT ($p = 0.061$).

The ML segmentation was also successful in the external dataset. The specific diagnoses for these patients were: cardiomyopathy ($n = 4$), normal anatomy ($n = 1$), repaired tetralogy of Fallot ($n = 1$), left PA stenosis ($n = 1$), anomalous pulmonary venous drainage ($n = 1$), repaired coarctation of the aorta with hypoplastic arch ($n = 1$), bicuspid aortic valve with severe AR and dilated aortic root ($n = 1$). The best, median and worst examples from the external test set are shown in Fig. 4.4. There was reasonable agreement between the Ext-ML and Ext-SO segmentations, with a median Dice score of 0.913 (0.889–0.927) for the aorta and 0.751 (0.728–0.797) for the PAs. Agreement was significantly lower than ML vs SO for the PAs ($p = 0.011$), but not for the aorta ($p = 0.089$). Similar findings were observed for IoU, HD and ASD (Fig. 4.5).

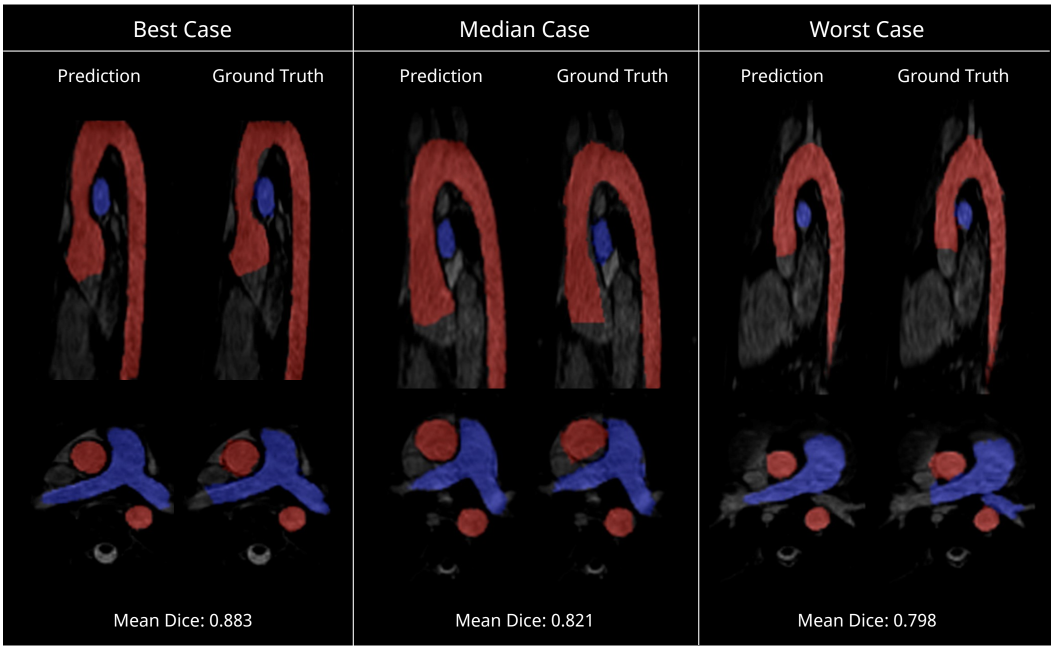

Figure 4.4: External test set segmentation overlays. Predicted and ground truth masks are overlaid over the original images for the best, median and worst test cases. Aorta and PA masks are shown in red and blue, respectively. Multiplanar reformats of the original 3D volume were manually selected on a case-by-case basis to be most informative. The best case had normal anatomy with gothic arch, the median case had bicuspid aortic valve and the worst case had mild left PA stenosis.

4.4.3 CFD metrics

There was overall good agreement in CFD metrics calculated using ML and GT segmentations (Fig. 4.2E–F). The median MAPE for pressure and velocity in the aorta were 10.1% (interquartile range: 8.5–15.7%) and 4.1% (3.1–6.9%) respectively, and for the PAs 14.6% (11.5–23.2%) and 6.3% (4.3–7.9%). PA MAPEs trended towards higher values compared to aortic MAPEs, but this did not reach statistical significance ($p = 0.081$ for pressure and $p = 0.093$ for velocity). However, pressure was more sensitive than velocity to different segmentations, with pressure MAPE being 2.5x greater than velocity MAPE ($***p < 0.001$).

Fig. 4.6 shows the surface meshes of test cases with the highest and lowest CFD MAPE, as well as pressure and velocities along the length of each vessel. Fig. 4.7 shows pressure and velocity fields calculated using both ML and GT manual segmentations. The main difference in the surface meshes (particularly for the

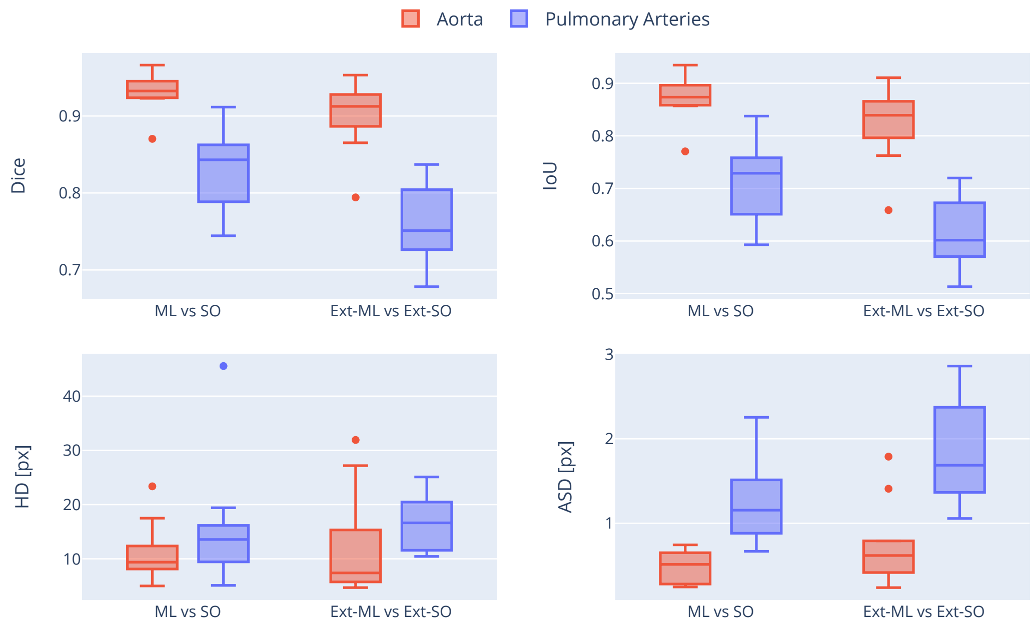

Figure 4.5: Segmentation metrics. The model's segmentation (ML) is compared with the observer's segmentation (SO) for two different datasets: the original data (ML vs SO) and an external test set from a different site and vendor (Ext-ML vs Ext-SO). (A–B): Confusion-based similarity metrics: Dice score and intersection-over-union (IoU). (C–D) Distance-based similarity metrics: Hausdorff distance (HD) and average surface distance (ASD), measured in pixels.

worst cases) were associated with the inlets and outlets (angle and size) and these differences propagated into pressure and velocity fields.

SO vs GT (inter-observer) and ML vs SO pressure and velocity MAPEs were of a similar magnitude to the errors from the ML segmentations (Fig. 4.2E–F, $p > 0.2$). When the clipping planes of the GT segmentations were used on the ML geometries, the median pressure and velocity MAPEs were reduced to 8.0/3.1% ($**p < 0.01$) for the aorta, and to 10.4/3.7% for the PA ($**p < 0.01$) (see Appendix A.3).

Fig. 4.8 illustrates the relationship between the segmentation metrics and the CFD errors on the ML vs GT comparison. No statistically significant correlations were found between any of the metrics, either for the aorta or the PAs, for either manual or equally clipped data.

Fig. 4.9 presents the relationship between outlet surface cross-sectional radius errors and angulation errors (ML vs GT) and flow errors (ML vs GT), com-

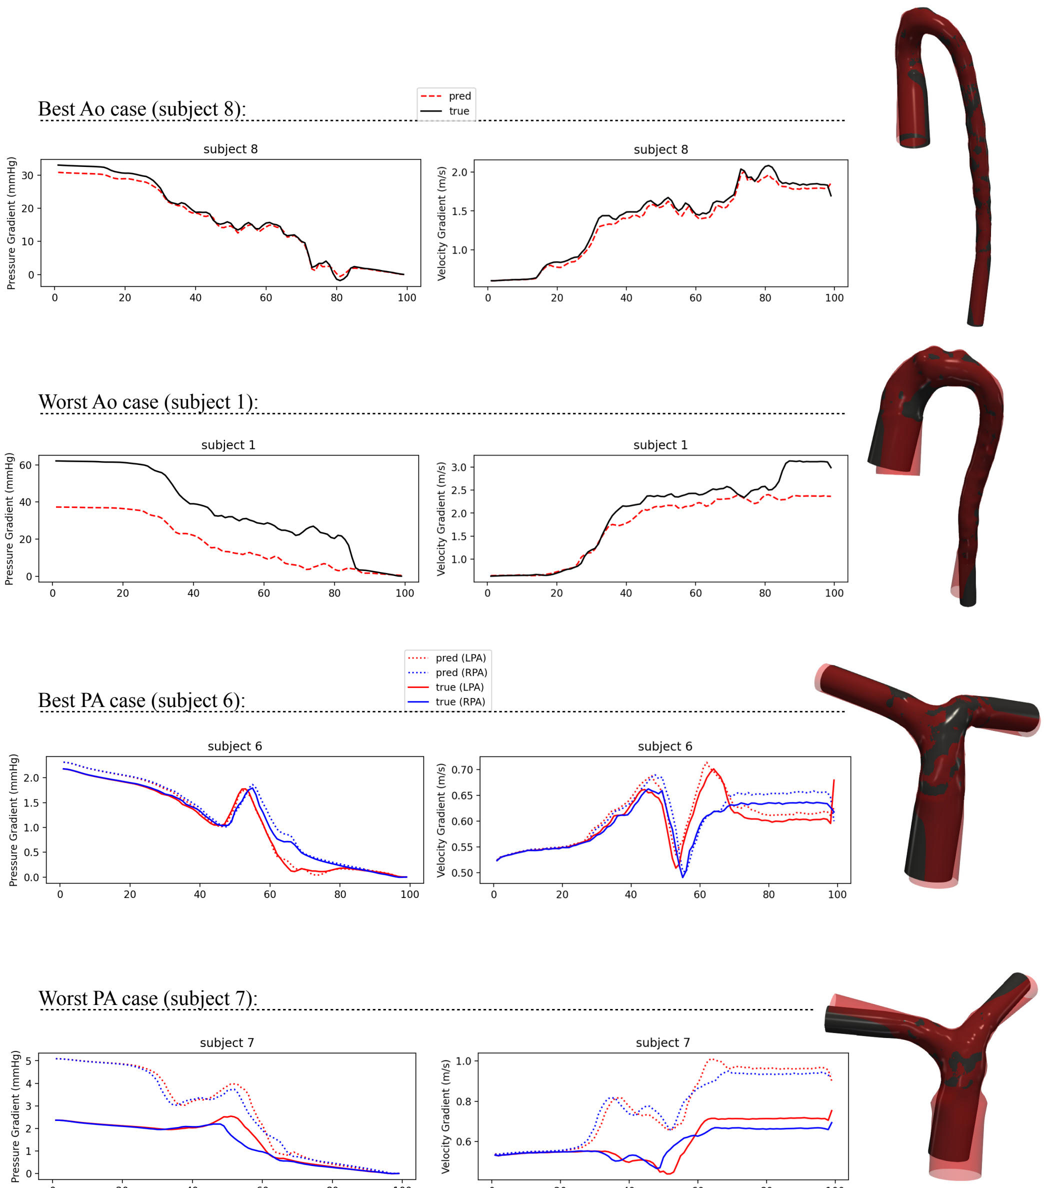

Figure 4.6: Best and worst cases for aortas and PAs. Graphs show planar-averaged metrics along the length of the vessels, starting from the inlet. Black geometries correspond to ground truth, whereas red correspond to predictions. The best aorta case had an atrial septal defect and the worst had a repaired double outlet right ventricle with right arch. The best PA has a bicuspid aortic valve with dilated aorta and unrepaired VSD, and the worst is a Marfan syndrome with dilated aorta.

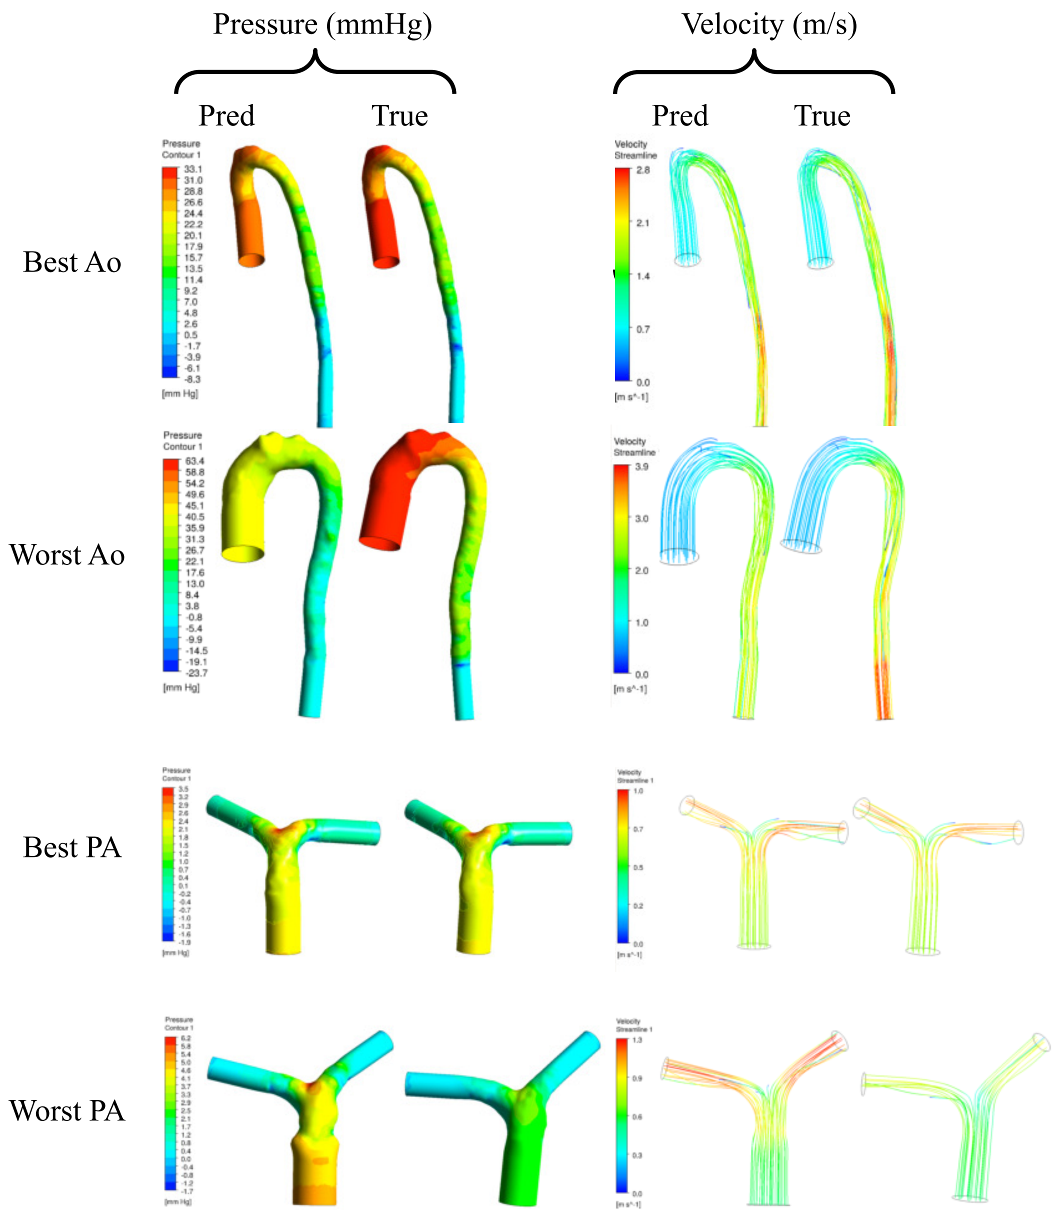

Figure 4.7: Best and worst aorta and PA predictions. Flow fields of pressure and velocity displayed, as contours and streamlines, respectively.

puted using the Pearson correlation coefficient. For the PAs, angulation/radius errors are the average result for both pulmonary branches. Only one statistically significant correlation was found, showing that pressure errors increased when the error in cross-sectional outlet radius increased, with respect to aortic simulations ($R = 0.86, **p < 0.01$). No other statistically significant correlations were found between outlet radius error/angulation error and pressure/velocity errors.

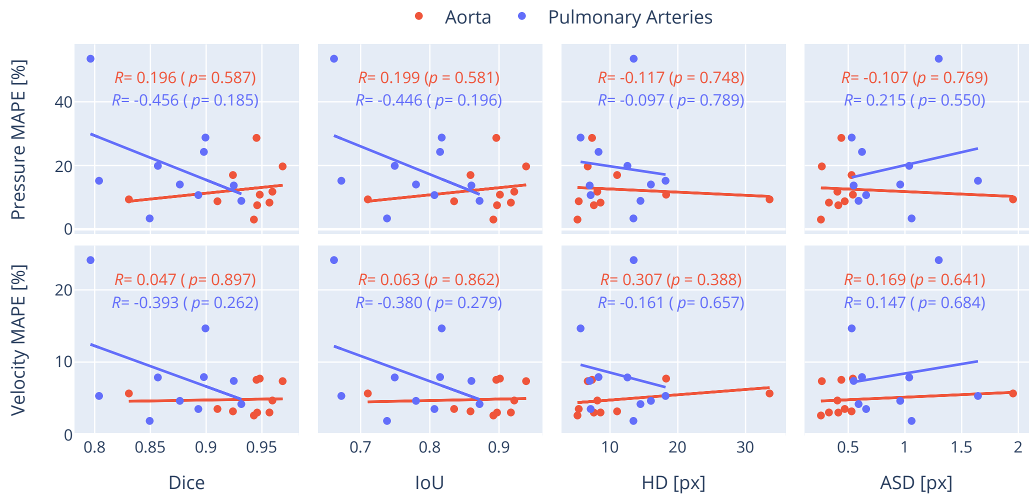

Figure 4.8: Flow errors against similarity metrics. The figure shows a scatter-plot matrix where each point corresponds to a subject. In the abscissas, two confusion-based metrics, Dice and IoU, and two distance-based metrics, the Hausdorff distance and the average surface distance, measured in pixels. In the ordinates, the pressure and velocity mean average percentage errors (MAPE). All values are for the ML vs GT comparison. Red and blue colours identify aorta and PA data, respectively. Trend lines are least-squares polynomial fits of degree 1. For Dice and IoU, higher is better (more similar). For Hausdorff distance, average surface distance and pressure and velocity MAPEs, lower is better.

4.5 Discussion

In this study, a DNN was trained to simultaneously segment the aorta and PAs from 3D CMR data. As its primary purpose was to provide patient specific anatomies for CFD models, accuracy was evaluated using conventional image-based segmentation metrics and resulting errors in CFD measures. The main findings were: (i) the proposed network achieved high performance in terms of image-based segmentation metrics, (ii) there was reasonable agreement between CFD models derived from

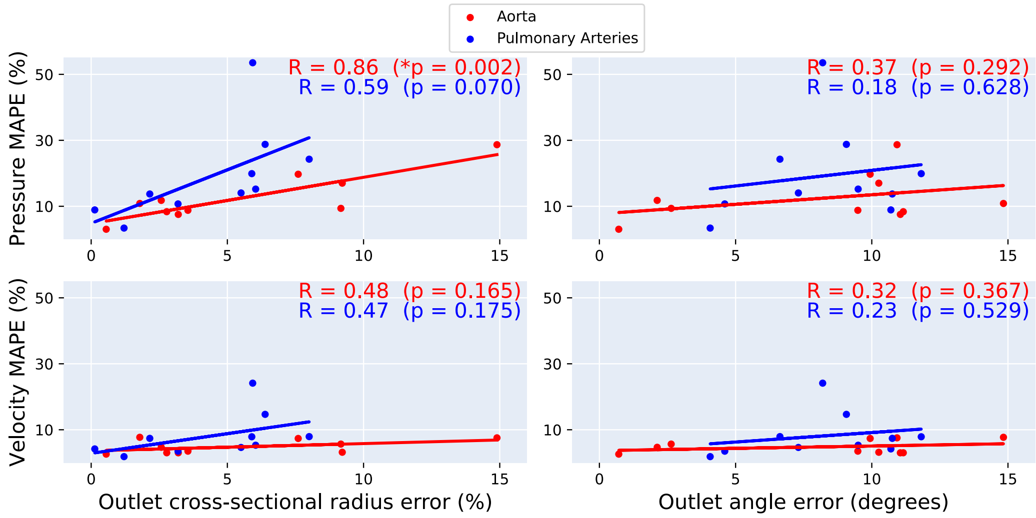

Figure 4.9: Flow errors compared to errors in the outlet cross-sectional radii or outlet angulation. Each point corresponds to a subject. Red and blue correspond to aortic/pulmonary artery errors, respectively. All values are for the ML vs GT comparison. Pearson correlation coefficients for each comparison are shown for aorta and pulmonary artery errors, with associated p values.

the ML and ground truth manual segmentation, (iii) these errors were similar in magnitude to those observed between two different manual segmentations, and (iv) there was no relationship between the segmentation metrics and the resulting CFD errors.

4.5.1 ML Segmentation

It was found that the segmentation model achieved comparable or better performance than previously reported 3D ML segmentation of the great vessels, including in patients with CHD [110, 179]. This suggests that the chosen network architecture and subsequent hyperparameter optimization were sufficient for accurate segmentation. Nevertheless, there were some differences between the GT and ML segmentations, and visual inspection reveals three main types of error. The first error was a tendency for ML to start and stop segmenting at slightly different points in the vessel compared to the ground truth. The second type of error was the presence of "bumps", due to the segmentation masks bleeding out at the locations of arterial branches, particularly in the aorta. Both these errors can be considered failures to properly demarcate vessel limits, rather than failures to correctly label blood pool pixels. The third type of error was inaccurate labelling of blood pixels at the vessel border, resulting in subtle differences in surface geometry. It should be noted that none of these patients had very abnormal pulmonary vascular or aortic anatomy, which was necessary to ensure that CFD models could be created from the segmentations. However, further testing on complex CHD is necessary if segmentation models are to be used more widely. Extension to complex disease may require further enhancements, and several strategies could potentially help improve the ML segmentation accuracy and generalizability. These include increasing the amount and heterogeneity of training data, or performing data augmentation, both of which improve generalizability and performance of ML models [180, 181]. Another interesting option might be the inclusion of statistical shape models [182, 183], which could help ensure that the segmented shapes conform to common patterns.

The model was also tested on 3D data acquired on a different vendor scanner. Although the type of sequence (3D whole-heart bSSFP) and imaging protocol were similar to the original data, there were visually apparent differences in image quality and characteristics. Nevertheless, a reasonable segmentation quality was observed. For the aorta, agreement with a human observer was only slightly lower than agreement with the same observer in the original data. However, there was a larger reduction in agreement for the PAs. This suggests there is scope for improving the generalisability of the model, especially with regards to more complex structures than the aorta, such as the PA. One of the best solutions for this is to include multi-site, multi-vendor data in the training set, but this would incur obvious labelling costs and potential data sharing difficulties. One possibility for bypassing data-sharing difficulties would be for institutions to build generative ML models for creating synthetic images [184]. Since synthetic images would follow the training data distribution, these could be used in cases where patient data sharing is not permissible. Indeed, Fernandez et al. showed that segmentation models could be effectively trained purely with synthetic data [185]. Other approaches to improve robustness to out-of-distribution data might be the use of data augmentation techniques (e.g. domain translation methods to generate multi-vendor datasets [186, 187]) and the use of strategies that incorporate additional domain knowledge [188]. Furthermore, data from different imaging modalities (CT, CMR) and different sequences, or at different resolutions could be included to make the models more robust and generalisable.

Segmentation is not a trivial task, and it was demonstrated that the agreement between two humans was similar to the agreement between ML and the GT human segmentation. This suggests that ML “errors” are approximately at the level of the inter-observer variability, and similar observations have previously been made for aortic segmentation [110]. Furthermore, there are significant advantages of ML over manual segmentation including very fast segmentation without user interaction and perfect reproducibility, due to its deterministic nature. This makes ML particularly useful for removing clinical bottlenecks and accelerating population-based research.

4.5.2 Relationship between CFD and Segmentation Errors

Reasonable agreement in pressure and velocity fields calculated from ML and manual segmentations was demonstrated. Importantly, the differences in CFD metrics using ML vs manual segmentations were of a similar magnitude to those between two independent manual segmentations. This suggests that ML can be successfully used to provide a starting point for CFD simulations, with accuracy similar to inter-observer variability.

However, there were some differences in CFD metrics between ML vs GT segmentations, particularly for pressure calculations. Pressure errors may be higher because local deviations in surface geometry tend to cause only local velocity field derangement, but have a global effect on upstream pressures. This effect is visible in the worst-case aorta, where a narrowing and angulation change in the ML aortic outlet (at the flow extension) resulted in localised flow acceleration, and significantly altered upstream pressures (worst Ao, Fig. 4.7 and Appendix B.7). This effect was analysed over the entire population in order to identify if it was a shared problem amongst all cases. The observation that larger outlet radius errors correlated with larger pressure errors was found (for aortas), supporting this hypothesis. However, no similar correlations existed for outlet angulation error, suggesting this was not influential (Fig. 4.9). No correlations could be identified for the pulmonary arteries. This is likely due to the presence of multiple outlets (flow splits) that introduce new sources of error, making the relationship between shape error and flow error more complex. In the future, the influence of different categories of segmentation/pre-processing errors (elongation/shortening, bumps, angulation changes, boundary errors) should be further explored, with particular emphasis on understanding how the magnitude of these errors affects the CFD flow fields.

Interestingly, no significant correlations between image-based segmentation metrics and errors in the pressure/velocity fields were found. This suggests that neither overlap-based (Dice, IoU) nor boundary distance-based (HD, ASD) metrics can accurately capture the features that ensure CFD accuracy. This may be because CFD models are highly sensitive to local geometric errors, while segmentation metrics are global and therefore may not fully capture these localised deviations. Another reason may be that differences in clipping (which were not accounted for by segmentation metrics) are responsible for some of the CFD errors, as shown by the analysis of equally clipped data. However, significant CFD errors remained after removing this confounding factor, and these errors were still not correlated with image-based segmentation metrics. Irrespective of the cause, the poor correlation between segmentation and CFD errors has some important implications. Specifically, in this application, it might be better to combine conventional global image-based losses with more CFD specific objective measures during training. For example, these could include losses which prioritise preserving accuracy of the domain boundaries, since it was observed that outlet radius errors are highly correlated with pressure errors in aortas (Fig. 4.9).

4.5.3 Limitations

One of the main limitations of this study was that a simplified CFD model was applied across all subjects (laminar, steady state, with no patient-specific parameters). This was done to better isolate the effect of segmentation differences on the resulting CFD model. However, it does limit the patient specific aspect of these comparisons and in the future, boundary conditions for each subject (such as velocity profiles taken from phase contrast CMR) could be incorporated into the model. Furthermore, now that good agreement has been demonstrated using simple CFD models, the utility of ML segmentation for more complex CFD models should be investigated. Another limitation is that the methods used for comparison do not necessary account for the full flow field. Plane-averaged pressures and velocities along the length of the centreline were used to quantitatively compare different CFD models. However, this averaging does lead to a loss of localised details in the flow fields. A node-based approach for comparing two distinct 3D flow fields could be useful for improved model evaluation, and will be presented in the following Chapter.

4.6 Summary and conclusion

A CNN was developed, optimised and trained for segmentation of the aorta and the PAs in 3D cardiovascular CMR. The segmentation network was validated for its primary purpose - the creation of CFD models and calculation of flow fields - with errors in the range of human inter-observer variability. The proposed method can help automate the workflow of clinical hemodynamic assessments and improve its robustness, as it provides 3D anatomical reconstructions nearly instantaneously after image acquisition. Assembling the specific patient CFD analysis and the related computational costs remain the next hurdle for a fully translatable tool in clinical practice. Thus, automation of the CFD setting-up and solver will be shown in the following Chapter, as the final stage for a fully automated CFD pipeline for cardiovascular applications.

Chapter 5

Design of a real-time computational fluid dynamics pipeline: a machine learning approach for predicting aortic hemodynamics

5.1 Introduction

Despite the potential benefits described in Chapter 2 and 3, CFD is still not integrated into routine clinical practice due to the need for an experienced engineer to set-up simulations, high computational resource requirements and long simulation times. Whilst the ML method presented in Chapter 4 can help speed up the conventional CFD simulation pipeline, many other steps still require manual inputs. ML models have previously shown to be suitable for accelerating and automating CFD pipelines [116, 117]. The main challenges of training ML models using CFD data include: (i) poor availability of clinical data, (ii) unstructured meshes without point correspondence, and (iii) large meshes and resultant flow fields.

5.2 Aims and objectives

The study presented in this Chapter has the following three primary objectives: (i) to create a suitably large synthetic cohort of 3D aortas based on real patients;

(ii) to train ML models to predict aortic/velocity pressure fields by representing unstructured/large data types with low-dimensional vectors; and (iii) to compare results between the ML-based CFD solution and the conventional CFD method.

5.3 Methods

5.3.1 Statistical shape modelling

The dataset used for development of the SSM consisted of cardiac and respiratory gated steady state free precession CMR images (N=67) from patients previously diagnosed with coarctation of the aorta (CoA). All patients were post-surgical repair, asymptomatic and underwent CMR imaging at a mean age of 22.4 ±6.2 years. Images for each subject were manually segmented and converted into surface meshes (Section 2.3). This was followed by remeshing and smoothing using functions from the VMTK, as described in Section 4.3.6. All geometries were aligned in the same local space and orientation through rigid registration using an iterative closest point algorithm in VMTK [189]. This ensured that shape modelling was not affected by any spatial misalignment. Surfaces were manually clipped above the aortic root for the inlet, and at the diaphragm for the outlet. An SSM was then built using these clipped aortic surfaces, using an approach previously described by Bruse et al [125].

The package Deformetrica 4 was used to build the SSM [123]. First, an average aortic shape (surface template) was computed, containing 2,541 nodes (Fig. 5.1). A volume template was also created by meshing the surface template with tetrahedral elements (29,000 nodes). Each subject could then be described as a non-linear deformation of 3D ambient space, relative to the template (Fig. 5.1). In this case, each deformation is fully parameterised by a paired set of 3D control points $(q_i)_{i=1,...n}$ and 3D momenta vectors $(\mu_i)_{i=1,...n}$ (Section 3.4.1) using a Gaussian kernel of width σ which was set as 10 mm. The number and location of control points were optimised (n=172) by initialising the model with a high resolution control point grid (n=500) and truncating points which were observed to have little influence on the deformation (low variance). The final computed 3D deformations for a given subject were represented by a deformation vector with 516 coefficients - the number of control points (172) multiplied by the number of deformation directions (3). The deformation vectors for the whole population were collected in the 2D matrix M [67, 516], which was decomposed using PCA in order to identify low-dimensional deformations which account for most of the variance. This required standardisation for each column in M (i.e. removal of the mean and scaling to unit variance). Following this, SVD was performed with M being decomposed into three matrices: $M = USV^T$. The V transpose matrix [516, 516] contained the principal component axes or PCA modes, S is a rectangular diagonal matrix [67, 516] where the singular values on the diagonal can be used to calculate the variance explained by the associated PCA modes, and the matrix product US [67, 516] contains the projection (weights) of each subject onto each PCA mode. It was found that the first 35 PCA modes were capable of approximating 99% of the variance in M. This meant that specific aortic shapes could be represented by a lower-dimensional deformation vector with 35 coefficients rather than 516 (almost 15 times reduction in the size).

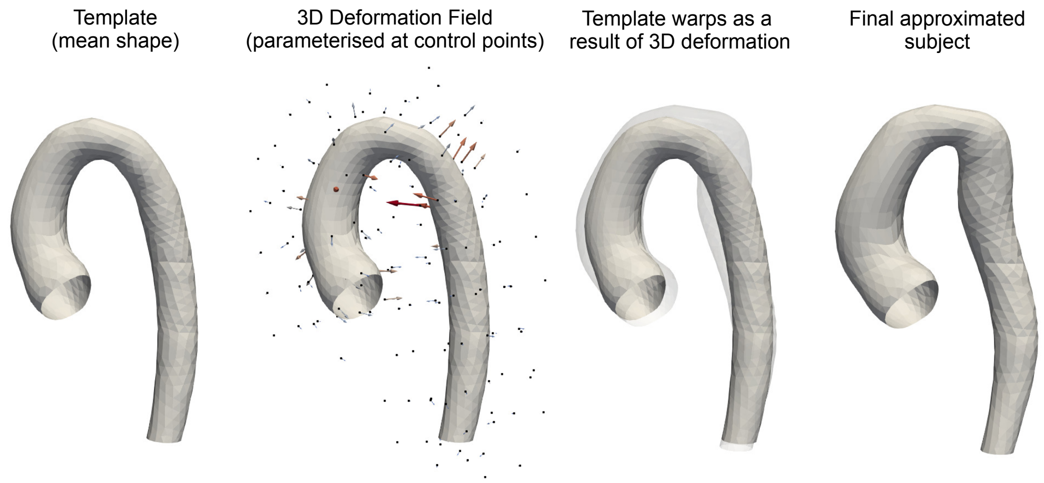

Figure 5.1: Mesh registration with SSM. An example aortic shape approximation using my SSM is shown. Individual surface or volumes can be reconstructed using a mean aortic shape and applied deformation field initialised on a set of control points (n=172).

Using the SSM, new synthetic aortic shapes could then be created using synthetic lower dimensional deformation vectors. Specifically, each of the 35 coefficients in a synthetic vector was generated by randomly sampling a Gaussian distribution (within 2 standard deviations) based on the distribution of weights in the US matrix. Following concatenation of all the lower dimensional deformation vectors, the matrix X [3000, 35] was transformed into a matrix L [3000, 516] by matrix multiplication ($L = XSV^T$), thus reversing PCA and mapping the matrix X onto the original axes. Standardisation was then reversed in all columns of L using the previously computed standard deviations and means in the M matrix. The deformation matrix L was then reshaped into a 3D momenta matrix [3000, 172, 3] and applied onto the aortic surface and volume template using Deformetrica, thus generating a surface and volume mesh for each new synthetic subject. Since all new meshes are derived from the same template, all synthetic aortas contained the same number of nodes/elements. Additionally, nodes can be thought to be lying within spatially correspondent locations within each aorta (see Appendix B.1). This was vital for enabling the dimensionality reduction of derived flow fields, as described in later sections.

The new synthetic population (n=3000) was compared to the original population (n=67) by computing geometric properties of the shapes based on a centreline approach. Mean centreline lengths and diameters were computed. Mean torsion is used to express how sharply the centreline is twisting in space. The parameter tortuosity describes the length ratio between the centreline and a rectilinear line between the endpoints. All parameters were computed using implementations within VMTK, as described by Piccinelli et al [190].

5.3.2 CFD pipeline

Volume meshes previously generated with the SSM were unsuitable for CFD computation. This was primarily because low mesh skewness could not be guaranteed, and remeshing was not an option since nodes were to be preserved in order to maintain point correspondence (see Appendix B.1). Therefore, separate meshes solely for CFD computation were built, starting from the surface of each aorta, and following the pipeline as described in Section 4.3.6. Firstly, each surface was extended by 40 mm at the inlet. This was done to produce a flat and circular inlet upon which a velocity profile could be uniformly applied. Extending the inlet further than 40 mm was avoided in order to reduce the likelihood of surface self-intersection. Following this, volume meshing with tetrahedral elements was performed [191] (∼400,000 cells on average). Element/node counts were deemed to be in satisfactory ranges after a mesh sensitivity analysis (see Appendix B.2).

CFD (Fluent, Ansys Technologies) was performed on all 3,000 synthetic cases. The same boundary conditions were applied to each simulation, as part of adopting a simple model which would reliably converge for all subjects. Laminar, steady-state flow conditions were enforced. An inlet velocity of 1.3 m/s, corresponding to an average ascending aortic flow rate at peak systole, was set [192, 193]. A velocity boundary condition was preferred to a volumetric flow rate, since it is invariant to any differences in inlet surface area between subjects. Outlet gauge pressure was fixed at 0 Pa. Standard non-slip conditions were applied at the wall, and the fluid was assumed to be Newtonian with density and dynamic viscosity equal to 1,060 kg/m^3 and 0.004 Pa·s, respectively [194]. The set-up and simulation of all 3,000 cases was fully automated.

5.3.3 Machine learning

5.3.3.1 Data interpolation

As CFD was performed on large unstructured meshes with inconsistent numbers of nodes/elements between cases, point correspondence had to be restored prior to PCA-based dimensionality reduction of the flow fields (needed for easier model training). This was done using the volume meshes previously generated with the SSM (by applying deformations to the volume template mesh). Since each of these 'SSM volume meshes' inherited its properties from the template, relative nodal positions were preserved. Consequently, pressure/velocity data for all subjects was resampled from unstructured CFD meshes onto SSM volume meshes using a Voronoi kernel (padding of 5 mm, grid resolution of 1,000,000 voxels) in the software Paraview (Fig. 5.2). This resulted in 3,000 newly resampled pressure/velocity fields, with all subjects containing 29,000 nodes in point correspondence. The data was concatenated into a feature vector and PCA was applied to reduce dimensionality. The aim was to capture 99% of variance with as few PCA modes as possible.

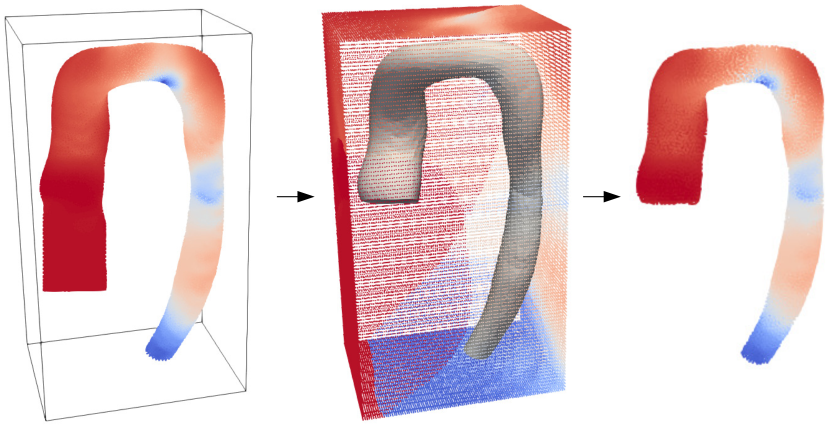

Figure 5.2: CFD data interpolation. CFD results are interpolated onto a point-correspondence mesh (generated by the SSM), thus restoring node concordance.

5.3.3.2 Deep neural network architecture

The architecture adopted was a standard sequential, fully-connected DNN with independent networks for pressure and velocity. The input for the model was the lower dimensional deformation vector, which is also referred to as a 'shape vector'. The outputs of the trainable part of the model were the pressure/velocity PCA scores (reduced order CFD field), referred to as a 'pressure/velocity vector'. A non-trainable inverse PCA layer (implemented in Keras using a lambda layer) serves to reconstruct the pressure/velocity vector into the full 3D flow field with 29,000 nodes (see Fig. 5.3). Rectified linear units (ReLU) were used in each hidden layer. Linear activation functions were set at the output. Model implementation was done using Keras and TensorFlow 2.0.

5.3.3.3 Deep neural network training

Models were built separately for predicting the static pressure and the velocity-magnitude. The loss function used for training was the mean absolute error (MAE), computed on the entire 3D flow field (i.e. after inverse PCA) rather than on the output pressure/velocity vector. This provides a more granular measure of error and effectively weights the importance of each PCA mode in the network according to the amount of variance it explains. Model optimisation was carried out using

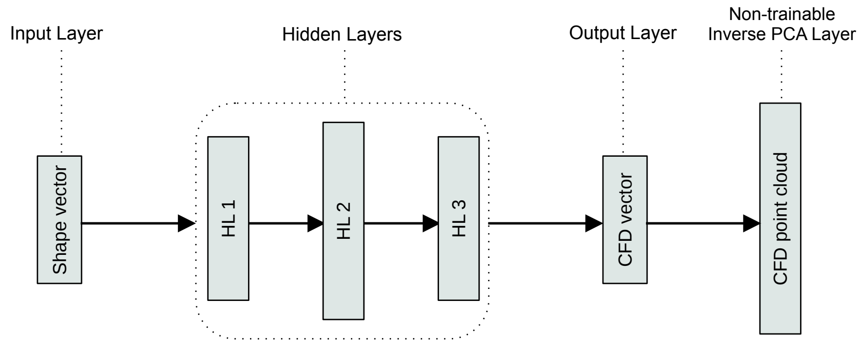

Figure 5.3: The general sequential, fully-connected DNN set-up used to build both pressure and velocity predictors ('CFD vector' can be either pressure or velocity PCA vectors).

the Adam optimiser [195]. Hyperparameter tuning was conducted with a 5-fold cross-validation process to find the best model settings. The number of hidden layers, number of hidden layer neurons and initial learning rate were all explored. This was repeated for 1,000 model configurations, sampled using a tree-structured Parzen estimator (TPE) algorithm [196]. Batch size and epochs were set at 32 and 50, respectively. Hyperband pruning was used to terminate early training rounds if the model was deemed to be poorly fitting the validation data.

After completing hyperparameter tuning using 5-fold cross-validation, the model was retrained on the entire training dataset. From 3,000 subjects, 2,800 were randomly selected to make up the training set (with 10% being used for validation). The remaining subjects (200) formed the test set. Model training lasted 1,000 epochs and training/validation loss was monitored. A workstation with an Nvidia GTX 1080Ti graphics card was used to perform training.

5.3.3.4 Model evaluation

Once trained, models were evaluated on the test set of synthetic aortas (n=200). Absolute errors were computed for every node in all test cases by comparing the prediction value (ML) to the ground-truth value (CFD). Errors were then normalised according to subject CFD data range, as detailed in Liang et al. [118]. Normalisation was necessary to enable direct comparison between individual cases and also between CFD metrics, since pressure/velocity ranges widely differed per subject. Equation 5.1 details how normalised absolute error (NAE) is computed for either pressure or velocity at a node i, belonging to a subject j. $True_{i,j}$ is the CFD nodal pressure/velocity value. $Pred_{i,j}$ is the ML nodal pressure/velocity value. $Range(True_j)$ is the difference between the maximum and minimum values in the CFD flow field for subject j.

$$NAE(i,\ j) = \frac{\left|True_{i,\ j}\ -\ Pred_{i,\ j}\right|}{Range(True_j)} \times 100\% \tag{5.1}$$

Mean node errors ($\mathrm{MNAE_N}$) were computed by averaging NAE values across the population for each node (n=29,000). These were then plotted on the template mesh points in order to better visualise the magnitude of these errors with respect to their location. However, since NAE values are absolute errors, this provides no insight regarding any systematic over or under-estimation during model inference. Therefore, a Bland-Altman plot was used to examine the bias and limits of agreement of the pressure and velocity DNNs. This was done for the overall aorta and for three separate regions; ascending aorta, transverse arch and descending aorta (anatomically defined).

Mean subject errors ($\mathrm{MNAE_S}$) were computed by averaging NAE values in each subject (n=200). Cases with the best, median and worse mean subject error values were compared. A single population error for both pressure and velocity was given by averaging all $\mathrm{MNAE_S}$ values. The relationship between shape mode scores and subject error ($\mathrm{MNAE_S}$) was investigated with scatter plots and assessed using Pearson R coefficients and p values ($p = 0.05$ considered significant).

In addition to evaluation of the models on the test set (n=200), the models were also tested on real, patient-specific aortas with previously repaired CoA (n=10), completely unseen from the SSM and the DNN. This was done in order to validate the robustness of the models for inferring accurate flow fields on real subjects outside the synthetic training/testing sets. CMR images of each patient were segmented. Surfaces were approximated by non-rigid registration (applying deformations on the template) using the SSM. Deformation matrices for each case were decomposed into PCA shape vectors and passed as inputs into the DNN models. Pressure and velocity-magnitude fields were inferred for each subject. Following this, all ten predictions were compared to CFD flow fields computed using both SSM derived geometries and real geometries.

Comparison between ML and SSM derived flow fields were performed using the $MNAE_S$ as previously described. However, it was not possible to compare the predicted and real CFD flow fields with node-based metrics due to the lack of node-to-node correspondence and exact surface matching. Therefore, a gradient-based approach was used to enable direct comparison of pressure/velocity flow fields without shape correspondence. Subject centrelines were used to calculate plane-averaged pressure/velocity gradients at 99 locations over the length of the aorta, as shown in Section 4.3.7. To compare gradients, the Fréchet distance (FD) was used. The FD is a measure of similarity between two point-sets of curves, taking into account the location and ordering of the curve coordinates. Intuitively, it can be thought of as the shortest possible distance between two observers traversing different paths while remaining connected. An advantage of using the FD is that it does not neglect sharp spikes or differences between gradients, which some other metrics may diminish through averaging. Additionally, FD is not a percentage error, therefore is stable around values close to 0 (such as at the very end of the descending aorta in pressure flow fields). An algorithmic implementation for computing FD as described by Eiter et al. was used to calculate this metric [197].

5.4 Results

5.4.1 Statistical shape modelling

PCA decomposition was performed on the deformation matrices (momenta) computed by statistical shape modelling. The first, second and third PCA modes captured 29.8%, 13.2% and 10.1% (total 53.1%) of the variability, respectively. Mode 1 relates to overall vessel size, mode 2 relates to ascending arch angulation/diameter, and mode 3 describes rounded versus triangular arches. After PCA decomposition, 99% of the variance in the momenta could be represented with the first 35 modes. Some examples of the 3,000 synthetic subjects produced by randomly sampling and combining 35 PCA mode scores are shown in Fig. 5.4. Anatomical characteristics of the synthetic aortas (length, diameter, tortuosity and torsion) were found to be statistically similar to those of the real patient cohort (see Table 5.1).

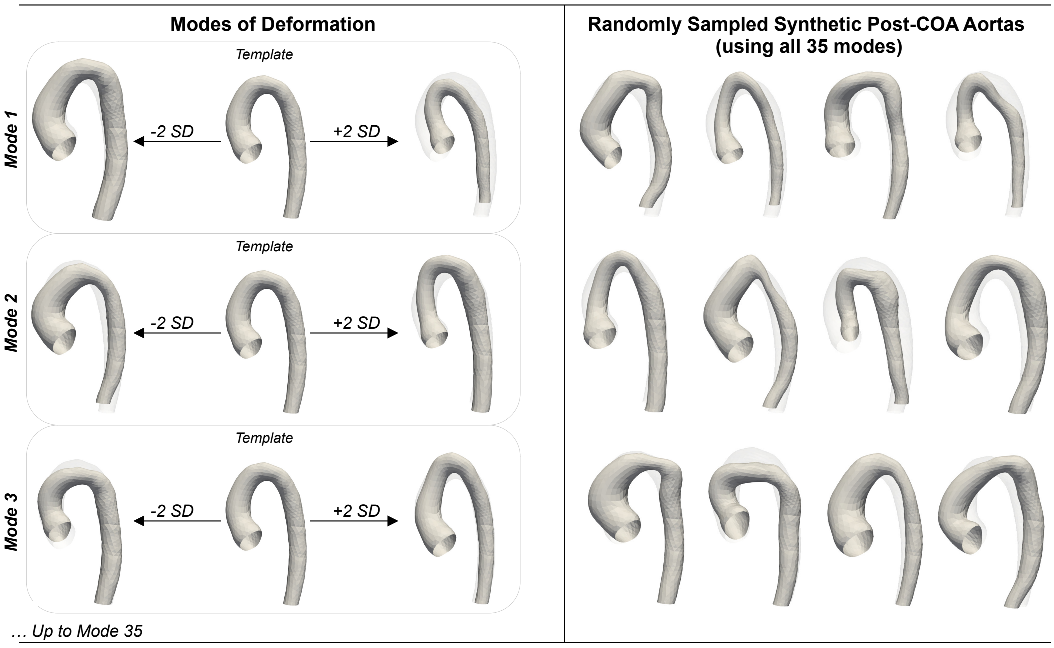

Figure 5.4: Modes of deformation. Left: first three modes of deformation from the SSM (SD = standard deviation). Right: examples of synthetic post-repair CoA aortas from the test set (using combinations of all 35 shape modes).

	Diameter (mm)	Length (mm)	Tortuosity	Mean torsion
Real cohort mean	19.74 ±1.29	257.3 ±29.88	2.21 ±0.39	0.0044 ±0.061
Synthetic cohort mean	20.12 ±1.32	258.5 ±23.08	2.19 ±0.35	0.0012 ±0.053
p value (Welch's t-test)	0.25	0.75	0.74	0.65

Table 5.1: Comparison of aorta dimensions between original real cohort (n=67) and synthetic cohort (n=3000).

5.4.2 Training data

After CFD was computed on all cases (n=3000), values were interpolated from high resolution meshes onto lower-resolution grids in point correspondence. The mean loss in accuracy due to interpolation was found to be 0.056% ±0.027 and 0.849% ±0.247 for pressure and velocity-magnitude, respectively. This was computed by calculating the mean percentage error in centreline pressure and velocity gradients for all cases (n=3000) and averaging the results. The data post-interpolation was used as the 'ground-truth' training and testing sets.

PCA decomposition of the pressure and velocity training data matrices ([2800, 29000] each) was then performed, following standardisation. After PCA decomposition, 99% of the standardised pressure variance could be captured with 20 modes. Only 87% of the standardised velocity variance could be captured with 55 modes, and it was felt that adding more modes to capture greater variance was not feasible due to massively diminishing returns. Subject errors resulting from PCA decomposition were tested on the 200 test cases (unseen by the PCA model). Average $\mathrm{MNAE_S}$ in pressure and velocity fields were found to be 1.46% $\pm$ 0.59 SD and 2.70% $\pm$0.49 SD. The reconstructed test cases with the highest $\mathrm{MNAE_S}$ for pressure and velocity (4.32% and 4.93%, respectively) are shown in Appendix B.3.

5.4.3 Model architecture

The model input layer size was set at 35 (number of shape modes). Output layer sizes were set at 20 and 55 for pressure and velocity, respectively (number of pressure/velocity modes). Hyperparameter tuning using cross-validation was performed 1,000 times to search for the optimal learning rate, number of neurons and number of layers, with tuning taking $\sim$4 hours per model. Pressure and velocity model architectures as a result of the optimisation process are shown in Appendix B.4.

5.4.4 Model predictive performance

Pressure and velocity-magnitude fields were computed on the test set (n=200) using the trained DNNs. Conventional CFD took $\sim$5 minutes per subject, whereas ML inference took $\sim$0.075 seconds, demonstrating a $\sim$4,000x speed-up.

5.4.4.1 Node errors

Average node prediction errors ($\mathrm{MNAE_N}$) were computed for all nodes (n=29,000). Fig. 5.5 (left) shows these values projected onto the template (average position of the nodes), allowing for the locations of the highest absolute errors to be assessed. Pressure errors were observed to be lower in the descending aorta, with the highest errors situated in the transverse arch. Velocity errors were notably more prevalent in the underside of the arch and descending aorta. The maximum $\mathrm{MNAE_N}$ was observed to be 12.46% and 14.86% for pressure and velocity, respectively.

Bland-Altman analysis showed negligible prediction biases for the overall aorta (Fig. 5.5, right) and within selected regions (Appendix B.5). Bland-Altman biases were found to be 0.19% and 0.28% for pressure and velocity, respectively. The limits of agreement were found to be marginally wider for pressure when compared to velocity (15.65% vs 12.89%, respectively).

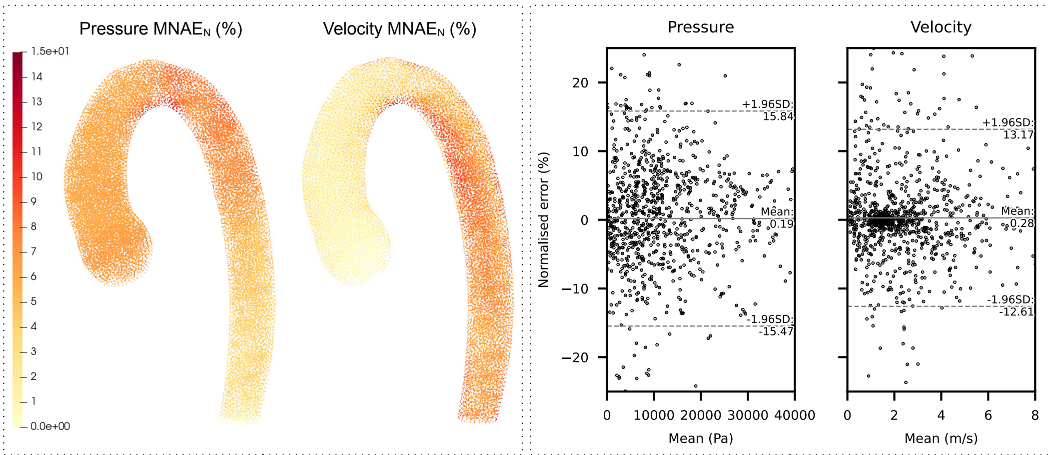

Figure 5.5: Nodal error analysis. Left: Distribution of mean nodal errors (MNAE$_N$), computed on the test set (n=200). Errors are absolute values and are projected on the template aorta. Right: Bland-Altman plots for the overall aorta. Normalised error (%) refers to the NAE of each node in every test case, without taking the absolute value (n=5,800,000). Only 1,000 randomly selected points were drawn to improve graph readability.

5.4.4.2 Subject errors

The population error for pressure and velocity was 6.01 $\pm$3.12% SD and 3.99 $\pm$0.93% SD, respectively. The test cases with the best, median and worst subject error (MNAE$_S$) are shown in Fig. 5.6 with corresponding pressure/velocity gradients. The maximum MNAE$_S$ for pressure and velocity were found to be 23.60% and 8.07%, respectively. The minimum MNAE$_S$ for pressure and velocity were found to be 1.54% and 1.91%, respectively.

The relationship between pressure/velocity MNAE$_S$ and shape mode coefficients is presented in Fig. 5.7. The second and third mode showed statistically significant correlations with velocity prediction MNAE$_S$, with the third shape mode showing the highest correlation (R=-0.31). No significant correlations were found between any other shape mode and MNAE$_S$.

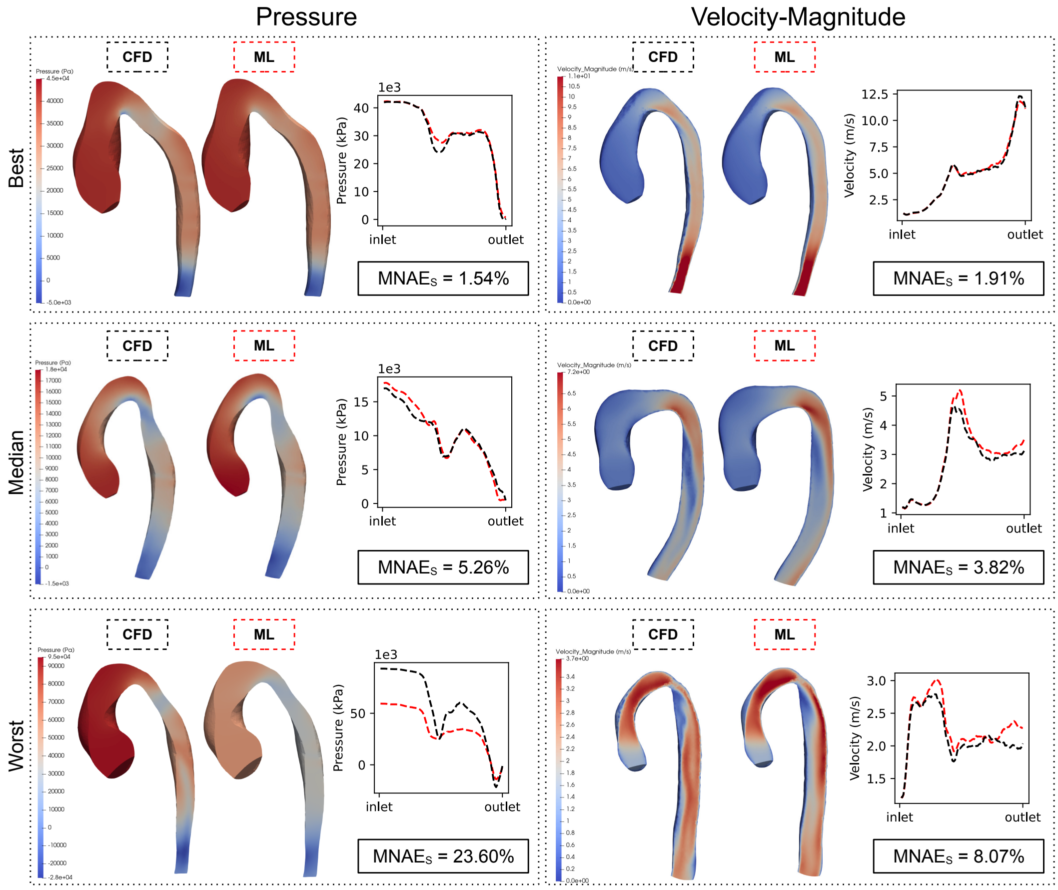

Figure 5.6: Best, median and worst test-set predictions. Comparisons between ground truth (CFD) and predicted (ML) in the test set (n=200). Best, median and worst cases for both pressure and velocity-magnitude are shown, ranked using the mean node-to-node error ($MNAE_S$). Pressure/velocity gradients are also displayed.

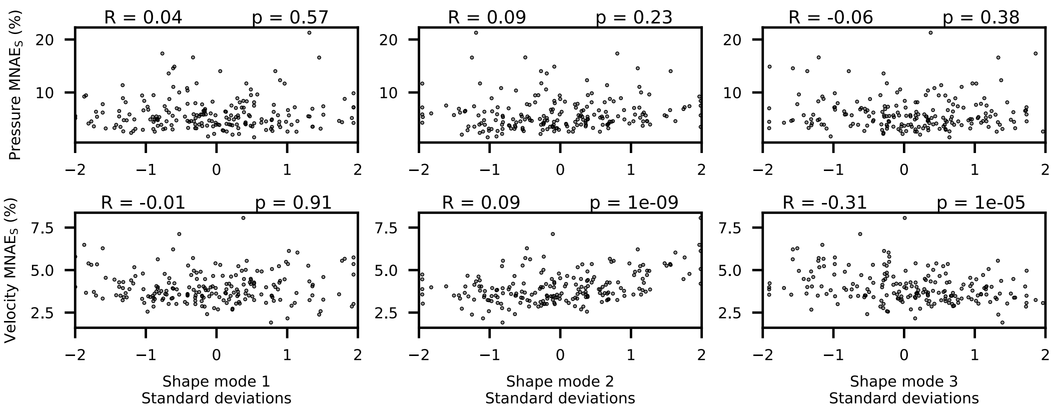

Figure 5.7: Shape modes vs. ML error. Scatter plot comparing the shape PCA mode values against subject error ($MNAE_S$) in the test set (n=200). Pearson R coefficients and p-values were computed for each subplot.

5.4.5 Validation with real patient data

Coefficients for the first 10 shape modes for all new subjects (n=10) were found to lie within the max-min ranges of the original PCA shape modes (n=67, see Appendix B.6). The average prediction error using the MNAE_S metric (ML vs CFD computed on SSM meshes) was 10.19% $\pm$10.41 and 4.47% $\pm$1.18 for pressure and velocity respectively. This corresponds to an increase in mean subject error by 4.18% for pressure and 0.48% for velocity, when compared to the population error in the test-set of synthetic cases (n=200).

Due to the lack of point correspondence and surface matching between the real and SSM shapes, Frechet distance of pressure and velocity gradients were used to compare the ML CFD results with the patient CFD data (computed using both SSM and real geometries). The FD computed between the ML and the CFD (SSM) gradients corresponded to the prediction error arising solely from the ML model (FD SSM). The FD computed between the ML and the CFD (real shape) gradients corresponds to the total error between the ML prediction and true CFD (FD real). The subjects with the best and worst FD (real) for pressure and velocity are shown in Fig. 5.8. The mean FD SSM was 1680 $\pm$629 Pa for pressure and 0.47 $\pm$0.17 m/s for velocity. This compares to 4583 $\pm$3210 Pa for pressure and 1.30 $\pm$0.37 m/s for velocity between the ML and real-CFD (FD real).

5.5 Discussion

The main findings of this study were: i) SSMs and PCA are suitable for creating synthetic training data and dimensionality-reduced representations of 3D shape and flow, ii) DNNs based on these dimensionality-reduced representations can predict pressure and velocity fields with high accuracy.

5.5.1 Synthetic data generation and dimensionality reduction

A key element of the approach presented was to use an SSM and PCA to both generate a synthetic training dataset (n=3,000), and to parameterise aortic shape/flow fields for simplifying DNN model training. It was demonstrated that volumetric meshes generated by the SSM were suitable for forcing point correspondence in the

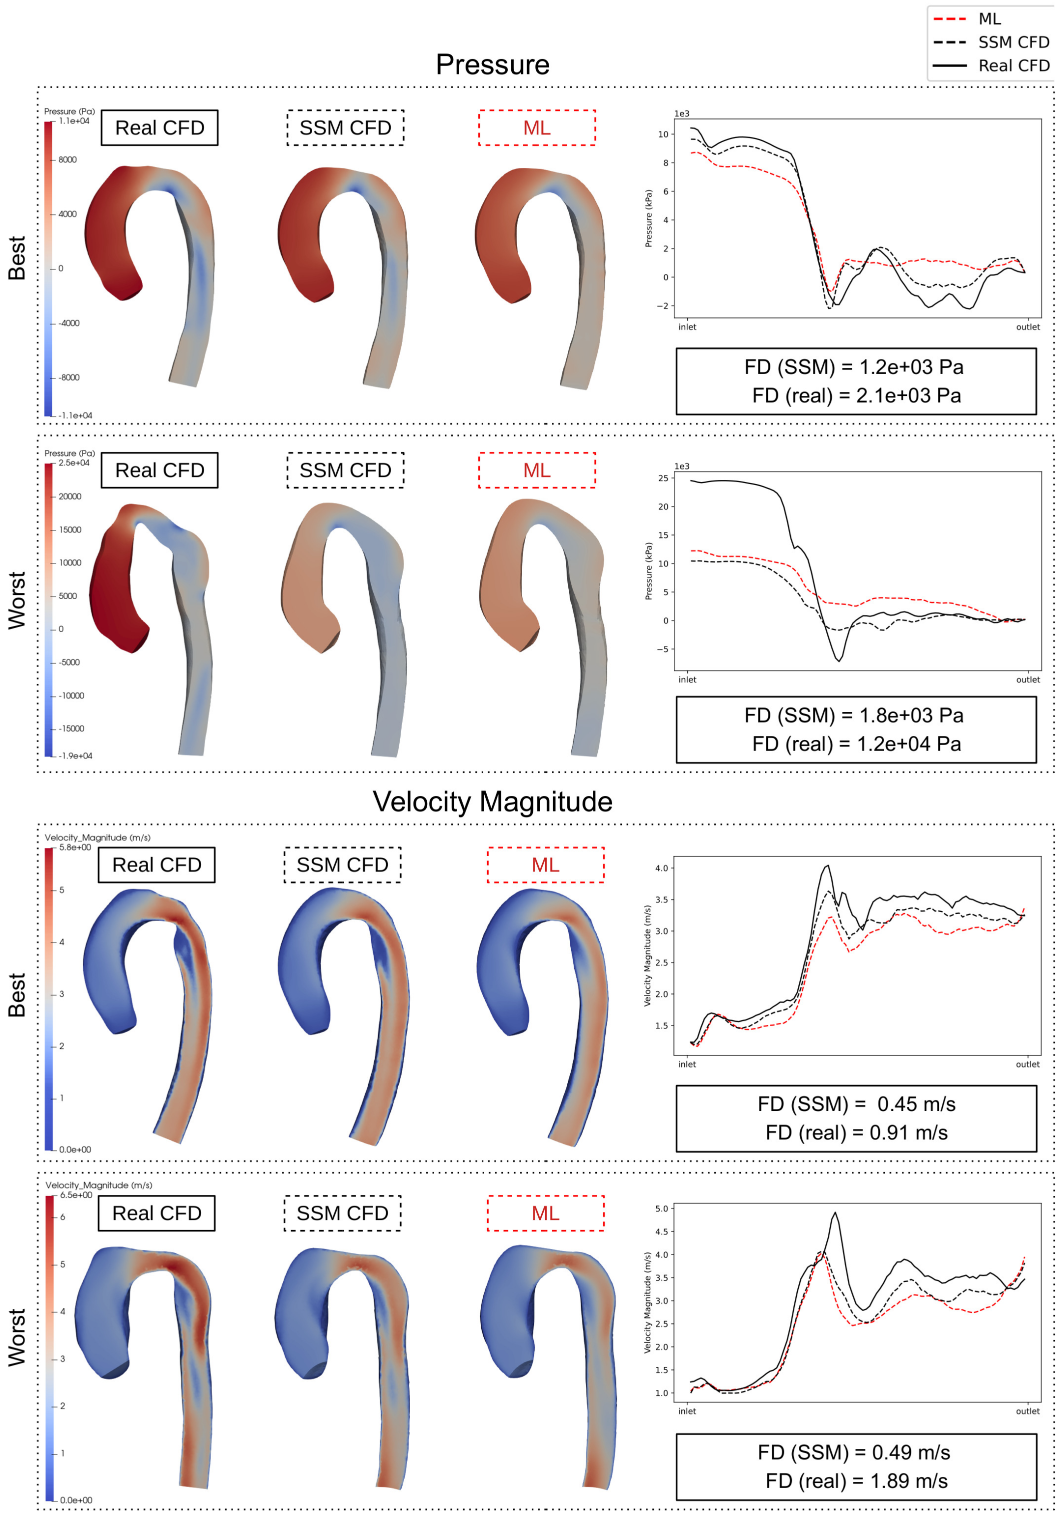

Figure 5.8: Best and worst pressure and velocity predictions on the real patient test cohort (n=10). ML predictions are performed on the SSM mesh representation. FD (SSM) is the error between the ML predicted (red) and SSM mesh CFD (dashed black) gradients. FD (real) is the error between the ML predicted (red) and true mesh CFD (solid black) gradients.

CFD dataset through interpolation of 3D aortic CFD flow fields. It was also shown that less than 60 PCA modes were needed to capture the majority of variance within aortic shape deformations and 3D pressure/velocity fields. For the shape deformations, it was found 99% of the total variance could be captured with the first 35 PCA modes. However, the selection of these modes was based purely on amount of variance captured, and potentially omits low-energy deformations which could still correlate with the CFD flow-fields. In the future, a sensitivity analysis could be conducted to assess how many and which shape/CFD modes should be incorporated in the DNN, and how this would affect model performance. From observation, mean PCA reconstruction errors of pressure/velocity were found to be low, but not insignificant in the worst observed test-set reconstructions (Appendix B.3). Geometric properties of the synthetic data were shown to be mostly close to the real cohort (Table 5.1). In the future, newer approaches for generating synthetic data and creating dimensionality-reduced representations of complex structures could be explored, notably deep-learning methods such as autoencoders [118] and generative adversarial networks (GANs) [198]. In some studies, autoencoders have been shown to be superior to PCA-based methods (e.g. for 3D facial surface reconstruction) [199, 200]. Additionally, GANs have shown promise for creating images of synthetic patients afflicted with CHD [201]. Finally, instead of sampling a Gaussian distribution to find combinations of shape PCA parameters, other methods for creating new DNN training data may be more appropriate, such as Latin hypercube sampling [202].

5.5.2 Model performance

DNN-based models were observed to predict point clouds of both pressure and velocity flow fields with good accuracy, while being approximately 4,000x faster than the presented conventional CFD method. Node errors for pressure were seen to be larger at the inlet, while velocity errors were more skewed towards the distal regions of the aorta (Fig. 5.5). This is most likely due to the CFD boundary conditions at the inlet and outlet constraining the pressure/velocity variability at these regions. The further away from the aortic inlet or outlet, the greater the variability in velocity or pressure, respectively. It should be noted that there were no significant biases in either pressure or velocity predictions, suggesting that there were no systematic errors with the models (Fig. 5.5). In testing, it was found that the mean pressure subject error was slightly higher than the mean velocity error. This is despite the pressure PCA model capturing more variance than the velocity model. A possible explanation for the higher pressure errors is that the association between shape and pressure is more complex than that between shape and velocity in aortic domains. This is supported by the observation that shape modes do not correlate with pressure errors (Fig. 5.7). Interestingly, there was a strong negative correlation between shape mode 3 and velocity error. This suggests that more 'gothic' aortas (characterised by a more triangular arch) were less prone to velocity prediction errors. A possible explanation may be that the gothic arch constrains downstream flow patterns (where most velocity errors occur), hence making it easier for the model to characterise flow features associated to this subset of aortic shapes.

An important element of this study was to apply trained DNN models to prospective, unseen cases in order to explore the feasibility of performing inference on real patient data (n=10). It was shown that the average node-based prediction error ($MNAE_S$) on the prospective cases did increase for pressure from 6.01% to 10.19%. Velocity errors increased only marginally in the prospective test cases (3.99% to 4.47%). This aligns with the previous observation that the relationship between shape and pressure may be more complex due to the lack of an observed correlation between the shape PCA modes and pressure prediction errors. ML predictions were also compared to the real CFD (performed on the raw segmentation mesh) for all cases. This enabled the proportion of the total error due to the SSM and DNN to be estimated. It was shown that $\sim$60% of the total gradient error in pressure/velocity was due to the SSM, which further strengthens the argument for improved shape parameterisation.

The approach of using DNNs to model 3D aortic pressure and velocity flow fields has been described in other works [117, 118]. However, an important limitation of this data-driven approach is that flow fields cannot be assumed to satisfy the

Navier-Stokes equations for incompressible flow. This may hinder the applicability of the approach for simulations where accurate flow field prediction is insufficient in of itself, and conservation of mass and momentum needs to be guaranteed. Future development should also include models for computing velocity x, y and z components, allowing visualisation of streamlines or possible derivation of parameters such as wall-shear stress. Other studies have reported the use of long short-term memory (LSTM) networks or specialised architectures such as PointNet to build CFD-based ML models [116, 203, 204]. LSTM networks in particular may be highly suited towards any ML transient flow applications, due to their inherent ability to learn temporal sequences of data. Other architectures such as probabilistic DNNs which output uncertainty intervals during inference should also be explored [203, 205, 206]. Additionally, alternative methods for sampling synthetic data (e.g. Latin hypercube sampling) may produce a more diverse dataset for model training. This may prevent the occurrence of outliers, such as the worst pressure case in the test-set ($MNAE_S$=23.6%).

5.5.3 Potential clinical utility

The fast computation of hemodynamics using the proposed method may have multiple potential clinical uses. However, this is a proof-of-concept study and further improvements are required prior to any clinical validation (particularly the inclusion of patient-specific boundary conditions and time varying flow fields - see limitations). Nevertheless, several clinical usages may be possible if this is achieved, such as for supporting the identification of patients who need an intervention and predicting the outcome. Specific to the post-repair CoA population, several studies have shown that CFD can be used to evaluate abnormal hemodynamics (particularly during stress) and predict normalisation of hemodynamics after stenting of coarctation. However, this approach is rarely used in the clinical environment because it is so time-consuming. The approach could possibly be extrapolated to evaluate stress hemodynamics by simulating each training dataset case under elevated cardiac stress conditions. A second application could involve a fast and automatic pipeline to predict post-stenting hemodynamics in aortas. This could be implemented in a two-step solution using: (i) a surrogate finite-element model for predicting an ideal post-op aortic shape following stenting, and (ii) a surrogate CFD model for predicting hemodynamics on the post-op aortic shape (following the approach presented in this study). Use of such models would bring forth a new level of precision medicine that is currently lacking in congenital heart disease.

In this proof-of-concept approach, the quantities of interest (full 3D pressure and velocity fields) were selected in order to assess the model's ability to predict hemodynamics over the entire aortic domain for individual subjects. However, for specific clinical applications, such as non-invasive estimation of transcoarctation pressure gradients, a different modelling approach may be more suitable than the one presented in this Chapter. This is because for such an application, sampling a pressure gradient from a 3D reconstructed flow field may be less accurate/reliable than using a model that is tailor-made to predict only a pressure gradient. Another advantage of this method is that clinically acquired pressure gradient data could be used to validate the ML model. However, the possibility to infer 3D pressure/velocity (including potentially wall shear stress) has greater importance when applied to clinical applications other than CoA (e.g. aortic aneurysms). Furthermore, the design of an interactive and visual 3D tool synergises better with the idea of a virtual stenting/surgery pipeline for pre-operative planning.

In addition to the current proposed approach, physics-informed neural networks (PINNs) are being increasingly used to solve systems of partial differential equations. These are especially useful in scenarios where the boundary conditions are uncertain or the governing fluid flow equations need to be conserved [207]. Other advantages of PINNs include 'mesh-free' network architectures which can integrate sparse data from multimodality sources into the learning process, by using an experimental data loss in addition to the Navier-Stokes loss. However, in contrast to the data-driven approach presented in this Chapter, PINNs take long amounts of time to optimise the network weights and can be challenging to train. Additionally, each model needs to be trained per patient, therefore the approach does not allow for development of a single, generalisable ML model.

5.6 Limitations

In order to translate the proposed DNN-based CFD approach to clinics, there are two main modelling limitations which need to be bypassed. The first is related to the loss in surface accuracy when using SSM representations of aortic shapes. The second is the current simplicity of the CFD approach, which needs to be further developed in order to generate more meaningful DNN training data.

5.6.1 Shape Parameterisation

It has been seen in previous studies and in Chapter 4 that aortic CFD flow fields are highly sensitive to geometric and topological variation [208]. For this reason, using shape vectors that are accurate descriptors of the aortic surfaces is of critical importance. In the future, the relationship between the SSM registration error and the resultant CFD flow fields should be further investigated. Where possible, augmentations to the shape vector should be trialled in order to see if DNN prediction errors for real subjects can be reduced. This may require including additional shape information as inputs in the DNN models to act as a form of regularisation, such as a registration error or a geometric feature (e.g. centreline diameters). Such complimentary shape descriptors may be used to better inform the network of important features not fully captured by the SSM shape vector alone. Additionally, multiple SSMs based upon templates other than the mean aortic shape could be generated. These would enable closer non-rigid registration for unique cases where the target aorta deviates significantly from the mean shape. Recently, Wiputra et al. demonstrated methods for augmenting SSMs to enable the inclusion of head and neck vessel geometry within the aortic shape parameterisation, while retaining high accuracy with low-dimensional PCA vectors [209]. A similar approach may be employed in future studies in order to be able to fully describe the patient-specific aorta with head and neck vessels. Of course, a simple initial improvement could be made by adding more subjects to the SSM to introduce more variability in the population.

5.6.2 Computational Fluid Dynamics

In this study, a simplified CFD pipeline was chosen in order to easily automate and ensure convergence for numerous simulations (n=3,000). A standard CFD solver set-up was used, assuming steady-state conditions and incompressible flow. In order to account for the pulsatility involved in aortic flow, a transient solver may be better suited to real-life applications. It was also assumed that the flow through the great arteries at peak systole was laminar [210]. In the future, the Reynolds number may be computed for individual cases to allow for the inclusion of turbulence modelling where necessary. However, for complex morphologies (such as CoA), Reynolds number has been seen to be an inconsistent measure of turbulence [211]. Feiger et al. proposed an alternative solution to turbulence modelling when performing CFD on a large scale, which involves meshing the domain with extremely high numbers of nodes [117].

Boundary conditions selected included a fixed, flat velocity inlet and zero pressure outlet condition for all cases. However, idealised inlet velocity profiles (flat, parabolic etc.) have been shown to be ineffective for producing clinically relevant data [212, 213]. Thamsen et al. showed that a synthetic aortic population could be created with realistic accompanying 4D CMR-derived vector flow profiles [214]. In the future, a similar approach could be taken, allowing for an additional velocity vector field input parameter into the ML model. Alternatively, a more accessible approach could involve the use of a parabolic velocity inlet condition and a patient-specific unsteady flow profile (derived from phase-contrast CMR) in conjunction with a transient solver for resolving the peak systolic flow field. In this study, it was decided to omit the head and neck vessels from the CFD model for simplicity, however this would be required when aiming to simulate realistic patient-specific aortic hemodynamics [208]. Indeed, Wiputra et al. showed that accurate modelling of the head and neck vessels is necessary for capturing local flow features in the arch and producing realistic downstream fluid flow forces [209]. Thus, the inclusion of head and neck vessels along with lumped parameter outlet models such as Windkessel models should be explored in the future [215,216], as modelling downstream resistance has been seen to produce more clinically meaningful results [217–219].

5.7 Summary and conclusion

In this proof-of-concept exploration, a pipeline for building ML-based models to perform repetitive vessel-based CFD tasks has been proposed. Generation of synthetic aortic training data by means of shape modelling allowed ML techniques to be used, even where data scarcity is an issue (n=67). Point correspondence was maintained between subject meshes in order to enable PCA. ML models were able to compute pressure and velocity flow fields much more rapidly (4,000x) than traditional CFD solvers, without large computational requirements or simulation setup. Comparison between predicted and ground truth test cases revealed good overall performance. Testing on prospective cases revealed that shape registration errors could produce misleading flow fields which deviated significantly from the 'real' CFD result, even in the presence of low ML errors. The approach described in this study is shape-driven and is applicable to any vascular structure which can be segmented from medical images. In the future, as further explained in Chapter 9, the models should be improved, so they can perform inference on prospective data from real patients. The requirements for this are two-fold; improving the accuracy of the shape representation methods while incorporating the head and neck vessels, and using a more realistic CFD pipeline for generating training data with the inclusion of patient-specific boundary conditions. Comparison of the ML models against clinically acquired data (such as catheter-based pressure drops) should be also performed in the future for validation purposes. Results from the CFD automated pipeline and from possible CHD simulated treatments could be visualised, together with the patient specific anatomy automatically segmented using the algorithm proposed in Chapter 4, in the VR platform that will be described in the following Chapter, to further enhance care in CHD.

Chapter 6

VheaRts: a virtual reality platform for supporting treatment and teaching of congenital heart disease

6.1 Introduction

Despite great potential, as discussed in Section 2.2.2 and Section 3.5.4, VR has not yet fully transitioned into clinical practice, neither for pre-operative planning of complex cases nor in educational settings. At the time of my project conceptualisation, VR software solutions bespoke for the visualisation of congenital abnormalities were sparse and not fully compliant with the needs identified at our clinical centre. Thus, in order to exploit the uses of VR for CHD, I worked on the development of a flexible in-house VR platform that could be fitted with tailored features to respond to the specific demand of care and education in CHD.

6.2 Aims and objectives

The aim of the VR platform was to provide clinicians and students with a user-friendly tool to enhance understanding of CHD. This was achieved by incrementally implementing specific features, taking into account the direct feedback from specialists in CHD surgery/clinical care and education/training. The overarching needs inspired the application to include: (i) easy to use interactive capabilities with patient-specific models via hand controllers, (ii) a collection of specialised tools for enhancing spatial and anatomical understanding of cardiac anatomy, and (iii) a multi-user feature implemented as a room-based system to enable participants to create/join online VR sessions, simultaneously.

6.3 Development tools

6.3.1 Unity

Unity is a conventional 3D game engine that has become popularised in recent years for the development of VR experiences, and follows an object-oriented programming (OOP) approach. Objects in this sense can contain data and/or functions, and can derive these properties from other objects or classes in a process known as *inheritence*. Since OOP is a highly modular approach by design, code and data can be easily reused within a project, allowing multiple objects in a scene to share the same or similar behaviours. In Unity, 'MonoBehaviour' is the base class which almost every script inherits from. MonoBehaviour contains a number of classes and functions that facilitate the development process such as: code for storing mesh information, positional information, computing physics interactions, collider detection and more.

Within Unity, assets such as images, videos, mesh files and more can be imported. The Unity Editor grants a graphical user-interface (GUI) for the developer to assemble a scene, change settings, write scripts and more. The majority of scripts within Unity are written in the language C#, and are appended onto objects in a scene (following OOP). Shaders (i.e. scripts that act in the rendering pipeline) are typically written using high-level shader language (HLSL). Following assembly of a scene (containing meshes, objects with behaviours, lighting sources etc.), the application can be started by simulating the game-loop. Physics, collider interactions, logic, rendering and more are updated once every frame, following a specific execution order.

6.3.2 Oculus software developer kit

Unity includes third-party software development kits for Oculus hardware (Section 3.5.2) in order to allow interfacing with Oculus devices. Some of the functionalities provided in the Oculus SDK include:

- integration of Unity with the device headset/controllers, in order to register inputs and send outputs
- scripts for managing VR-based interaction with objects in the Unity scene
- scripts for raycasting and creating VR-compatible GUIs
- prefabricated assets (prefabs) which allow for visualising the controllers/hands
- scripts to manage audio and voice interaction

6.4 3D data pre-processing

To generate 3D heart meshes for viewing within VheaRts, cross-sectional patient images require pre-processing prior to importing. The first step involves segmentation of the 3D CT or CMR dataset, the choice of which is dependent on the existing imaging availability for the patient (see Section 2.3.5). Due to the higher spatial resolution offered, CT is generally preferred. For segmentation of CMR, a 3D whole-heart SSFP sequence is typically used. All images are acquired in DICOM format. Software used to generate the masks of the images includes ScanIP (Synopsys Simpleware) and Mimics (Materialise), and segmentation is typically performed using semi-automatic threshold functions. Manual editing is often necessary, e.g. when ensuring structures such as the atria/ventricles are properly separated from one another. For most applications, the blood pool of the images is segmented, omitting the myocardium unless specifically requested for clinical assessment. Surface meshes typically require further editing, which is performed in Meshmixer (Autodesk), and includes common operations such as re-meshing, removal of artefacts and cropping of unnecessary sections. In most cases, a negligible and artificial outer thickness of 0.1 mm is added to surface meshes by duplicating the mesh in the outward direction of the surface normals. This operation is typically done for 3D printing, but in VheaRts also improves the visibility of mesh boundaries when using the clipping tool (see Section 6.5.2), although this is optional. For compatibility with Unity, mesh files required conversion from .stl format to .obj, which also supports multi-object meshes, allowing for hearts composed of subcomponents in the application. As explained in the later sections, the file size of heart meshes is typically kept under $\sim$50 MB for files imported within Quest 2 devices, although HMDs tethered to a workstation can support meshes of sizes over 1 GB (depending on processing power).

6.5 Core functionalities of VheaRts

The in-house VR platform, referenced in the context of this thesis as 'VheaRts', underwent a process of continuous development and functional optimisation over the timeline of my PhD project. Core functionality developed for VheaRts was used to form two main applications of the software, one tailored for clinical decision-making of CHD and the other tailored towards educational needs in CHD. Features were developed to grant users a wider range of possible interactions with the 3D heart models. Most of the following tools/functions were accessible across both educational and clinical versions of VheaRts, along with Android/Desktop builds for Oculus Quest/Rift (Section 3.5.2). [1]

6.5.1 Model grabbing and handling

VR model handling is the most intuitive and important of all interactions, being fundamental in almost any modern VR experience which utilises hand-tracked controllers. The functionality of this feature is provided within Oculus' Unity SDK, and its implementation relies on two main scripts from the library: *OVRGrabber* and *OVRGrabbable*, which distinguish target colliders (object) from user colliders (the hands) and manage interactions between the two. Objects can be grabbed with either hand (by holding down the trigger button), and also translated and rotated intuitively. This allows for natural manipulation of heart models.

[1] Video of primary VheaRts functionalities: `https://vimeo.com/831949370`

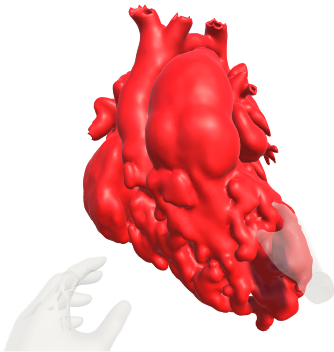

Figure 6.1: Grabbing/handling of a patient-specific heart model using two individual VR hand controllers.

6.5.2 Shader-based graphical clipping tools

Real-time object clipping allows users to dynamically crop 3D models for an unrestricted investigation of intracardiac anatomy. The tool was implemented with shader code (HLSL) which allows Unity to dictate how an object is rendered by the GPU. With the clipping tool, the user can dynamically hide sections of the model by passing a planar or spherical 'clipping object' into the target mesh. The clipping tool can be translated/rotated using the hand controllers. Sphere size can be adjusted using the user-interface. The position/orientation of the clipping plane/sphere respective to the heart can be fixed, allowing the user to 'lock' specific cuts in place.

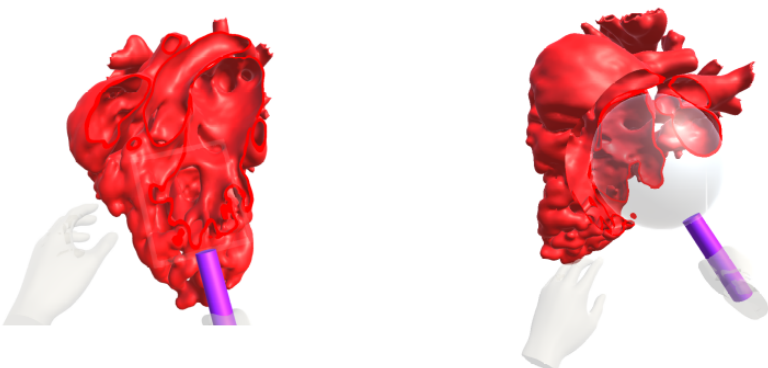

Figure 6.2: Left: plane-based clipping. Right: sphere-based clipping. Bright red areas indicate gaps in the mesh wall where 'false thickness' has been rendered.

The shader for the plane clipper runs by looping through each vertex in the mesh, and calculates the dot product of the vertex-plane vector and the plane's normal vector. If the result is positive, the vertex is discarded from the render process and not displayed. If the result is less than or equal to zero, this is rendered. For the sphere slicer, a similar approach is taken, except the vertex is discarded if the vertex-sphere distance is greater than the sphere radius, since it implies the vertex is not within the bounds of the sphere.

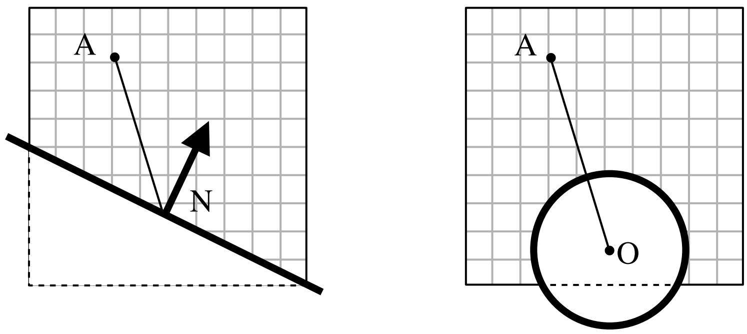

Figure 6.3: Left: plane-clipping. Dot products between plane normal (N) and each vertex-plane vector determine if a vertex is rendered (e.g. vertex A is rendered). Right: sphere-clipping. Vertices outside the bounds of the sphere radius are rendered (e.g. vertex A).

In some heart meshes, the myocardium or other types of wall thicknesses may be present. In order to visualise a 'thickness' within the gap between the two shells, two render passes in the shader are required. The first pass renders only the back-facing triangles as a flat, unlit colour with no illumination or shadows. The second renders only the front-facing triangles (with normal illumination). This combined effect results in the mesh appearing as it normally would, until the user clips the model and views the back-faces (which are unlit), giving the impression of 'thickness' or muscle mass.

6.5.3 Virtual echocardiogram probe

A virtual echocardiogram tool enables users to practice generating ultrasound-like 2D slices from a 3D cardiac object. The VR probe was designed by attaching an additional camera in the scene to an interactive/grabbable probe handle.

During runtime, the camera's output is converted into a texture. The texture is updated constantly and overlaid onto a two-dimensional conical sprite, in order to emulate the projection of an ultrasound. The camera has an infinitesimally short depth for creating a flat 2D projection. Post-processing filters including noise,

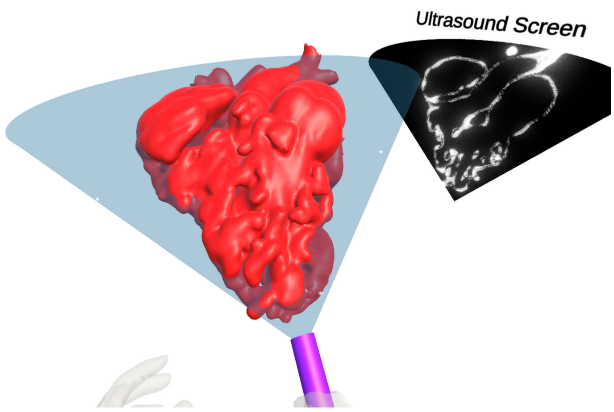

Figure 6.4: Echocardiogram tool, with projected four chamber view on the right.

monochrome and bloom were applied to create a stronger likeness to conventional ultrasound imagery.

6.5.4 Marking and measuring

Functionality for inserting markers and measuring distances allows users to highlight structures or assess diameters, lengths and sizes of cardiac structures.

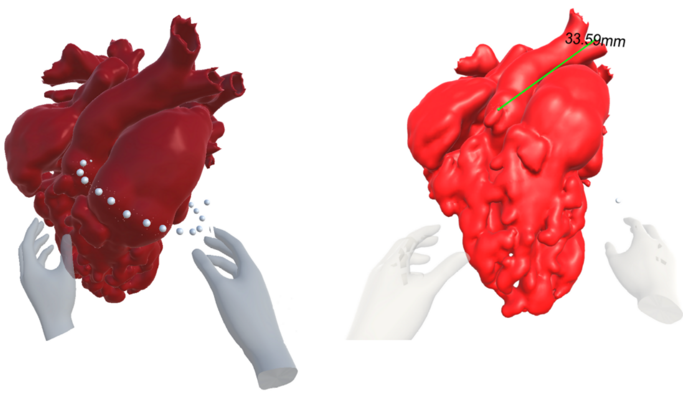

Figure 6.5: Left: markers placed on a heart model. Right: linear distance between two points positioned on the aortic arch of the patient-specific model.

Marker placing was implemented by instantiating primitives (in-built Unity meshes) at runtime. Markers attached to heart objects would remain in the same respective positions. Previously placed markers could be deleted. Distances between two markers could be measured using a function that converts Unity's standard distance units into millimetres, while accounting for any scaling performed to the mesh within the application.

6.5.5 Mesh slicing

The ability to cross-cut and section the heart in any plane enables users to subdivide the heart with any desired configuration of slices. The new meshes can be then grabbed/interacted with, similarly to regularly imported objects.

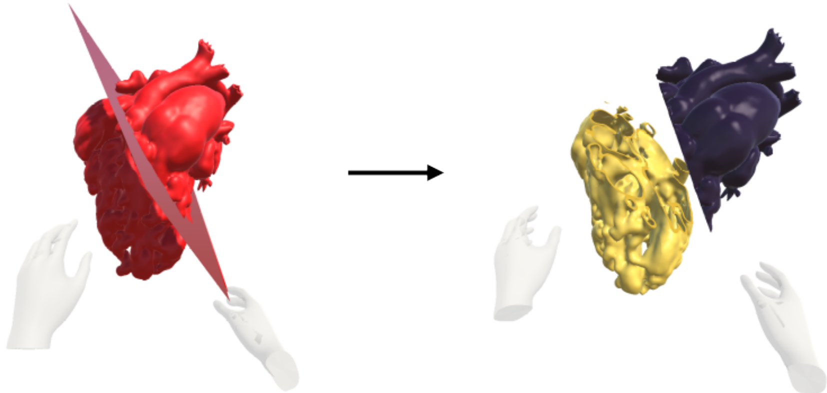

Figure 6.6: Left: heart model (pre-cut) with user deciding plane orientation. Right: heart model (post-cut) split into two separate meshes.

The slicing algorithm involves a number of steps. The first is to define the desired plane using three points in space. These points are recorded by the user with a laser pointer, setting the plane position (Fig. 6.6). Following this, a for loop checks the vertices of each triangle and identifies whether it is left/right-sided from the plane, depending on the computed signed distance (dot product). Triangles which are identified to lie on the axis of the plane have their intersection points computed. This is performed for each triangle edge with a plane-line intersection formula [137]. For a plane with normal n and offset D from the origin, the intersection vector I on a line given by two points AB is given as:

$$I = A + \left(\frac{D - n \cdot A}{n \cdot (B - A)} \right) \times (B - A) \tag{6.1}$$

Following identification of a plane-triangle intersection, the polygon is split and/or sorted depending on the orientation of the intersection. Five main types of intersections were identified, with only two cases necessitating re-triangulation (Fig. 6.7). To maintain consistency, the vertices are always passed into the retriangulation function in clockwise winding order. In fringe situations where the cut lies extremely close to a vertex (intersection point is $<1\%$ edge length), the triangle may be sorted instead of undergoing splitting and re-triangulation, following cases 3 or 4 (Fig. 6.7). Once the new triangles are constructed, they are sorted into the new mesh collections depending on what side of the plane they lie on. Limitations of this approach include the possibility of generating highly skewed triangles if the intersection line is very close to the border.

Algorithm 4 ComputePlane

```
% User draws a plane P in 3D with the laser pointer
% Save any three points (p1, p2, p3) that lie on the plane P
p1, p2, p3 = points[0], points[1], points[2]
% Compute two vectors that lie on the plane
p3p1 = p3 - p1
p2p1 = p2 - p1
% Compute the cross product of p3p1 and p2p1 for the plane normal n
n = Cross(p3p1, p2p1)
% Compute offset of plane (d) using any point on the plane
d = Dot(n, p1)
return n, d
```

Algorithm 5 GetSideOfTriangle

```
% Get the signed distances of all triangle vertices from the plane n
for Vertex v in triangle t do
    signedDistance = Dot((v-n), n)
    distances.Add(signedDistance)
end for
return distances
```

Algorithm 6 GetIntersectionPoints

```
% Compute the locations of the triangle/plane intersection points (if they exist)
for Edge e in triangle t do
    % Use line-plane intersection formula (Eq. 6.1)
    i = IntersectionEquation(e, n, d)
    if intersection is found then
        break
    end if
end for
return i
```

Algorithm 7 Retriangulate

```
% Build new triangles from vertices and intersection points (Fig. 6.7)
if Case 1 then
    % Preserve winding order of triangles
    newTri1, newTri2, newTri3 = [A, AB, AC], [AB, C, AC], [AB, B, C]
    % Return as left/right sided depending on signed distance
    return newTri1, newTri2, newTri3
end if
if Case 2 then
    newTri1, newTri2 = [A, B, BC], [A, BC, C]
    return newTri1, newTri2
end if
if Case 3 then
    % Does not require retriangulation
    return t
end if
if Case 4 then
    % Does not require retriangulation
    return t
end if
if Case 5 then
    % Does not require retriangulation - duplicate in both left/right arrays
    return t, t
end if
```

Algorithm 8 Mesh cutting algorithm

```
n, d = ComputePlane()
for triangle t in mesh do
    distances = GetSideOfTriangle(T)
    if t is entirely left-sided then
        leftTris.Add(t)
    end if
    if t is entirely right-sided then
        rightTris.Add(t)
    end if
    if t not entirely right or left sided then
        % Determine t intersection points
        i = GetIntersectionPoints(t, n)
        % define new triangles and add them to corresponding left/right arrays
        newTrisLeft, newTrisRight = Retriangulate(t, i, n, d)
        leftTris.Add(newTrisLeft)
        rightTri.Add(newTrisRight)
    end if
end for
```

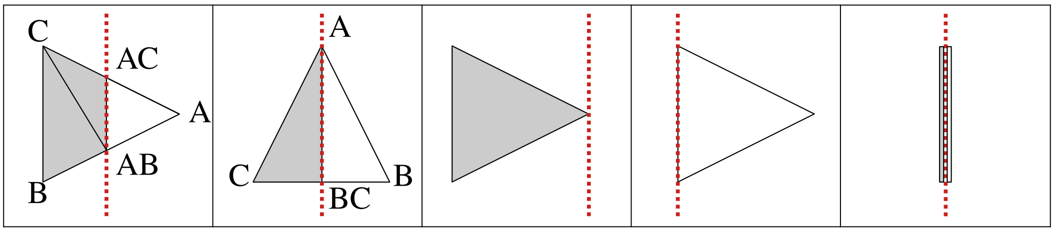

Figure 6.7: The five main triangle-plane intersection conditions. The red dotted line indicates the cutting plane. The intersection conditions from left to right are: (i) two edges, (ii) one vertex and one edge, (iii) one vertex, (iv) one edge, and (v) all three vertices.

6.5.6 Mesh surface painting

Surface mesh painting to grant users the ability to highlight and/or delete specific areas on the surface of the mesh in VR, thus providing further mesh-editing functionality. Newly generated meshes can be exported in an .stl format. A selection of different 'brush sizes' are accessible with the tool.

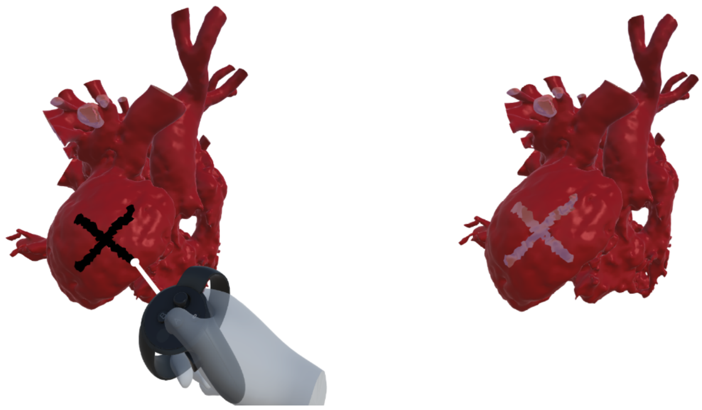

Figure 6.8: Left: user selecting and painting mesh surface elements by raycasting with laser pointer. Right: heart after selected elements were deleted.

Implementation of this tool is mesh-based rather than shader-based. During initialisation, a duplicate mesh of the target heart is generated with empty triangles (not visible). With the use of a laser pointer, a raycast 'hit' is registered on the target mesh. Following this, a function identifies which triangles were selected and activates them in the secondary mesh. Since the secondary mesh has a higher priority in the rendering queue, triangles appear to overlay the original mesh, acting as a mask.

When the user decides to delete any highlighted triangles from the mesh, active

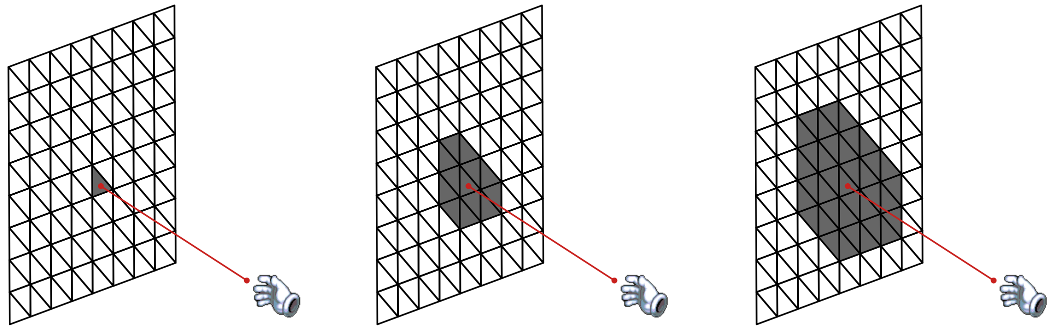

Figure 6.9: Mesh surface element painting with increasing brush-sizes. Triangle connectivity is used to determine element selection patch.

triangles in the secondary mesh are removed from the original mesh. The user has the ability to change the 'paint brush' size by parsing through a mesh connectivity dictionary and highlighting nearby triangles (Fig. 6.9).

6.5.7 Miscellaneous features

In addition to the aforementioned tools, other functionalities have been developed for VheaRts. These included:

- tutorial videos for each tool which can be viewed in-app
- scaling (or resizing) of imported 3D models using a GUI
- generating screenshots from the VR user perspective
- recording videos from the VR user perspective
- annotation and colouring of meshes with multiple subcomponents
- interactivity with individual mesh subcomponents for creating an 'exploded' view
- importing of patient ultrasound or CT/CMR images
- animations for cases with multiple mesh frames, acquired typically from 4D-CT sequences

6.6 Networking implementation

Thus far, all VheaRts described features are accessible in the offline, single-user educational and clinical platform versions. In order to enable online and multi-user interaction while preserving the same functionality, modifications and changes were necessary.

6.6.1 Photon unity networking

Photon unity networking (PUN) is a package for adding multi-user functionality into Unity applications. It comes pre-built with scripts and behaviours that enable individual players to access shared rooms, where objects are synced over the network. PUN offers servers so that players can connect over the cloud, removing the need for managing end-to-end connections. Players connect to a master server containing a list of identifiers for various regions. When PUN determines the region of choice (or lowest latency region by default), the user is guided via the region's master server into a game server. Creation/joining of individual rooms is conducting and hosted on the game servers.

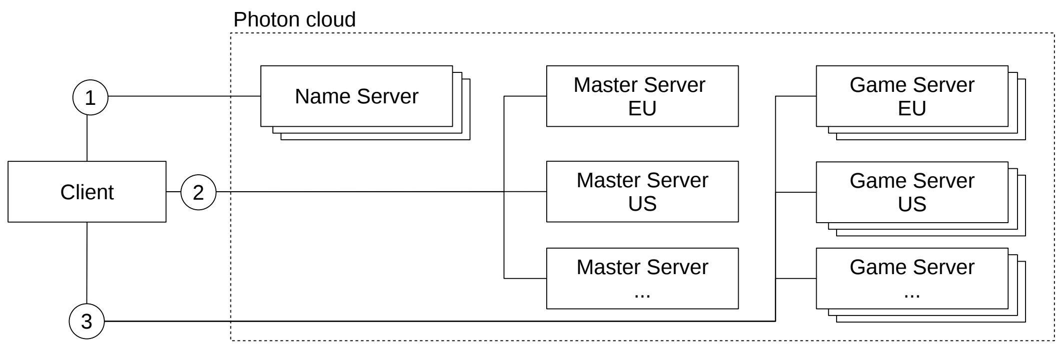

Figure 6.10: PUN cloud connectivity system.

An application programming interface (API) within PUN provides implementations of a number of networking functions, allowing for much of the base functionality of multiplayer experiences to be quickly built-up (e.g. creating rooms, joining rooms). Generic objects in the scene can be converted into networked objects with the use of a 'Photon View' script. Ownership of objects can be transferred to other players, allowing more than a single person to control/interact with an object. Remote procedure calls (RPCs) can be conducted over the network, allowing users to execute commands on other clients' applications via the server. This enables pre-existing functionality to be modified for supporting multi-user capabilities. For example, actions such as toggling a model on/off could be performed using an RPC so that the scene is synchronised between clients. Audio transfer and in-built avatars are also included within the PUN package.

6.6.2 Online rooms

VheaRts applications were fitted with multiplayer functionality in order to allow users access to shared rooms. Two classifications of online lobbies were assigned: a classroom-type lobby (one tutor and several students), or a clinical lobby (all with equal permissions). Before entering a classroom-type lobby, setting the level of permissions (student or tutor) was required. This allowed only one tutor to be active in the scene, which was important in classroom contexts, since it granted only the course coordinate access to the GUI for toggling tools and models. This streamlined the teaching process by preventing other users from modifying settings. The flowchart for entering/creating rooms when a user loads into a classroom-type lobby can be viewed in Fig. 6.11. For clinical lobbies, all users were granted full permissions.

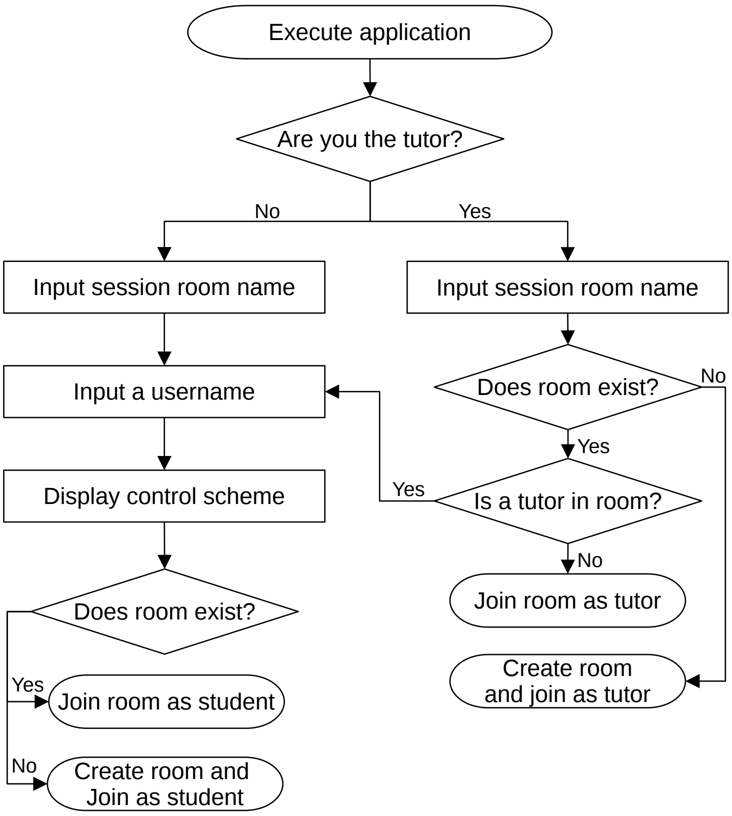

Figure 6.11: The flowchart for users accessing the online component in an online classroom-type lobby.

Each user is granted a digital avatar upon loading into a room, composing of head and hand models. Usernames can be entered and are displayed over each avatar. Voice communications using the headset in-built mic and headphones were enabled. Each tool previously described in Section 6.5 had online-compatibility implemented using a combination of RPCs and pre-fabricated PUN scripts to ensure functions could be correctly synced over the network. Transfer of ownership for some objects (e.g. ultrasound probe, clipper) was made possible to allow users with lower permissions (students) to also interact with tools and grab heart models. Locomotion was enabled for each player, allowing users to reposition themselves in the room with the controller joysticks. All users could toggle a laser pointer to highlight or indicate structures.

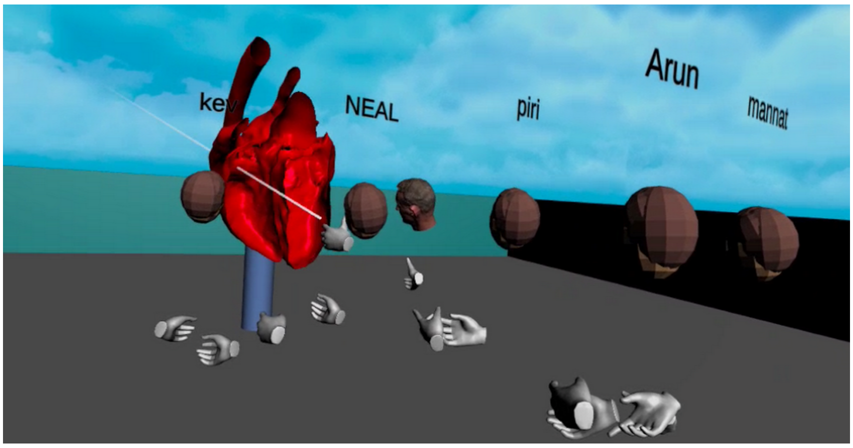

Figure 6.12: Example session of a VR multiplayer environment in the educational application. The tutor (pointing at heart) can be seen using a clipped model to teach cardiac anatomy in a classroom setting.

6.7 VheaRts for clinical decision-making

VheaRts for clinical decision-making was built by assembling the previously described tools for an application in which 3D models of the heart could be loaded during runtime for VR assessment. All core and networking functionalities were included except for those not relevant to clinical usage (e.g. labelling of mesh subcomponents). Builds for both PC (Windows) and Android (Quest devices) were developed, in order to be able to target tethered Oculus headsets (Rift) and standalone headsets (Quest). Functions for the serialisation and deserialisation of meshes were included. This enabled models to be saved in a persistent library, avoiding the need to re-import models each time the application is halted.

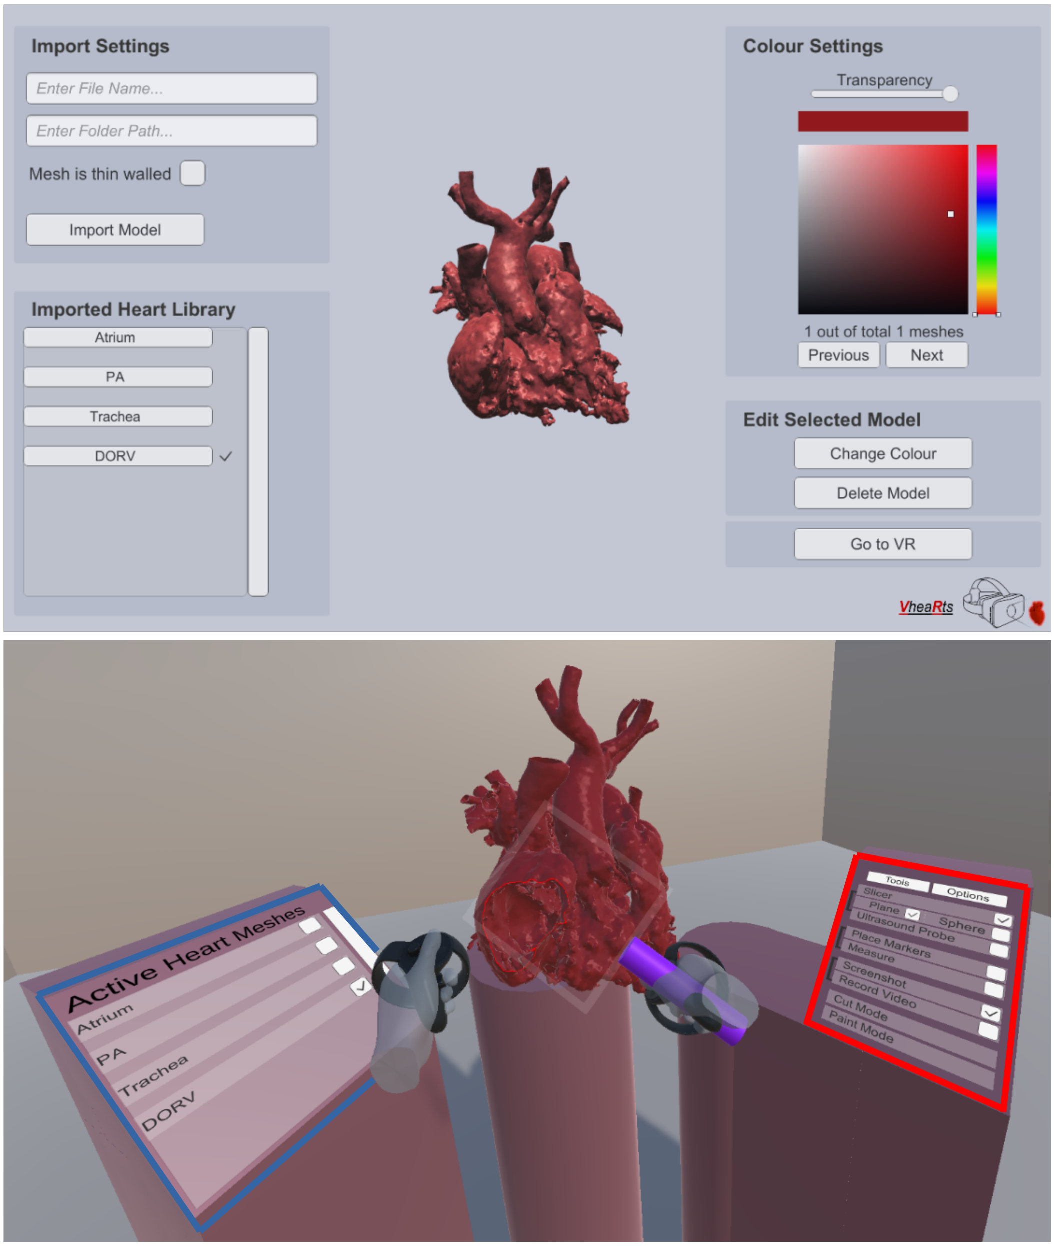

Figure 6.13: Top: Desktop GUI for importing models in the clinical application (a VR-based GUI is available for the Quest build). Bottom: User-perspective during clipping tool interaction in the VR environment (blue: GUI for toggling models, red: GUI for toggling tools/settings).

Fig. 6.13 presents the desktop version of the VheaRts clinical application. After initialisation, the user may alter settings or import new models via a GUI (Fig. 6.13, top). In the desktop version this is performed with mouse/keyboard, whereas for the standalone Quest build the user uses a laser pointer to make selections in an in-VR GUI. Features of the GUI included: (i) text fields to be able to point to a file location (ii) a button for executing the import command (iii) a list to be able to select models, (iv) a button to delete selected meshes, and (v) a canvas for selecting the colour of individual submeshes. An option for flagging the imported file as a thick-walled mesh was needed in order for assigning shader properties so that the gaps between shells are rendered as full (e.g. the myocardium). The shader was discussed in greater depth in Section 6.5.2. From the desktop GUI, access to VR is granted using a UI button in the bottom-right corner, with online functionalities accessible in-VR via the settings menu (Fig. 6.13).

6.8 VheaRts for education

In the educational version of VheaRts, interactions with models, user-interfaces and tools were mostly unchanged from the clinical application. The most significant differentiating feature was the replacement of runtime mesh importing with a selected library of curated CHD models for teaching. The selection of models is described in Chapter 8. Additionally, some specific features required implementation in the educational application. These functions included: the possibility to toggle on/off mesh labels and the ability to have dynamic, ‘animated’ hearts by using a script which loops through multiple meshes in a patient case. Networking capabilities were also included for the educational app version. However, the implementation differed from the clinical version in order to more closely simulate a conventional anatomy lab through the use of ‘classroom-type lobbies’ (see Section 6.6.2).

6.9 Summary and conclusion

A framework for a VR platform (VheaRts) consisting of a specific set of features for teaching and planning treatment of CHD has been outlined. The general software architecture and development approach was described, along with the implementation of individual tools, including the scripting approach and any relevant algorithms used to create the feature. The configuration of the application to support networking capabilities through server-hosted online rooms was described. The final assembly of two individual applications - clinical and educational - was outlined and examples of their adoption in the respective areas will be presented in the next two chapters, together with studies carried out to gain evidence on the advantages of the use of VR for clinical care and education in CHD.

Chapter 7

VheaRts for clinical planning

7.1 Introduction

The overall potential of VR for clinical applications in CHD was discussed in Section 3.5.4, including the remaining challenges of translating VR into actual clinical practice as recognised from the literature. In this Chapter, I report the overall experience related to the integration of VheaRts (see Chapter 6) into clinical pre-operative planning for CHD at Great Ormond Street Hospital for Children and at an external centre, Bristol Royal Hospital for Children, from 2019 to date. The first part of the Chapter lists the CHD cases where VheaRts was employed for planning complex treatments (Section 7.2), the second section focuses on three examples of a specific lesion, DORV (Section 7.3), the third part describes a retrospective study to assess the advantages of VR compared to other image modalities (Section 7.4), and the final part describes the use of VheaRts beyond CHD (Section 7.5). The aims of this part of the project were to: (i) explore the possibilities for supporting complex surgical pre-operative planning within CHD and (ii) capture the potential benefits and disadvantages of VR vs other 3D imaging modalities.

7.2 Overview of VheaRts clinical cases of CHD

The use of VheaRts was requested on a case-by-case basis by clinical practitioners from Great Ormond Street Hospital for Children and the Bristol Royal Hospital for Children for pre-operative planning of a complex CHD cases. Following the creation of the 3D anatomical reconstruction for each individual case by conventional segmentation techniques (Section 2.3) of routine diagnostic imaging, users were able to interact with the patient specific 3D models using the tools described in Chapter 6. VR interaction was typically self-guided, and cross-sectional medical images (CT/CMR) were available on a separate screen as part of the VR patient assessment. During the period 2019-2023, a total of 48 patients were evaluated using VheaRts at Great Ormond Street Hospital and 8 patients at Bristol Royal Hospital for Children. Thus, a total of 56 cases have been reviewed (combining both clinical centres) using VheaRts, with primary diagnosis shown in Table 7.1.

Case	Primary diagnosis	Case	Primary diagnosis
1	DORV	29	Multiple associated complex lesions
2	DORV	30	AOCA
3	DORV	31	AOCA (intramural)
4	DORV	32	AOCA (intramural)
5	DORV	33	Complex heterotaxy
6	DORV	34	Complex heterotaxy
7	DORV	35	Double aortic arch
8	DORV, HLHS	36	Hypoplastic aortic arch
9	DORV (TGA-type)	37	Truncus arteriosus
10	DORV (TGA-type)	38	Aortic valve dysplasia
11	DORV (TGA-type)	39	Shone complex
12	DORV (TOF-type)	40	Complex LVOTO
13	DORV (TOF-type)	41	PDA (type B)
14	DORV (TOF-type)	42	PPVI stent implantation
15	DORV (complex)	43	Ebstein's anomaly
16	Cardiac tumour	44	Double mitral valve dextrocardia
17	Cardiac tumour	45	Double inlet left ventricle
18	Cardiac tumour	46	TGA
19	Cardiac tumour	47	HLHS
20	Cardiac tumour	48	CCTGA
21	Cardiac tumour	49	CCTGA
22	AVSD	50	CCTGA
23	AVSD	51	CCTGA
24	VSD	52	HLHS
25	VSD	53	HLHS
26	ASD and sinus venosus	54	MAPCAs
27	Multiple associated complex lesions	55	Right atrial isomerism, TAPVC
28	Multiple associated complex lesions	56	Obstructed TAPVD repair

Table 7.1: Compiled list of cases (n=56) reviewed with VR. The last 8 cases (highlighted) were managed and performed at the external clinical centre.

Overall, cases with conotruncal lesions (see Section 2.1.3) as primary diagnosis formed 39% of the referrals. The most frequently requested lesion (Chapter

2) was DORV (27%), followed by cardiac tumour/fibroma (11%); TGA-type and TOF-type DORVs each constituted 20% of the total DORV cases. Cases with a primary query related to a septal defect (AVSD, ASD, VSD) formed 9% of the referrals, although these defects were much more common as an associated lesion in other diseases (e.g. in all DORV cases). Cases were deemed to be with 'multiple associated complex lesions' if 5+ abnormalities were present with no clear primary diagnosis. Arterial lesions as a primary diagnosis (see Section 2.1.4) constituted 7% of the referrals. CCTGA was observed to be the most frequently requested lesion at the external centre (38%). Notably, no DORV or cardiac tumour cases were requested at the external centre. The only lesions requested in both centres were CCTGA (internal n=1, external n=3) and hypoplastic left heart syndrome (HLHS) (internal n=1, external n=2). Prolonged testing of the VheaRts clinical application in the external centre will be needed in order to compare both populations in greater detail. Finally, cases recorded included some patients with highly unique lesions, such as anomalous origin of the coronary arteries (AOCA) or patients with cardiac tumors (Fig. 7.1). Almost all cases were referred from surgeons/cardiologists with the intention to evaluate the anatomical arrangement of the patient cardiac structures prior to surgical intervention. In these instances, model handling interactions and clipping tools were most frequently used.

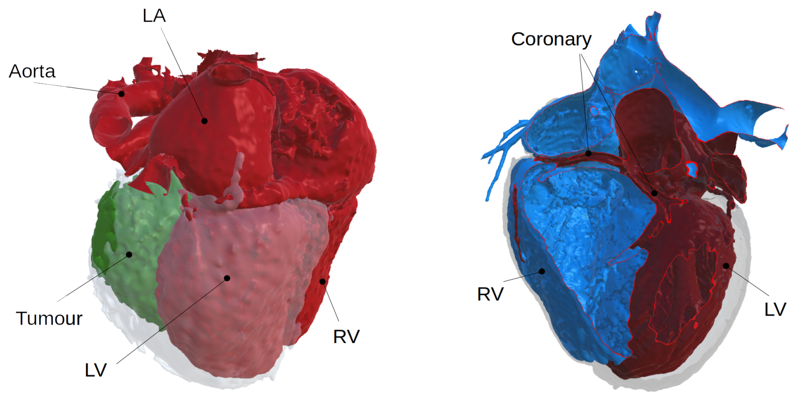

Figure 7.1: Left: rare case of a tumor embedded within the myocardium in VR (red: heart, green: tumor, translucent: myocardium). Right: VR model of a patient with an anomalous intramural coronary course (blue: right side of heart, red: left side).

One request only was received from an interventional cardiologist to plan a case of percutaneous pulmonary valve implantation (PPVI) in the presence of possible risks of compressing the nearby coronary artery (Fig. 7.2). This resulted in a unique outcome where VheaRts was used as a method to visualise a set of structural simulation results. The implantation of two stents (SAPIEN, Edwards and Melody, Medtronic) was simulated at different levels of expansion and at different locations using the software Abaqus (Dassault Systèmes, SIMULIA corp). The results were visualised in VheaRts and used to better understand how the impingement of the stented PA root on the coronary artery varied with the stent type/placement/expansion.

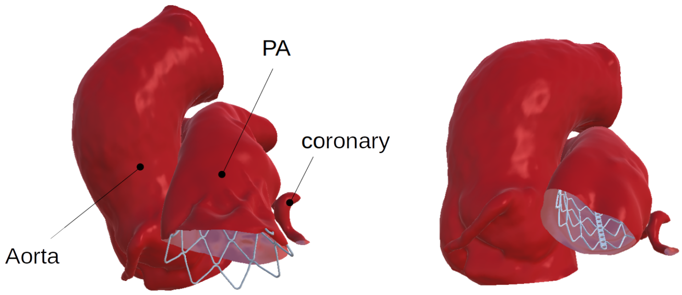

Figure 7.2: Left: PPVI simulation with a Melody stent expanded at a diameter of 18 mm. Right: PPVI simulation with a SAPIEN stent expanded at a diameter of 17 mm. All models were rendered in VheaRts.

At the end of the first year of VheaRts deployment and testing (2019-2020), clinical information related to the patients that were pre-operatively assessed with VR (n=14) was retrospectively collected and analysed, together with feedback from the clinicians/surgeons who had used the application (n=6) via a short survey. No early mortality was recorded, and mean operative time, cardiopulmonary bypass time, aortic cross clamp time, ICU stay, ventilation time and hospital stay length were similar to those reported from historical data without the use of VheaRts. Further prospective analyses on larger cohorts will be required in order to establish the possible value of VheaRts for surgical planning. In terms of clinical feedback for the use of VheaRts, respondents had an average experience of 17.5$\pm$9.6 years of specialisation within their field. The majority of users (67%) had tried VR prior to VheaRts. Following VheaRts testing, 83% agreed strongly or very strongly that the application and tools were intuitive. Feedback indicated that model handling was observed to be the most intuitive and useful tool, shortly followed by clipping. All users stated strong interest in using VheaRts again for clinical pre-operative planning of CHD. Examples with more details and specific surgical approaches are reported in the following Section.

7.3 VheaRts for planning DORV repair: case studies

As described in Section 7.2, DORV was the most commonly referred lesion for VR assessment (27%, n=56). From 2019-2022, the surgical team reported three patient cases in detail where VheaRts was used for pre-operative assessment. These are presented with the aim of showing: (i) the general approach that was taken using VR, (ii) how VR contributed to the formulation of the pre-operative plan, and (iii) the potential applicability of VR in DORV.

7.3.1 Case report 1: planning biventricular repair

The patient presented with a postnatal diagnosis of DORV with side-by-side great arteries, large non-committed VSD and a small muscular apical VSD. Previous palliation included a PA banding. Cardiac CT confirmed the presence of a significant subpulmonary obstruction with RV hypertrophy. The 3D model was reconstructed from cardiac CT.

As part of the pre-operative planning phase, the surgeon used VheaRts to assess the intracardiac anatomy of the patient from both conventional (axial, sagittal, coronal) and non-conventional views. The graphical clipping slicer was used to evaluate the severity of the subpulmonary obstruction and the feasibility of a biventricular repair. The spatial relationship between the VSDs and the great arteries was evaluated.

The patient underwent successful intracardiac biventricular repair at 10 months-of-age with VSD enlargement, division of mid-cavity RV muscle bundles and intraventricular tunnelling through the VSD to the aorta. Postoperative echocardiography showed a patent LV to aorta tunnel, with laminar flow and no residual VSD. Postoperative CT was performed and segmented, allowing additional evaluation of the surgical outcomes in 3D (Fig. 7.3, right).

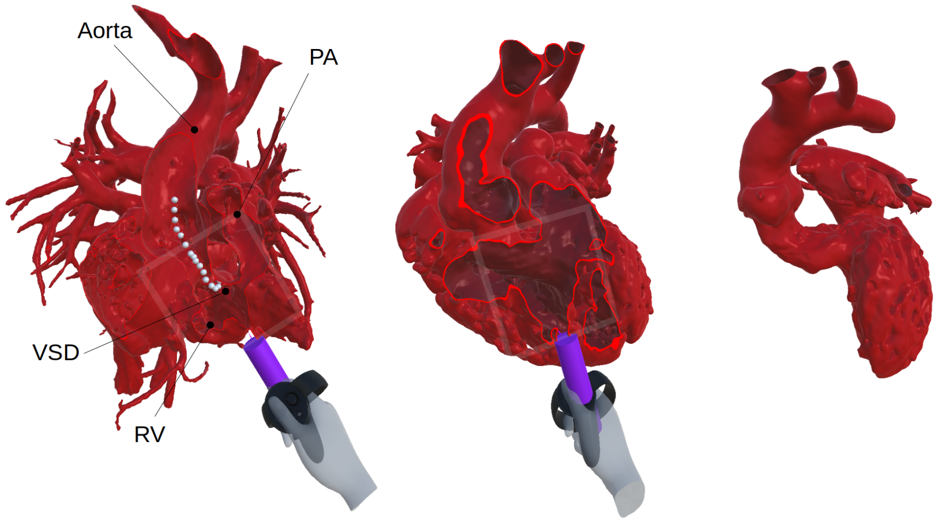

Figure 7.3: Left: pre-surgical patient model, shown using the planar clipping tool. The markers delineate the planned LV-aorta baffle path. Centre: Post-surgical patient model, with newly reconstructed LVOT shown. Right: the same post-surgical patient model with only the isolated LV to aorta baffle route shown.

7.3.2 Case report 2: planning biventricular repair with arterial switch

The patient presented with DORV, previously palliated with a PA band. Due to an identified tricuspid valve (TV)/mitral valve (MV) fibrous continuity, the interventricular communication was suspected to be perimembranous, and the conduction axis expected at its posterior-inferior margin. The 3D model was reconstructed from cardiac CT images.

Using VheaRts, assessment of the intracardiac anatomy was performed. The relationship between the VSD and the great arteries was inspected in a similar approach to case 1 using the clipping tools. The course of the suture line and expected 3D conduction system route were mapped by placing markers and measuring distances (Section 6.5.4). The diagnosis of the VSD phenotype was confirmed.

Following VR assessment, interventricular tunnelling of the LV to the PA with arterial switch was deemed the most suitable approach. At 10 months-of-age, the patient successfully underwent PA de-banding, VSD enlargement, intracardiac routing of the LV to the pulmonary root, arterial switch and re-positioning of the neo-PA to the right of the neo-aorta.

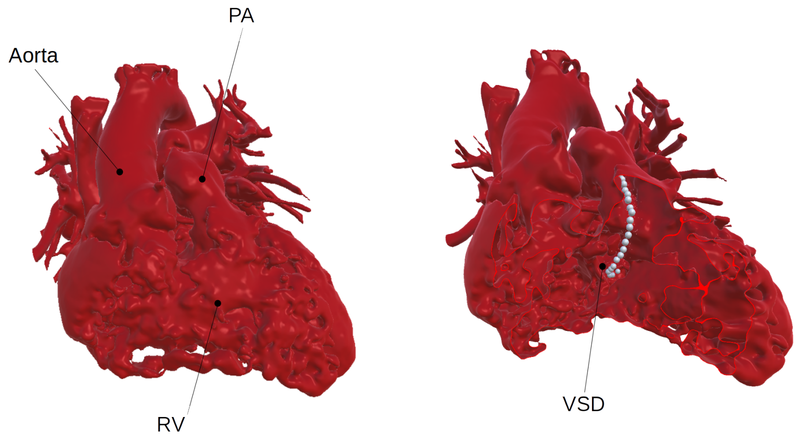

Figure 7.4: Pre-surgical patient blood pool model. Left: exterior view, right: intracardiac view with the expected LV-PA baffle route plotted using markers.

7.3.3 Case report 3: planning complex biventricular repair

A patient with complex DORV was referred following a series of interventions in other centres. The original anatomy presented with aortic coarctation, arch hypoplasia and DORV. As a neonate, the patient underwent aortic arch repair and PA banding. At 10 months-of-age, the patient underwent the Glenn and Damus-Kaye-Stansel (DKS) procedure [220], planning for an eventual total cavopulmonary connection (TCPC). A 3D model was reconstructed from CT images.

Using VheaRts, a small LV cavity, narrow VSD and thickened interventricular septum could be clearly identified. The fibrous continuinity between the TV/MV implied a perimembraneous VSD. With VheaRts, the 3D conduction system route and suture lines were mapped by measuring and using markers and the clipping tools.

The patient was operated with a one and one half ventricle repair approach with VSD enlargement. The postoperative course was complicated, with bilateral pleural effusion and a failed attempt at extubation. The superior vena cava pressure was measured at 20mmHg, with the inferior vena cava at 10mmHg. Seven days after the one and one half ventricle repair, the Glenn was taken down and a biventricular repair was successfully performed. Three weeks after the repair, cardiac catheterisation showed a patent SVC-RA connection, open RV-PA conduit, good LV function and an unobstructed LV-DKS channel and LVOT.

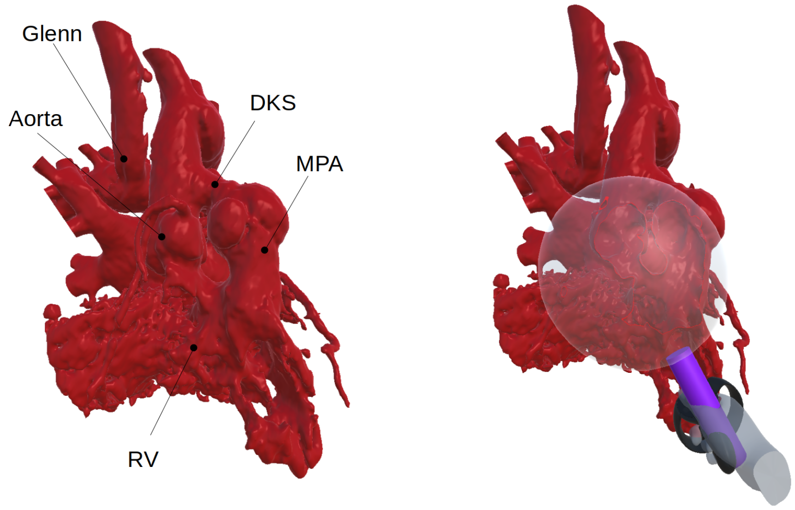

Figure 7.5: Pre-surgical patient blood pool model. Left: exterior view, right: intracardiac view with the interior DKS revealed using the clipping tool.

7.4 VheaRts for planning DORV repair: a retrospective pilot study to assess the benefits of VR

7.4.1 Introduction

As described in Sections 2.1.3 and 3.5.4, DORV represents a wide range of anatomical configurations that, together with associated abnormalities, often result in unique anatomies requiring individualised surgical repair approaches. In many cases, deciding on a biventricular/univentricular approach can be challenging, with potentially significant impacts on patient outcomes. Since the introduction of 3D options for exploring intracardiac anatomy, there has been a lack of evidence for favouring one 3D modality over the other, or against conventional cross-sectional imaging (CT/CMR). This study reports a first attempt to identify evidence regarding the potential usefulness of VR compared to other image modalities in the pre-operative planning of DORV.

7.4.2 Aims and objectives

This study aimed to evaluate the role of different 3D modelling techniques (3D PDF, printed models, VR) and conventional cross-sectional imaging in the surgical planning of patient-specific DORV cases. The main objectives were to: (i) evaluate how the surgeon approach differs when exposed to various 3D imaging modalities, and (ii) investigate any potential advantages for using VR as a supplementary 3D visualisation tool during the pre-operative planning of complex DORV.

7.4.3 Methods

7.4.3.1 Patient population and image data

Ten consecutive patients with complex DORV who successfully underwent biventricular repair with intracardiac baffle (without arterial switch), between August 2015 and March 2018 at Great Ormond Street Hospital for Children were retrospectively selected. CT images were acquired using dual source multidetector CT scan (Siemens Somatom Force, Siemens Healthineers, Erlangen, Germany), as contrast enhanced non ECG-gated datasets. Cardiac MR imaging were acquired with contrast enhanced three-dimensional balance steady state free precession (3D whole heart) in mid-diastole with respiratory navigator at 1.5 Tesla (Siemens Avanto). Relevant clinical information, including post-operative follow-up and cross-sectional imaging data (CT/CMR) was collected for all patients.

Volumetric images were post-processed (see Section 2.3) to reconstruct the atria, ventricles, great vessels and, where possible, valvular structures. For 3D printing, two planar cuts were performed on the 3D meshes, one across the RV free wall and one across the LV posterior wall to expose the intracardiac anatomy and VSD. The cuts were indicated by the operating surgeons as optimal to evaluate the patient-specific anatomies. A 1 mm uniform thickness was added to each cardiac surface for printing.

7.4.3.2 Three-dimensional visualisation tools

Following image reconstruction, the 3D PDF file of each patient model (with planar cuts) was created from ScanIP (Synopsis, U.S.A.). All models were then printed at 1:1 scale, in rigid white nylon (EOS PA2200 Nylon 12) using selective laser technology (EOS P100). The 3D reconstructions (as .obj files) were imported into VheaRts. The target HMD was the Oculus Rift system (Section 3.5.2). In the VR environment, the user could freely move, rotate and interact with each patient-specific model in virtual space, as detailed in Chapter 6.

7.4.3.3 Evaluation of patient model and surgical strategy

The ten cases were independently evaluated by two experienced paediatric cardiac surgeons from different centres, each with more than 15 years of experience as a first operator in paediatric cardiac surgery. Neither surgeons were involved in the original intervention nor with the clinical care of the patients, and were completely blinded to the actual surgical repair and outcome. Each surgeon was individually asked to provide a surgical plan for all ten retrospective cases, using four different visualization tools. The decision on the type of repair (biventricular vs univentricular) was recorded at the end of each stage of the analysis, noting whether arterial switch would be required. The first assessment of patients was based on clinical history and conventional imaging modalities such as echocardiography and CT/CMR images. The review of the images was guided by a senior cardiologist with extensive experience in cardiovascular imaging. Following this, the surgeons were presented sequentially with patient-specific three-dimensional models of each case in the form of: (i) a 3D PDF, (ii) a physical 3D printed model, and, (iii) the VR setup (Fig. 7.6).

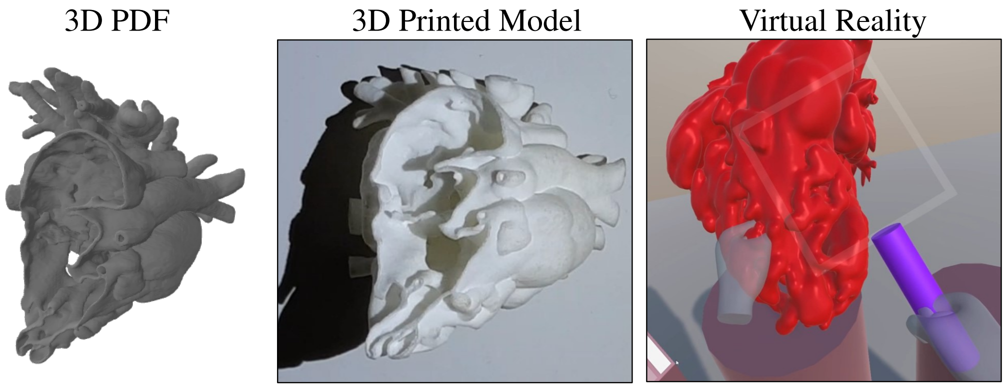

Figure 7.6: Three-dimensional visualisation tools used by surgeons A and B for an example case.

At each step, the surgeons were asked to confirm or change their potential surgical approach. Each decision was compared to the actual strategy performed on each patient (i.e. the choice of reference) in order to evaluate the accuracy of each type of 3D modelling modality in planning complex surgical repairs. The actual performed surgical repair was considered successful for all patients, due to reported positive outcomes within their follow-up time.

7.4.4 Results

7.4.4.1 Patient population

Patient baseline characteristics are summarized in Table 7.2. Nine patients had previous palliation at a median of 0.6 months (ranging from 0.1 - 3) including PA banding (PAB, n=6), arch repair and PAB (n=2) and bilateral PAB and PDA stent (n=1). In accordance with the plan agreed by the multidisciplinary team, all patients underwent biventricular repair with intracardiac tunnelling ± VSD enlargement at a median of 8.4 months (ranging from 2.6 - 13.8). In three cases (case 1, case 6, case 7), ASO was performed. Details of each procedure and LVOT gradients at discharge are recorded in Table 7.3.

Case	**Imaging modality**	**Location of interventricular communication**	**Arrangement of great arteries**	**Additional findings**
1	CT	Multiple (non committed+apical)	Parallel, aorta anterior and to right	Atrial septal defect, aortic arch hypoplasia, aortic coarctation
2	CMR	Subpulmonary	Parallel, side by side, aorta to right	Patent foramen ovale, aortic coarctation, RCA from left facing sinus
3	CT	Non committed	Parallel, side by side, aorta to right	N/A
4	CT	Non committed	Parallel, side by side, aorta to right	Large VSD split in two by large muscular bridge
5	CT	Double committed with inlet extension	Parallel, side by side, aorta to right	Sub aortic narrowing, anomalous left anterior descending from RCA
6	CT	Non committed	Parallel, side by side, aorta to right	Large conal branch running on the anterior wall of RV
7	CT	Multiple (non committed + small muscular)	Parallel, aorta anterior and to right	N/A
8	CT	Non committed	Parallel, side by side, aorta to right	Subpulmonary and main pulmonary artery stenosis
9	CT	Non committed	Parallel, aorta anterior and to right	ASD, chordal attachment of tricuspid valve to interventricular septum
10	CT	Non committed	Parallel, aorta anterior and to right	Aortic arch hypoplasia, aortic coarctation

Table 7.2: Baseline characteristics for the 10 DORV patients.

The median follow-up was 31 months (ranging from 10.2 - 44.6). There were no mortalities during follow up. Two patients underwent further surgical interventions. Case 1 had surgical closure of residual apical VSDs due to the persistence of significant left to right shunt 24 months after the initial repair. Case 7 underwent

Case	Type of biventricular repair	Previous palliation	Age at palliation (months)	Age at repair (months)	Discharge LVOT (m/s)
1	ASO, intraventricular tunnel and arch repair	Bilateral PAB and PDA stent	0.7	6	1.2
2	Intraventricular tunnel	Arch repair, PAB	0.1	11	1.1
3	Intraventricular tunnel	PAB	1.4	7	1.4
4	Intraventricular tunnel	PAB	0.5	5	2.2
5	Intraventricular tunnel	PAB	2	3	1.2
6	ASO and intraventricular tunnel	PAB	2.6	10	1.4
7	ASO and intraventricular tunnel	PAB	0.4	13	2.5
8	Intraventricular tunnel	n/a	n/a	5	1.3
9	Intraventricular tunnel	PAB	3	11	1.2
10	Intraventricular tunnel	Arch repair, PAB	0.1	10	1.1

Table 7.3: Procedures performed for each DORV patient. CPB = Cardio-pulmonary bypass, ASO = arterial switch operation, PAB = pulmonary artery banding.

two reoperations for recurrent LVOT obstruction. A resection of a fibro-muscular shelf was performed at 8 months after repair, followed by a later replacement of the VSD patch at 23 months. One more patient (i.e. case 5) was scheduled for reoperation at the time of writing, with an LVOTO gradient of 4.1 m/sec 40 months after the biventricular repair. Follow-up data such as LVOT gradients are summarised in Appendix C.1.

7.4.4.2 Surgical planning

The choices of the two surgical evaluators were compared to the actual operation performed (Table 7.4). After reviewing the CT or CMR data, a biventricular repair strategy was correctly proposed in 9/10 cases by surgeon A and in 6/10 cases by surgeon B.

Following review of the 3D PDF models, the feasibility of any biventricular repair (with or w/o ASO) did not change for surgeon A (9/10) or surgeon B (6/10). Additionally, no individual approaches drastically changed (i.e. from univentricular to biventricular or vice versa).

After viewing 3D printed models, the total feasibility of biventricular repair remained unchanged for surgeon A (9/10), and increased to 7/10 for surgeon B. However, the pre-operative plan for some individual cases was changed, for both

Case	Repair		CT/CMR	3D PDF	3D print	VR
1	Bi-V +ASO	surgA	Bi-V+ASO	Bi-V+ASO	Uni-V repair	Bi-V+ASO
		surgB	Uni-V repair	Uni-V repair	Uni-V repair	Bi-V
2	Bi-V	surgA	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
		surgB	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
3	Bi-V	surgA	Bi-V+ASO	Bi-V+ASO	Bi-V	Bi-V
		surgB	Bi-V	Bi-V	Bi-V+ASO	Bi-V
4	Bi-V	surgA	Bi-V	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
		surgB	Bi-V+ASO	Bi-V+ASO	Bi-V	Bi-V+ASO
5	Bi-V	surgA	Bi-V+ASO	Bi-V	Bi-V	Bi-V
		surgB	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
6	BiV +ASO	surgA	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
		surgB	Bi-V+ASO	Bi-V+ASO	Uni-V repair	Uni-V repair
7	BiV +ASO	surgA	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO	Bi-V+ASO
		surgB	Not sure	Not sure	Bi-V+ASO	Bi-V+ASO
8	Bi-V	surgA	Bi-V	Bi-V	Bi-V	Bi-V
		surgB	Bi-V	Bi-V	Bi-V	Bi-V
9	Bi-V	surgA	Uni-V repair	Uni-V repair	Bi-V	Bi-V
		surgB	Uni-V repair	Uni-V repair	Not sure	Bi-V
10	Bi-V	surgA	Bi-V	Bi-V	Bi-V	Bi-V+ASO
		surgB	Not sure	Not sure	Bi-V	Bi-V
Bi-V or Bi-V+ASO approach			**15/20 (75%)**	**15/20 (75%)**	**16/20 (80%)**	**19/20 (95%)**
Approach matches performed surgery			**9/20 (45%)**	**9/20 (45%)**	**11/20 (55%)**	**12/20 (60%)**

Table 7.4: Agreement with chosen surgical strategy according to different 3D tools. ASO = arterial switch operation, Bi-V = biventricular repair, Uni-V = univentricular repair. Green: suggested biventricular approach exactly matches the repair, Yellow: suggested approach is biventricular but need for ASO not correctly identified, Red: approach disagrees with performed repair (i.e. is univentricular).

surgeons A and B. For example, both surgeon A and B incorrectly switched to univentricular repairs for one case each after 3D printed model evaluation (1 and 6, respectively).

Following the 3D printed model step, VR was then performed for all cases. The perceived feasibility of biventricular repair increased to 10/10 for surgeon A and 9/10 for surgeon B. No approaches were changed into univentricular repairs following VR, although for patient 6, surgeon B maintained a univentricular approach even after VR.

Overall, the agreement between the surgeon evaluators and the actual repair changed from 75% after cross-sectional imaging review, to 95% following all 3D modalities. The inter-observer agreement in surgical strategies changed from 70% after cross-sectional imaging review, to 90% after VR. Compared to traditional cross-sectional imaging, the 3D PDF did not increase the perceived feasibility of biventricular repair or the agreement with the surgical approach. The perceived feasibility of biventricular repair improved marginally with the use of the physical 3D printed model (by 5%) and more significantly using the VR setup (by 15%). The agreement between the actual repair and proposed approach improved by 10% with 3D printing (45% to 55%) and by 5% with VR (55% to 60%).

7.4.5 Discussion

The main findings are:

- VR assessment following 3D PDF and 3D printing resulted in a 95% accuracy in determination of biventricular repair feasibility and a 60% accuracy in determining ASO suitability.
- VR assessment showed the most significant improvement for determining biventricular repair feasibility (increase of 15%), whereas 3D printing showed the most significant improvement for identifying the suitability of ASO (increase of 10%).
- There was no evidence to support that the use of a 3D model on screen (a 3D PDF in this study) provided extra assistance for determining the feasibility of biventricular repair or ASO, when compared to conventional cross-sectional imaging.
- 3D printing and VR were found to be able to drastically change the opinion on the surgical plan (univentricular to biventricular or vice versa), whereas 3D PDF did not.

According to the results, screen-based 3D reconstruction such as a 3D PDF had a negligible benefit for identifying the suitability of biventricular repair. This may be explained by the limited range of interactions 3D PDF allows, when compared to 3D printing or VR. Additionally, viewing 3D models on a screen does not provide the same level of depth perception as VR or 3D printing does.

Rigid white nylon 3D-printed models were used in this study. Physical 3D printed models in a 1:1 scale showed a marginal overall improvement when compared to the use of conventional imaging and 3D PDF. However, in two individual cases (one for each surgeon), the exposure to 3D printed models led to an incorrect decision to perform univentricular repair. In a real case scenario, this may have lead to suboptimal surgical outcomes.

VR proved to be the best tool for the two evaluators to identify the feasibility of biventricular repair and the need for arterial switch operation. This is likely due to the immersive capabilities of VR, which have shown to be highly suitable for the study of complex anatomy [153, 221]. VheaRts enabled evaluation of the patient-specific models from conventional and non-conventional surgical views. Clipping tools enabled exploration of the intracardiac anatomy from angles not possible with other imaging modalities.

7.4.6 Limitations

Although small, the cohort represents a carefully selected and homogenous population of DORV patients with complex interventricular communications. In the future, the study should be extended to include a larger and heterogeneous cohort of patients and conditions. In this study, the visualisation modalities were presented in consecutive order to the surgeons (3D PDF, 3D printing, VR). This strategy was decided since it better mimics a real-world scenario where cross-sectional imaging is likely followed up by 3D modalities in order of most commonplace to least commonplace. A limitation is that in the current approach, bias is possibly introduced in each case due to previous exposure from a different imaging modality. In the future, the order of 3D modalities should be randomised to better nullify this effect and enable each modelling technique to be assessed in isolation.

A further limitation is that testing was done by two senior evaluators only. However, it was felt that given the complexity of the assessed conditions and the experience of the actual operating surgeon, additional testing by less experienced evaluators may not offer as conclusive results. In this study, the choices of the two evaluating surgeons were compared to the actual operation performed. As summarised in Table 7.3 and Appendix C.1, excellent short-term results from all repairs were achieved. At Great Ormond Street Hospital, biventricular repair is generally the preferred option for patients with complex DORV and remote VSDs. However, ultimately it is the long-term re-intervention free survival that determines the best approach to this challenging subgroup of patients, and good results are possible with various surgical strategies. By including a diverse set of patients in the future with longer-term successful outcomes, the assumption of ‘correctness’ for the procedure will become more reliable.

All 3D modelling techniques used in this study have significant limitations related to their capacity for assessment of the atrioventricular valves. This is due to the source imaging modality (CT/CMR) which cannot display valve structures with sufficient resolution for accurate 3D reconstruction. Since the insertion of the tricuspid valve represents a crucial part of the surgical correction of DORV, it is recommended that 3D models fuse information from multimodality imaging such as 3D echocardiography.

7.5 VheaRts for planning complex surgery beyond CHD

VheaRts was used in supporting some highly complex surgical procedures, outside the scope of CHD, in conjoined twins separations. In these rare cases, the challenging anatomy often poses several questions in surgical planning, requiring a multidisciplinary approach and the involvement of specialists from various medical disciplines. Thus, VR represents a useful tool to assist the multidisciplinary team in interpreting these complicated cases and in guiding the decision-making process.

In 2019, the first complex case that involved the use of VheaRts outside CHD applications, was a set of twins with total vertical craniopagus, an exceptionally rare condition where the crania of twins are joined, with varying degrees of fusion between the brains and parts of the cranial vasculature.

Anatomical reconstructions from multi-modality imaging (CT and MR) of all structures of interest (skin, bone, brains, dura, vasculature) from both twins were imported into VheaRts, where they could be studied. As part of the operation to separate the twins, the complex network of arteries/veins and their relationship with the parenchyma need to be identified. VR was used to intuitively navigate the anatomies and better understand the three-dimensional location of the vessels within the brains by activating/deactivating different layers of structures and leveraging transparency tools. In total, three VR sessions were organised where the lead neurosurgeon spent time navigating the 3D reconstructed models and identifying vascular points of interest (Fig. 7.7, left). Following a planning phase involving several cross-sectional imaging (CT/MR) data, 3D printing models and VR sessions, the patients were separated successfully in a series of complex surgeries [222].

Similarly, VheaRts was used to support a successful second separation of craniopagus twins at Great Ormond Street Hospital for Children in 2020 [223], and a third set, in 2022, this time with a different group of surgeons from the Instituto Estadual do Cerebro Paulo Niemeyer in Brazil [1]. For this case, a VR session was held remotely between the lead neurosurgeons from Great Ormond Street Hospital for Children and the equivalent surgical counterpart from the Brazilian pediatric centre using the multi-user VheaRts application to discuss the best surgical approach. Both surgeons were connected to the same VR room using Meta Quest 2 headsets, and had shared control of the twins 3D-reconstructed models within the scene. Various approaches of surgical separation were discussed, with specific challenging anatomical regions highlighted during the consultation. The twins were later separated successfully in a lengthy procedure in Brazil [224].

A final set of omphalopagus twins, joined at the abdomen and with two functional lower limbs, were separated at Great Ormond Street Hospital for Children in 2022. This case required a wide range of surgical and clinical specialists to take part in both the planning and surgical separation, due to the complex arrangement and fusion of multiple organs including livers, vasculatures, kidneys, urinal tract systems and more. The 3D reconstructed models were shown to the main surgeons during the pre-operative planning phase (Fig. 7.7, right). In addition, a multi-user

[1] Video of UK-Brazil online VR surgical plan for craniopagus twin separation: `https://vimeo.com/831952801`

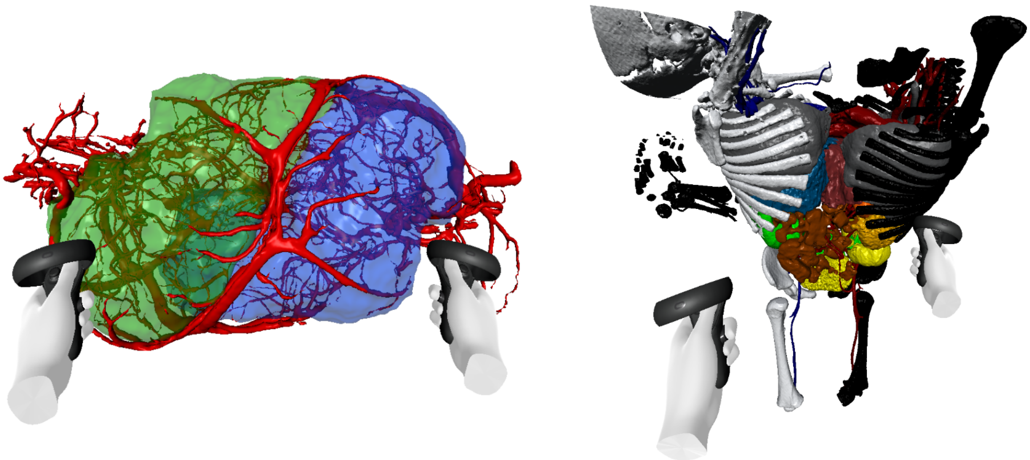

Figure 7.7: Left: model of the first set of craniopagus twins, with brains in separate colours and vessels shown intermingling. Right: conjoined omphalopagus twins, with full skeletal structure and internal organ 3D models reconstructed.

version of the application was built on multiple Quest 2 devices and used for the pre-operative team briefing. This facilitated an active discussion of the case, since the acting surgeons joined the virtual room and reviewed the surgical plan in 3D, enhancing multidisciplinary team communication. Additionally, the VR scene was projected on monitors, allowing the rest of the surgical team (15+ members) to observe and contribute to the discussion. Following the consultation, the separation of the twins was successfully performed [225].

7.6 Future integration of VR in clinics

In this Chapter, a model for the early adoption of a VR platform in CHD pre-operative planning has been shown, with possible applications and benefits outlined. In the future, greater efforts must be made in order to establish the level of benefit that clinical VR assessment provides. This is necessary to overcome one of the main hurdles preventing wider integration of VR in clinics - a lack of extensive evidence quantifying if patient outcomes are improved with VR planning and if so, by how much. A solution could involve recording patient/surgical metrics in a randomised clinical trial, where participants have access to VR for pre-operative planning over a prolonged period of time. The intervention group (imaging + VR) could be compared to a control group (only imaging). Measures such as success rates, aortic cross-clamp time, cardiopulmonary bypass time, rates of re-intervention, post-op complications and more could be recorded to assess the level of improvement that VR brings to CHD surgery. More granular and qualitative measures could also be extracted, by interviewing surgeons pre- and post-surgery in order to assess if the anatomical insight provided by VR correlated with reality (e.g. I observed that the VSD was large and perimembranous as seen in the VR model). This could be done for various types of operations to assess the accuracy of VR for providing a realistic expectation of the 3D anatomical configuration (and hence the surgical plan). Other measures such as a self-perceived improvement in confidence, understanding of 3D patient anatomy and more could be gathered as part of the questionnaires. Data could be analysed to assess how the impact depends on the experience level of the surgeon, and whether the case is performed at a high volume centre or not. Results from prospective studies and clinical trials are pivotal for providing an improved understanding on the clinical adoption of VR on a global scale.

As demonstrated in this Chapter, the VR application can be distributed amongst multiple centres and in some cases consultation can be done between multiple hospitals using VR. These capabilities could also be further studied in order to examine how they can impact current conventions for planning treatment of challenging CHD cases. Risks involved in clinical implementation of VR relate primarily to the quality of the 3D mesh segmentation. In order to establish a routine protocol for mitigating the error, this would ideally require that each 3D segmentation is conducted and checked by a cardiac imaging specialist to confirm it accurately represents the source data. Since segmentations discard information from the images for the purpose of creating a binary model, they should only be used as a supplement and not as a replacement to the imaging review. Additionally, VR assessment should take place with the presence of the imaging expert (with the cross-sectional images) to ensure that any doubts about the model can be addressed during the VR review.

7.7 Summary and conclusion

The application of VheaRts in various clinical settings for supporting decision-making has been outlined. Individual case studies of the most commonly assessed lesion (DORV) were presented, including a retrospective pilot study which investigated any potential benefits of VR for deciding on DORV optimal surgical approach. Finally, an insight into the scope of VR beyond CHD was given, in complex cases of conjoined twins. In conclusion, VheaRts shows promise as a clinical visualisation tool for supporting the complex repair of CHD and other complex procedures when integrated into a pediatric centre. Integration of new technologies, such as VheaRts, in care settings can be fostered by early adoption in education and training of healthcare professionals. Examples of this with VheaRts are reported in the following Chapter.

Chapter 8

VheaRts for education

8.1 Introduction

VheaRts, the VR platform I developed during my PhD (see Chapter 6), not only finds applications within clinical decision-making of complex CHD (Chapter 7), but also responds to the rising needs for a robust visualisation tool with specialised functions for teaching and learning CHD (Section 2.2.2). As discussed, heart specimens are considered the 'gold standard' for learning cardiac morphology, yet are not widely available. Therefore, more accessible resources are required to support the growing need for effective CHD education. Due to overall advances in portable headset technology and software capabilities, it is becoming increasingly affordable and accessible for academic/healthcare institutions to acquire multiple HMD devices, thus creating opportunities for teaching entirely in a virtual ecosystem (Section 3.5.4). Additionally, the COVID-19 pandemic created further demand for VR applications with online features that were previously underutilised. However, integration of VR into CHD courses for students at varying stages of education remains widely unexplored. Lack of quantitative feedback, showing tangible improvement in knowledge acquisition, is the biggest barrier to wider scale adoption. Additionally, the suitability of online VR rooms for teaching CHD is still relatively new. In the context of exploiting VR for CHD education, this Chapter presents how I: (i) created a curated dataset of CHD models for teaching (Section 8.2), (ii) implemented applications of VR for teaching CHD in different settings including undergraduate, postgraduate and highly specialised courses (Section 8.3), and (iii) investigated the potential benefits that VR can bring over conventional anatomical teaching methods for improving understanding of CHD and for supporting highly specialised anatomical education (Sections 8.4, 8.5 and 8.6).

8.2 Virtual anatomy lab

The development of the core functionalities in the VheaRts educational application was shown in Chapter 6. In order to create a repository of hearts for teaching CHD, the application was populated with 3D heart models. These were incrementally added over the period 2019-2022, with the most up-to-date collection including the 30 models detailed in Table 8.1.

Models were curated to address the requirements of the senior cardiac morphologists and clinicians, who were also responsible for lecturing and delivering the courses to the recipients. The heart 3D meshes were generated from three main sources: (i) CT/CMR patient images, retrieved from a single pediatric centre (Great Ormond Street Hospital for Children), (ii) micro-CT or Synchrotron scans of anatomical specimens, and (iii) idealised and artistic models of the heart, publicly available from online repositories. All DICOM images were segmented using commercially available software ScanIP (Simpleware, Synopsis). For the anatomical samples, the myocardium was reconstructed, whilst for the clinical images only the blood pool was segmented.

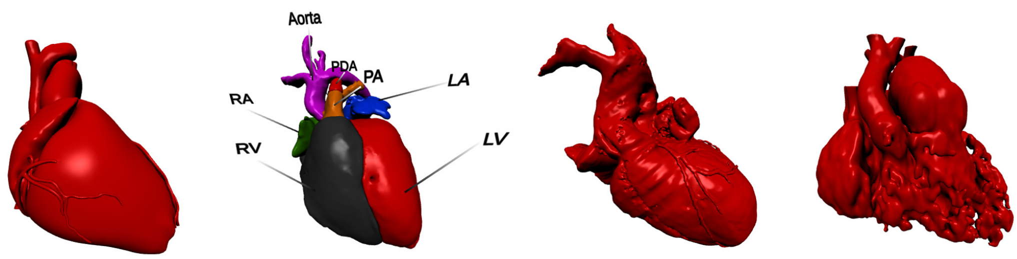

Figure 8.1: Examples of 4 models from different imaging modalities or sources. From left to right: (i) idealised artistic model (case 1), (ii) anatomical model (micro-CT, case 3), (iii) anatomical model (Synchrotron, case 21), and (iv) patient blood pool model (CT, case 10).

Real patient-specific cases constituted 70% of the library models, with DORV

Case	Description	Source	Type	Submeshes
1	Normal	online repository	myocardial	✓
2	Normal	online repository	myocardial	✓
3	Normal	specimen	myocardial	✓
4	Normal	patient	blood pool	×
5	ASD	online repository	myocardial	✓
6	VSD	online repository	myocardial	✓
7	VSD	patient	blood pool	×
8	AVSD	online repository	myocardial	✓
9	AVSD	specimen	myocardial	✓
10	DORV 1	patient	blood pool	×
11	DORV 1	patient	myocardial	×
12	DORV 1 repaired	patient	blood pool	×
13	DORV 2	patient	blood pool	×
14	DORV 3	patient	blood pool	×
15	DORV 4	patient	blood pool	×
16	DORV 5	patient	blood pool	×
17	TOF	specimen	myocardial	✓
18	TOF	patient	blood pool	×
19	TOF repair 1	patient	blood pool	×
20	TOF repair 2	patient	blood pool	×
21	TGA	specimen	myocardial	✓
22	TGA repair	patient	blood pool	×
23	CCTGA 1	patient	blood pool	×
24	CCTGA 2	patient	blood pool	×
25	DORV + TGA	patient	blood pool	×
26	HLHS	patient	blood pool	×
27	Mitral stenosis	patient	blood pool	✓
28	Sinus venosus	patient	blood pool	×
29	Sinus venosus repaired	patient	blood pool	×
30	PDA	patient	blood pool	×

Table 8.1: Up-to-date collection of 3D models available in the online anatomy lab. Models from online repositories are idealised/artistic creations. Patient cases are reconstructed from CT/CMR. Myocardial models have the muscle reconstructed instead of the blood pool. Submeshes indicates the model is broken down into multiple components.

being the most frequent lesion (24% of patient-derived models), followed by idealised/artistic models (17%) and anatomical specimens (13%). Most of the meshes derived from clinical cases were blood pool reconstructions (80%), with the remaining including the myocardium. Where possible, images were segmented into separate individual structures to enable highlighting/labelling of each subcomponent. Out of all anatomical specimen cases (n=4), only case 21 (Fig. 8.1, centre-right) was scanned using X-ray phase contrast Synchrotron imaging technology (PB-XPCI, Paul Scherrer Institut, Switzerland). The other anatomical specimens (n=3) were scanned using micro-CT (Nikon Metrology HMX ST 225, spatial resolution of 5-125 μm). Lastly, two patients (cases 25 & 27) which were previously scanned using 4D-CT sequences had each frame of the cardiac cycle segmented individually to create dynamic models of the heart (animations) in addition to static models.

8.3 Overview of VheaRts for teaching CHD

VheaRts was implemented in six different courses/workshops at UCL and Great Ormond Street Hospital to support the education and training of cardiac anatomy and CHD. The breakdown of these courses is shown in Table 8.2. From 2019 to 2022, more than 240 students were taught using VheaRts by Prof. Andrew Cook, Professor of Cardiac Morphology at UCL.

Course	Level	Total attendees	Years held	Multi-user
iBSc	undergraduate	71	2020-2022	✓
iNUGSC workshop	undergraduate	52	2022	✓
MSc	postgraduate	44	2019-2022	mixed
Cardiac morphology workshop	professional	58	2018-2020	×
Echocardiography workshop	professional	14	2021-2022	✓
DORV workshop	professional	5	2022	✓

Table 8.2: Overview of the different courses VheaRts was implemented in at UCL and Great Ormond Street Hospital. Multi-user means that sessions were conducted with multiple headsets connecting to a shared VR room.

8.3.1 iBSc

At UCL, the integrated BSc (iBSc) in Cardiovascular Science offers medical students a year out of their Medicine programme to study foundational cardiovascular science. An optional CHD module within the course covers the anatomy of the normal heart and a range of complex lesions. The module is delivered via traditional lectures together with anatomical case reviews using pathological specimens.

VheaRts was integrated into the delivery of this CHD module for the first time in 2020-2021, when iBSc students (n=40) were taught on-campus over the course of four weeks in groups, due to COVID-19 restrictions (Fig. 8.2) [1].

Figure 8.2: An iBSc VR session (n=5) during COVID-19 restrictions (2020-2021 cohort).

During each session, the normal heart anatomy and/or a selection of lesions matching the progression of the course were presented in multi-user VR rooms. In the following academic year 2021-2022, the end of COVID-19 restrictions allowed for iBSc students (n=31) to attend longer sessions in larger groups. Three one-hour VR sessions were held in-person, with each session being repeated twice (group sizes of ∼15).

8.3.2 iNUGSC

In 2022, a multi-user VR workshop on cardiac morphology was developed for UCL medical students (n=52), hosted at the International Undergraduate and Foundation Surgery conference (iNUGSC 2022) organised by the UCL Surgical Society. Four sessions were held over the course of a day (average of 13 participants in each session). All participants were provided with Meta Quest 2 headsets. In VR, after a short demonstration, a brief of the normal anatomy was first delivered, followed by an explanation of a patient-specific anatomical model with TGA (case 21, 8.1).

[1] Example session (VR view) of a VR session in 2020: `https://vimeo.com/831951658`

The description of the TGA case also included some proposed surgical solutions to the lesion. Participants were assessed post-workshop and feedback was collected to investigate the effectiveness of a VR-only based approach for teaching CHD. Further details are presented in Section 8.4.

8.3.3 MSc

A CHD module similar to that delivered for the iBSc course is an available option for post-graduate students who undertake the UCL MSc in Cardiovascular Science. Students come from a range of multidisciplinary backgrounds. The module programme starts from the fundamentals of normal cardiac anatomy. Teaching and examination are heavily focused on the identification and understanding of CHD morphological abnormalities. In the first year of testing (2019-2020), the VheaRts educational platform was made available to students with two models. The first was an idealistic model of the heart, derived from the Living Heart Project [226] (Table 8.1, case 1). This model represents an average adult male heart at 70% ventricular diastole. Nineteen individual anatomical structures were subdivided in various colours and labelled. The second model was a TOF post-natal specimen reconstructed from micro-CT images (Table 8.1, case 17). Through a combination of mesh cuts, the model was subdivided artificially. Eight anatomical areas of interest (specific to TOF) were labelled in addition to the basic cardiac structures (LA, LV, PA and more).

In the first year of testing (2019-2020), VheaRts was used to investigate the effectiveness of VR for improving knowledge acquisition of CHD. Students (n=22) were divided into two groups (control and intervention) to investigate how the addition of VR affected assessment scores (see Section 8.5 for details). Following this, VheaRts was fully integrated into the MSc CHD module (2020-2021) with multi-user functionality and the full suite of CHD models (Table 8.1). However, COVID-19 lockdown restrictions meant that headsets had to be posted to the students (n=11) in order to conduct VR anatomy lab sessions. Using this approach, three VR sessions (one hour each) were held remotely over the duration of the CHD programme, where the normal heart anatomy and a selection of lesions were

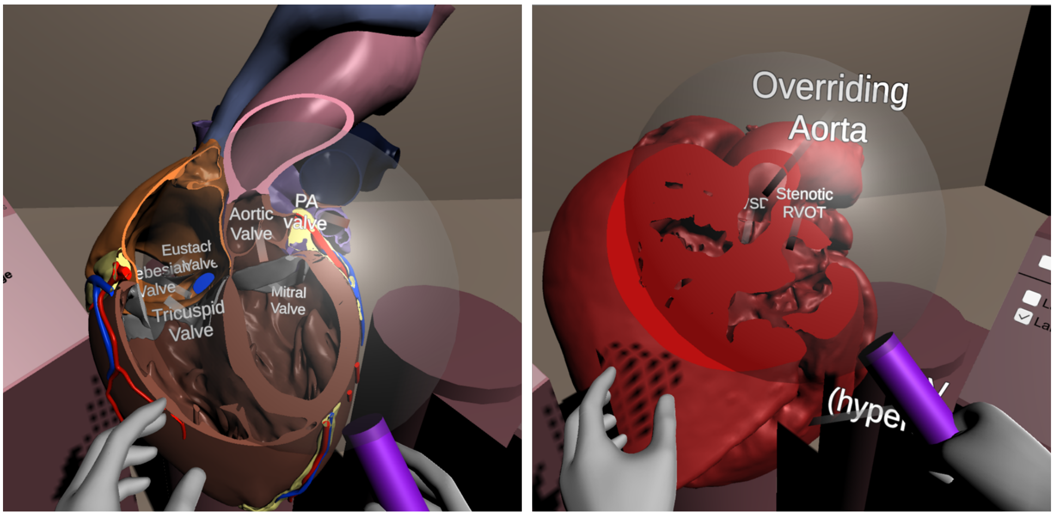

Figure 8.3: VheaRts normal (left) and TOF (right) models with labels and the sphere clipping tool active.

explored collectively yet remotely in VR. In the following academic year (2021-2022), normal in-person VR sessions resumed, with three VR sessions held over the length of the programme. Due to a smaller cohort (n=11), all students could be taught in a single group [2].

Figure 8.4: An MSc virtual anatomy lab session with a group size of 11 (2021-2022 cohort).

[2]Example MSc VR session in 2021: `https://vimeo.com/831950425`

8.3.4 Professional cardiac morphology workshop

VheaRts was tested in the UCL ‘hands-on’ Cardiac Morphology Course. This 3-day short course is focused on teaching CHD morphology to healthcare professionals from a mix of backgrounds with the aid of anatomical specimens, all under expert guidance [227]. Participants (n=58) to the courses, which ran between October 2018 and January 2020 before the COVID-19 pandemic, were invited to additionally test VheaRts over a dedicated 2-hour session. Six patient-specific models were included, one normal and a selection of 5 CHDs (Table 8.3).

Lesion	Case	Specimen	Image modality	Age
None (normal)	3	✓	Micro-CT	16 weeks gestation’
AVSD	9	✓	Micro-CT	16 weeks gestation’
TOF	10	✓	Micro-CT	post-natal
TGA	17	✓	Synchrotron	16 weeks gestation’
DORV	21	×	CT	4 months
PDA	30	×	CT	12 months

Table 8.3: All image datasets converted into 3D models and used in this study. The case number of the model is with respect to the final collection of hearts (Table 8.1).

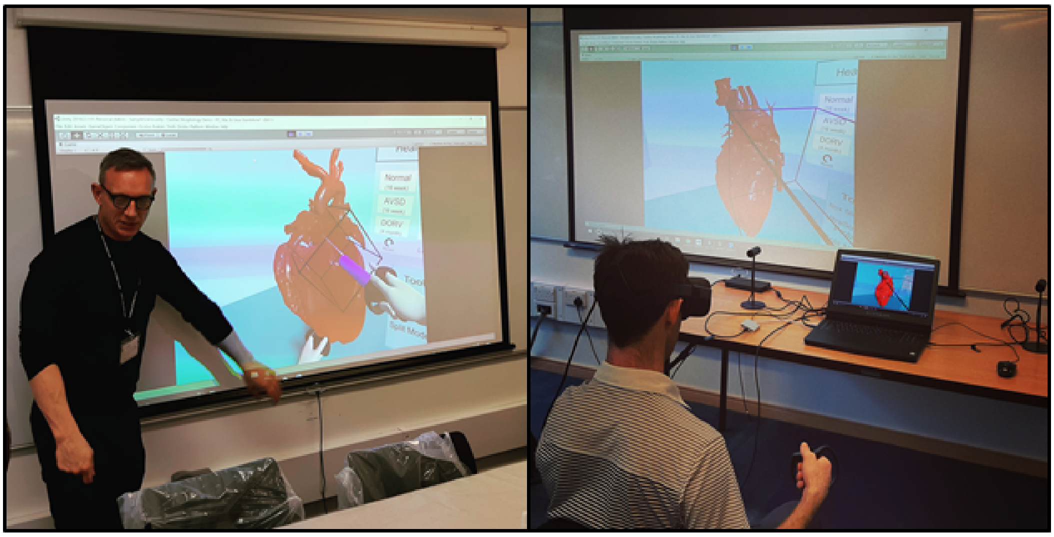

Figure 8.5: The typical VR set-up used in the VR 2-hour session slot for clinical professionals training in CHD. Left: instructor, right: course attendee.

The VR setup constituted of a single Oculus Rift connected to an Alienware 17 R5 laptop. Users had access to the following functionalities: (i) model handling,

(ii) clipping tools, (iii) labelled structures, (iv) VR ultrasound probe and (v) measuring/placing markers. After a short adaptation period, users independently explored the contents and tools of the VR demo. The perspective of the participants was streamed live and guidance was offered by the course's instructor (Fig. 8.5). User feedback was collected to assess the feasibility/effectiveness of VR for teaching healthcare professionals (see Section 8.6 for details).

8.3.5 Professional echocardiography workshop

In collaboration with the Great Ormond Street Hospital Learning Academy (GLA), VheaRts was adapted to support the delivery of a specialised echocardiography simulation course for professional cardiologists and surgeons from Great Ormond Street Hospital. Two pilot sessions were organised (in 2021 and 2022) of one hour each. A total of 14 participants attended the courses. At each session, the participants were split in two groups: one group followed a 30-minute lecture in VheaRts (which included CHD clinical case studies) while the other trialled a physical echocardiography simulator. The groups then swapped for the following 30 minutes. Echocardiography training was delivered by a Consultant Cardiologist from Great Ormond Street Hospital. Three models were used for the course teaching: (i) CCTGA, (ii) VSD, (iii) and Sinus venosus (pre/post repair) corresponding to cases 23, 7, and 28+29 from Table 8.1, respectively. Echocardiography video clips were retrieved for each patient case and added into the VR environment (Fig. 8.6). This allowed participants to emulate echo views using the VR ultrasound probe and compare 2D projections to the real echo images and the 3D model anatomy.

Feedback was collected from the participants (n=5) after the first pilot workshop had ended. Questions related to the experience were presented in a five-point Likert scale format. Only one attendee had prior experience in VR. All participants agreed or strongly agreed that: (i) VR is useful for understanding cardiac structures, (ii) VR enhances medical imaging and (iii) VR is useful for reviewing the surgical procedure. All respondents answered 'agree' in response to 'The VR headsets were easy to use'. However, four out of five attendees noted that they had experienced some levels of discomfort, with one expressing strong discomfort due to 'neck pain

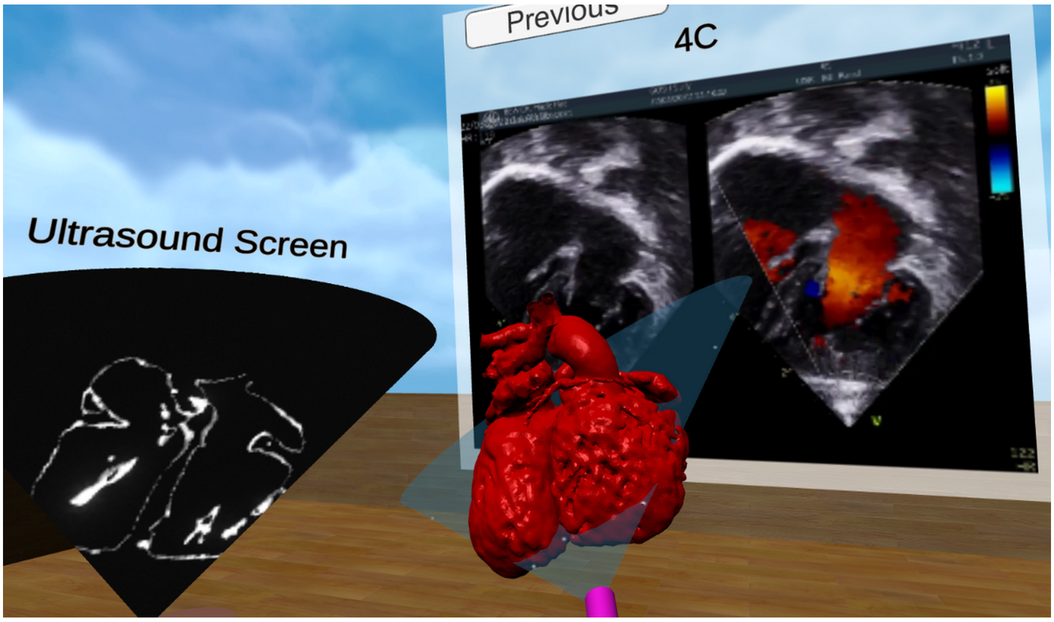

Figure 8.6: In-person view of VR ultrasound probe compared to actual four-chamber echocardiography view of the same CCTGA patient.

and headaches'. A common request was to include more post-operative cases and patients with 'complex 2D anatomy'.

8.3.6 Professional online DORV workshop

VheaRts was used to explore the possibility of conducting anatomical workshops remotely in VR. This was supported by a small grant from the UCL Africa and Middle East Teaching Fund Initiative. The workshop involved a group of South African pediatric cardiologists and cardiac surgeons (n=5) from the Nelson Mandela Children Hospital. The session was centred around complex DORV anatomy, and how the anatomy drove repair. Almost all DORV models from the virtual anatomy lab were used and included in the workshop (Cases 10-15, Table 8.1). Attendees connected with their Quest 2 devices to a virtual classroom hosted on a server in South Africa. After a short familiarisation with the VR environment, participants collaboratively analysed each of the DORV case studies in the application, with the course coordinator managing the presentation of each patient case. Anonymised clinical notes were visible for each patient in VR. After the workshop was completed, feedback was collected from all the participants.

Out of all respondents (n=5), only one had previously tried VR. All partici-

Figure 8.7: VR user from Johannesburg, South Africa attending the DORV workshop, with the course coordinator based in London, UK.

pants stated that VheaRts allowed them to explore DORV anatomy in greater detail and better understand the spatial relationships between structures. Two participants reported discomfort/VR sickness. Three participants stated VR was intuitive and easy to use, while two responses were neutral. All participants stated that they were interested in using VheaRts for clinical applications, such as pre-operative planning. Requested lesions for VR assessment included pulmonary atresia, AVSD and major aortopulmonary collateral arteries (MAPCAs). Requested improvements to the application included ”more realistic textures” and ”different colours for all structures”. Overall, feedback was positive, and the application was well received by the participants.

8.4 Assessment of VR for teaching CHD

8.4.1 Introduction

As explained in Section 8.3.2, at iNUGSC 2022, the multi-user VheaRts environment was trialled as a primary method for delivering CHD education to undergraduate medical students. In this context, my aims were: (i) to investigate the feasibility of delivering a CHD workshop solely through VR, (ii) to measure the level of engagement through user feedback, and (iii) to assess post-intervention knowledge acquisition in the participants via a questionnaire.

8.4.2 Methods

Before and after the VR session, students were asked to fill out a questionnaire using a web-based mobile application. Demographic questions were collected through text-based prompts. Qualitative feedback questions were answered by selecting a value from 0-100 on an interactive GUI slider. As part of the student evaluation, attendees also had to fill out an assessment composed of ten multiple-choice questions related to the taught content (see Appendix D.1).
Pre-workshop questions:

- How much experience/knowledge do you have with the topics and instructional material presented in the class? (0=none, 100=extensive)
- How much experience/knowledge do you have with the application of Virtual Reality (VR) in the medical domain? (0=none, 100=extensive)
- Are you motivated to attend this class? (0=not at all, 100=extremely)

Post-workshop questions:

- How much mental workload did the activities in the class impose on you? (0=none, 100=excessive)
- How successful were you in the class? In other words, how satisfied were you with your gain in understanding? (0=not at all, 100=extremely satisfied)
- Were you alerted during the class? In other words, were you sleepy/tired or fully alert/awake? (0=not at all alert, 100=very alert)

In addition to computing means and standard deviations of all measured metrics, Pearson correlation coefficients with associated p values were calculated to examine correlations between metrics. Values were computed using data from all participants.

8.4.3 Results

After the course, data from the collected feedback (n=52) was compiled. The average number of students per group was 13 ± 1.8. The average VR session time was 16.8 ±1.8 minutes excluding the initial demonstration and set-up. The average year of study per participant was 1.9±1.1. Mean scores for all categories and the assessment are shown in Table 8.4.

	Mean ± standard deviation						
G	**Cardiac experience**	**VR experience**	**Motivation**	**Mental workload**	**Learning effectiveness**	**Alertness**	**Test score**
1	43.0±22.2	16.0±13.8	74.9±23.5	50.5±15.9	60.1±19.7	62.9±22.9	53.1±14.4
2	51.3±22.1	19.8±17.8	85.9±17.9	54.4±16.4	73.5±18.1	64.5±24.9	42.5±27.7
3	51.5±24.0	24.3±11.6	80.7±19.1	62.8±7.56	55.3±21.2	71.8±17.0	50.8±18.8
4	56.5±20.4	41.1±22.4	67.7±26.0	45.4±16.3	61.2±21.7	45.5±21.7	42.7±17.9
M	50.4±22.1	24.4±18.7	78.1±22.0	53.4±15.5	63.3±20.7	61.8±23.3	47.1±20.9

Table 8.4: Results from feedback and assessment for all groups (G=group, M=mean). All questions are out of 100 (maximum). The first three columns are pre-workshop questions, the last three are post-workshop questions.

Findings showed that there were no statistically significant correlations for any of the feedback metrics with assessment score (all $p > 0.05$). Self-reported prior VR experience showed a strong negative correlation with post-workshop mental workload ($R = -0.36,^{*} p < 0.01$). Self-reported prior VR experience also showed a strong negative correlation with post-workshop reported alertness levels ($R = 0.39,^{*} p < 0.01$). Pre-workshop motivation levels showed a strong positive correlation with post-workshop satisfaction levels ($R = 0.44,^{*} p < 0.001$). No other significant correlations between metrics were found.

8.4.4 Discussion

This pilot study demonstrated the possibility to deliver short VR-based cardiac anatomy workshops for educational and training purposes. Average participant satisfaction was high (63.3±20.7 out of 100), although lower than the self-reported motivation levels pre-session (78.1±22.0). No significant correlations between feedback scores and test scores were found. Therefore, in this study, the effectiveness of VR (evaluated by assessment) was not correlated to the participants' opinions regarding the experience. Students expressing lower levels of experience with VR were found to report higher mental workloads and lowered alertness levels. However, there was no correlation with prior VR experience and satisfaction levels/test scores, which may suggest that teaching effectiveness in a short session (mean = 16.8±1.8 minutes) is maintained even for those with little VR experience. In a future study, participants should be tested with a pre-workshop assessment in order to better quantify the levels of knowledge uptake following VR usage. Additionally, groups should be homogenised according to the year of study, and repeated at numerous intervals in order to be able to retrieve more statistically meaningful correlations. In conclusion, the virtual anatomy lab was well-received by participants and found to be suitable for teaching groups (n<20) cardiac anatomy, although a prolonged study measuring both pre- and post-intervention assessment levels would provide stronger evidence for/against supporting the use of VR-only anatomical workshops.

8.5 Assessment of VR for improving knowledge of cardiac anatomy

8.5.1 Introduction

During the 2019/2020 MSc congenital heart disease module (see Section 8.3.3), a study was set up to (i) evaluate levels of improvement in student understanding of CHD following VR usage, and (ii) record participant feedback regarding the functionality of VheaRts.

8.5.2 Methods

Participants were randomly divided into two groups (A and B). Each student was assigned a unique identifier. Firstly, all participants were required to complete a demographic survey and a pre-class assessment on the knowledge of normal cardiac anatomy, in order to evaluate baseline levels. In this assessment, two different photographs of a normal heart were presented, requiring labelling. The first showed the anterior view of a normal heart specimen, displaying the chambers and great vessels. The second presented a posterior view, showing structures requiring more advanced anatomical knowledge (Fig. 8.8).

Following the assessment, students attended an introductory lecture and anatomical lab session. After teaching had concluded, group A was granted the

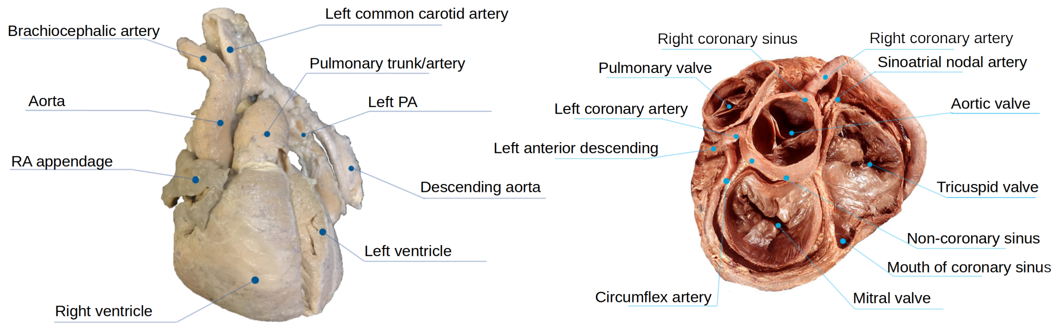

Figure 8.8: Both sides of the normal heart assessment, with all answers shown (n=21). Left: anterior view of normal heart specimen. Right: posterior view of a normal dissected heart.

opportunity to use VheaRts to inspect the normal virtual heart model. Each user was allowed ∼7 minutes to use VR, due to time considerations. Participants from group A and B were then asked to fill in the same post-class test on normal heart anatomy, with group B having completed only the lecture and lab session (without VR).

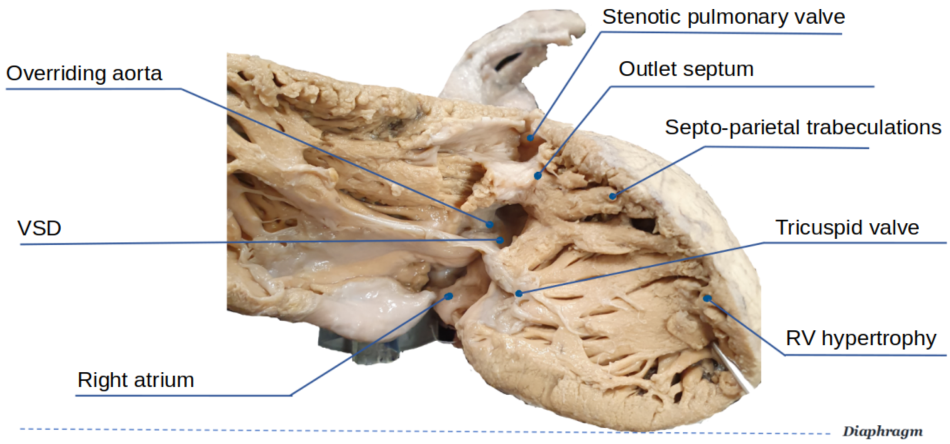

Figure 8.9: TOF survey with lesion-specific questions already answered (n=8).

At a later date, the students attended a second lecture and lab session focused on the anatomy and physiology of TOF. This time, participants from group B only were granted access to the VR TOF model for ∼7 minutes each. A post-class survey consisting of a four-chamber view of a TOF specimen was then given to both groups. Students were asked to identify the structures.

After completing post-class surveys, participants were asked to complete an evaluation feedback survey. Users were asked about their prior experience with VR.

Three questions using a five-point Likert scale asked participants to state: (i) how helpful they found VR for learning cardiac anatomy, (ii) how intuitive VR was to them, and (iii) how keen they would be to use VR continually through the congenital heart disease module. Finally, an empty section was given for respondents to write down any added thoughts relating to VR.

Normality of data was assessed using a Shapiro-Wilk's test. Unless stated otherwise, an independent t-test was performed to compare the means between two groups. For non-normally distributed paired data, a Wilcoxon signed-rank test was used. Otherwise, paired t-tests were used to compare results from individual participants 'pre' and 'post' teaching. A p value of $<$0.05 was regarded as statistically significant.

8.5.3 Results

Twenty-two students enrolled in the congenital heart disease module registered to participate in the study, with 13 in group A (59%) and 9 in group B (41%). The backgrounds of the participants in the two groups are summarised in Table 8.5. In group A, 77% of students reported previous clinical experience, whereas only 23% of students in group B reported clinical experience. Students reporting prior cardiovascular field experience were 62% and 56% for group A and B, respectively.

MSc Student Background	Frequency in group A	Frequency in group B
Biomedical Science BSc	5	4
Current medical student	4	0
Practising physician	2	2
Pharmacology BSc	0	2
Cardiovascular Technology BSc	1	0
Human Genetics BSc	1	0
Medical Science BSc	0	1

Table 8.5: Demographics of MSc students in the two groups.

In all assessments, a correct answer was equal to one mark. No marks were awarded for partially correct responses. Results from all assessments were normally distributed (Shapiro-Wilk test $*p < 0.05$). At baseline (pre-class), average scores for group A and B were 6.85$\pm$4.12 and 2.78$\pm$2.11, respectively ($*p < 0.05$). Overall, the most frequently identified structures were the left ventricle (group

A=12, group B=5), the right ventricle (group A=11, group B=6) and the aorta (group A=11, group B=5). Students with prior clinical experience (7.08±4.23) attained higher marks compared to those without clinical experience (2.90±1.97, $*p < 0.05$).

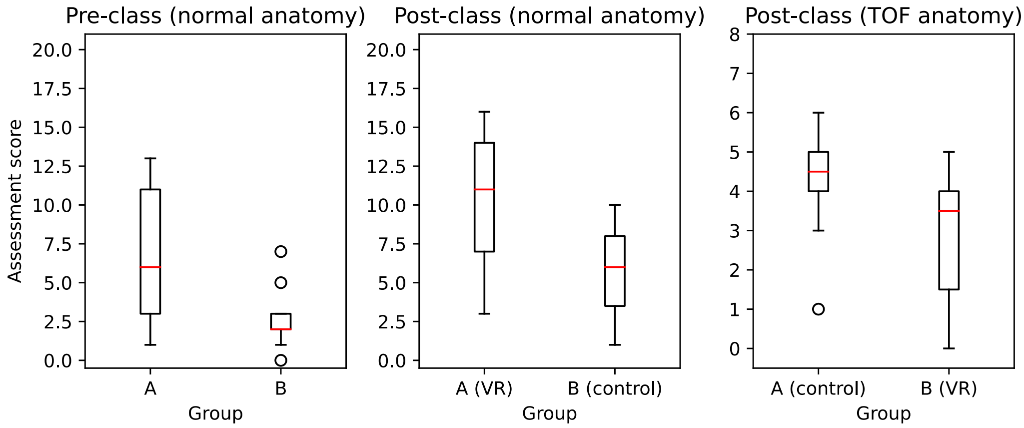

Figure 8.10: Results from all assessments for both intervention and control groups. Pre-class (normal) is also referred to as the baseline test, as no VR takes place.

In the post-class normal anatomy test, group A was exposed to the additional VR session (intervention). The average post-class score for group A and group B was (10.08±4.21) and (5.71 ±3.40) out of a possible 21, respectively ($*p < 0.05$). A paired t-test was performed to compare pre- and post-class results, resulting in both groups showing statistically significant improvement post-class ($*p < 0.05$).

In the post-class TOF anatomy test, group B had the additional VR session (intervention). The total number of possible marks was 8. Average results for group A and B were 4.19±1.28 and 3.00±1.83, respectively. An independent t-test showed no statistically significant difference between the results of cohorts A and B ($p > 0.05$).

Regarding previous experience, only 4 students reported prior experience with VR technology. Fifteen students (68.2%) reported the application to be ‘useful’ or ‘very useful’ for learning CHD. The remainder of respondents were neutral (n=6) or rated the application ‘not useful’ (n=1). VheaRts was deemed to be ‘intuitive’ or ‘very intuitive’ by 16 users (72.7%). Finally, there were 13 students (59.1%) that desired to use the application more frequently throughout the course programme, with the others being neutral (n=6) or not keen (n=3).

8.5.4 Discussion

In the normal heart anatomy assessment, it was found that the majority (95%) of students improved their understanding of anatomical structures (post-class) when compared to the baseline. This improvement was independent of whether or not users were exposed to the VR session. Post-class, students in group A (exposed to VR) performed better in the normal anatomy test than those in group B ($^*p < 0.05$), suggesting that VR may improve knowledge acquisition. However, for the TOF examination, group B (exposed to VR) showed no difference in post-class assessment scores when compared to group A ($p > 0.05$). This is possibly due to the smaller sample size of TOF assessment questions (n=8) when compared to the normal anatomy questions (n=21). Also, another explanation may lie in the fact that group A performed better on the baseline test than group B ($^*p < 0.05$), suggesting group A contained higher performing students. This is supported by the observation that group A was mostly composed of students with prior clinical experience (75.9%), which were also shown to perform better than those without clinical experience ($^*p < 0.05$). Therefore, it can be argued that the benefits of VR (for group B) in the TOF assessment were negated by the inherently better-performing students of group A, thus equalising the two groups. A final explanation for the lack of statistical significance in the TOF comparison is that the TOF VR model was derived from a specimen, with structures less clearly delineated than the normal anatomical VR model (which was an idealised creation). This may have made it more complicated to use as a reference for understanding TOF anatomy, thus hindering its usefulness.

A number of limitations were highlighted in this pilot study. The first was that groups had to be selected based on teaching hours, preventing a homogeneous and equal distribution of students/experience levels. Secondly, this study was focused only on two models of hearts (one normal and one defected). In order to properly understand the extent to which VR could support CHD teaching, more lesions should have been introduced. Each assessment should follow an identical structure in order to make them more directly comparable. In this study, a pre-class assessment of TOF was not possible due to the scheduling, but should be enforced in a future attempt to allow for more useful statistical analysis. Finally, in the future, an idealistic TOF model should be used since the one used was reconstructed from micro-CT images of an anatomical specimen, and thus more difficult to inspect than the normal heart model.

8.6 Assessment of VR for teaching CHD to clinical professionals

8.6.1 Introduction

In addition to undergraduate/postgraduate teaching, it is critical to explore the utility of VR also for teaching clinical professionals. VheaRts was tested during a regular dedicated session in the UCL 'hands-on' Cardiac Morphology Course between 2018 and 2020 (see Section 8.3.4 for details) to support CHD training for professionals. User feedback regarding VR was collected from the participants using metrics such as perceived usefulness and ease-of-use.

8.6.2 Methods

Following the VR demo, users (n=58) were asked to provide feedback via a short questionnaire. Participants were required to provide information related to their professional background and clinical area of expertise. Next, the respondent was asked if they had previous experience with VR, and, if so, to elaborate further on what technologies they had previously tried. Three feedback questions were scored using a five-point Likert scale and asked users about: i) the ease of interaction with the content/tools in VR; ii) the perceived usefulness of VR for improving their understanding of morphology; iii) their willingness to implement a VR experience in their respective working environment. Finally, an open-ended question concluded the survey to capture any further comments. Responses were analysed by calculating average feedback scores and exploring correlations with subgroups.

8.6.3 Results

In total, 58 participants tested VheaRts over seven sessions (average = 8.3 participants per session) for ∼10 minutes each. Twenty participants listed their current profession to be strongly related to cardiology (i.e. cardiology trainee/registrar/fellow, cardiac sonographer, cardiac intensive care fellow and cardiac surgeon). Other commonly reported disciplines included roles such as paediatrician (n=6) and anaesthetist (n=5). From the 58 participants, 11 had previously tried VR before (Fig. 8.12).

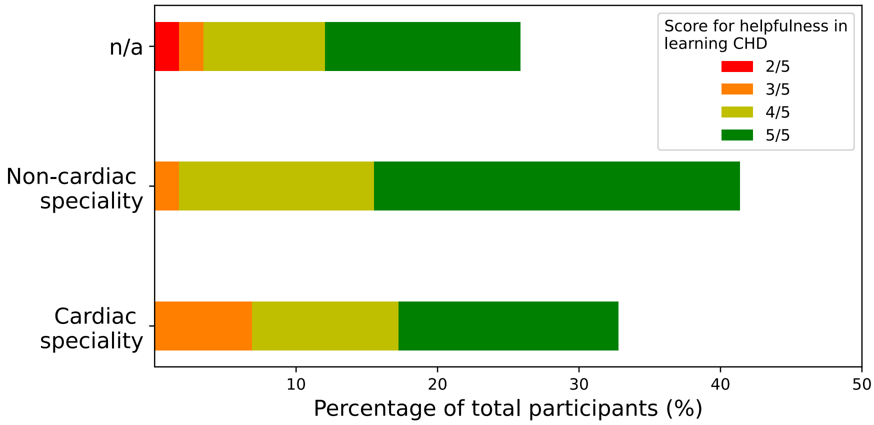

Figure 8.11: Comparing the spread of scores for 'helpfulness' between groups with VR experience and groups without VR experience (n=58).

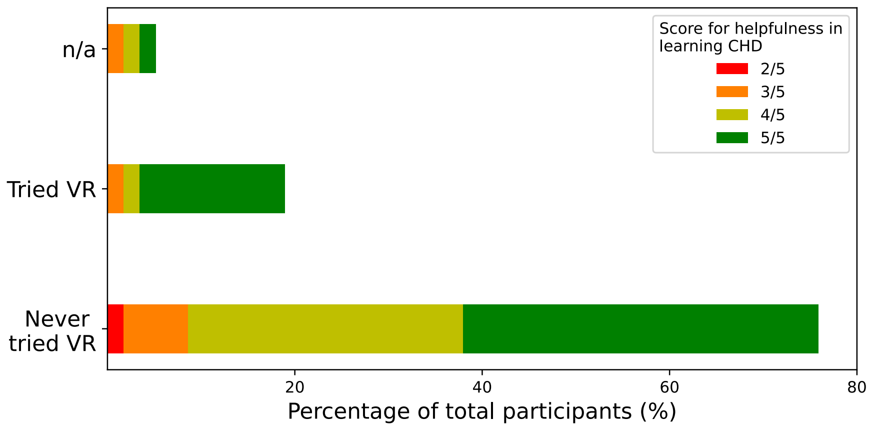

Figure 8.12: Comparing feedback regarding 'helpfulness' of the app between cardiac vs. non-cardiac speciality groups (n=58).

VheaRts was found to be very intuitive and easy to use by 93% of the total participants, with an average feedback score of 4.6 ± 0.8 SD. The slicing tool was the most commonly used feature. The average ‘perceived helpfulness of VR for learning CHD’ score was 4.4 ± 0.6 SD. Out of the total respondents, 88% viewed VR as ‘helpful’ or ‘very helpful’. VR was deemed ‘helpful’ or ‘very helpful’ by 75% of participants from cardiac backgrounds, and 95% from non-cardiac backgrounds (Fig. 8.12). VR was rated as ‘helpful’ or ‘very helpful’ by 91% of participants with prior VR experience, and 89% without prior VR experience. Over 89% of users declared their willingness to implement VR in their clinical practice. This result is observed to correlate with the ’helpfulness’ scores recorded (see Appendix D.2).

A total of 47 participants answered the open-ended question. Participants repeatedly reported that the application was useful (n=14), provided clear understanding of heart anatomy (n=10) or was intuitive and easy to use (n=9). From the 47 who added comments, 28 elaborated further to make suggestions for improvements on the software. Some future requests included: incorporating ultrasound images (n=5) and adding dynamic/beating 3D models (n=5), both of which were implemented in later VheaRts versions.

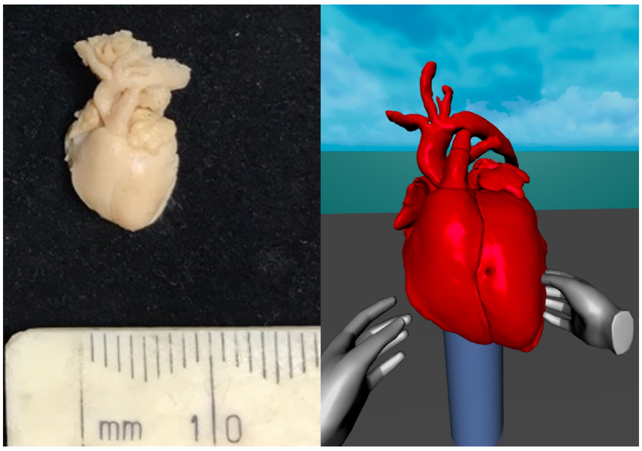

Figure 8.13: The normal 16-week specimen with roughly 1 cm volumetric bounds, compared against its 3D model in VheaRts which can be scaled up freely.

The ability to scale up models was noted to be particularly advantageous for fetal heart specimens, which are difficult to fully inspect with the naked eye (Fig. 8.13). In addition, other discovered benefits of VR (some for specific lesions) were observed during the sessions (Fig. 8.14). For example, it was found that sphere clipping in DORV could allow simultaneous observation of ventriculo-arterial connections and septal defects on the curved cutting plane (Fig. 8.14, top-left). This was more intuitive and easy to accomplish in VheaRts than in conventional cross-sectional imaging (e.g. CT/CMR). Secondly, VheaRts was found to be useful for mapping out the electrical conduction system in the heart with markers, particularly in specimens with uncommon or multiple septal defects such as in the AVSD subject (Fig. 8.14, top-right). After tracing the pathway, the heart could be hidden to leave only the conduction system tracks visible, enabling isolated assessment of 3D conduction pathways. Finally, VheaRts demonstrated that it could be used for precisely assessing the origin of the great arteries. In the TOF case, it was discovered that the aortic root was found to mostly originate from the right ventricle, suggesting a possible association with DORV (Fig. 8.14, bottom).

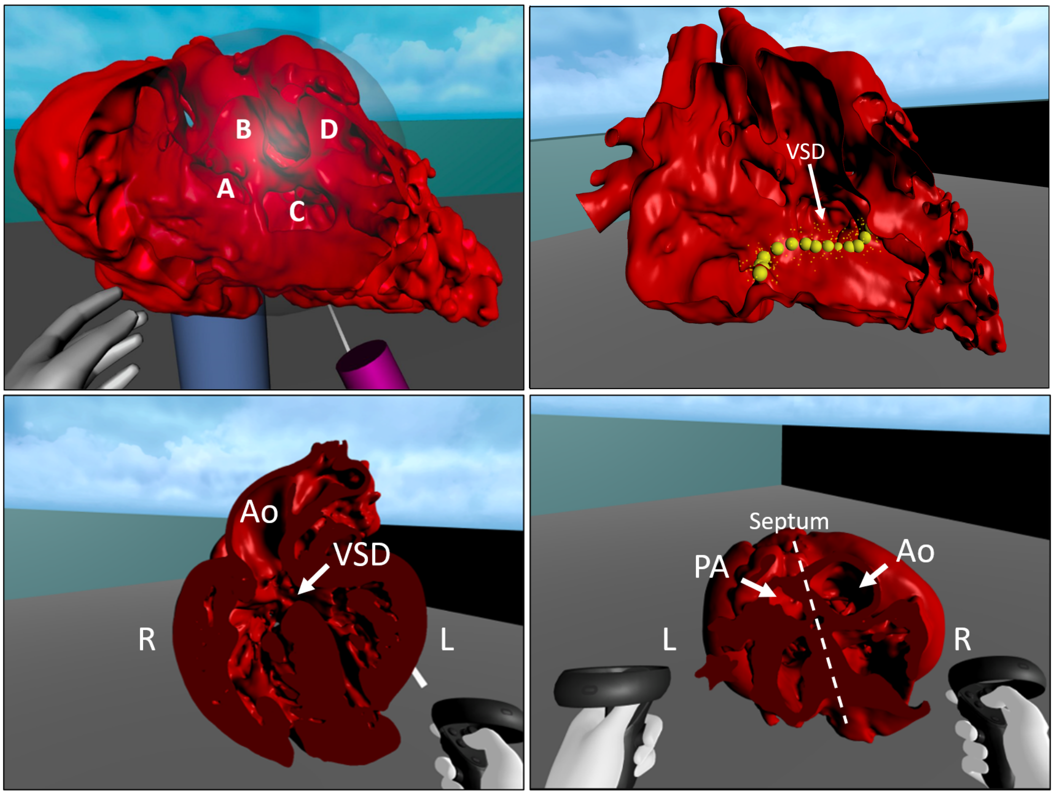

Figure 8.14: Top left: sphere clipper for unique right-sided view in DORV. A = ASD, B = aortic root, C = VSD and D = RVOT and PA root. Top right: conduction system mapping using markers in DORV. Bottom: VR evaluation of VA connections in TOF. Four-chamber and short-axis views display the degree of aortic override.

8.6.4 Discussion

This study explored the feasibility of combining VR with a library of patient-specific heart models for enhancing the cardiac morphology education for professionals. Early findings indicate that the addition of VR to conventional teaching modalities is both viable and extremely well-received by professional clinicians studying CHD.

Averages for each feedback metric were highly positive (all above 4.4 out of 5). 'Intuitiveness' recorded the best result (4.6 $\pm$0.6), despite 76% of attendees having had no prior experience with VR technologies. 'Helpfulness for learning CHD' displayed the lowest mean and largest standard deviation (4.4 $\pm$0.8), likely due to the highly mixed participant demographics, resulting in a wider range of knowledge in CHD. From Fig. 8.11, 82% of users with prior VR experience voted 5 for 'helpfulness', whereas only 50% of users with no prior VR experience voted 5 for 'helpfulness'. This may suggest that users with VR experience are better prepared to interpret the content. This could be due to reduced/removed novelty and unfamiliarity with VR, as described by Moro et al. [228]. From Fig. 8.12, it can be seen that non-cardiac attendees reported the application to be slightly more helpful on average, suggesting that the perceived helpfulness of VR is increased in those with lowered experience levels.

Some important limitations related to the study were noted. Firstly, the survey was designed to be concise due to time constraints in the course. In the future, the feedback forms should be expanded to include more specific questions (e.g. related to individual tools). More precise demographic data (e.g. years of experience, specific profession) should also be gathered. Assessments may also be performed in the future to evaluate participant knowledge 'pre' and 'post' VR [229]. Finally, by including VR surgical/clinical simulators, the users could potentially be assessed in the future by correlating perceived levels of improvement to tangible clinical metrics measured through simulation [138].

8.7 Summary and conclusion

The implementation of VheaRts in various educational settings related to CHD has been described. VheaRts was used for teaching more than 240 students over the course of its implementation at UCL and Great Ormond Street Hospital. This spanned a range of demographics and experience levels. The details of the virtual anatomy lab were presented, including applications in teaching, undergraduate and postgraduate students, and in professional settings (such as echocardiography simulation and DORV workshops). During the COVID-19 pandemic remote learning, the online application supported teaching to small to medium-sized groups. A few pilot studies measured the quantitative benefits of using VR for learning cardiac anatomy and showed promise for enhancing conventional anatomical teaching. User feedback was overall positive, with the VheaRts platform being well-received. In conclusion, VR shows significant promise as a tool for enhancing the understanding of complex cardiac anatomy over a wide range of disciplines and background experience levels.

Chapter 9

Conclusions and future developments

9.1 Contribution summary

This thesis aimed to explore pathways for translating potentially high-impact computational tools into CHD standard clinical/teaching applications. In the introduction, it was stated that the main goals were two-fold: (i) to develop a more clinically suitable pipeline for CFD by accelerating and automating processes, and (ii) to develop and implement a VR application in clinical and educational settings. In this thesis, it was shown that ML-based segmentation/CFD models were capable of approximating the ground-truth data with reasonable accuracy. Future steps towards clinical implementation should involve the use of more realistic simulation data and improved shape parameterisation methods. This should be followed by validation against acquired clinical metrics, such as pressure gradients. Also, as part of this thesis, the development and implementation of a VR platform for CHD teaching/clinics was demonstrated. This early exploration showed that VR was well-received in multiple applications, such as for the planning of complex DORV repair and for teaching cardiac anatomy to students of various experience levels. Future efforts should involve a multi-centre approach for evaluating the usefulness of VR, especially in pre-operative settings, through the collection of clinical data over a prolonged period. Overall, I have shown that the translation of CFD and VR towards clinical applications is possible, and with necessary improvements, may contribute to the teaching and clinical care of CHD on a wider scale.

9.2 Contributions from each chapter

In Chapter 4, the development of a fast, automatic image segmentation model for the great arteries was presented. The model was built using a deep neural network architecture (U-net), and was evaluated using conventional imaging metrics. The use of this automatic segmentation model for producing 'CFD-suitable' reconstructions was explored. CFD was performed on a selection of aortas and pulmonary arteries inferred by the model, followed by comparison against ground truth data. The model produced surfaces with CFD accuracy equivalent to human inter-observer variability when performing manual segmentation. Therefore, outcomes from this chapter suggest fast automatic segmentation models for the great arteries are suitable for clinical CFD, thus accelerating the overall CFD pipeline greatly by removing the initial segmentation step.

In Chapter 5, a framework for fully automatic and accelerated CFD of aortic geometries was presented to complement the fast automatic segmentation model. The two-step approach included a statistical shape model and PCA for parameterising shape/flow, followed by a deep neural network for regressing pressure and velocity CFD characteristics. The model was trained using a synthetically generated population of realistic aortas. Results showed that reasonable CFD pressure/velocity flow fields could be inferred for unseen geometries, although errors in shape parameterisation were significant. It was demonstrated that with this approach, large computational requirements or simulation setups can be avoided. For this reason, this pipeline is highly suitable for clinical implementation and addresses the primary barriers hindering current CFD methods, although further improvements are needed before full utilisation.

Chapter 6 presented VheaRts, a Unity-based VR platform for assessing 3D patient heart models or learning the anatomy of CHD. Technical details related to the development were discussed, including the software architecture and individual tool designed. The integration of networking in the application was shown. Specific features were implemented for CHD which enabled the support of both clinical and educational activities within university and hospital.

In Chapter 7, clinical applications of VheaRts were shown. An overview of the patient cases planned pre-operatively with VheaRts was shown, including specific case studies. A pilot retrospective study into the impact of VR for determining the feasibility of biventricular repair in DORV was conducted. Findings showed that adding VR to other 3D imaging modalities increased the likelihood for the surgeon to decide on the correct surgical approach. Finally, the use of VheaRts was extended beyond planning in CHD to the surgical planning of three complex cases of conjoined twins.

Chapter 8 presented the implementation of VheaRts in educational settings at university and hospital courses. An overview of the 3D library of heart models currently available in the virtual anatomy lab was given. Teaching with VheaRts included undergraduate and postgraduate students; an assessment of VR for improving the structural knowledge of cardiac anatomy suggested that VR usage notably improved knowledge acquisition of the basic heart anatomy, yet there was inconclusive evidence regarding CHD lesions. VheaRts showed great potential also in the education of clinical professionals, with applications including CHD morphology teaching, echocardiography training and above all the possibility of online delivery of specialised courses to units around the world.

9.3 Limitations

A common limitation between Chapter 4 and 5 was the use of simplified CFD models. The model included laminar, steady-state flow without patient-specific boundary conditions. This enabled easier convergence for large-scale CFD simulation. Additionally, it could be argued that simplified CFD made it easier to isolate the effect of segmentation differences (Chapter 4) or correlate aortic shape features to flow features (Chapter 5). However, whilst the current approach may have been sufficient for demonstrating the feasibility of both applications, more realistic CFD is necessary in order to prove the suitability of the approaches in direct clinical applications. Some other limitations include the methods performed for evaluation. For example, in both Chapter 4 and 5, pressure gradient data was computed by averaging cross-sections along the length of vessels. In Chapter 5, a more granular method for comparing CFD volumes was possible, due to point-to-point correspondence. In Chapter 5, the accuracy of the ML model was dependent on the accuracy of the SSM surface registration; therefore, improved shape parameterisation was necessary to compute more realistic flow fields.

In Chapter 7 and 8, much of the work performed was exploratory and investigated the feasibility of implementing VR in various clinical/educational settings, using mostly a combination of surveys and qualitative assessments. The study of VR into planning the repair of DORV in Section 7.4 was limited by the assumption that the performed repair was always the optimal choice. Additionally, since observers were always exposed to 3D imaging modalities in a fixed order, bias may have been introduced and resulted in over-inflated results favouring VR visualisation. In Chapter 8, qualitative feedback was the primary method for evaluating user experience in VR, with short surveys typically used due to imposed time-constraints in classroom environments. Since questionnaires varied amongst the studies, this made results difficult to compare and aggregate since the data was not consistent. In Section 8.5, the effectiveness of VR for improving CHD knowledge was assessed. However, the short-term aspect of the study and inability to test multiple lesions over the period of the course made it more challenging to establish strong supporting evidence for/against the use of VR for teaching CHD.

9.4 Future work

In the future, the study from Chapter 4 should be repeated using more realistic CFD models which include the head and neck vessels. Similarly, a more realistic CFD model should be used to create the training data in 5, including synthetic boundary conditions for aortas generated by patient-specific flow data. Due to the increased complexity and number of parameters, more synthetic aortas will need to be generated. An improved SSM should be trialled, including additional shape features for capturing the geometric properties not described by the current SSM. This may include centreline-based data, which can also act as regularisation during model training. Transient CFD models should be tested with the use of LSTM DNN architectures. ML models for other CFD parameters of interest such as wall shear stress should also be built. Training datasets with 'exercise boundary conditions' should also be created, in order to be able to simulate hemodynamics under stress in patients which are high risk.

Once reasonable performance is demonstrated on aortas, the pipeline should be extended to the pulmonary arteries and to more complex surgical structures such as TCPC. Finally, an end-to-end image to CFD report framework should be assembled, by aligning together: (i) automatic segmentation of the great arteries, (ii) mesh pre-processing VMTK pipeline, (iii) deep neural network CFD model and (iv) a post-processing GUI tool for visualisation of the CFD results. This would enable rapid approximation of a full 3D CFD flow field, requiring only patient inlet conditions and a 3D image dataset. Following this, the automatic CFD tool could be deployed in clinical situations and researched to assess its potential for improving the clinical management of patients. Possible developments (in the more distant future) could include extending this approach to predict flows in post-intervention aortas, using a two-step approach that involves: (i) an ML model for predicting the ideal post-op aortic shape, and (ii) a CFD ML model for predicting the post-op hemodynamics.

Future work for VheaRts includes the development and improvement of new features and tools, in particular to update networking capabilities and visualisation options. Restructuring for scalability (for more easily implementing in multiple centres) is first needed, including adapting the VheaRts code to different HMD brands and models. In terms of specific features, new shaders should be developed to better render transparency in objects. The post-processing stack in the echocardiography probe could be improved to produce a more realistic ultrasound effect. A greater suite of 3D mesh editing tools should be implemented, such as dynamic cutting of mesh and coupling with biomechanical tissue simulators. New options for students in the virtual anatomy lab could be implemented, such as being able to interact with a personal, offline version of the currently active heart during classroom sessions. This would create less passive teaching environments where students could remain engaged for longer periods of time.

Research into the clinical benefit of VR could be conducted by performing larger and prospective studies, with additional questions and measured parameters for extracting qualitative and quantitative data. Specific information related to the surgical questions should be gathered (with follow-ups), in order to isolate the potential ramifications on the surgery. For example: pre-VR: "is the RV too small for biventricular repair?", post-VR: "in VR, I was satisfied that the RV is insufficiently sized for biventricular repair", post-surgery: "the RV cavity was confirmed to be too small in the operating theatre". Additionally, there should be a greater effort to assess the effectiveness of VR with quantitative measures, for example by recording surgical information (cross-clamp time, bypass time) from patients assessed pre-op with VR and comparing the data to a similar control population without VR access.

Similarly, future research into the educational aspects of VR in CHD could be improved by repeating studies on larger cohorts and with more granular questionnaires. The study in Section 8.5 should be re-run over a longer time-frame, testing students on multiple lesions and also evaluating the retention of knowledge by repeating examinations. As an alternative to labelling structures in 2D (paper-based examination), 3D examinations could be implemented in VR (by marking structures). To support the deployment of VR in wider-educational settings, strong quantitative evidence would be required which demonstrated faster knowledge uptake and improved performance longer-term. Outside purely anatomical education, simulators (echocardiography, surgical) and more tools could be implemented, in order to improve trainee knowledge. Exploration of the potential benefits using similar feedback and assessment methods could be conducted to further establish the usefulness of VR in CHD teaching/training.

9.5 Final remarks

This thesis showed the feasibility of translating two potentially high-impact computational tools (CFD and VR) for the clinical management/education of CHD. I strongly believe that with more work, CFD and VR could be fully implemented for applications in the clinical practice, management and education of CHD. The potential in an automatic and accelerated clinical CFD model in CHD care is significant, primarily for reasons such as risk stratification and decision-making. In the more distant future, I foresee that ML-based approaches could be used to design interactive physics-derived tools for applications such as virtual intervention (e.g. aortic stenting). I also believe that VR has a future significance in CHD, both for surgical planning and teaching via virtual classrooms. With improved software, hardware and accessibility of HMD devices, it is highly likely that the uptake of VR in clinical and teaching institutions will increase. For this reason, it is of paramount importance to properly research the usefulness of VR, and in particular where and how it should be applied. Most importantly, there are no practical restrictions preventing the approaches presented in this thesis from being utilised into other healthcare (non-cardiac) applications. At present, the prospect of using ML, CFD and VR to tackle complex problems in healthcare is an exciting one, and is gathering more and more traction and momentum each year.

Appendix A

Chapter 4

A.1 Sensitivity Analysis: Laminar vs Turbulent

It was found that an element count of 300,000 was well-suited to capture the flow details in both the aorta and the PA, and for yielding a stable solution. Iterations of 300 and 750 were sufficient for the aorta and PA respectively, in order to reach convergence and accurate results.

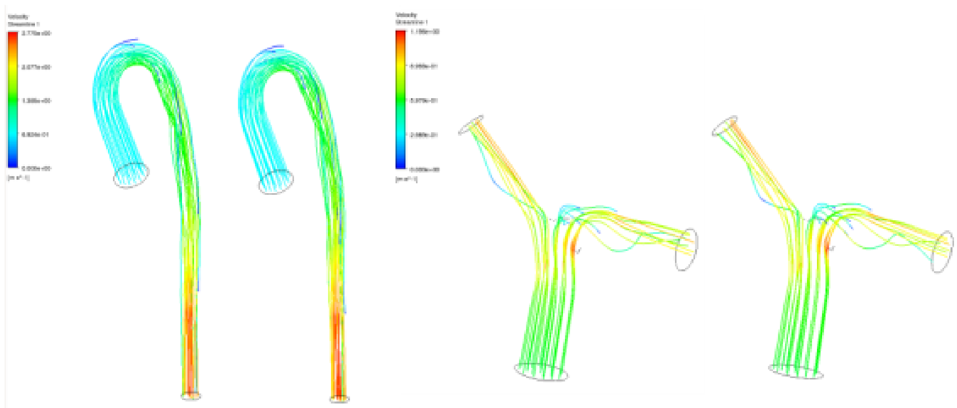

Figure A.1: Difference velocity fields in Ao and PA when applying a k-omega turbulence model(left = laminar, right = turbulent).

The gradient error (MAPE) for the aorta and PA (∼300,000 elements) vs the highest resolution aorta and PA (1,000,000+ elements) were 2.0% and 2.15% for pressure, respectively. The error for the velocity was 3.4% and 3.23%, respectively. Lastly, the Reynold's number was computed on both test cases, and resulted in 3221.8 and 4250.5 for the aorta and pulmonary artery, respectively. The same simulations were carried out with a k-omega turbulence model (with default parameters), which showed no difference when compared to the laminar model.

A.2 Hyperparameter Optimization

Training, including hyperparameter optimization, took 24 h, during which a total of 124 hyperparameter configurations were sampled by the Hyperband iterations. The top 10 performing configurations in terms of mean validation Dice score are reported in Table A.1 and Fig. A.2.

Table A.1: Top 10 hyperparameter configurations.

#	Scales	Layers per block	Initial filters	Learning rate	Batch size	Loss function	Dice
1	3	2	64	$3.46 \cdot 10^{-4}$	2	Focal Tversky	0.946
2	3	4	32	$1.99 \cdot 10^{-4}$	2	Tversky	0.944
3	4	3	32	$3.40 \cdot 10^{-4}$	4	Dice	0.944
4	4	3	32	$1.37 \cdot 10^{-4}$	4	IoU	0.943
5	3	3	64	$6.77 \cdot 10^{-4}$	4	IoU	0.939
6	3	3	64	$3.27 \cdot 10^{-4}$	4	Dice	0.938
7	4	3	32	$1.81 \cdot 10^{-4}$	2	Dice	0.938
8	3	4	32	$1.83 \cdot 10^{-3}$	2	Focal Tversky	0.937
9	2	3	32	$1.24 \cdot 10^{-3}$	4	Jaccard	0.935
10	3	2	32	$4.68 \cdot 10^{-3}$	4	Dice	0.934

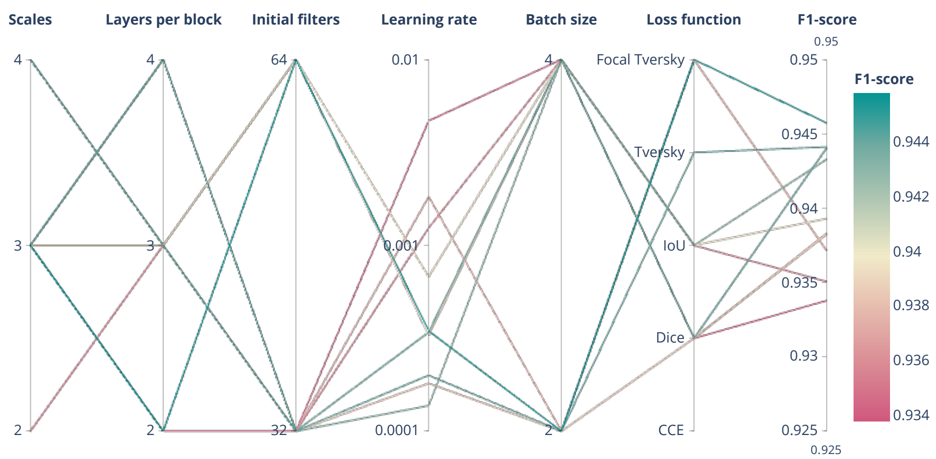

Figure A.2: Parallel coordinates view of the top 10 hyperparameter configurations. Each hyperparameter as well as the mean validation Dice score is shown on its own axis. Each coloured line represents a hyperparameter combination, with vertices at the corresponding values on the parallel axes.

The best performing configuration was as follows: scales = 3, layers per block = 2, initial filters = 64, learning rate = $3.46 \cdot 10^{-4}$, batch size = 2, and loss function = focal Tversky. This model was selected and used in all further experiments.

The top 8 performing configurations have 3 or 4 scales, which suggests that models with 2 scales may have too limited representational power or deep layer receptive fields. Models with at least 3 layers per block also tended to perform better, with the notable exception of the best performing model, which had 2. A relatively small learning rate also seems to be advantageous, with 6 of the top 7 configurations below 0.0005. Confusion-based losses seem to outperform cross-entropy, which does not appear on any of the top 10 configurations, but none of those appears to consistently outperform the rest.

A.3 Manual and Equal Clipping

In all cases, the equally clipped data produces better agreement between the ML and GT CFD simulations (Fig. A.3). The median pressure and velocity error in aortas is reduced through equal clipping by 2.1 and 1.0 percentage points, respectively. The median pressure and velocity error in PAs is reduced through equal clipping by 4.2 and 2.7 percentage points, respectively. The magnitude of outliers is also reduced through equal clipping.

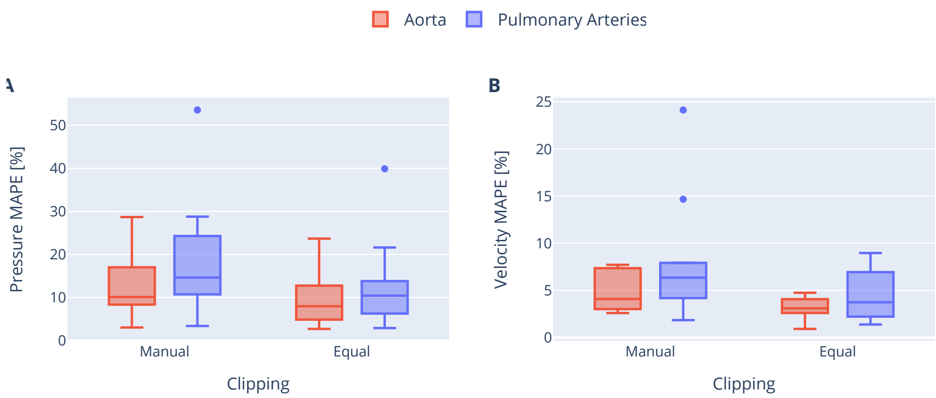

Figure A.3: Pressure and velocity mean average percentage errors (MAPE) for manual and equally clipped aortas and pulmonary arteries.

Appendix B

Chapter 5

B.1 Template volume mesh deformation

Presented is an example of template volume mesh deformation when generating the volume mesh for a new subject. In Deformetrica, the smoothness of ambient space deformations preserves the relative spatial positions of nodes in the meshes, however this affects mesh skewness which makes these meshes unsuitable for CFD simulation.

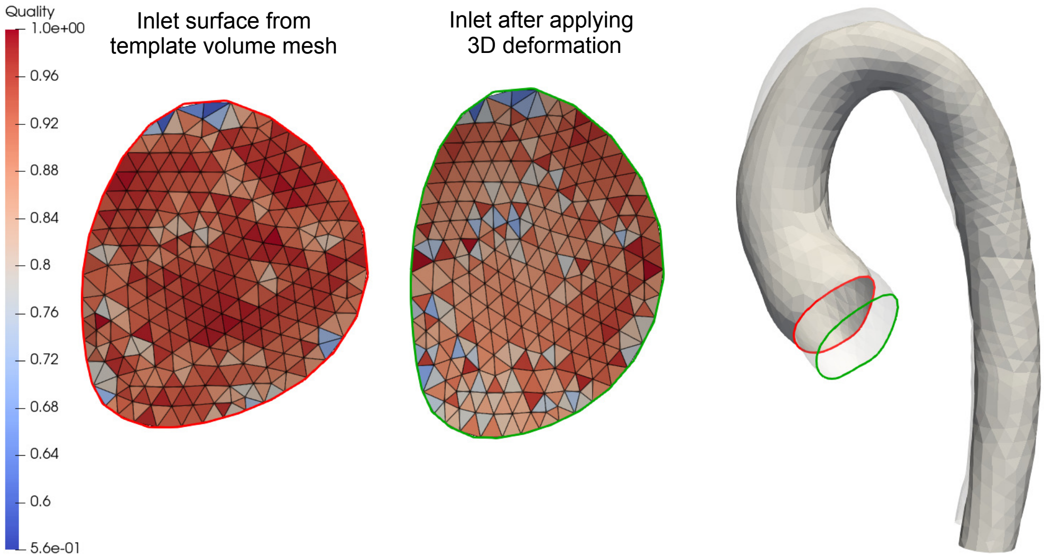

Figure B.1: Mesh skewness due to SSM. The inlet surface of the template volume mesh pre-deformation (left, red boundary), and the inlet post-deformation (right, green boundary) are presented. Mesh skewness (Jacobian) is plotted, with a result of 1 implying optimal triangle quality.

B.2 CFD sensitivity study

A test aortic shape with a sharp arch angulation was chosen to be used in a mesh sensitivity study. Pressure was monitored at a plane in the arch, as well as along the entire centerline. The mean percentage error between the highest resolution solution (∼1e6 cells) and our chosen mesh resolution (∼3e5 cells) was 2.0% and 3.4% for pressure and velocity, respectively. We determined that for this study, this level of convergence was a suitable compromise in order to prove the feasibility of our approach and reduce the computational burden during training data generation. During batch CFD simulation, no divergence of errors was observed. Error residuals (momentum and continuity) were all observed to fall under the 1e-3 threshold set. Pressure was seen to converge more closely at the outlet, and velocity was seen to converge more closely at the inlet. This is assumed to be due to the proximity to the boundary conditions (velocity inlet, pressure outlet).

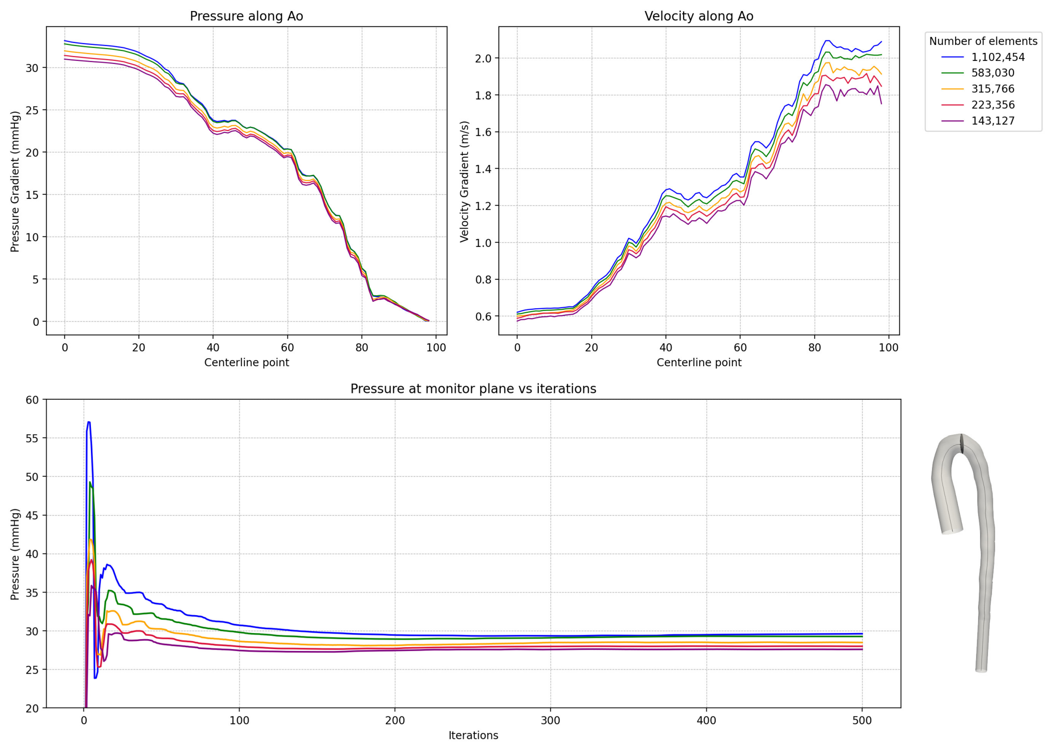

Figure B.2: Mesh sensitivity CFD plots. Top left: pressure along the centreline for 5 mesh resolutions. Top right: velocity-magnitude along the centreline for the 5 mesh resolutions. Bottom: Pressure monitored over 500 iterations. Iterations = computations of the steady-state solver (no time-dependent terms).

B.3 Best and worst PCA errors

The comparison in flow fields is shown in the worst PCA reconstructions (performed on the test-set, n=200). The data is previously unseen from the PCA model. PCA is reversed by undoing the SVD computation described in the Methods section. The subject with highest errors in pressure after reconstruction with 20 PCA modes had an error of MNAE_S=4.32%. The subject with highest errors in velocity after reconstruction with 55 PCA modes had an error of MNAE_S=4.93%. In Fig. B.3, negative pressure values are respective to the reference gauge pressure set during simulations (atmospheric pressure, 100kPa).

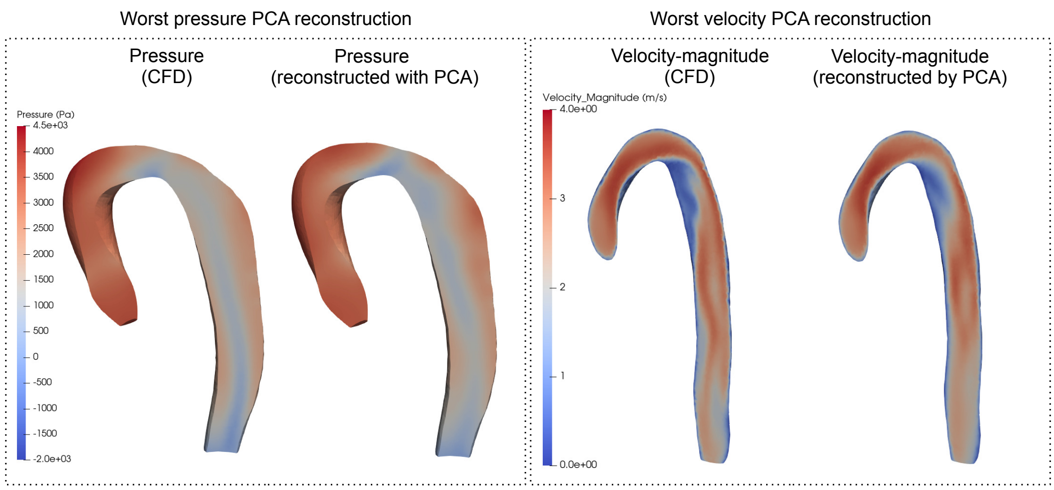

Figure B.3: Cases with worst PCA reconstruction errors. Left: subject with highest pressure MNAE_S after reconstruction with 20 PCA modes (4.32%). Right: subject with highest velocity MNAE_S after reconstruction with 55 PCA modes (4.93%).

B.4 Pressure and velocity DNN architectures

The final pressure and velocity DNN architectures. In both networks, 4 layers were found to be the optimal configuration. Rectified linear units (ReLU) are used in each hidden layer, with linear units for the output layer.

B.5 Further Bland-Altman analysis

Regional Bland-Altman analysis. This was conducted on three different regions of the aorta (ascending, transverse arch, descending). Normalised error is equivalent

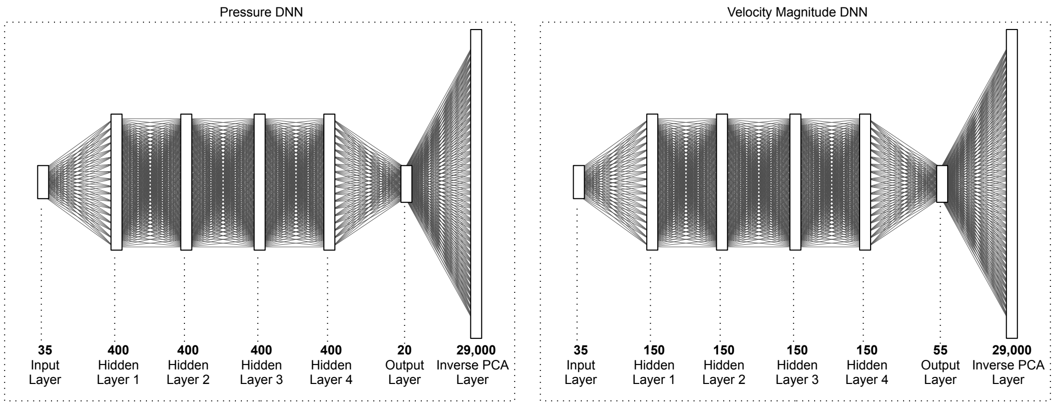

Figure B.4: DNN model diagrams. The architecture and parameters of both pressure and velocity-magnitude networks is shown after hyperparameter optimisation.

to the NAE of each node (without taking the absolute value). The percentage of nodes lying in the ascending aorta, transverse arch and descending aorta are 38%, 22% and 40%, respectively. It was observed that there was no evidence of any significant systematic over-estimation or under-estimation for both models in all three regions.

B.6 Projection of new shapes on PCA axes

PCA projections of real cases. All ten new prospective real patient cases were registered using the shape model and had their corresponding deformations decomposed using the shape PCA model. The first 35 shape mode projections were taken for each subject and plotted against the range of the original dataset scores (n=67). It can be seen that in almost all cases, the 10 subjects fall within the interquartile or maximum-minimum range of the original cohort.

B.7 Worst aortic pressure prediction

Shown are additional angles of the worst aortic pressure case from the test set (n=200). Of particular note is the descending aorta flow extension, which has a visibly reduced diameter and angulation in the ground truth segmentation. This constriction is expected to be the cause of the local flow acceleration and resulting pressure elevation in the ascending aortic domain.

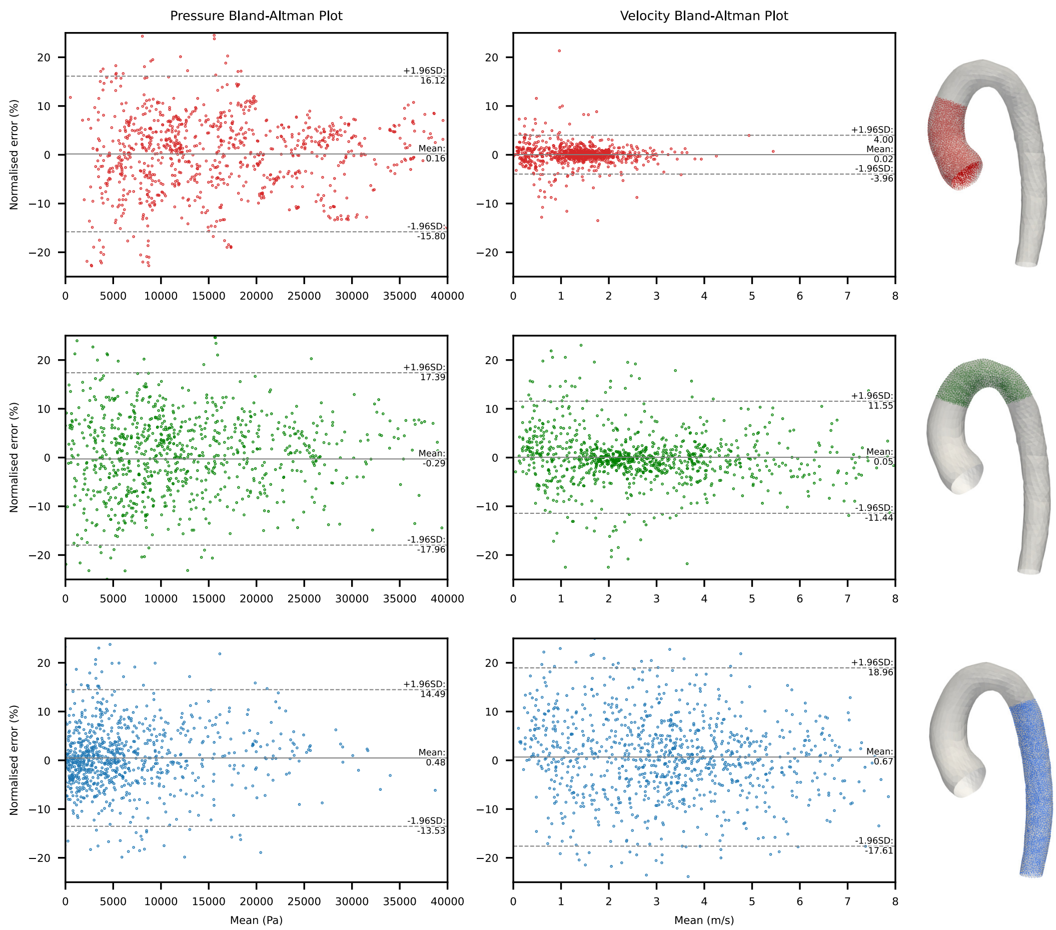

Figure B.5: Further Bland-Altman analysis. Errors in the three primary anatomical sections of the aorta are inspected. The ascending, transverse arch and descending aortic regions are all shown.

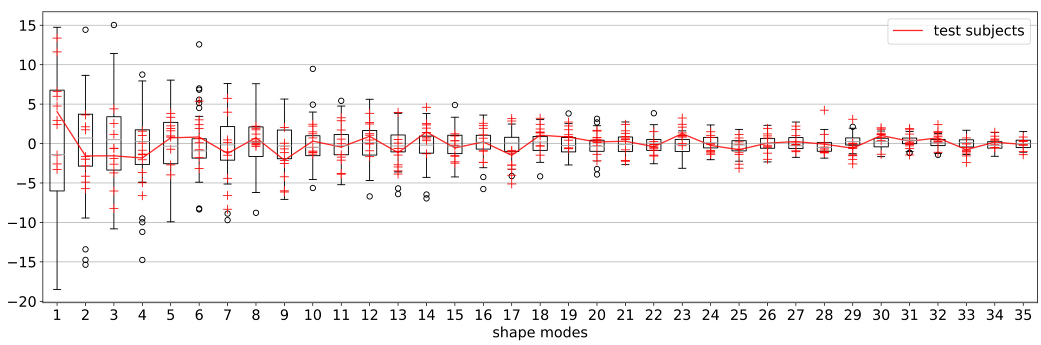

Figure B.6: Prospective test cases shape PCA scores. The 35 shape mode projections for all 10 new cases are shown. Boxplots represent the range of PCA mode scores in the original cohort (n=67). Whiskers are 1.5 × the interquartile range. Red markers are the PCA scores for the new subjects (n=10).

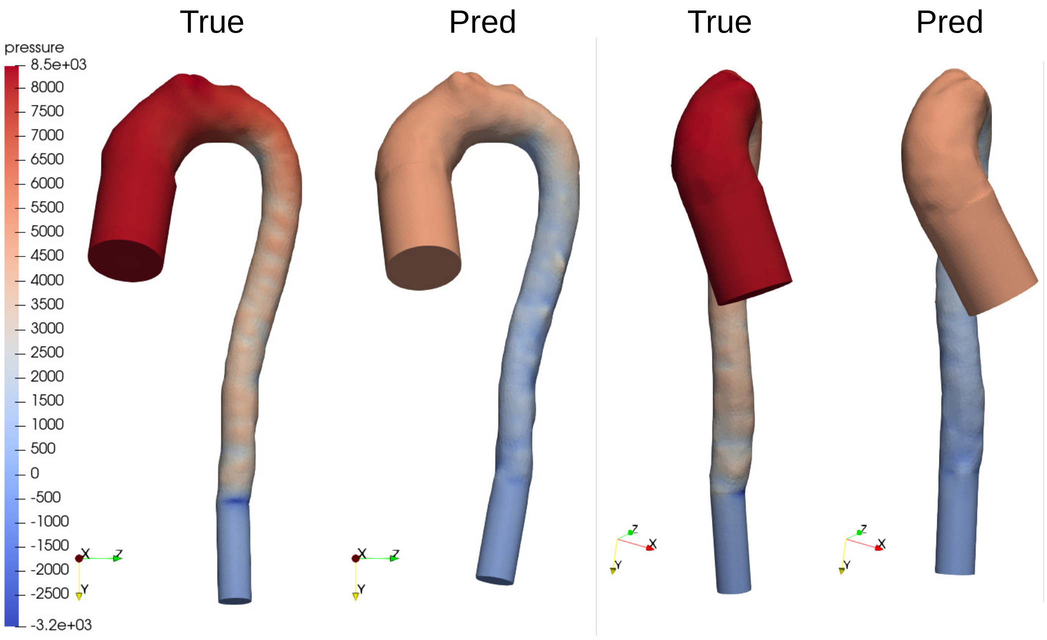

Figure B.7: Worst aorta pressure prediction. Shown are two different angles (from the side and front) of the worst aortic pressure prediction ($\mathrm{MNAE_S}$ = 23.6%)

Appendix C

Chapter 7

C.1 Follow-up data

The table presented here includes all the relevant follow-up information for all 10 DORV patients in the study. Included is the time of follow-up, information regarding re-operations during follow-up, LVOT max discharge velocity and RVOT max discharge velocity. All patients had been repaired with a biventricular approach.

Patient case	Time of follow-up (months)	Reoperations during follow-up	LVOT max velocity (m/s)	RVOT max velocity (m/s)
1	44.6	Closure of multiple apical VSDs	1.5	1.7
2	33.6	No	1.5	1.5
3	40.5	No	3.3	2.1
4	28.6	No	2.2	2.4
5	39.9	No	4.1	1.9
6	43.8	No	1.5	1.5
7	23	LVOTO relief (2x)	1.5	1.3
8	10.2	No	1.3	2
9	24.2	No	1.3	1.7
10	16	NO	1.3	1.5

Table C.1: Follow-up data for all DORV cases.

Appendix D

Chapter 8

D.1 Assessment questions for undergraduate CHD VR workshop

The questions that were posed to attendees of the workshop are presented here. Only one answer is correct for each multiple-choice question (shown in bold).

1. In the normal heart, the arterial roots: (A) are in parallel arrangement, (B) are inferiorly located, (C) are both on the right-side, **(D) cross one-another,** or (E) vary in position.

2. In regular transposition (TGA), the aorta is most commonly: (A) left-sided, (B) centrally positioned, (C) posterior and to the left, (D) anterior and to the left, or **(E) anterior and to the right.**

3. In a typical arterial switch operation (ASO), the following manoeuvre is performed: (A) Leiden, **(B) Lecompte,** (C) Namagashi, (D) Metras, or (E) REV.

4. After a typical arterial switch (ASO), the pulmonary artery branches are: **(A) anterior to the aorta,** (B) posterior to the aorta, (C) to the left-side of the aorta, (D) to the right-side of the aorta, or (E) variable in position.

5. In transposition (TGA), the left-hand facing sinus typically gives rise to: (A) all of the coronary arteries, (B) the left coronary artery, **(C) the right coro-**

nary artery, (D) the circumflex coronary artery, or (E) both the left and circumflex coronary arteries.

6. Regular transposition (TGA) can be defined as: (A) separation of aorta from the mitral valve, (B) separation of the pulmonary trunk from the mitral valve, (C) aorta and pulmonary trunk arising from the LV, (D) aorta and pulmonary trunk arising from the RV, **(E) aorta arising from the RV and pulmonary trunk from the LV.**

7. In terms of segmental connections, regular transposition(TGA) is defined as: (A) discordant AV + concordant VA connections, (B) discordant AV + discordant VA connections, (C) concordant AV + variable VA connections, **(D) concordant AV + discordant VA connections,** or (E) concordant AV + concordant VA connections.

8. In the arterial switch operation (ASO), the coronary arteries are: (A) ligated prior to transfer, (B) divided and sutured circumferentially, **(C) removed on buttons of aortic wall,** (D) not a consideration for ASO, or (E) prepared as the first step of the ASO.

9. In the typical arterial switch operation (ASO), the 'switch' refers to the: (A) arterial valves, **(B) arterial trunks**, (C) outflow tracts, (D) coronary arteries, or (E) the branch pulmonary arteries.

10. A typical arterial switch operation can be prevented by: **(A) severe left ventricular outflow obstruction,** (B) a bicuspid aortic valve, (C) an intramural coronary artery, (D) aortic coarctation, or (E) a large ventricular septal defect.

D.2 Willingness of participants to implement VR

This graph details the willingness of participants to implement VR in their practice, with respect to how useful/helpful they perceived VR during their session. It can be seen that the large majority of users who found VR useful were also very willing to implement it in their practice. One user reported a very high interest to implement VR in their practice, despite reporting a 2/5 score for usefulness of VR in understanding CHD.

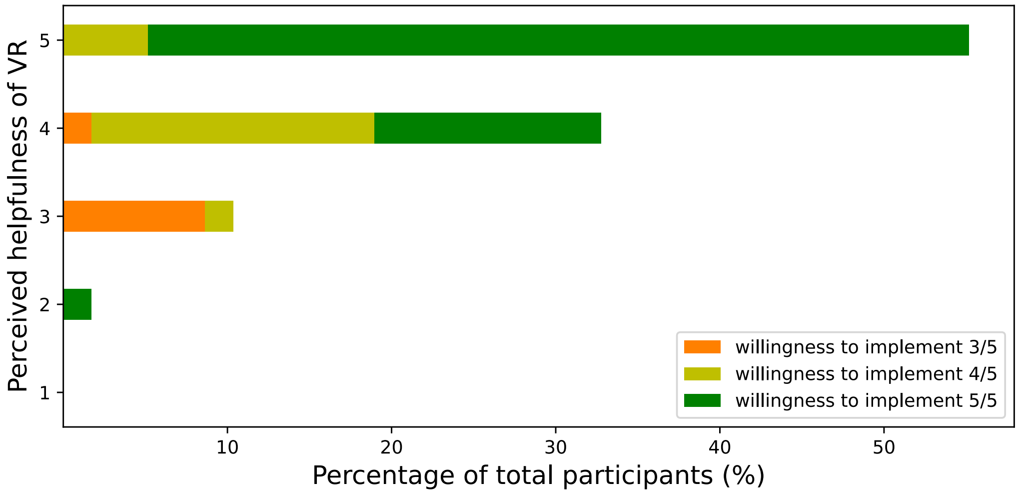

Figure D.1: Relationship between scores for 'helpfulness' and 'willingness to implement' (n=58).

Bibliography

[1] Mayo Clinic. Chd anomalies pictured on wesbite.

[2] ecgwaves.com. Two dimensional echocardiography.

[3] radiopaedia.org. Cardiac angiogram of coarctation of the aorta.

[4] Julie M Harris. Binocular vision: moving closer to reality. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 362(1825):2721–2739, 2004.

[5] Oculus VR LLC, Boyd Drew Allin, and Robin Michael Miller. Oculus vr patent: System for discharging heat out of head-mounted display based on hybrid fan and heat pipe u.s. 2018/0307282 a1, 2018.

[6] Oculus VR LLC, Peter Wesley Bristol, Yi-Yaun Chen, and Glen Jason Tompkins. Oculus vr patent: Hand-held controller using led tracking ring u.s. 2018/0311575 a1, 2018.

[7] Oculus VR LLC, Mark Alan Tempel, Jared I. Drinkwater, Robin Michael Miller, and Joseph Patrick Sullivan. Oculus vr patent: Dual-screen head-mounted displays u.s. 2016/036699 a1, 2016.

[8] Rajesh Puranik, Vivek Muthurangu, David S Celermajer, and Andrew M Taylor. Congenital heart disease and multi-modality imaging. *Heart, Lung and Circulation*, 19(3):133–144, 2010.

[9] Yingjuan Liu, Sen Chen, Liesl Zühlke, Graeme C Black, Mun-kit Choy, Ningxiu Li, and Bernard D Keavney. Global birth prevalence of congeni-

tal heart defects 1970–2017: updated systematic review and meta-analysis of 260 studies. *International Journal of Epidemiology*, 48(2):455–463, 02 2019.

[10] Denise Van Der Linde, Elisabeth EM Konings, Maarten A Slager, Maarten Witsenburg, Willem A Helbing, Johanna JM Takkenberg, and Jolien W Roos-Hesselink. Birth prevalence of congenital heart disease worldwide: a systematic review and meta-analysis. *Journal of the American College of Cardiology*, 58(21):2241–2247, 2011.

[11] SC Mitchell, SB Korones, and HW Berendes. Congenital heart disease in 56,109 births incidence and natural history. *Circulation*, 43(3):323–332, 1971.

[12] Benoit G Bruneau. The developmental genetics of congenital heart disease. *Nature*, 451(7181):943–948, 2008.

[13] Giulia Ottaviani and L Maximilian Buja. Congenital heart disease: pathology, natural history, and interventions. In *Cardiovascular pathology*, pages 223–264. Elsevier, 2022.

[14] Kimberly A Holst, Sameh M Said, Timothy J Nelson, Bryan C Cannon, and Joseph A Dearani. Current interventional and surgical management of congenital heart disease: specific focus on valvular disease and cardiac arrhythmias. *Circulation research*, 120(6):1027–1044, 2017.

[15] John T Berger, Richard Holubkov, Ron Reeder, David L Wessel, Kathleen Meert, Robert A Berg, Michael J Bell, Robert Tamburro, J Michael Dean, Murray M Pollack, et al. Morbidity and mortality prediction in pediatric heart surgery: physiological profiles and surgical complexity. *The Journal of thoracic and cardiovascular surgery*, 154(2):620–628, 2017.

[16] Kate E Best and Judith Rankin. Long-term survival of individuals born with congenital heart disease: a systematic review and meta-analysis. *Journal of the American Heart Association*, 5(6):e002846, 2016.

[17] Maria M Ossa Galvis, Rupal T Bhakta, Abdulla Tarmahomed, and Magda D Mendez. Cyanotic heart disease. In *StatPearls [Internet]*. StatPearls Publishing, 2022.

[18] Tim Leiner, Jan Bogaert, Matthias G Friedrich, Raad Mohiaddin, Vivek Muthurangu, Saul Myerson, Andrew J Powell, Subha V Raman, and Dudley J Pennell. Scmr position paper (2020) on clinical indications for cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 22(1):1–37, 2020.

[19] Francois Lacour-Gayet, Edward Bove, Viktor Hraška, Victor Morell, and Thomas L Spray. *Surgery of conotruncal anomalies*. Springer, 2016.

[20] Sonia Stefanovic, Heather C Etchevers, and Stéphane Zaffran. Outflow tract formation—embryonic origins of conotruncal congenital heart disease. *Journal of Cardiovascular Development and Disease*, 8(4):42, 2021.

[21] Jin-Chung Shih, Shu-Chien Huang, Chia-Hui Lin, Tzu-Hung Lin, Ying-Ning Su, Shin-Yu Lin, En-Ting Wu, Ming-Kwang Shyu, and Chien-Nan Lee. Diagnosis of transposition of the great arteries in the fetus. *Journal of Medical Ultrasound*, 20(2):65–71, 2012.

[22] Henry L Walters III, Constantine Mavroudis, Christo I Tchervenkov, Jeffrey P Jacobs, François Lacour-Gayet, and Marshall L Jacobs. Congenital heart surgery nomenclature and database project: double outlet right ventricle. *The Annals of Thoracic Surgery*, 69(3):249–263, 2000.

[23] Kathy J Jenkins, Kimberlee Gauvreau, Jane W Newburger, Thomas L Spray, James H Moller, and Lisa I Iezzoni. Consensus-based method for risk adjustment for surgery for congenital heart disease. *The Journal of thoracic and cardiovascular surgery*, 123(1):110–118, 2002.

[24] Olubunmi Oladunjoye, Breanna Piekarski, Christopher Baird, Puja Banka, Gerald Marx, J Pedro, and Sitaram M Emani. Repair of double outlet right

ventricle: midterm outcomes. *The Journal of Thoracic and Cardiovascular Surgery*, 159(1):254–264, 2020.

[25] Shoujun Li, Kai Ma, Shengshou Hu, Zhongdong Hua, Keming Yang, Jun Yan, and Qiuming Chen. Surgical outcomes of 380 patients with double outlet right ventricle who underwent biventricular repair. *The Journal of Thoracic and Cardiovascular Surgery*, 148(3):817–824, 2014.

[26] Timothy J Bradley, Tara Karamlou, Alex Kulik, Bojana Mitrovic, Trisha Vigneswaran, Salima Jaffer, Patrick D Glasgow, William G Williams, Glen S Van Arsdell, and Brian W McCrindle. Determinants of repair type, reintervention, and mortality in 393 children with double-outlet right ventricle. *The Journal of Thoracic and Cardiovascular Surgery*, 134(4):967–973, 2007.

[27] Olivier Villemain, Emre Belli, Magalie Ladouceur, Lucile Houyel, Zakaria Jalal, Virginie Lambert, Mohamed Ly, Pascal Vouhé, and Damien Bonnet. Impact of anatomic characteristics and initial biventricular surgical strategy on outcomes in various forms of double-outlet right ventricle. *The Journal of Thoracic and Cardiovascular Surgery*, 152(3):698–706, 2016.

[28] Joanna G Escalon, Lorna P Browne, Tami J Bang, Carlos S Restrepo, Daniel Ocazionez, and Daniel Vargas. Congenital anomalies of the pulmonary arteries: an imaging overview. *The British journal of radiology*, 92(1093):20180185, 2019.

[29] Adriana DM Villa, Eva Sammut, Arjun Nair, Ronak Rajani, Rodolfo Bonamini, and Amedeo Chiribiri. Coronary artery anomalies overview: The normal and the abnormal. *World journal of radiology*, 8(6):537, 2016.

[30] Elles J Dijkema, Tim Leiner, and Heynric B Grotenhuis. Diagnosis, imaging and clinical management of aortic coarctation. *Heart*, 103(15):1148–1155, 2017.

[31] Lan Nguyen and Stephen C. Cook. Coarctation of the Aorta. Strategies for Improving Outcomes. *Cardiology Clinics*, 33(4):521–530, 2015.

[32] Cristina Basso, Domenico Corrado, and Gaetano Thiene. Congenital coronary artery anomalies as an important cause of sudden death in the young. *Cardiology in review*, 9(6):312–317, 2001.

[33] Julie A Davis, Frank Cecchin, Thomas K Jones, and Michael A Portman. Major coronary artery anomalies in a pediatric population: incidence and clinical importance. *Journal of the American College of Cardiology*, 37(2):593–597, 2001.

[34] Hussam Suradi and Ziyad M Hijazi. Current management of coarctation of the aorta. *Global Cardiology Science and Practice*, 2015(4):44, 2015.

[35] Evan J Zucker. Cross-sectional imaging of congenital pulmonary artery anomalies. *The International Journal of Cardiovascular Imaging*, 35(8):1535–1548, 2019.

[36] Sarv Priya, Richard Thomas, Prashant Nagpal, Arun Sharma, and Michael Steigner. Congenital anomalies of the aortic arch. *Cardiovascular diagnosis and therapy*, 8(Suppl 1):S26, 2018.

[37] Julien IE Hoffman. The challenge in diagnosing coarctation of the aorta. *Cardiovascular journal of Africa*, 29(4):252–255, 2018.

[38] Abraham Rothman. Coarctation of the aorta: an update. *Current problems in pediatrics*, 28(2):37–60, 1998.

[39] Margarida Q Dias, António Barros, Adelino Leite-Moreira, and Joana O Miranda. Risk factors for recoarctation and mortality in infants submitted to aortic coarctation repair: a systematic review. *Pediatric Cardiology*, 41(3):561–575, 2020.

[40] Judith Anne Adriane Ellen Cuypers, Maarten Witsenburg, Denise van der Linde, and Jolien W Roos-Hesselink. Pulmonary stenosis: update on diagnosis and therapeutic options. *Heart*, 99(5):339–347, 2013.

[41] Meghan S Zimmerman, Alison Grace Carswell Smith, Craig A Sable, Michelle Marie Echko, Lauren B Wilner, Helen Elizabeth Olsen, Hagos Tasew Atalay, Ashish Awasthi, Zulfiqar A Bhutta, Jackie LeeAnne Boucher, et al. Global, regional, and national burden of congenital heart disease, 1990–2017: a systematic analysis for the global burden of disease study 2017. *The Lancet Child & Adolescent Health*, 4(3):185–200, 2020.

[42] Julien IE Hoffman. The global burden of congenital heart disease. *Cardiovascular journal of Africa*, 24(4):141–145, 2013.

[43] Derek P Chew, Craig Juergens, John French, Will Parsonage, Matthew Horsfall, David Brieger, Stephen Quinn, Predict Study Investigators, et al. An examination of clinical intuition in risk assessment among acute coronary syndromes patients: observations from a prospective multi-center international observational registry. *International journal of cardiology*, 171(2):209–216, 2014.

[44] Nicholas A Christakis, Julia L Smith, Colin Murray Parkes, and Elizabeth B Lamont. Extent and determinants of error in doctors' prognoses in terminally ill patients: prospective cohort studycommentary: Why do doctors overestimate? commentary: Prognoses should be based on proved indices not intuition. *Bmj*, 320(7233):469–473, 2000.

[45] Julien Hoffman. The challenge in diagnosing coarctation of the aorta. *Cardiovascular Journal of Africa*, 29(4):252–255, 2018.

[46] Zahra Keshavarz-Motamed, Farhad Rikhtegar Nezami, Ramon A Partida, Kenta Nakamura, Pedro Vinícius Staziaki, Eyal Ben-Assa, Brian Ghoshhajra, Ami B Bhatt, and Elazer R Edelman. Elimination of transcoarctation pressure gradients has no impact on left ventricular function or aortic shear stress after intervention in patients with mild coarctation. *JACC: Cardiovascular Interventions*, 9(18):1953–1965, 2016.

[47] Alexander C Egbe, Thomas G Allison, and Naser M Ammash. Mild coarctation of aorta is an independent risk factor for exercise-induced hypertension. *Hypertension*, 74(6):1484–1489, 2019.

[48] Phalla Ou, David S Celermajer, Elie Mousseaux, Alain Giron, Yacine Aggoun, Isabelle Szezepanski, Daniel Sidi, and Damien Bonnet. Vascular remodeling after "successful" repair of coarctation: impact of aortic arch geometry. *Journal of the American College of Cardiology*, 49(8):883–890, 2007.

[49] Jan L Bruse, Abbas Khushnood, Kristin McLeod, Giovanni Biglino, Maxime Sermesant, Xavier Pennec, Andrew M Taylor, Tain-Yen Hsia, Silvia Schievano, Sachin Khambadkone, et al. How successful is successful? aortic arch shape after successful aortic coarctation repair correlates with left ventricular function. *The Journal of thoracic and cardiovascular surgery*, 153(2):418–427, 2017.

[50] Edward J Hickey, Christopher A Caldarone, Eugene H Blackstone, Gary K Lofland, Thomas Yeh Jr, Christian Pizarro, Christo I Tchervenkov, Frank Pigula, David M Overman, Marshall L Jacobs, et al. Critical left ventricular outflow tract obstruction: the disproportionate impact of biventricular repair in borderline cases. *The Journal of thoracic and cardiovascular surgery*, 134(6):1429–1437, 2007.

[51] Olivier Villemain, Damien Bonnet, Lucile Houyel, Mathieu Vergnat, Magalie Ladouceur, Virginie Lambert, Zakaria Jalal, Pascal Vouhé, and Emre Belli. Double-outlet right ventricle with noncommitted ventricular septal defect and 2 adequate ventricles: is anatomical repair advantageous? In *Seminars in Thoracic and Cardiovascular Surgery*, volume 28, pages 69–77. Elsevier, 2016.

[52] Antonio F Corno, Saravanan Durairaj, and Gregory J Skinner. Narrative review of assessing the surgical options for double outlet right ventricle. *Translational Pediatrics*, 10(1):165, 2021.

[53] Sara K Pasquali, Dylan Thibault, Sean M O'Brien, Jeffrey P Jacobs, J William Gaynor, Jennifer C Romano, Michael Gaies, Kevin D Hill, Marshall L Jacobs, David M Shahian, et al. National variation in congenital heart surgery outcomes. *Circulation*, 142(14):1351–1360, 2020.

[54] M Hauser. Exercise blood pressure in congenital heart disease and in patients after coarctation repair. *Heart*, 89(2):125–126, 2003.

[55] Jorge Corral-Acero, Francesca Margara, Maciej Marciniak, Cristobal Rodero, Filip Loncaric, Yingjing Feng, Andrew Gilbert, Joao F Fernandes, Hassaan A Bukhari, Ali Wajdan, et al. The 'digital twin'to enable the vision of precision cardiology. *European heart journal*, 41(48):4556–4564, 2020.

[56] John C McLachlan and Debra Patten. Anatomy teaching: ghosts of the past, present and future. *Medical education*, 40(3):243–253, 2006.

[57] Jon Cornwall and Mark D Stringer. The wider importance of cadavers: Educational and research diversity from a body bequest program. *Anatomical Sciences Education*, 2(5):234–237, 2009.

[58] John C McLachlan, John Bligh, Paul Bradley, and Judy Searle. Teaching anatomy without cadavers. *Medical education*, 38(4):418–424, 2004.

[59] Giovanni Biglino, Claudio Capelli, Despina Koniordou, Di Robertshaw, Lindsay-Kay Leaver, Silvia Schievano, Andrew M Taylor, and Jo Wray. Use of 3d models of congenital heart disease as an education tool for cardiac nurses. *Congenital heart disease*, 12(1):113–118, 2017.

[60] Ivan Lau and Zhonghua Sun. Three-dimensional printing in congenital heart disease: A systematic review. *Journal of medical radiation sciences*, 65(3):226–236, 2018.

[61] Jason Yuen. What is the role of 3d printing in undergraduate anatomy education? a scoping review of current literature and recommendations. *Medical Science Educator*, 30(3):1321–1329, 2020.

[62] John M Simpson and Owen Miller. Three-dimensional echocardiography in congenital heart disease. *Archives of cardiovascular diseases*, 104(1):45–56, 2011.

[63] Michael D Puchalski, George K Lui, Wanda C Miller-Hance, Michael M Brook, Luciana T Young, Aarti Bhat, David A Roberson, Laura Mercer-Rosa, Owen I Miller, David A Parra, et al. Guidelines for performing a comprehensive transesophageal echocardiographic: examination in children and all patients with congenital heart disease: recommendations from the american society of echocardiography. *Journal of the American Society of Echocardiography*, 32(2):173–215, 2019.

[64] Julio Garcia, Romain Capoulade, Florent Le Ven, Emmanuel Gaillard, Lyes Kadem, Philippe Pibarot, and Éric Larose. Discrepancies between cardiovascular magnetic resonance and doppler echocardiography in the measurement of transvalvular gradient in aortic stenosis: the effect of flow vorticity. *Journal of Cardiovascular Magnetic Resonance*, 15(1):1–9, 2013.

[65] Neil Patel, Melissa Dodsworth, and John F Mills. Cardiac output measurement in newborn infants using the ultrasonic cardiac output monitor: an assessment of agreement with conventional echocardiography, repeatability and new user experience. *Archives of Disease in Childhood-Fetal and Neonatal Edition*, 96(3):F206–F211, 2011.

[66] Shankar Sridharan, Gemma Price, Oliver Tann, Marina Hughes, Vivek Muthurangu, and Andrew M Taylor. *Cardiovascular MRI in congenital heart disease: an imaging atlas*. Springer Science & Business Media, 2010.

[67] Annelies E van der Hulst, Arno AW Roest, Jos JM Westenberg, Lucia JM Kroft, and Albert de Roos. Cardiac mri in postoperative congenital heart

disease patients. *Journal of Magnetic Resonance Imaging*, 36(3):511–528, 2012.

[68] Gareth T Archer, Alaa Elhawaz, Natasha Barker, Benjamin Fidock, Alexander Rothman, RJ Van Der Geest, Rod Hose, Norman Briffa, Ian R Hall, Ever Grech, et al. Validation of four-dimensional flow cardiovascular magnetic resonance for aortic stenosis assessment. *Scientific reports*, 10(1):1–10, 2020.

[69] Baiyan Zhuang, Arlene Sirajuddin, Shihua Zhao, and Minjie Lu. The role of 4d flow mri for clinical applications in cardiovascular disease: current status and future perspectives. *Quantitative Imaging in Medicine and Surgery*, 11(9):4193, 2021.

[70] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.

[71] David A Field. Laplacian smoothing and delaunay triangulations. *Communications in applied numerical methods*, 4(6):709–712, 1988.

[72] Henk Kaarle Versteeg and Weeratunge Malalasekera. *An introduction to computational fluid dynamics: the finite volume method*. Pearson education, 2007.

[73] Chaoqun Liu, Jiyuan Tu, and Guan Heng Yeoh. *Computational fluid dynamics : a practical approach*. Elsevier, third edition edition, 2018.

[74] Richard A Malinauskas, Prasanna Hariharan, Steven W Day, Luke H Herbertson, Martin Buesen, Ulrich Steinseifer, Kenneth I Aycock, Bryan C Good, Steven Deutsch, Keefe B Manning, et al. Fda benchmark medical device flow models for cfd validation. *Asaio Journal*, 63(2):150–160, 2017.

[75] Giovanni Biglino, Daria Cosentino, Jennifer A. Steeden, Lorenzo De Nova, Matteo Castelli, Hopewell Ntsinjana, Giancarlo Pennati, Andrew M. Taylor,

and Silvia Schievano. Using 4D Cardiovascular Magnetic Resonance Imaging to Validate Computational Fluid Dynamics: A Case Study. *Frontiers in Pediatrics*, 3(December):1–10, 2015.

[76] Yulei Zhu, Rui Chen, Yu Hsiang Juan, He Li, Jingjing Wang, Zhuliang Yu, and Hui Liu. Clinical validation and assessment of aortic hemodynamics using computational fluid dynamics simulations from computed tomography angiography. *BioMedical Engineering Online*, 17(1):1–12, 2018.

[77] Robert Ascuitto, Nancy Ross-Ascuitto, Martin Guillot, and Carey Celestin. Computational fluid dynamics characterization of pulsatile flow in central and sano shunts connected to the pulmonary arteries: importance of graft angulation on shear stress-induced, platelet-mediated thrombosis. *Interactive Cardiovascular and Thoracic Surgery*, 25(3):414–421, 2017.

[78] Keiichi Itatani, Kagami Miyaji, Yi Qian, Jin Long Liu, Tomoyuki Miyakoshi, Arata Murakami, Minoru Ono, and Mitsuo Umezu. Influence of surgical arch reconstruction methods on single ventricle workload in the norwood procedure. *The Journal of thoracic and cardiovascular surgery*, 144(1):130–138, 2012.

[79] Namheon Lee, Michael D Taylor, and Rupak K Banerjee. Right ventricle-pulmonary circulation dysfunction: a review of energy-based approach. *Biomedical engineering online*, 14:1–20, 2015.

[80] Friso M Rijnberg, Mark G Hazekamp, Jolanda J Wentzel, Patrick JH De Koning, Jos JM Westenberg, Monique RM Jongbloed, Nico A Blom, and Arno AW Roest. Energetics of blood flow in cardiovascular disease: concept and clinical implications of adverse energetics in patients with a fontan circulation. *Circulation*, 137(22):2393–2407, 2018.

[81] Pouya Youssefi, Alberto Gomez, Taigang He, Lisa Anderson, Nick Bunce, Rajan Sharma, C Alberto Figueroa, and Marjan Jahangiri. Patient-specific computational fluid dynamics—assessment of aortic hemodynamics in a

spectrum of aortic valve pathologies. *The Journal of thoracic and cardiovascular surgery*, 153(1):8–20, 2017.

[82] Andres F Osorio, Ruben Osorio, Andres Ceballos, Reginald Tran, William Clark, Eduardo A Divo, I Ricardo Argueta-Morales, Alain J Kassab, and William M DeCampli. Computational fluid dynamics analysis of surgical adjustment of left ventricular assist device implantation to minimise stroke risk. *Computer methods in biomechanics and biomedical engineering*, 16(6):622–638, 2013.

[83] Rabin Gerrah and Stephen J Haller. Computational fluid dynamics: a primer for congenital heart disease clinicians. *Asian Cardiovascular and Thoracic Annals*, 28(8):520–532, 2020.

[84] M Conijn and GJ Krings. Computational analysis of the pulmonary arteries in congenital heart disease: A review of the methods and results. *Computational and Mathematical Methods in Medicine*, 2021:1–10, 2021.

[85] John F LaDisa Jr, Charles A Taylor, and Jeffrey A Feinstein. Aortic coarctation: recent developments in experimental and computational methods to assess treatments for this simple condition. *Progress in pediatric cardiology*, 30(1-2):45–49, 2010.

[86] John F LaDisa, C Alberto Figueroa, Irene E Vignon-Clementel, Hyun Jin Kim, Nan Xiao, Laura M Ellwein, Frandics P Chan, Jeffrey A Feinstein, and Charles A Taylor. Computational simulations for aortic coarctation: representative results from a sampling of patients. *Journal of Biomedical Engineering*, 2011.

[87] Atefeh Razavi, Shagun Sachdeva, Peter C Frommelt, and John F LaDisa Jr. Patient-specific numerical analysis of coronary flow in children with intramural anomalous aortic origin of coronary arteries. In *Seminars in Thoracic and Cardiovascular Surgery*, volume 33, pages 155–167. Elsevier, 2021.

[88] Charlotte Schubert, Jan Brüning, Leonid Goubergrits, Anja Hennemuth, Felix Berger, Titus Kühne, and Marcus Kelm. Assessment of hemodynamic responses to exercise in aortic coarctation using MRI-ergometry in combination with computational fluid dynamics. *Scientific Reports*, 10(1):1–12, 2020.

[89] Petter Frieberg, Nicolas Aristokleous, Pia Sjöberg, Johannes Töger, Petru Liuba, and Marcus Carlsson. Computational fluid dynamics support for fontan planning in minutes, not hours: the next step in clinical pre-interventional simulations. *Journal of Cardiovascular Translational Research*, 15(4):708–720, 2022.

[90] Yue Qiu, Jiarong Wang, Jichun Zhao, Tiehao Wang, Tinghui Zheng, and Ding Yuan. Association between blood flow pattern and rupture risk of abdominal aortic aneurysm based on computational fluid dynamics. *European Journal of Vascular and Endovascular Surgery*, 64(2-3):155–164, 2022.

[91] William M DeCampli, I Ricardo Argueta-Morales, Eduardo Divo, and Alain J Kassab. Computational fluid dynamics in congenital heart disease. *Cardiology in the Young*, 22(6):800–808, 2012.

[92] Giancarlo Pennati, Chiara Corsini, Tain-Yen Hsia, Francesco Migliavacca, and Modeling of Congenital Hearts Alliance (MOCHA) Investigators. Computational fluid dynamics models and congenital heart diseases. *Frontiers in pediatrics*, 1:4, 2013.

[93] Tain-Yen Hsia, Richard Figliola, Modeling of Congenital Hearts Alliance (MOCHA) Investigators, Edward Bove, Adam Dorfman, Andrew Taylor, Alessandro Giardini, Sachin Khambadkone, Silvia Schievano, Tain-Yen Hsia, et al. Multiscale modelling of single-ventricle hearts for clinical decision support: a leducq transatlantic network of excellence, 2016.

[94] Lucian Itu, Puneet Sharma, Kristóf Ralovich, Viorel Mihalef, Razvan Ionasec, Allen Everett, Richard Ringel, Ali Kamen, and Dorin Comaniciu.

Non-invasive hemodynamic assessment of aortic coarctation: validation with in vivo measurements. *Annals of biomedical engineering*, 41(4):669–681, 2013.

[95] Giovanni Biglino, Claudio Capelli, Jan Bruse, Giorgia M. Bosi, Andrew M. Taylor, and Silvia Schievano. Computational modelling for congenital heart disease: How far are we from clinical translation?, 1 2017.

[96] Claudio Capelli, Emilie Sauvage, Giuliano Giusti, Giorgia M Bosi, Hopewell Ntsinjana, Mario Carminati, Graham Derrick, Jan Marek, Sachin Khambadkone, Andrew M Taylor, et al. Patient-specific simulations for planning treatment in congenital heart disease. *Interface Focus*, 8(1):20170021, 2018.

[97] Aurélien Géron. *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow*. " O'Reilly Media, Inc.", 2022.

[98] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

[99] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[100] Renato Cuocolo, Teresa Perillo, Eliana De Rosa, Lorenzo Ugga, and Mario Petretta. Current applications of big data and machine learning in cardiology. *Journal of geriatric cardiology: JGC*, 16(8):601, 2019.

[101] Andreas Hauptmann, Simon Arridge, Felix Lucka, Vivek Muthurangu, and Jennifer A Steeden. Real-time cardiovascular mr with spatio-temporal artifact suppression using deep learning–proof of concept in congenital heart disease. *Magnetic resonance in medicine*, 81(2):1143–1156, 2019.

[102] Jennifer A Steeden, Michael Quail, Alexander Gotschy, Kristian H Mortensen, Andreas Hauptmann, Simon Arridge, Rodney Jones, and Vivek Muthurangu. Rapid whole-heart cmr with single volume super-resolution. *Journal of Cardiovascular Magnetic Resonance*, 22(1):1–13, 2020.

[103] Tim Leiner, Daniel Rueckert, Avan Suinesiaputra, Bettina Baeßler, Reza Nezafat, Ivana Išgum, and Alistair A Young. Machine learning in cardiovascular magnetic resonance: basic concepts and applications. *Journal of Cardiovascular Magnetic Resonance*, 21:1–14, 2019.

[104] Rima Arnaout, Lara Curran, Yili Zhao, Jami C Levine, Erin Chinn, and Anita J Moon-Grady. An ensemble of neural networks provides expert-level prenatal detection of complex congenital heart disease. *Nature medicine*, 27(5):882–891, 2021.

[105] Ali Madani, Ramy Arnaout, Mohammad Mofrad, and Rima Arnaout. Fast and accurate view classification of echocardiograms using deep learning. *NPJ digital medicine*, 1(1):1–8, 2018.

[106] C Wang, H Frimmel, and O Smedby. Fast level-set based image segmentation using coherent propagation. *Medical physics*, 41(7), 2014.

[107] Jürgen Weese, Michael Kaus, Christian Lorenz, Steven Lobregt, Roel Truyen, and Vladimir Pekar. Shape Constrained Deformable Models for 3D Medical Image Segmentation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 2082, pages 380–387. Springer, Berlin, Heidelberg, 2001.

[108] Jochen Peters, Olivier Ecabert, Carsten Meyer, Hauke Schramm, Reinhard Kneser, Alexandra Groth, and Jürgen Weese. Automatic Whole Heart Segmentation in Static Magnetic Resonance Image Volumes. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelli-*

gence and Lecture Notes in Bioinformatics), 4792 LNCS(PART 2):402–410, 2007.

[109] Chen Chen, Chen Qin, Huaqi Qiu, Giacomo Tarroni, Jinming Duan, Wenjia Bai, and Daniel Rueckert. Deep Learning for Cardiac Image Segmentation: A Review. *Frontiers in Cardiovascular Medicine*, 0:25, 3 2020.

[110] Haben Berhane, Michael Scott, Mohammed Elbaz, Kelly Jarvis, Patrick McCarthy, James Carr, Chris Malaisrie, Ryan Avery, Alex J. Barker, Joshua D. Robinson, Cynthia K. Rigsby, and Michael Markl. Fully automated 3D aortic segmentation of 4D flow MRI for hemodynamic analysis using deep learning. *Magnetic Resonance in Medicine*, 84(4):2204–2218, 10 2020.

[111] Xiahai Zhuang, Lei Li, Christian Payer, Darko Štern, Martin Urschler, Mattias P. Heinrich, Julien Oster, Chunliang Wang, Orjan Smedby, Cheng Bian, Xin Yang, Pheng-Ann Heng, Aliasghar Mortazi, Ulas Bagci, Guanyu Yang, Chenchen Sun, Gaetan Galisot, Jean-Yves Ramel, Thierry Brouard, Qianqian Tong, Weixin Si, Xiangyun Liao, Guodong Zeng, Zenglin Shi, Guoyan Zheng, Chengjia Wang, Tom MacGillivray, David Newby, Kawal Rhode, Sebastien Ourselin, Raad Mohiaddin, Jennifer Keegan, David Firmin, and Guang Yang. Evaluation of algorithms for Multi-Modality Whole Heart Segmentation: An open-access grand challenge. *Medical Image Analysis*, 58:101537, 12 2019.

[112] Saeed Karimi-Bidhendi, Arghavan Arafati, Andrew L Cheng, Yilei Wu, Arash Kheradvar, and Hamid Jafarkhani. Fully-automated deep-learning segmentation of pediatric cardiovascular magnetic resonance of patients with complex congenital heart diseases. *Journal of cardiovascular magnetic resonance*, 22(1):1–24, 2020.

[113] Wenjia Bai, Matthew Sinclair, Giacomo Tarroni, Ozan Oktay, Martin Rajchl, Ghislain Vaillant, Aaron M. Lee, Nay Aung, Elena Lukaschuk, Mihir M. Sanghvi, Filip Zemrak, Kenneth Fung, Jose Miguel Paiva, Valentina Cara-

pella, Young Jin Kim, Hideaki Suzuki, Bernhard Kainz, Paul M. Matthews, Steffen E. Petersen, Stefan K. Piechnik, Stefan Neubauer, Ben Glocker, and Daniel Rueckert. Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance 2018 20:1*, 20(1):1–12, 9 2018.

[114] Benjamin Böttcher, Ebba Beller, Anke Busse, Daniel Cantré, Seyrani Yücel, Alper Öner, Hüseyin Ince, Marc-André Weber, and Felix G. Meinel. Fully automated quantification of left ventricular volumes and function in cardiac MRI: clinical evaluation of a deep learning-based algorithm. *The International Journal of Cardiovascular Imaging 2020 36:11*, 36(11):2239–2247, 7 2020.

[115] Wouter Huberts, Stefan GH Heinen, Niek Zonnebeld, Daniel AF van den Heuvel, Jean-Paul PM de Vries, Jan HM Tordoir, D Rodney Hose, Tammo Delhaas, and Frans N van de Vosse. What is needed to make cardiovascular models suitable for clinical decision support? a viewpoint paper. *Journal of computational science*, 24:68–84, 2018.

[116] Pavlo Yevtushenko, Leonid Goubergrits, Lina Gundelwein, Arnaud Setio, Tobias Heimann, Heiko Ramm, Hans Lamecker, Titus Kuehne, Alexander Meyer, and Marie Schafstedde. Deep learning based centerline-aggregated aortic hemodynamics: An efficient alternative to numerical modelling of hemodynamics. *IEEE Journal of Biomedical and Health Informatics*, 2021.

[117] Bradley Feiger, John Gounley, Dale Adler, Jane A Leopold, Erik W Draeger, Rafeed Chaudhury, Justin Ryan, Girish Pathangey, Kevin Winarta, David Frakes, et al. Accelerating massively parallel hemodynamic models of coarctation of the aorta using neural networks. *Scientific reports*, 10(1):1–13, 2020.

[118] Liang Liang, Wenbin Mao, and Wei Sun. A feasibility study of deep learning for predicting hemodynamics of human thoracic aorta. *Journal of Biomechanics*, 99, 2020.

[119] M. J.M.M. Hoeijmakers, I. Waechter-Stehle, J. Weese, and F. N. Van de Vosse. Combining statistical shape modeling, CFD, and meta-modeling to approximate the patient-specific pressure-drop across the aortic valve in real-time. *International Journal for Numerical Methods in Biomedical Engineering*, 36(10):1–18, 2020.

[120] Yanji Qu, Xinlei Deng, Shao Lin, Fengzhen Han, Howard H Chang, Yanqiu Ou, Zhiqiang Nie, Jinzhuang Mai, Ximeng Wang, Xiangmin Gao, et al. Using innovative machine learning methods to screen and identify predictors of congenital heart diseases. *Frontiers in cardiovascular medicine*, 8:2087, 2022.

[121] Gerhard-Paul Diller, Aleksander Kempny, Sonya V Babu-Narayan, Marthe Henrichs, Margarita Brida, Anselm Uebing, Astrid E Lammers, Helmut Baumgartner, Wei Li, Stephen J Wort, et al. Machine learning algorithms estimating prognosis and guiding therapy in adult congenital heart disease: data from a single tertiary centre including 10 019 patients. *European heart journal*, 40(13):1069–1077, 2019.

[122] Joshua Cates, Shireen Elhabian, and Ross Whitaker. Shapeworks: particle-based shape correspondence and visualization software. In *Statistical Shape and Deformation Analysis*, pages 257–298. Elsevier, 2017.

[123] Alexandre Bône, Maxime Louis, Benoît Martin, and Stanley Durrleman. Deformetrica 4: an open-source software for statistical shape analysis. In *International Workshop on Shape in Medical Imaging*, pages 3–13. Springer, 2018.

[124] Gene H Golub and Christian Reinsch. Singular value decomposition and least squares solutions. In *Linear algebra*, pages 134–151. Springer, 1971.

[125] Jan L Bruse, Kristin McLeod, Giovanni Biglino, Hopewell N Ntsinjana, Claudio Capelli, Tain-Yen Hsia, Maxime Sermesant, Xavier Pennec, Andrew M Taylor, and Silvia Schievano. A statistical shape modelling framework to extract 3d shape biomarkers from medical imaging data: assessing arch morphology of repaired coarctation of the aorta. *BMC medical imaging*, 16(1):1–19, 2016.

[126] Froso Sophocleous, Alexandre Bône, Andrew IU Shearn, Mari Nieves Velasco Forte, Jan L Bruse, Massimo Caputo, and Giovanni Biglino. Feasibility of a longitudinal statistical atlas model to study aortic growth in congenital heart disease. *Computers in Biology and Medicine*, 144:105326, 2022.

[127] Sarah Kollar, Elias Balaras, Laura J Olivieri, Yue-Hin Loke, and Francesco Capuano. Statistical shape modeling reveals the link between right ventricular shape, hemodynamic force, and myocardial function in patients with repaired tetralogy of fallot. *American Journal of Physiology-Heart and Circulatory Physiology*, 323(3):H449–H460, 2022.

[128] Xenia Alba, Marco Pereañez, Corné Hoogendoorn, Andrew J Swift, Jim M Wild, Alejandro F Frangi, and Karim Lekadir. An algorithm for the segmentation of highly abnormal hearts using a generic statistical shape model. *IEEE transactions on medical imaging*, 35(3):845–859, 2015.

[129] Fanwei Kong, Nathan Wilson, and Shawn Shadden. A deep-learning approach for direct whole-heart mesh reconstruction. *Medical image analysis*, 74:102222, 2021.

[130] Sang-Wook Lee, Luca Antiga, J David Spence, and David A Steinman. Geometry of the carotid bifurcation predicts its exposure to disturbed flow. *Stroke*, 39(8):2341–2347, 2008.

[131] Anupama Goparaju, Ibolya Csecs, Alan Morris, Evgueni Kholmovski, Nassir Marrouche, Ross Whitaker, and Shireen Elhabian. On the evaluation and

validation of off-the-shelf statistical shape modeling tools: a clinical application. In *Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings*, pages 14–27. Springer, 2018.

[132] Andrew J Parker. Binocular depth perception and the cerebral cortex. *Nature Reviews Neuroscience*, 8(5):379–391, 2007.

[133] Ralf Dörner, Wolfgang Broll, Paul Grimm, and Bernhard Jung. *Virtual and augmented reality (VR/AR) Foundations and Methods of Extended Realities (XR)*. Springer, 2022.

[134] Nisal Menuka Gamage, Deepana Ishtaweera, Martin Weigel, and Anusha Withana. So predictable! continuous 3d hand trajectory prediction in virtual reality. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pages 332–343, 2021.

[135] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.

[136] RE Stevens, DP Rhodes, A Hasnain, and P-Y Laffont. Varifocal technologies providing prescription and vac mitigation in hmds using alvarez lenses. In *Digital Optics for Immersive Displays*, volume 10676, pages 142–158. SPIE, 2018.

[137] Christer Ericson. *Real-time collision detection*. Crc Press, 2004.

[138] Hyun Woo Goo, Sang Joon Park, and Shi-Joon Yoo. Advanced medical use of three-dimensional imaging in congenital heart disease: augmented reality, mixed reality, virtual reality, and three-dimensional printing. *Korean journal of radiology*, 21(2):133–145, 2020.

[139] Kanwal M Farooqi, James C Nielsen, Santosh C Uppu, Shubhika Srivastava, Ira A Parness, Javier Sanz, and Khanh Nguyen. Use of 3-dimensional printing to demonstrate complex intracardiac relationships in double-outlet

right ventricle for surgical planning. *Circulation: Cardiovascular Imaging*, 8(5):e003043, 2015.

[140] Israel Valverde, Gorka Gomez-Ciriza, Tarique Hussain, Cristina Suarez-Mejias, Maria N Velasco-Forte, Nicholas Byrne, Antonio Ordoñez, Antonio Gonzalez-Calle, David Anderson, Mark G Hazekamp, et al. Three-dimensional printed models for surgical planning of complex congenital heart defects: an international multicentre study. *European Journal of Cardio-Thoracic Surgery*, 52(6):1139–1148, 2017.

[141] Giovanni Biglino, Sarah Moharem-Elgamal, Matthew Lee, Robert Tulloh, and Massimo Caputo. The perception of a three-dimensional-printed heart model from the perspective of different stakeholders: a complex case of truncus arteriosus. *Frontiers in Pediatrics*, 5:209, 2017.

[142] Hong-Chang Guo, Yang Wang, Jiang Dai, Chang-Wei Ren, Jin-Hua Li, and Yong-Qiang Lai. Application of 3d printing in the surgical planning of hypertrophic obstructive cardiomyopathy and physician-patient communication: a preliminary study. *Journal of thoracic disease*, 10(2):867, 2018.

[143] Silvia Schievano, Francesco Migliavacca, Louise Coats, Sachin Khambadkone, Mario Carminati, Neil Wilson, John E Deanfield, Philipp Bonhoeffer, and Andrew M Taylor. Percutaneous pulmonary valve implantation based on rapid prototyping of right ventricular outflow tract and pulmonary trunk from mr data. *Radiology*, 242(2):490–497, 2007.

[144] Matthew Lee, Sarah Moharem-Elgamal, Rylan Beckingham, Mark Hamilton, Nathan Manghat, Elena Giulia Milano, Chiara Bucciarelli-Ducci, Massimo Caputo, and Giovanni Biglino. Evaluating 3d-printed models of coronary anomalies: a survey among clinicians and researchers at a university hospital in the uk. *BMJ open*, 9(3):e025227, 2019.

[145] Deane Yim, Andreea Dragulescu, Haruki Ide, Mike Seed, Lars Grosse-Wortmann, Glen Van Arsdell, and Shi-Joon Yoo. Essential modifiers

of double outlet right ventricle: revisit with endocardial surface images and 3-dimensional print models. *Circulation: Cardiovascular Imaging*, 11(3):e006891, 2018.

[146] Shi-Joon Yoo and Glen S Van Arsdell. 3d printing in surgical management of double outlet right ventricle. *Frontiers in pediatrics*, 5:289, 2018.

[147] Shujie Deng, Gavin Wheeler, Nicolas Toussaint, Lindsay Munroe, Suryava Bhattacharya, Gina Sajith, Ei Lin, Eeshar Singh, Ka Yee Kelly Chu, Saleha Kabir, et al. A virtual reality system for improved image-based planning of complex cardiac procedures. *Journal of imaging*, 7(8):151, 2021.

[148] Amir H Sadeghi, Sulayman El Mathari, Djamila Abjigitova, Alexander PWM Maat, Yannick JHJ Taverne, Ad JJC Bogers, and Edris AF Mahtab. Current and future applications of virtual, augmented, and mixed reality in cardiothoracic surgery. *The Annals of Thoracic Surgery*, 2020.

[149] Jennifer NA Silva, Michael Southworth, Constantine Raptis, and Jonathan Silva. Emerging applications of virtual reality in cardiovascular medicine. *JACC: Basic to Translational Science*, 3(3):420–430, 2018.

[150] Ana Mendez, Tarique Hussain, Amir-Reza Hosseinpour, and Israel Valverde. Virtual reality for preoperative planning in large ventricular septal defects. *European heart journal*, 40(13):1092–1092, 2019.

[151] Animesh Tandon, Barbara EU Burkhardt, Maria Batsis, Thomas M Zellers, Mari Nieves Velasco Forte, Israel Valverde, Ryan P McMahan, Kristine J Guleserian, Gerald F Greil, and Tarique Hussain. Sinus venosus defects: anatomic variants and transcatheter closure feasibility using virtual reality planning. *JACC: Cardiovascular Imaging*, 12(5):921–924, 2019.

[152] Elena Giulia Milano, Endrit Pajaziti, Emilie Sauvage, Andrew Cook, Silvia Schievano, Kristian H Mortensen, Andrew M Taylor, Jan Marek, Martin

Kostolny, and Claudio Capelli. Taking surgery out of reality: A virtual journey into double outlet right ventricle. *Circulation: Cardiovascular Imaging*, 12(7):e009297, 2019.

[153] Francesca Raimondi, Vladimiro Vida, Charlotte Godard, Francesco Bertelli, Elena Reffo, Nathalie Boddaert, Mohamed El Beheiry, and Jean-Baptiste Masson. Fast-track virtual reality for cardiac imaging in congenital heart disease. *Journal of Cardiac Surgery*, 36(7):2598–2602, 2021.

[154] Reena M Ghosh, Christopher E Mascio, Jonathan J Rome, Matthew A Jolley, and Kevin K Whitehead. Use of virtual reality for hybrid closure of multiple ventricular septal defects. *Case Reports*, 3(14):1579–1583, 2021.

[155] Victor Manuel Caicedo Ayerbe, Martha Lucía Velasco Morales, Carlos Javier Latorre Rojas, and María Lucía Arango Cortés. Visualization of 3d models through virtual reality in the planning of congenital cardiothoracic anomalies correction: an initial experience. *World Journal for Pediatric and Congenital Heart Surgery*, 11(5):627–629, 2020.

[156] Jianzheng Cen, Rong Liufu, Shusheng Wen, Hailong Qiu, Xiaobin Liu, Xiaokun Chen, Haiyun Yuan, Meiping Huang, and Jian Zhuang. Three-dimensional printing, virtual reality and mixed reality for pulmonary atresia: early surgical outcomes evaluation. *Heart, Lung and Circulation*, 30(2):296–302, 2021.

[157] Mohamed Estai and Stuart Bunt. Best teaching practices in anatomy education: A critical review. *Annals of Anatomy-Anatomischer Anzeiger*, 208:151–157, 2016.

[158] Loren D Sacks and David M Axelrod. Virtual reality in pediatric cardiology: hype or hope for the future? *Current Opinion in Cardiology*, 35(1):37–41, 2020.

[159] Chin Siang Ong, Aravind Krishnan, Chen Yu Huang, Philip Spevak, Luca Vricella, Narutoshi Hibino, Juan R Garcia, and Lasya Gaur. Role of virtual

reality in congenital heart disease. *Congenital Heart Disease*, 13(3):357–361, 2018.

[160] Byeol Kim, Yue-Hin Loke, Paige Mass, Matthew R Irwin, Conrad Capeland, Laura Olivieri, and Axel Krieger. A novel virtual reality medical image display system for group discussions of congenital heart disease: Development and usability testing. *JMIR cardio*, 4(1):e20633, 2020.

[161] Ivan Lau, Ashu Gupta, and Zhonghua Sun. Clinical value of virtual reality versus 3d printing in congenital heart disease. *Biomolecules*, 11(6):884, 2021.

[162] Kuberan Pushparajah, Ka Yee Kelly Chu, Shujie Deng, Gavin Wheeler, Alberto Gomez, Saleha Kabir, Julia A Schnabel, and John M Simpson. Virtual reality three-dimensional echocardiographic imaging for planning surgical atrioventricular valve repair. *JTCVS techniques*, 7:269–277, 2021.

[163] W Ye, X Zhang, T Li, C Luo, and L Yang. Mixed-reality hologram for diagnosis and surgical planning of double outlet of the right ventricle: A pilot study. *Clinical Radiology*, 76(3):237–e1, 2021.

[164] Tiffany R Lim, Hunter C Wilson, David M Axelrod, David K Werho, Stephanie S Handler, Sunkyung Yu, Katherine Afton, Ray Lowery, Patricia B Mullan, James Cooke, et al. Virtual reality curriculum increases paediatric residents' knowledge of chds. *Cardiology in the Young*, pages 1–5, 2022.

[165] Neil Patel, Anthony Costa, Stephen P Sanders, and David Ezon. Stereoscopic virtual reality does not improve knowledge acquisition of congenital heart disease. *The International Journal of Cardiovascular Imaging*, 37:2283–2290, 2021.

[166] Jennifer A. Steeden, Michael Quail, Alexander Gotschy, Kristian H. Mortensen, Andreas Hauptmann, Simon Arridge, Rodney Jones, and Vivek Muthurangu. Rapid whole-heart CMR with single volume super-resolution. *Journal of Cardiovascular Magnetic Resonance*, 22(1):56, 8 2020.

[167] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9351, pages 234–241. Springer Verlag, 2015.

[168] Lisha Li, Kevin Jamieson, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.

[169] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, Savannah, GA, November 2016. USENIX Association.

[170] Javier Montalt-Tordera. Tensorflow mri: A library of tensorflow operators for computational magnetic resonance imaging., July 2022.

[171] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1026–1034. Institute of Electrical and Electronics Engineers Inc., 2015.

[172] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, 12 2015.

[173] W Silversmith. cc3d: Connected Components on Multilabel 3D Images, 2021.

[174] Luca Antiga, Marina Piccinelli, Lorenzo Botti, Bogdan Ene-Iordache, Andrea Remuzzi, and David A Steinman. An image-based modeling framework for patient-specific computational hemodynamics. *Medical & biological engineering & computing*, 46(11):1097–1112, 2008.

[175] F. J.H. Gijsen, F. N. Van De Vosse, and J. D. Janssen. The influence of the non-Newtonian properties of blood on the flow in large arteries: steady flow in a carotid bifurcation model. *Journal of Biomechanics*, 32(6):601–608, 6 1999.

[176] A. D. Caballero and S. Laín. A Review on Computational Fluid Dynamics Modelling in Human Thoracic Aorta. *Cardiovascular Engineering and Technology 2013 4:2*, 4(2):103–130, 4 2013.

[177] Mirko Bonfanti, Stavroula Balabani, John P. Greenwood, Sapna Puppala, Shervanthi Homer-Vanniasinkam, and Vanessa Díaz-Zuccarini. Computational tools for clinical support: a multi-scale compliant model for haemodynamic simulations in an aortic dissection based on multi-modal imaging data. *Journal of The Royal Society Interface*, 14(136), 11 2017.

[178] I.T. Gabe, J.H. Gault, J. Ross, D.T. Mason, C.J. Mills, J.P. Schillingford, and E. Braunwald. Measurement of instantaneous blood flow velocity and pressure in conscious man with a catheter-tip velocity probe. *Circulation*, 40(5):603–614, 1969.

[179] Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Multi-label Whole Heart Segmentation Using CNNs and Anatomical Label Configurations. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 10663 LNCS:190–198, 2017.

[180] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2017-October, pages 843–852. Institute of Electrical and Electronics Engineers Inc., 12 2017.

[181] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, 25, 2012.

[182] Riddhish Bhalodia, Shireen Y. Elhabian, Ladislav Kavan, and Ross T. Whitaker. DeepSSM: A Deep Learning Framework for Statistical Shape Modeling from Raw Images. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11167 LNCS:244–257, 9 2018.

[183] Ashwin Raju, Shun Miao, Dakai Jin, Le Lu, Junzhou Huang, and Adam P Harrison. Deep implicit statistical shape models for 3d medical image delineation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2135–2143, 2022.

[184] Gerhard-Paul Diller, Julius Vahle, Robert Radke, Maria Luisa Benesch Vidal, Alicia Jeanette Fischer, Ulrike MM Bauer, Samir Sarikouch, Felix Berger, Philipp Beerbaum, Helmut Baumgartner, et al. Utility of deep learning networks for the generation of artificial cardiac magnetic resonance images in congenital heart disease. *BMC Medical Imaging*, 20(1):1–8, 2020.

[185] Virginia Fernandez, Walter Hugo Lopez Pinaya, Pedro Borges, Petru-Daniel Tudosiu, Mark S Graham, Tom Vercauteren, and M Jorge Cardoso. Can segmentation models be trained with fully synthetically generated data? In *Simulation and Synthesis in Medical Imaging: 7th International Workshop, SASHIMI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Proceedings*, pages 79–90. Springer, 2022.

[186] Jun Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *Proceedings of the IEEE International Conference on Computer Vision*, 2017-October:2242–2251, 3 2017.

[187] Wenjun Yan, Yuanyuan Wang, Shengjia Gu, Lu Huang, Fuhua Yan, Liming Xia, and Qian Tao. The domain shift problem of medical image segmentation and vendor-adaptation by unet-gan. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 623–631. Springer, 2019.

[188] Lu Liu, Jelmer M. Wolterink, Christoph Brune, and Raymond N.J. Veldhuis. Anatomy-aided deep learning for medical image segmentation: a review. *Physics in Medicine & Biology*, 66(11):11TR01, 5 2021.

[189] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. Spie, 1992.

[190] Marina Piccinelli, Alessandro Veneziani, David A Steinman, Andrea Remuzzi, and Luca Antiga. A framework for geometric analysis of vascular structures: application to cerebral aneurysms. *IEEE transactions on medical imaging*, 28(8):1141–1155, 2009.

[191] Si Hang. Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw*, 41(2):11, 2015.

[192] Julio Garcia, Roel LF van der Palen, Emilie Bollache, Kelly Jarvis, Michael J Rose, Alex J Barker, Jeremy D Collins, James C Carr, Joshua Robinson, Cynthia K Rigsby, et al. Distribution of blood flow velocity in the normal aorta: effect of age and gender. *Journal of Magnetic Resonance Imaging*, 47(2):487–498, 2018.

[193] AJ Powell, SE Maier, T Chung, and T Geva. Phase-velocity cine magnetic resonance imaging measurement of pulsatile blood flow in children

and young adults: in vitro and in vivo validation. *Pediatric cardiology*, 21(2):104–110, 2000.

[194] Mirko Bonfanti, Stavroula Balabani, John P Greenwood, Sapna Puppala, Shervanthi Homer-Vanniasinkam, and Vanessa Díaz-Zuccarini. Computational tools for clinical support: a multi-scale compliant model for haemodynamic simulations in an aortic dissection based on multi-modal imaging data. *Journal of The Royal Society Interface*, 14(136):20170632, 2017.

[195] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[196] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, 24, 2011.

[197] Thomas Eiter and Heikki Mannila. Computing discrete fréchet distance. *Technical Report CD-TR 94/64*, 1994.

[198] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018.

[199] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 704–720, 2018.

[200] Yasi Wang, Hongxun Yao, and Sicheng Zhao. Auto-encoder based dimensionality reduction. *Neurocomputing*, 184:232–242, 2016.

[201] Gerhard-Paul Diller, Julius Vahle, Robert Radke, Maria Luisa Benesch Vidal, Alicia Jeanette Fischer, Ulrike MM Bauer, Samir Sarikouch, Felix Berger, Philipp Beerbaum, Helmut Baumgartner, et al. Utility of deep learning networks for the generation of artificial cardiac magnetic resonance images in congenital heart disease. *BMC Medical Imaging*, 20(1):1–8, 2020.

[202] Michael Stein. Large sample properties of simulations using latin hypercube sampling. *Technometrics*, 29(2):143–151, 1987.

[203] Asad Abbas, Ashkan Rafiee, Max Haase, and Andrew Malcolm. Geometric convolutional neural networks–a journey to surrogate modelling of maritime cfd. In *The 9th Conference on Computational Methods in Marine Engineering (Marine 2021)*, 2022.

[204] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CVPR 2017 Open Access*, 2016.

[205] Nobuyuki Umetani and Bernd Bickel. Learning Three-Dimensional Flow for Interactive Aerodynamic Design regression prediction for new shape. *ACM Trans. Graph.*, 37(4), 2018.

[206] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.

[207] Amirhossein Arzani, Jian-Xun Wang, and Roshan M D'Souza. Uncovering near-wall blood flow from sparse data with physics-informed neural networks. *Physics of Fluids*, 33(7):071905, 2021.

[208] Kwong Ming Tse, Rong Chang, Heow Pueh Lee, Siak Piang Lim, Sudhakar Kundapur Venkatesh, and Pei Ho. A computational fluid dynamics study on geometrical influence of the aorta on haemodynamics. *European Journal of Cardio-Thoracic Surgery*, 43(4):829–838, 2013.

[209] Hadi Wiputra, Shion Matsumoto, Jessica E Wagenseil, Alan C Braverman, Rochus K Voeller, and Victor H Barocas. Statistical shape representation of the thoracic aorta: accounting for major branches of the aortic arch. *Computer methods in biomechanics and biomedical engineering*, pages 1–15, 2022.

[210] ANDRES D Caballero and SJCE Laín. A review on computational fluid dynamics modelling in human thoracic aorta. *Cardiovascular Engineering and Technology*, 4(2):103–130, 2013.

[211] Jonas Lantz, Tino Ebbers, Jan Engvall, and Matts Karlsson. Numerical and experimental assessment of turbulent kinetic energy in an aortic coarctation. *Journal of biomechanics*, 46(11):1851–1858, 2013.

[212] Pouya Youssefi, Alberto Gomez, Christopher Arthurs, Rajan Sharma, Marjan Jahangiri, and C Alberto Figueroa. Impact of patient-specific inflow velocity profile on hemodynamics of the thoracic aorta. *Journal of biomechanical engineering*, 140(1), 2018.

[213] Umberto Morbiducci, Raffaele Ponzini, Diego Gallo, Cristina Bignardi, and Giovanna Rizzo. Inflow boundary conditions for image-based computational hemodynamics: impact of idealized versus measured velocity profiles in the human aorta. *Journal of biomechanics*, 46(1):102–109, 2013.

[214] Bente Thamsen, Pavlo Yevtushenko, Lina Gundelwein, Arnaud AA Setio, Hans Lamecker, Marcus Kelm, Marie Schafstedde, Tobias Heimann, Titus Kuehne, and Leonid Goubergrits. Synthetic database of aortic morphometry and hemodynamics: overcoming medical imaging data availability. *IEEE Transactions on Medical Imaging*, 40(5):1438–1449, 2021.

[215] Nico Westerhof, Jan-Willem Lankhaar, and Berend E Westerhof. The arterial windkessel. *Medical & biological engineering & computing*, 47(2):131–141, 2009.

[216] Rodrigo M Romarowski, Adrien Lefieux, Simone Morganti, Alessandro Veneziani, and Ferdinando Auricchio. Patient-specific cfd modelling in the thoracic aorta with pc-mri–based boundary conditions: A least-square three-element windkessel approach. *International journal for numerical methods in biomedical engineering*, 34(11):e3134, 2018.

[217] Andrea S Les, Shawn C Shadden, C Alberto Figueroa, Jinha M Park, Maureen M Tedesco, Robert J Herfkens, Ronald L Dalman, and Charles A Taylor. Quantification of hemodynamics in abdominal aortic aneurysms during rest and exercise using magnetic resonance imaging and computational fluid dynamics. *Annals of biomedical engineering*, 38(4):1288–1313, 2010.

[218] Sudharsan Madhavan and Erica M Cherry Kemmerling. The effect of inlet and outlet boundary conditions in image-based cfd modeling of aortic flow. *Biomedical engineering online*, 17(1):1–20, 2018.

[219] S Pirola, Z Cheng, OA Jarral, DP O'Regan, JR Pepper, T Athanasiou, and XY Xu. On the choice of outlet boundary conditions for patient-specific analysis of aortic flow using computational fluid dynamics. *Journal of biomechanics*, 60:15–21, 2017.

[220] Doff B McElhinney, V Mohan Reddy, Norman H Silverman, and Frank L Hanley. Modified damus-kaye-stansel procedure for single ventricle, subaortic stenosis, and arch obstruction in neonates and infants: midterm results and techniques for avoiding circulatory arrest. *The Journal of thoracic and cardiovascular surgery*, 114(5):718–726, 1997.

[221] Loren D Sacks and David M Axelrod. Virtual reality in pediatric cardiology: hype or hope for the future? *Current Opinion in Cardiology*, 35(1):37–41, 2020.

[222] Karen A Eley, Maria Camilla Rossi-Espagnet, Silvia Schievano, Antonio Napolitano, Juling Ong, Aurelio Secinaro, Luca Borro, David Dunaway, Carlo Efisio Marras, Adam Rennie, et al. Multiparametric imaging for presurgical planning of craniopagus twins: the experience of two tertiary pediatric hospitals with six sets of twins. *Radiology*, 298(1):18–27, 2021.

[223] NHS Great Ormond Street Hospital. Separating yiḡit and derman.

[224] Shiona McCallum. Conjoined twins separated with the help of virtual reality.

[225] NHS Great Ormond Street Hospital. Conjoined twins annie and issie return home after separation surgery.

[226] Brian Baillargeon, Nuno Rebelo, David D Fox, Robert L Taylor, and Ellen Kuhl. The living heart project: a robust and integrative simulator for human heart function. *European Journal of Mechanics-A/Solids*, 48:38–47, 2014.

[227] UCL Institute of Cardiovascular Science (ICS). Short courses, hands-on cardiac morphology course.

[228] Christian Moro, Zane Štromberga, Athanasios Raikos, and Allan Stirling. The effectiveness of virtual and augmented reality in health sciences and medical anatomy. *Anatomical sciences education*, 10(6):549–559, 2017.

[229] HS Maresky, A Oikonomou, I Ali, N Ditkofsky, M Pakkal, and B Ballyk. Virtual reality and cardiac anatomy: Exploring immersive three-dimensional cardiac imaging, a pilot study in undergraduate medical anatomy education. *Clinical Anatomy*, 32(2):238–243, 2019.