



POLITECNICO
MILANO 1863

Data-driven approaches for quantification of complex biomarkers from medical imaging

Doctoral Dissertation of:
Simone Saitta

Supervisor:
Prof. Alberto Redaelli

Tutor:
Prof. Dario Gastaldi

The Chair of the Doctoral Program:
Prof. Gabriele Dubini

DEPARTMENT OF ELECTRONICS, INFORMATION AND BIOENGINEERING
DOCTORAL PROGRAM IN BIOENGINEERING - 35TH CYCLE

1	Introduction	1
1.1	Background and significance	1
1.2	Machine learning	3
1.3	Structure of the thesis	4
1.3.1	Addressed applications	4
1.3.2	Aim of the thesis	8
I	Supervised approaches for morphological assessment	9
2	A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography	11
2.1	TEVAR preprocedural planning	13
2.2	Building an automatic pipeline from CT data to complex geometric characterization of the thoracic aorta	14
2.2.1	Data Collection	14
2.2.2	Image Pre-analysis and Manual Segmentation	14
2.2.3	Neural Network Training	15
2.2.4	Pipeline Implementation	15
2.2.5	Statistical Analysis	22
2.2.6	Comparison with semi-automated measurements	22
2.3	Performance Evaluation	22
2.3.1	Test set results	22
2.4	Comparison with state of the art and solutions	28
2.5	Conclusions	31

3	A CT-based deep learning system for automatic assessment of aortic root morphology for TAVI planning	33
3.1	TAVI planning and state of the art	35
3.2	Development of an automatic method for aortic root characterization . .	36
3.2.1	Data collection and manual annotation	36
3.2.2	Analysis of inter-operator variability	38
3.2.3	Image preprocessing: aortic root region detection	38
3.2.4	Automatic segmentation	39
3.2.5	Pipeline implementation	40
3.2.6	Statistical analysis	43
3.3	Results	43
3.3.1	Evaluation of segmentation performance	43
3.3.2	Comparison vs. manual measurements	45
3.4	Discussion and comparison with existing approaches	48
3.5	Final remarks	50
4	Mixed Reality and Deep Learning for External Ventricular Drainage Placement: A Fast and Automatic Workflow for Emergency Treatments	51
4.1	Clinical problem and current solutions	53
4.2	Approach	53
4.2.1	Automatic Segmentation	54
4.2.2	Automatic Registration	55
4.2.3	Registration Accuracy Assessment	56
4.3	Result assessment	58
4.3.1	Automatic segmentation	58
4.3.2	Registration Accuracy Assessment	59
4.4	Conclusions	60
II	Unsupervised approaches for hemodynamic assessment	63
5	Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta	65
5.1	Significance and state of the art	67
5.2	Statistical shape model of aortic velocity profile	68
5.2.1	Image data	68
5.2.2	Data preprocessing	68
5.2.3	Statistical shape modeling of inlet velocity profiles	70
5.2.4	Flow morphology descriptors	72
5.2.5	Acceptance criteria	73
5.2.6	Statistical analysis	74
5.3	Results	74
5.3.1	Statistical shape model analysis	74

5.3.2	Associations of velocity profile modes with flow morphology descriptors	78
5.4	Results impact and comparison with previous work	80
5.5	Conclusions	82
6	Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI	83
6.1	Background and significance	85
6.1.1	Related Work	86
6.2	Implicit neural representations of 4D flow MRI	89
6.2.1	Problem setting	89
6.2.2	Training a SIREN in 4D	90
6.2.3	SIREN evaluation	90
6.2.4	Error quantification	91
6.2.5	Wall shear stress analysis	91
6.2.6	Case 1: synthetic 4D flow magnetic resonance imaging (MRI)	92
6.2.7	Case 2: <i>in vivo</i> 4D flow MRI	95
6.3	Experiment results	99
6.3.1	Case 1: quantification of image degradation	99
6.3.2	Case 1: hyperparameter tuning	102
6.3.3	Case 1: comparison with existing methods	108
6.3.4	Case 2: velocity field assessment	119
6.3.5	Case 2: wall shear stress field assessment	121
6.4	Discussion of SIRENs for 4D flow MRI	121
6.5	Final considerations	123
7	Final remarks	125

List of Figures

1.1	Graphical overview of the methods, applications and anatomical districts covered in this thesis.	4
1.2	Graphical outline of the thesis. The present work encompasses a variety of biomedical applications; the addressed clinical problem, employed methods and solutions are outlined for each chapter.	7
2.1	Schematic representation of the implemented automatic pipeline. In steps 1-3, red domains and surfaces represent the thoracic aorta; green domains and surfaces represent the pulmonary artery.	16
2.2	Representation of the algorithms for automatic detection of centerlines seed and target points. a) the first normalized eigenvector of the Laplace-Beltrami operator used to find seed (yellow circle) and target (grey circle) points for the main aortic centerline. b) seed (green cross) and targets (red crosses) are manually annotated on the reference surface, which is morphed into the processed surface. From the morphed surface, left and right pulmonary arteries centerlines are computed. c) the normalized heat kernel signature (HKS) is computed on the aortic surface and recursively thresholded to isolate supraaortic branch ends; tube surfaces are generated for the aorta (grey), brachiocephalic trunk (green), left carotid artery (red) and subclavian artery (blue)	18
2.3	Automatic computation of geometric descriptors. A standard arch configuration (left) exemplifying centerline curvature radius (RC) and outer curvature radius (RO) (in blue). A CILCA arch (right) reporting β and tortuosity (τ) angles. N.B.: distances not to scale for display purposes . .	21

2.4	Comparison of manual (ground truth, GT) segmentation in bright red (aorta, Ao) and bright green (pulmonary arteries, PA), vs. automatic segmentation obtained by the trained neural network (convolutional neural network (CNN)) in dark red (aorta, Ao) and dark green (pulmonary arteries, PA) on a test case. Four axial slices with color fill are shown on the left. A sagittal and a coronal slices with dashed contours are shown on the right	23
2.5	Aortic arches with thoracic aortic aneurysm from the test set with centerlines colored by landing zone. Six arches present a standard configuration with the presence of zone 1; all other arches present a common origin of the innominate and left carotid artery (CILCA) arch configuration without zone 1. β and tortuosity (τ) angles for zones longer than 20 mm are reported, with subscripts indicating the zone they refer to. . . .	26
2.6	Representation of a 60 mm centerline tract covering the landing zones (left), and maximum diameter plotted against the centerline abscissa (right). Values for standard (orange) and CILCA (blue) cases are plotted as mean (solid lines) with 95% confidence intervals (transparent bands)	27
3.1	Schematic representation of the adopted workflow. Dataset A includes manual segmentations of the aortic root, annulus and sino-tubular junction (STJ); it was used to train and validate two neural networks (model 1 and model 2) for automatic segmentation. Dataset B included manual annotations of aortic root measurements, used to validate the developed morphological analysis pipeline.	37
3.2	Template-matching approach for automatic identification of the region of interest (ROI) from 3D computed tomography (CT). A template image (X) for which aortic segmentation (S) is known, is affinely registered to the processed image (Y). The transformation mapping X to Y (T_{α^*}) is applied to S and the bounding box of $T_{\alpha^*}(S)$ is used to crop the ROI from Y	39
3.3	Schematic representation of the implemented automatic pipeline. The ROI is detected from the input CT scan. Model 1 infers the segmentation of the aortic root (AR) with the left ventricle outflow tract (LVOT), shown as the red label (left side). Model 2 infers the segmentation of the aortic annulus and STJ (in green and blue, respectively) (right side). The aortic surface processing step computes the second eigenvector of the Laplace-Beltrami operator (LBO) and its contours (left side). The annulus and STJ processing step performs refinement of model 2 segmentation. In the AR analysis step (bottom), the AR is isolated, and the Sinuses of Val-salva are detected. Anatomical measurements are computed together with calcium volume (shown in white, bottom panel).	41

3.4	steps implemented for aortic annulus and STJ plane detection refinement. a) the point sets yielded by model 2 for the annulus (green) and STJ (blue), visualized with the aortic segmentation produced by model 1 (transparent gray). b) the best fitting planes for the two structures (Π_{ann} and Π_{STJ}), together with the region inside which their center of mass is allowed to move during the refinement procedure (magenta spheres). c) the resulting planes (Π_{ann}^* and Π_{STJ}^*) after the optimization procedure. d) visualization of the refined planes in dark gray and the initialization planes obtained by model 2. e) the isolated AR with the contours of the second eigenvalue of the LBO (in black) and the plane identifying the sinuses of Valsalva (yellow).	43
3.5	Comparison of manual (ground truth, GT) segmentation in bright red (aorta), bright green (annulus) and light blue (STJ) vs. automatic segmentation obtained by the trained neural networks (CNN) in dark red (aorta), dark green (annulus) and dark blue (STJ) on four test cases: the best (top row) and the worst (bottom row) for the two models. Coronal slices with color fill and with dashed contours are shown.	44
3.6	Bland–Altman plots of the proposed algorithm versus expert manual measurements. Mean differences are shown as continuous horizontal lines, while 95% limits of agreement are shown as dashed horizontal lines. Critical cases are highlighted with a red arrow and identification number.	46
3.7	Results of the implemented automatic pipeline for 3 cases for which the discrepancies between expert and automatic measurements were largest. The first column shows the segmentation inferred by model 1 and bounded to the AR region (in red). The second column shows the 3D anatomy reconstructions together with either the annulus (green), STJ (blue) or sinuses plane (yellow). The third column shows an interpolation of the image onto the plane identified by either the annulus, STJ or sinuses planes. Segmentation contours are show in green (annulus), blue (STJ) and yellow (sinuses), together with D_{max} (dashed arrows, black) and D_{min} (dashed arrows, white).	47
4.1	Schematic representation of the developed workflow: the 3D CT images are automatically segmented to obtain the brain structures that will be sent to HoloLens2 (H2) device through an internet connection protocol. The user is then able to visualize the holographic models and set a path between the Foramen of Monro (FoM) and an entry point on the skin layer. The point cloud (PtC) of the skin surface of the patient is acquired using H2 depth camera and this data is used to estimate the transformation matrix T_{CT}^{H2}	54

4.2	A. Detailed representation of the registration workflow for the mixed reality (MR) environment with B. the considered coordinate systems and transformation matrices.	56
4.3	A. Experimental setup showing the hardware and physical components needed to test the registration accuracy; user executing the first test is evaluating the position of the 4 marker on the printed phantom B., and on the registered holographic model C.; D. phantom used by neurosurgeons for the second test while E. trying to insert the optical tracker (OT) probe to match the holographic trajectory.	57
4.4	Comparison of manual (ground truth, GT) segmentation in bright red (lateral ventricles, LV) and bright green (third ventricle, TV), vs. automatic segmentation obtained by the trained neural network (NN) in dark red (lateral ventricles, LV) and dark green (third ventricle, TV) on a test case.	59
4.5	Experimental results from evaluation tests. A: hologram-to-phantom accuracy evaluated on the 4 markers; B: quantitative targeting accuracy from the tests done by 7 neurosurgeons; C: hologram to phantom accuracy evaluated over the whole model surface; D: visualization of catheter tip positions and target from the external ventricular drain (EVD) simulations.	60
5.1	Schematic representation of the adopted workflow. All 4D flow acquisitions go through a preprocessing pipeline for the extraction of velocity profiles. The SSM process consists in a common alignment and spatiotemporal resampling of the profiles and then a combination of principal component analysis (PCA) modes to generate new ones. Only the profiles that meet specific acceptance criteria are included in the final dataset.	69
5.2	Method for ascending aortic plane selection. The pulmonary artery bifurcation point (P_B , green) is selected from the PC-MRA images. The corresponding point on the aortic centerline (P_{BA0}) is identified (pink, bottom right), and the selected plane (red, bottom right) is chosen at a 30 mm arc length upstream along the centerline and locally normal to the centerline.	70

5.3	a) Representation of an inlet plane coordinate set from a representative subject in the clinical cohort (red), as obtained upon rigid roto-translation, and of the corresponding point set in the fixed disk domain (green). b) A velocity profile from the clinical cohort (top), resampled to the fixed reference disk using radial basis functions (RBFs) (bottom). c) Exemplification of the computed flow descriptors. FDI (top): $area^{top15\%}$ is highlighted in black; FJA (middle): angle between plane normal (black) and mean velocity direction (magenta); SFD (bottom): ratio between the normal component (black) and in-plane component (cyan) of the mean velocity vector (magenta).	71
5.4	a. Individual and cumulative variance associated with each principal component. b. Orientation of the displayed planes and profiles in with respect to the subject in 2D and 3D. c. Mean velocity profile ($\bar{\mathbf{V}}$) colored by velocity magnitude in 2D (top row) and 3D (bottom row) at early systole (ES), peak systole (PS) and late systole (LS). d. 2D and 3D visualizations of velocity profiles deformed towards minimum and maximum for the first 4 modes and colored by velocity magnitude; color scale shown in panel b.	76
5.5	Box plots showing distributions of time-averaged flow descriptors for real and synthetic cohorts. Similar medians and ranges were observed. No significant differences ($p \geq 0.05$) were found except for PPV ($p = 0.040$). Whiskers indicate 1.5 interquartile ranges; diamonds highlight outliers.	77
5.6	Plots of normalized flow rates over time computed for the first four modes by varying their corresponding coefficient $b^{(m)}$. Mean flow rate over time corresponding to $\bar{\mathbf{V}}$ is shown in red for all subplots.	80
6.1	a) Computational domain. b) Time-varying 3-directional velocity profiles prescribed at as inlet boundary condition. c) Flow waveform imposed at Γ_i	92
6.2	Visualization of velocity magnitude colormaps on a sagittally oriented 2D slice (bottom left) corresponding to the different steps to create synthetic 4D flow data from computational fluid dynamics (CFD) results. CFD results (top row) are temporally averaged (second row). Temporally average velocity fields defined on unstructured meshes are resampled to a fine 3D Cartesian grid (third row). Fine cartesian grids are downsampled in k -space and complex Gaussian noise is added to produce synthetic flow measurements (bottom row).	94
6.3	Left: segmented 3D geometry (red) superimposed on a slice representation of phase-contrast magnetic resonance angiography (PC-MRA) images. Right: fine 3D mesh for spatial super-resolution	96

6.4	Sagittally oriented 2D slice within the aortic aneurysm (bottom left), together with colormaps of left-to-right (first and fourth rows), posterior-to-anterior (second and fifth rows) and foot-to-head (third and sixth rows) velocity components. Rows 4 to 6 show velocities sampled on a 3D mesh by linear interpolation. Columns from left to right correspond to increasing time points.	97
6.5	Axially oriented 2D slice within the aortic aneurysm (bottom left), together with colormaps of left-to-right (first and fourth rows), posterior-to-anterior (second and fifth rows) and foot-to-head (third and sixth rows) velocity components. Rows 4 to 6 show velocities sampled on a 3D mesh by linear interpolation. Columns from left to right correspond to increasing time points.	98
6.6	Velocity magnitude colormaps on a sagittal slice (bottom left corner) corresponding to: CFD solution (top row), images with <i>mild</i> added noise (second row), images with <i>medium</i> added noise (third row), images with <i>extreme</i> added noise (bottom row). Columns from left to right correspond to increasing time points.	100
6.7	Detailed view of velocity magnitude colormaps on a sagittal slice (magenta rectangle in bottom left corner) corresponding to: CFD solution (top row), images with <i>mild</i> added noise (second row), images with <i>medium</i> added noise (third row), images with <i>extreme</i> added noise (bottom row).	101
6.8	Sagittal view of the effect of different SIREN architectures on synthetic images with <i>extreme</i> noise level. Rows 1 and 5: images with <i>extreme</i> noise added; rows 2 and 6: CFD solution; rows 3 and 7: results of a SIREN with 20 layers and 500 neurons per layer; rows 4 and 8: results of a SIREN with 20 layers and 300 neurons per layer.	106
6.9	Axial view of the effect of different SIREN architectures on synthetic images with <i>extreme</i> noise level. Rows 1 and 5: images with <i>extreme</i> noise added; rows 2 and 6: CFD solution; rows 3 and 7: results of a SIREN with 20 layers and 500 neurons per layer; rows 4 and 8: results of a SIREN with 20 layers and 300 neurons per layer.	107
6.10	Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with <i>mild</i> noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.	109
6.11	Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with <i>medium</i> noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.	110

6.12	Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with <i>extreme</i> noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.	111
6.13	Results corresponding to <i>mild</i> noise levels. wall shear stress (WSS) colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	113
6.14	Results corresponding to <i>mild</i> noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	114
6.15	Results corresponding to <i>medium</i> noise levels. WSS colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	115
6.16	Results corresponding to <i>medium</i> noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	116
6.17	Results corresponding to <i>extreme</i> noise levels. WSS colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	117
6.18	Results corresponding to <i>extreme</i> noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).	118
6.19	Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on a sagittally oriented 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.	120
6.20	Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on an axially oriented 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.	120

6.21 WSS magnitude colormaps (row 1) and WSS vectors (rows 2) computed from SIREN results.	121
---	-----

List of Tables

2.1	Whole arch-based geometric parameters calculated by the automatic tool for standard and CILCA arch cases. Median [min, max] or mean \pm standard deviations are reported together with the p values obtained by the statistical analysis	24
2.2	Zone arc lengths in mm calculated by the developed automatic tool for standard and CILCA arch cases. Median [min, max] are reported together with p values obtained from the statistical analyses. ^a Standard vs. CILCA. ^b Anova among landing zones.	24
2.3	Zone maximum diameters in mm calculated by the developed automatic tool for standard and CILCA arch cases. Median [min, max] are reported together with p values obtained from the statistical analyses. ^a Standard vs. CILCA. ^b Anova among landing zones.	25
2.4	Zone β angles in degrees calculated by the developed automatic tool for standard and CILCA arch cases. Mean \pm standard deviations are reported together with p values obtained from the statistical analyses. ^a Standard vs. CILCA. ^b Anova among landing zones.	25
2.5	Zone tortuosity angles in degrees calculated by the developed automatic tool for standard and CILCA arch cases. Mean \pm standard deviations are reported together with p values obtained from the statistical analyses. ^a Standard vs. CILCA. ^b Anova among landing zones.	25
2.6	Comparison between geometric measurements obtained with the commercial software 3Mensio vs. our tool. Mean of differences and 95% limits of agreement are reported. Zone 1 metrics refer to standard arches only, while all other metrics refer to both standard and CILCA arches . .	28
3.1	Comparison between automatic segmentation approaches and ground truth manual segmentations for the AR, annulus and STJ. Mean value and [min, max] ranges of Dice score and mean surface distance (MSD) are reported.	44

3.2	Inter-operator variability analysis results. Dice scores, mean surface distance (MSD) and physical volume differences (ΔV) are computed for each pairwise comparison. Median [minimum, maximum] values are reported.	45
5.1	Comparison of time-averaged flow descriptors between real and synthetic cohorts. Normally distributed variables are expressed as mean \pm standard deviation; non-normally distributed variables are expressed as median [min, max]. [†] indicates statistical significance $p \leq 0.05$	75
5.2	PCA modes correlations with velocity profile flow descriptors at peak systole. Results of Pearson correlation analyses are reported. $p < 0.05$ indicates statistical significance.	79
6.1	Degradation parameters for generating synthetic measurements from the reference velocity field.	93
6.2	Errors introduced by the implemented degradation steps.	99
6.3	Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with <i>mild</i> noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.	103
6.4	Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with <i>medium</i> noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.	104
6.5	Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with <i>extreme</i> noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.	105
6.6	Velocity field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.	108
6.7	WSS field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.	112
6.8	Quantitative measurements for Case 2 on raw flow MR measurements and on velocity fields reconstructed with our method.	119

Acronyms

AI artificial intelligence.

AR aortic root.

ATAA ascending thoracic aortic aneurysms.

BAV bicuspid aortic valve.

BC boundary condition.

CFD computational fluid dynamics.

CILCA common origin of the innominate and left carotid artery.

CNN convolutional neural network.

CT computed tomography.

DE direction error.

DF-RBF divergence-free radial basis function.

DFW divergence-free wavelets.

EVD external ventricular drain.

FDI flow dispersion index.

FJA flow jet angle.

FoM Foramen of Monro.

GNN graph neural network.

GUI graphic-user interface.

H2 HoloLens2.

HKS heat kernel signature.

HMD head mounted display.

HU Hounsfield units.

ICP iterative closest point.

INR implicit neural representation.

L-BFGS limited memory Broyden-Fletcher-Goldfarb-Shanno.

LBO Laplace-Beltrami operator.

LDDMM large deformation diffeomorphic metric mapping.

LITP linear interpolation.

LVOT left ventricle outflow tract.

LZ landing zone.

ML machine learning.

MLP multi-layer perceptron.

mNRSME magnitude normalized root mean square error.

MR mixed reality.

MRA magnetic resonance angiography.

MRI magnetic resonance imaging.

N-S Navier-Stokes.

NN neural network.

OT optical tracker.

PA pulmonary artery.

PC-MRA phase-contrast magnetic resonance angiography.

PC-MRI phase-contrast magnetic resonance imaging.

PCA principal component analysis.

PDE partial differential equation.

PINN physics informed neural network.

PLZ proximal landing zone.

PPV positive peak velocity.

PtC point cloud.

RBF radial basis function.

RFI retrograde flow index.

ROI region of interest.

SFD secondary flow degree.

SNR signal to noise ratio.

SR super-resolution.

SSM statistical shape model.

STJ sino-tubular junction.

TAA thoracic aortic aneurysms.

TAV tricuspid aortic valve.

TAVI transcatheter aortic valve implantation.

TEVAR thoracic endovascular aortic repair.

VENC velocity encoding.

VMTK vascular modeling toolkit.

vNRMSE velocity normalized root mean square error.

WSS wall shear stress.

1.1 Background and significance

Over the past few decades, the paradigm of personalized medicine has revolutionized the healthcare industry. At the cornerstone of this transformation is the widespread adoption of medical imaging techniques, such as computed tomography (CT) and magnetic resonance imaging (MRI). These advanced imaging systems, in conjunction with efficient image reconstruction algorithms, offer high-resolution 3D anatomical visualizations and functional tissue information. In recent years, medical imaging has progressed beyond its traditional role as a diagnostic tool, and it has now become a prevalent component in the realm of clinical sciences, central in patient care and treatment planning. The integration of medical imaging into clinical practice has opened up new avenues for personalized and precise medical treatment, allowing for the individualized evaluation of each patient's unique anatomy and physiology.

The utilization of medical imaging goes beyond mere visual inspection, it provides a vast amount of data that can be analyzed for quantifiable information. The field of radiomics has seen the development of computer-aided tools that are capable of analyzing digital medical images and extracting meaningful quantitative features. These features, which can be based on attributes such as object intensity, shape, texture and wavelets, provide valuable information that can be used to improve healthcare outcomes in a multitude of applications [1]. Furthermore, the extraction of more complex quantitative features and biomarkers from medical images holds immense potential for advancing the field of personalized medicine. The ability to extract and analyze such information from medical images can significantly enhance the accuracy of clinical decision making, facilitate early disease detection and characterization, and provide more individualized treatment options for patients.

When designing an intervention, medical professionals must assess the internal

anatomical structures of the patient, in order to formulate a plan that optimizes the procedure's strategy and steps. Medical imaging provides a visual representation of the targeted area, allowing clinicians to identify the optimal approach. Moreover, an integrated use of medical imaging becomes critical in the context of minimally invasive procedures, where the field of view is restricted during the intervention [2]. In such scenarios, pre-procedural planning based on medical images becomes indispensable. A detailed analysis of a patient's pre-procedural scans is essential when the treatment involves the implantation of prosthetic devices. The proper sizing of these devices is contingent upon the patient's anatomy, and as such, a thorough analysis and evaluation of medical images is mandatory in order to make an informed decision. In essence, the integration of medical imaging into the decision-making process is of utmost importance and plays a pivotal role in ensuring the success of the intervention.

Medical imaging is also used during minimally invasive procedures as it serves as a visual aid for physicians to accurately place instruments and devices within the body of the patient [3]. The use of medical imaging enables the identification of key anatomical landmarks and structures to target, supporting the surgeon in making more informed decisions during the intervention. This leads to an improvement in the success rate of the procedure and reduction in the associated risks. Overall, medical imaging plays an integral role in facilitating minimally invasive procedures, contributing to improved patient outcomes and a safer medical environment.

Medical images have been widely utilized to create and inform computational models for non-invasive functional characterization of tissue with an emphasis on precision. These image-based numerical simulations leverage patient-specific image data to construct high-fidelity digital replicas of the patient, referred to as "digital twins" [4]. This approach enables the development of large-scale *in silico* trials, which provide a cost-effective alternative for validating medical devices. The integration of medical images in the formation of these digital twins enables a more accurate representation of a patient's anatomy and physiology, providing a robust foundation for simulating the functional properties of tissues. These simulations are critical in advancing the field of medical technology, as they allow for the evaluation of medical devices under a wide range of conditions without the need for costly and time-consuming physical trials. This approach has the potential to significantly reduce the costs associated with medical device validation, while also accelerating the development and implementation of new technologies [5, 6].

Medical imaging has become an indispensable tool for functional assessment of patients, particularly in the context of risk stratification [7]. Different functional imaging modalities can provide valuable insights into the function of various organs and systems in the body, including oxygen consumption of a tissue, blood flow dynamics, electrophysiology, enabling medical professionals to make better informed decisions regarding patient care.

The advancement of clinical care has led to a significant increase in the variety and amount of biomarkers and features that can be extracted from modern medical imaging. This has resulted in an increasingly complex and task-specific landscape for the

analysis of medical imaging data. Traditional radiomics analysis tools are not well-suited to address this issue, creating a gap between the image data and the physician responsible for making informed treatment and care decisions. To bridge this gap, there is a need for the development of new tools that can improve upon four key characteristics: accuracy, automation, speed, and consistency. Machine learning (ML) and data-driven approaches have the potential to greatly enhance these properties and provide improved analysis of medical imaging data.

1.2 Machine learning

ML is a discipline within the field of artificial intelligence (AI) that focuses on the development of algorithms that can learn from data and make predictions or decisions without being explicitly programmed. The goal of ML is to map inputs to outputs, and to do so, ML models extract features from the input data that are relevant to the specific task at hand. The performance of a model in ML is defined as a loss (or cost) function that measures the discrepancy between the model's predictions and the ground truth. The parameters of the ML model are then adjusted in order to minimize the loss function, with the most commonly used approach being gradient-based optimization algorithms. Based on the nature of the training data, ML can be divided into two main approaches: supervised and unsupervised learning.

In supervised learning, the model is trained using a labeled dataset where the ground truth outputs are known. Supervised models learn by being explicitly exposed to a large number of inputs-outputs pairs, known as training set. Validation is then carried out by evaluating the trained models on new, unseen data samples. One of the most successful applications of supervised learning in medical image analysis is automatic semantic segmentation. The introduction of convolutional neural networks (CNNs) has enabled automatic segmentation of a wide range of anatomical structures in medical images [8]. The main building blocks of a CNN are convolutional layers, which apply a series of filters to the input image to extract features. These filters slide over the input, generating feature maps that are progressively aggregated within the deep layers of the model. The application of non-linear activation functions on the feature maps allows a CNN to learn complex relations between input and output. During training, a CNN processes a large number of sample images and compares the automatically segmented anatomy with the corresponding ground truth segmentation. After training, the CNN can be validated on a test set, which consists of images that were not used during training, to assess its performance. However, it is important to note that semantic segmentation only provides direct quantitative information on the volume of the segmented object and does not fully bridge the gap between images and the quantification of features that are directly involved in the task at hand.

In unsupervised learning, the model is trained using an unlabeled dataset, and the goal is to learn a useful representation of the data. Unsupervised learning approaches enable the extraction of information from data without prior knowledge or explicit supervision. This approach is designed to uncover underlying patterns and relationships

within the data, ultimately learning to map inputs to outputs through the identification of informative features. The learning process is accomplished without the use of a ground truth-dependent loss function, which is a hallmark of supervised learning. Instead, unsupervised methods often involve clustering, dimensionality reduction, or generative models to create a useful representation of the data. One of the key advantages of unsupervised learning is its ability to identify patterns and structures in the data that are representative of the data. Successful applications of unsupervised learning also include denoising approaches based on neural networks.

1.3 Structure of the thesis

The present thesis is highly interdisciplinary in its nature, and various biomedical applications are at the center of this work. In each chapter, a specific application is addressed. The overarching rationale behind the development of this thesis was to devise novel solutions to challenges that emerge in precision medicine. In particular, we focus on improving the extraction of biomedical quantities from medical image data. Two imaging modalities are herein taken into account: 3D CT (for quantification of morphological biomarkers, Part I) and 4D flow MRI (for quantification of hemodynamic biomarkers, Part II). To do so, we employ data-driven methods, including both supervised (Part I) and unsupervised techniques (Part II). A graphical overview of the imaging modalities, methods, computed biomarkers and addressed applications is shown in Figure 1.1.

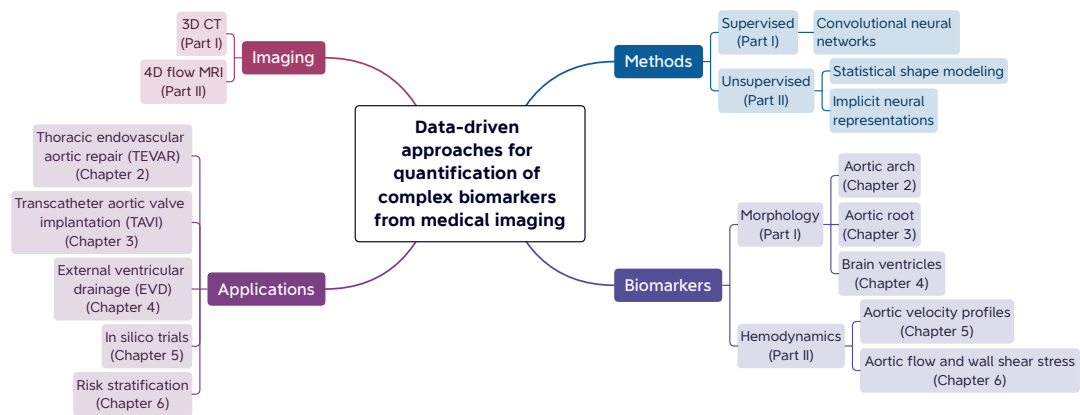


Figure 1.1: Graphical overview of the methods, applications and anatomical districts covered in this thesis.

1.3.1 Addressed applications

This thesis is divided in two parts, and each part is divided into chapters. An outline of the biomedical applications addressed in this work is shown in Figure 1.2

Part I presents three distinct applications of supervised learning techniques for accurately quantifying anatomical features to facilitate intervention planning and execution.

Chapter 2 focuses on the morphological characterization of the aortic arch for pre-procedural planning of thoracic endovascular aortic repair (TEVAR), a minimally invasive procedure used to repair a thoracic aortic aneurysm or dissection. During TEVAR, a stent is fit inside the damaged part of the aorta to provide reinforcement. Feasibility assessment and planning of TEVAR require CT-based analysis of geometric aortic features to identify adequate landing zones (LZs) for endograft deployment. However, taking the necessary measurements is complex, time-consuming and suffers from low repeatability. This chapter describes the development of a fully automatic pipeline processing CT images of the thorax, performs automatic segmentation of the aorta and pulmonary artery via a CNN with a 3D Unet architecture, and operates a sequence of differential geometry post-processing steps to extract aortic morphological features that are relevant for planning TEVAR.

Chapter 3 follows the development of Chapter 2 and concerns the automatic anatomical analysis of the aortic root from CT images, aimed at improving transcatheter aortic valve implantation (TAVI) pre-procedural planning. TAVI is a minimally invasive heart procedure to replace a stenotic aortic valve. During this intervention, a new valve is inserted without removing the old, damaged valve. Choosing the correct device and understanding the patient-specific anatomy is crucial for the success of the procedure, and a proper planning is based on morphological measurements of the aortic root. However, performing a meticulous pre-procedural assessment is time-consuming and affected by inter-operator variability. This chapter describes a devised solution that includes two CNNs for segmentation of the aortic root and robust landmark detection, allowing for quantification of aortic root metrics involved in planning of TAVI.

Chapter 4 describes a methodology to support neurosurgeons during external ventricular drain (EVD). EVD is the procedure of choice for treating hydrocephalus, a neurological disorder caused by an abnormal buildup of cerebrospinal fluid in the ventricles (cavities) deep within the brain. During EVD, the neurosurgeon inserts a draining catheter into the brain aiming at the Foramen of Monro (FoM), a small area at the turn of the brain ventricles. Given the extremely limited visibility of the operating field and the small dimensions of the procedural target, EVD is associated to a high chance of complications. This chapter described a workflow for EVD support based on a CNN for automatic segmentation of the brain ventricles, detection of procedural target, and integration within a mixed reality application for holographic visualization of brain structures.

Part II focuses on extraction of hemodynamic characteristics in the ascending aorta using unsupervised data-driven methods.

Chapter 5 concerns the development of computational models for *in silico* trials and addresses the scarcity of realistic flow conditions for building digital twins. In fact, computational models can drastically accelerate the development of new drugs and medical devices, to reduce, refine and replace tests on animals and clinical trials, and

to improve the safety of medical products. To obtain reliable results from numerical simulations of blood flow, using realistic inflow boundary conditions plays a crucial role. Nonetheless, such data are difficult to find and limited in their number. This chapter describes the usage of principal component analysis (PCA) to build a statistical shape model of blood flow velocity profiles in thoracic aortic aneurysm, suitable to be used for realistic numerical simulations and *in silico* trials.

Chapter 6 concerns the quantification of hemodynamics from 4D flow MRI, an imaging technique that provides time-resolved, 3-directional blood flow velocity field measurements in an anatomical district of interest. Velocity fields measured by 4D flow are affected by noise and suffer from low spatio-temporal resolution, dramatically limiting the applicability of this imaging modality for near-wall quantity (e.g. wall shear stress (WSS)) extraction and hemodynamics-based risk stratification. This chapter describes the development of a neural network-based approach for denoising and super-resolution of blood flow velocity fields measured by 4D flow MRI; the described method is a novel solution for rapid and accurate quantification of hemodynamic biomarkers in thoracic aortic aneurysm.

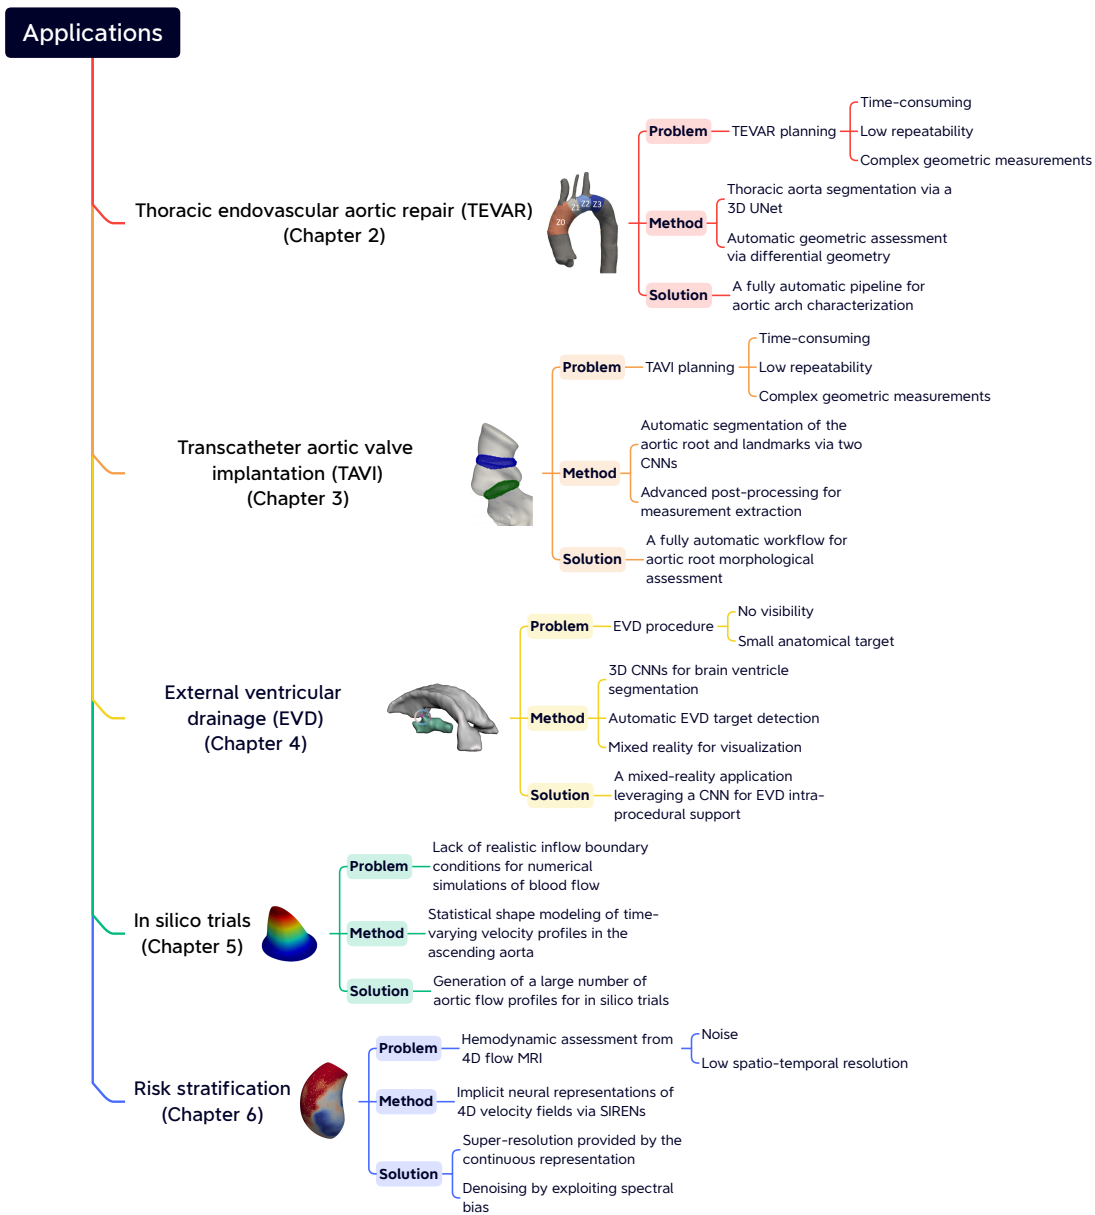


Figure 1.2: Graphical outline of the thesis. The present work encompasses a variety of biomedical applications; the addressed clinical problem, employed methods and solutions are outlined for each chapter.

1.3.2 Aim of the thesis

The overarching objective of this thesis is to develop data-driven and ML-based tools for analyzing medical images, with the aim of extracting quantitative features that are quickly and effectively utilized for a specific biomedical application. The approaches devised in the present work aimed to provide automatic assessment of both anatomical (Part I) and functional (Part II) features and biomarkers. We employ supervised learning approaches to address anatomical or morphological assessment to either support pre-procedural planning for minimally invasive cardiovascular interventions (Chapters 1 and 2), or assist neurosurgeons during procedures (Chapter 3). Then, we make use of unsupervised learning methods to extract functional characteristics, namely, hemodynamic features in the ascending aorta, with the aim to boost accurate development of digital twins for *in silico* trials (Chapter 4), and compute biomarkers for more accurate risk stratification (Chapter 5).

Part I

Supervised approaches for morphological assessment

CHAPTER 2

A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography

Based on:

Saitta, S., Sturla, F., Caimi, A., Riva, A., Palumbo, M. C., Nano, G., Votta, E., Della Corte, A., Glauber, M., Chiappino, D., Marrocco-Trischitta, M., & Redaelli, A. (2022). A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography. Journal of Digital Imaging, 35(2), 226-239.

Chapter Summary

Feasibility assessment and planning of TEVAR require CT-based analysis of geometric aortic features to identify adequate LZs for endograft deployment. However, no consensus exists on how to take the necessary measurements from CT image data. We trained and applied a fully automated pipeline embedding a CNN, which feeds on 3D CT images to automatically segment the thoracic aorta, detects proximal landing zones (PLZs), and quantifies geometric features that are relevant for TEVAR planning. For 465 CT scans, the thoracic aorta and pulmonary arteries were manually segmented; 395 randomly selected scans with the corresponding ground truth segmentations were used to train a CNN with a 3D U-Net architecture. The remaining 70 scans were used for testing. The trained CNN was embedded within computational geometry processing pipeline which provides aortic metrics of interest for TEVAR planning. The resulting metrics included aortic arch centerline radius of curvature, PLZ maximum diameters, angulation and tortuosity. These parameters were statistically analyzed to compare standard arches vs. arches with a common origin of the innominate and left carotid artery (CILCA). The trained CNN yielded a mean Dice score of 0.95 and was able to generalize to 9 pathological cases of thoracic aortic aneurysm, providing accurate segmentations. CILCA arches were characterized by significantly greater angulation ($p=0.015$) and tortuosity ($p=0.048$) in PLZPLZ 3 vs. standard arches. For both arch configurations, comparisons among PLZs revealed statistically significant differences in maximum zone diameters ($p<0.0001$), angulation ($p<0.0001$) and tortuosity ($p<0.0001$). Our tool allows clinicians to obtain objective and repeatable PLZs mapping, and a range of automatically derived complex aortic metrics.

2.1 TEVAR preprocedural planning

TEVAR represents the preferred approach for the treatment of pathologies of the descending thoracic aorta, including aortic aneurysm and dissection [9]. Advancements in endograft engineering and design have allowed the application of this minimally invasive technique also for the treatment of aortic disease involving the proximal tract of this vessel, namely the aortic arch. However, despite the favorable short-term outcomes with respect to surgery, TEVAR of the aortic arch still carries a significant risk of medium- and long-term complications, including endoleak, endograft migration and collapse [10, 11], likely due to its peculiar geometric and fluid-dynamic features [12]. As of today, no standard risk assessment tool exists for arch TEVAR, but a thorough preoperative decision-making process is considered crucial for an uneventful endograft deployment, and a durable clinical success [13, 14, 15]. Feasibility of TEVAR and correct endograft sizing are based on measurements of Ishimaru's proximal landing zones [16], i.e., the aortic wall sections corresponding to potential endograft apposition sites, during pre-procedural planning. A minimum length of 20 mm, and maximum diameter of 40 mm are required for the proximal and distal landing zones to ensure adequate endograft apposition, and a variable stent graft diameter oversizing is recommended based on aortic disease [13, 14]. These criteria, however, neglect additional geometric parameters, such as angulation and tortuosity, which proved to be associated with an increased risk of endograft failure after repair, and therefore allow to identify hostile landing zones for TEVAR [11, 17, 18]. Moreover, the presence of anomalous aortic arch configurations may further complicate TEVAR planning. Of note, arches with a common origin of the innominate and left carotid arteries (CILCA) were shown to have a higher risk of developing thoracic aorta disease [19], and to be associated with dismal clinical outcome after TEVAR [20]. Interestingly, both findings are apparently related to the consistent and peculiar geometric and fluid-dynamic pattern of proximal landing zones of the CILCA arch [21, 22]. Contrast-enhanced CT represents the preferred imaging modality [13, 23] to assess the pre-procedural aortic geometry, but no consensus nor detailed guidelines exist on how to extract the necessary measurements from CT image data. Such an analysis requires manual detection and definition of anatomical landmarks [23], and in some cases measurements are manually taken directly on 2D CT image slices, thus preventing a true and accurate 3D characterization [24]. In fact, the inherent 3D nature and complexity of the anatomical data, together with the high inter-operator variability, can significantly affect the measurements, thus hindering a repeatable geometry-based pre-procedural planning for TEVAR. Nonetheless, volumetric CT-based segmentation of the thoracic aorta can provide a 3D model on which the relevant geometric parameters can be measured objectively and systematically. However, this is a time-consuming process which is rarely carried out in clinical practice. Recent advances in deep learning algorithms for computer vision have allowed for automating complex tasks involved in medical image analysis, including fully automatic segmentation [25]. In particular, CNNs have achieved state-of-the-art results on a wide range of anatomical structures, including brain tumor [26], left atrium [27], liver [28] and oth-

ers belonging to the Medical Segmentation Decathlon Dataset [29]. A CNN can learn to perform segmentation through a training set, i.e., by iteratively processing a large amount of sample images, comparing the automatically segmented anatomy vs. the corresponding ground truth segmentation, and refining the tuning of the parameters to minimize the observed mismatch. After training, a CNN can be validated with a test set, i.e., by processing completely new images unused during training. In this study, we developed and applied a fully automated pipeline that feeds on 3D CT images and provides a comprehensive geometric analysis of the aortic arch, for both standard and CILCA arch configurations. Our pipeline allows clinicians to obtain automated segmentation of the thoracic aorta, automated proximal landing zones mapping, and a range of automatically derived aortic metrics. Our tool is also suitable to be used for unbiased quantitative analyses of geometric parameters in population studies.

2.2 Building an automatic pipeline from CT data to complex geometric characterization of the thoracic aorta

2.2.1 Data Collection

We retrospectively retrieved 3D CT scans of 515 subjects, acquired between 2010 and 2020, from three Italian referral center for aortic disease. Acquisitions with artifacts due to movement, presence of metal devices, or with slice thickness greater than 1.5 mm were excluded. Accordingly, 465 subjects were eventually included in this study and, among them, 219 subjects presented a CILCA arch. Within the entire dataset, 9 cases had thoracic aortic aneurysm (TAA). For the collected CT scans, pixel spacing ranged from $0.29 \times 0.29 \text{ mm}^2$ to $0.976 \times 0.976 \text{ mm}^2$, while slice thickness ranged between 0.29 and 1.5 mm. The study was approved by the local ethics committee and informed consent was waived because of the retrospective nature of the study and the analysis of anonymized data.

2.2.2 Image Pre-analysis and Manual Segmentation

To standardize the dataset, images were cropped to exclude regions below the celiac trunk, when present. Subsequently, all the CT volumes were manually segmented by five independent and experienced operators, guaranteeing that each acquisition was segmented by at least one operator to reduce single operator bias and ensuring an almost even/balanced split of the dataset among them. To avoid potential bias due to the choice of a specific segmentation software, manual segmentations were performed either using ITK-snap [30], 3D Slicer [31] or Materialise Mimics Medical v.22.0 (Materialise, Leuven, Belgium), and using gold standard segmentation methods such as level-sets, region growing and thresholding.

2.2.3 Neural Network Training

The dataset was subsequently divided into training and test sets: 395 (85%) scans of healthy subjects with their corresponding ground truth segmentations were randomly selected and used for neural network training (training set), while the remaining 70 scans (15%), including 61 healthy subjects and 9 patients, were kept out for testing (test set). The training set included both standard and CILCA aortic arches in random amounts. To investigate the network’s generalization capability in real clinical scenarios, all the 9 pathological cases were included in the test set. Herein, the trained CNN was based on the 3D U-Net architecture proposed in [32], which is composed of a input layer of 16 filters, encoder and decoder branches of 5 resolution levels each, a bottleneck block and skip connections via concatenation. Training was accomplished through a patch/batch-based strategy, which consists in subdividing the image volume into fixed-size subvolumes (patches), and processing groups of N patches (batches) separately [33]. Patch size was set to 128x128x128 pixels, while batch size (N) was set equal to 2. A data augmentation routine was used to increase dataset diversity. The random transformations applied to the processed patches included rotation, mirroring and elastic deformation [33]. The network was implemented in Pytorch and trained in parallel on 2 NVIDIA V100 GPUs over 300 epochs, each one defined as an iteration over 300 batches. A combination of Lovsz-Softmax loss [34] and focal loss [35] was used to trained the network end-to-end to perform multi-structure segmentation of the thoracic aorta and pulmonary arteries.

2.2.4 Pipeline Implementation

We implemented a fully automated pipeline, embedding the trained CNN, for the geometrical characterization of the aortic arch anatomy (Figure 2.1). The pipeline consists of the following main processing steps: i) segmentation, ii) detection of vessel ends and extraction of centerlines, iii) landing zone mapping and v) parameter computation.

Segmentation

The trained CNN receives a stack of 3D CT images and automatically segments the thoracic aorta and the pulmonary trunk and arteries, yielding a multi-label mask image.

Vessel end detection and centerline extraction

Triangulated surface meshes for the aorta and pulmonary vessels are obtained from the multi-label segmentation image using a marching cubes algorithm [36]. Centerlines of the aorta and of the pulmonary arteries are extracted from the corresponding surface meshes using the vascular modeling toolkit (VMTK) library [37], according to which vessel centerlines are computed as the locus of the centers of maximal inscribed spheres [38]. Centerline computation requires definition of a seed and target points on the surface of interest, which are automatically found by taking advantage of the

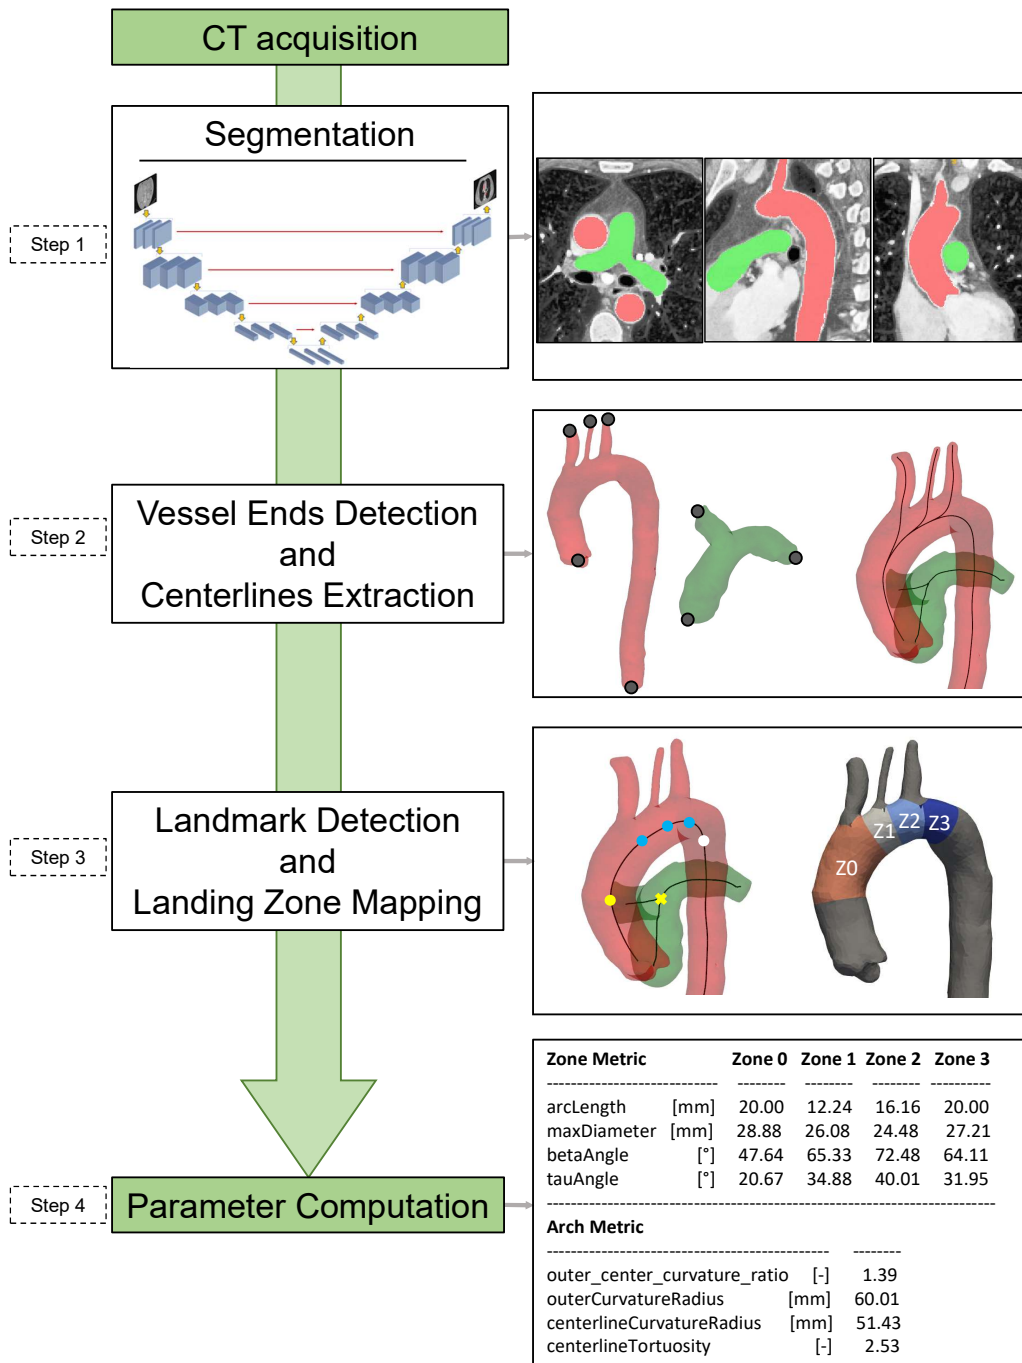


Figure 2.1: Schematic representation of the implemented automatic pipeline. In steps 1-3, red domains and surfaces represent the thoracic aorta; green domains and surfaces represent the pulmonary artery.

Laplace-Beltrami operator (LBO) [39]. Eigenfunctions of LBO provide a set of bases for continuous functions defined on the surface manifold, allowing one to form a basis which is descriptive of the analyzed shape [40].

Detection of seed and target points for the aortic centerline. The LBO of the triangulated aortic surface is computed. For tubular surfaces with a dominant longitudinal direction, such as the thoracic aorta, the first eigenvector of the LBO provides for each point on the surface a scalar measure of its proximity to the left ventricle (Figure 2.2). The points on the aortic surface associated with the minimum and maximum of this shape descriptor were used as centerline seed and target.

Algorithm 1: Detection of seed and target points for pulmonary vessel centerlines

Data: Reference surface (S_r), processed surface (S_p)

Result: Coordinates and indices of seed and target points on the processed surface

Previous manual annotation of:

$s_{idx} \leftarrow$ seed index;

$l_{idx} \leftarrow$ left pulmonary artery target index;

$r_{idx} \leftarrow$ right pulmonary artery target index;

for $idx \in [s_{idx}, l_{idx}, r_{idx}]$ **do**

 Find the point's transformed coordinate on S_r ;

 Find the closest point on S_p ;

 Store point coordinates and index;

end

Detection of seed and target points for pulmonary vessels centerlines. In this case, the bifurcated nature of the pulmonary trunk and arteries does not allow for a straightforward identification of seed and target points by computing the LBO eigenfunctions. The automation of this step required prior processing of a reference pulmonary artery surface that was randomly selected from the dataset. For the reference geometry, centerline seed and targets points are manually identified, and their indices stored. Then, for every processed case, the reference surface is morphed into the processed surface using a previously implemented large deformation diffeomorphic metric mapping (LD-DMM) [41]. After registration, a soft point-to-point correspondence is achieved between the processed and the reference surfaces. The three points corresponding to seed and targets are then defined on the processed surface as the nearest to the previously stored point indices on the reference surface (Figure 2.2b, Algorithm 1).

Detection of Target Points for Supra-Aortic Branch Centerlines Computation The definition of the LBO on the aortic surface is used to solve a heat equation and to compute, for every point on the processed surface, the HKS descriptor [42]. The computed

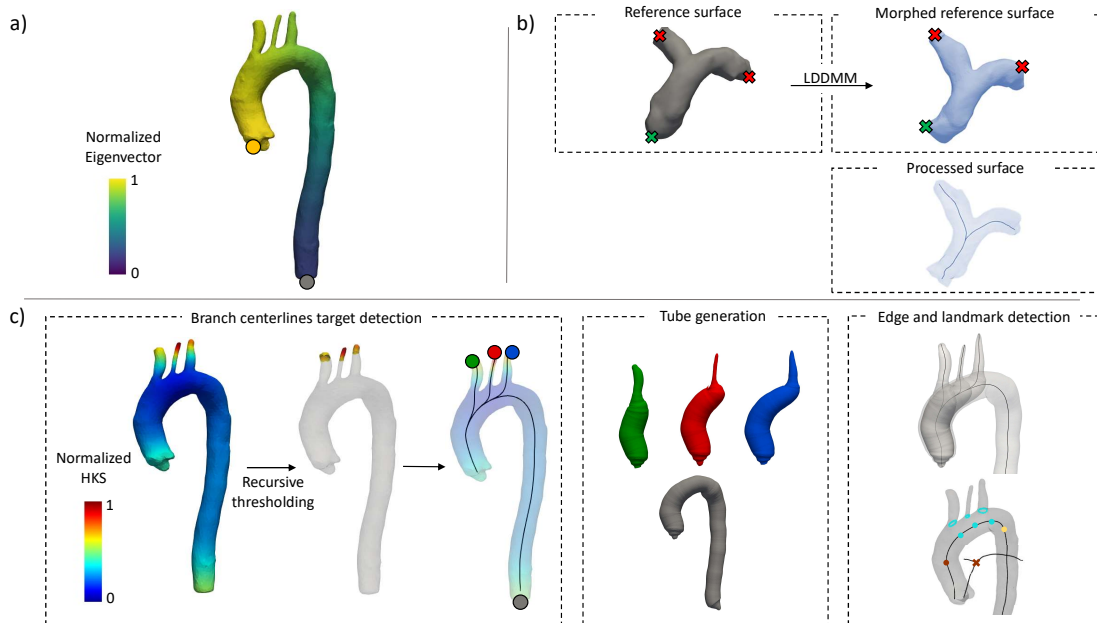


Figure 2.2: Representation of the algorithms for automatic detection of centerlines seed and target points. a) the first normalized eigenvector of the Laplace-Beltrami operator used to find seed (yellow circle) and target (grey circle) points for the main aortic centerline. b) seed (green cross) and targets (red crosses) are manually annotated on the reference surface, which is morphed into the processed surface. From the morphed surface, left and right pulmonary arteries centerlines are computed. c) the normalized HKS is computed on the aortic surface and recursively thresholded to isolate supraaortic branch ends; tube surfaces are generated for the aorta (grey), brachiocephalic trunk (green), left carotid artery (red) and subclavian artery (blue)

HKS is normalized with respect to its maximum value. The HKS gives a scalar approximation of the local curvature of the surface by measuring the amount of heat that remains after a specified time at each point after application of a unit heat impulse. The normalized HKS can be visualized on the surface mesh (left panel of Figure 2.2c). A recursive thresholding approach is used to identify supra-aortic branch extremities. Briefly, a thresholding filter is initialized to extract the surface regions with an associated normalized HKS between 1 and $1 - \delta$, with $\delta = 0.01$. The value of δ is then recursively increased by a factor of 2 until the number of connected regions extracted after thresholding is equal to 3. To filter out undesired points characterized by a high curvature, candidate target regions are constrained to lie above the main aortic centerline. Finally, for each of the 3 extracted regions, the point with the highest associated normalized HKS is identified as a branch centerline target and 3 centerlines are finally computed (Figure 2.2c, Algorithm 2).

Algorithm 2: Detection of seed and target points for supra-aortic branches centerlines

Data: Triangulated surface manifold of the thoracic aorta (M)

Result: Coordinates and indices of target points on the aortic surface

HKS descriptor \leftarrow solution of the heat equation;

$HKS \leftarrow HKS / \max(HKS)$;

$N_c \leftarrow 0$;

$\delta \leftarrow 0.01$;

while $N_c < 3$ **do**

 Find subregions of M where $HKS > 1 - \delta$;

$N_c \leftarrow$ number of connected extracted regions;

$\delta \leftarrow 2\delta$;

end

for all extracted subregions do

 Find point with highest HKS ;

 Store point coordinates and index;

end

Landing zone mapping

Once calculated, centerlines are separated according to the vessel to which they correspond and a tube surface inscribed in the vessel is generated for each centerline [37]. Herein, a tube is defined as a centerline envelope whose cross-sections are circles with radius locally equal to the radius of the maximal inscribed sphere within the vessel. Intersections between tubes corresponding to supra-aortic branches and the tube of the main aortic centerline are computed (Figure 2.2c, right). In this way, three bounded lines silhouetting branch origins on the arch are generated (cyan rings in Figure 2.2c, right). For each edge, the landing zone landmark on the aortic centerline is identified as the projection of the most distal point in the edge onto the aortic centerline. Two

additional landmarks (brown and yellow circles, Figure 2.2c, right) are automatically placed on the centerline to delimit zone 0, proximally, and zone 3, distally. To automatically detect the landmark delimiting zone 0 proximally, the bifurcation point of the pulmonary vessels is identified (brown cross, Figure 2.2c, bottom right) and projected onto the aortic centerline, keeping the axial coordinate fixed (brown circle, Figure 2.2c, bottom right). The landmark delimiting zone 3 distally is automatically placed on the aortic centerline so as to fix the length of this zone to 20 mm, which is considered as the minimum length to identify a stable landing zone for TEVAR [23] (yellow circle, Figure 2.2c, bottom right).

Parameter computation

In this step (Figure 2.1 - step 4), two kinds of metrics are automatically quantified to characterize the aortic arch morphology.

Whole-arch metrics For quantitative characterization of the whole arch, two points are defined (Figure 2.3). Point A is identified as the most proximal point in the vessel centerline, falling within the ascending aorta. Point B is the point of the aortic centerline that lies on the same CT image axial plane passing through point A. The portion of the centerline bounded between A and B is considered. The outer curvature path corresponding to this centerline tract is found as the diametrically opposed path to the Dijkstra path [43] with respect to the aortic centerline. For both the centerline tract and the outer curvature path, the radius of curvature is computed in two steps; first, by projecting the 3D curve onto its best-fitting plane found by singular value decomposition and, second, by finding the radius of the best-fitting circle to the curve in a least-square sense (Figure 2.3, left). Additionally, arch centerline tortuosity is calculated as the ratio of the curvilinear length between A and B to their straight-line distance.

Zone-based metrics Arc length along the centerline, maximum diameter, β angle and tortuosity angle are calculated for each zone according to [23] (Figure 2.3, right). Maximum zone diameters are calculated as the diameters of the largest inscribed spheres within each zone [37]. Briefly, for each zone, the most proximal point (p_0) is identified together with its tangent unitary vector. Two other points, p_{20} and p_{40} , located 20 mm and 40 mm distal to p_0 along the centerline, respectively, are selected with their tangent unitary vectors. β is the angle formed by the tangent unitary vectors to p_0 and p_{40} . The tortuosity angle (referred to as τ in Figure 2.3, right) is defined as the angle formed by the vectors connecting p_0 to p_{20} and p_{20} to p_{40} .

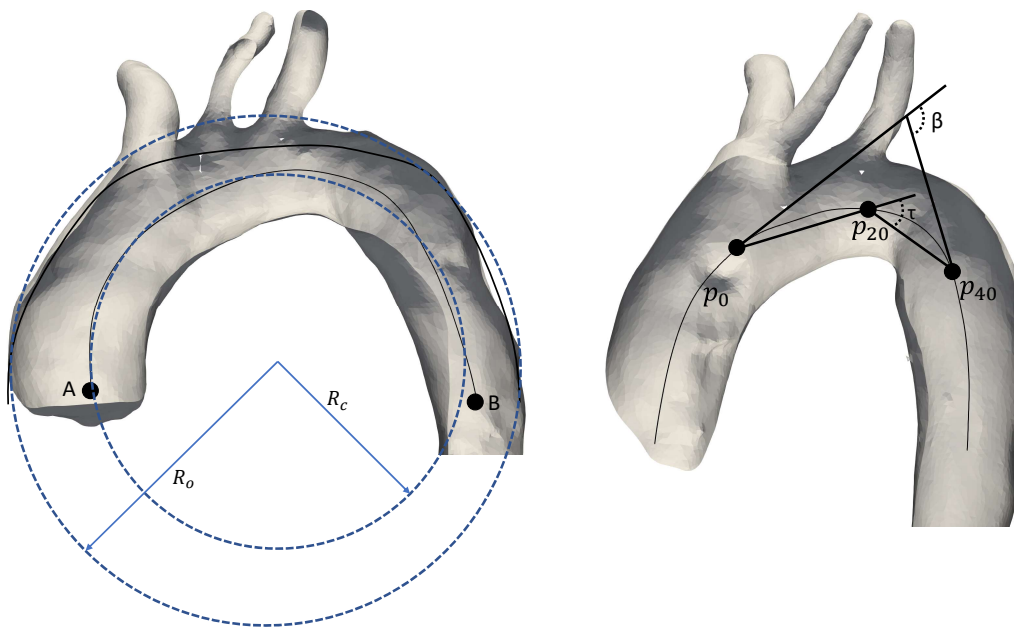


Figure 2.3: Automatic computation of geometric descriptors. A standard arch configuration (left) exemplifying centerline curvature radius (RC) and outer curvature radius (RO) (in blue). A CILCA arch (right) reporting β and tortuosity (τ) angles. N.B.: distances not to scale for display purposes

2.2.5 Statistical Analysis

Statistical analyses were performed using the Python scipy 1.5.1 statistics library. Data normality was determined using the Shapiro-Wilk test. Normally distributed variables are expressed as mean \pm standard deviation; non-normally distributed variables are expressed as median [min, max]. Comparison between standard and CILCA arches was accomplished using unpaired t-tests for normally distributed variables and Mann-Whitney U tests for non-normally distributed data. Comparison among zone-based geometric parameters for standard and CILCA arches were carried out using one-way ANOVA for normally distributed data and Kruskal-Wallis tests for non-normally distributed data. A p value lower than 0.05 was considered statistically significant.

2.2.6 Comparison with semi-automated measurements

Out of the 70 test cases, 30 were randomly selected and processed by an experienced operator using the semi-automated commercial software 3Mensio Vascular 8.0 (3Mensio Medical Imaging B.V., Bilthoven, The Netherlands). The results of the geometric analysis obtained with the proposed methodology were compared against the corresponding measurements taken through the commercial software.

2.3 Performance Evaluation

2.3.1 Test set results

The trained CNN was employed to run inference on the test set made of 70 CT scans, 9 of which presented TAA. The automatic segmentation algorithm proved able to accurately extract the thoracic aorta and pulmonary vasculature anatomies for all the subjects of the test subjects (Figure 2.4). CNN performance was quantitatively assessed, in terms of Dice coefficients and Hausdorff distances [44], with respect to manual ground truth segmentations. For the thoracic aorta, Dice score and Hausdorff distance were equal to 0.954 [0.873, 0.999] and 11.97 [1.96, 68.83] mm, respectively. For the pulmonary trunk and arteries, mean Dice score and Hausdorff distance were equal to 0.906 [0.710, 0.999] and 11.55 [0.94, 29.86] mm. The CT acquisition phase was chosen to specifically visualize aortic structures via the injection of a contrast-enhancing agent. The reduced presence of contrast agent in the pulmonary vasculature limited the relative brightness of these vessels with respect to the background. Thus, the trained neural network was considerably better at segmenting the thoracic aorta as compared to the pulmonary arteries. Nonetheless, the worse performance in pulmonary artery segmentation did not invalidate the automated detection of the bifurcation landmark, relevant for the identification of zone 0. The automatic pipeline effectively identified Ishimaru's landing zones for all the 70 test subjects, including the 9 pathological cases, as detailed in Figure 2.5.




Figure 2.4: Comparison of manual (ground truth, GT) segmentation in bright red (aorta, Ao) and bright green (pulmonary arteries, PA), vs. automatic segmentation obtained by the trained neural network (CNN) in dark red (aorta, Ao) and dark green (pulmonary arteries, PA) on a test case. Four axial slices with color fill are shown on the left. A sagittal and a coronal slices with dashed contours are shown on the right

Geometric and Statistical analysis

The developed automated geometric analysis (steps 2 to 4 of the pipeline) was applied to the entire dataset of manually segmented CT scans ($n = 465$).

Comparison of CILCA vs. standard arches Statistically significant differences were observed for whole arch metrics in centerline curvature radius ($p=0.001$), outer curvature radius ($p=0.004$) and centerline tortuosity ($p=0.021$), with CILCA arches exhibiting lower values (Table 2.1). Focusing on zone-based metrics, zone 0 was found to be significantly longer in CILCA than standard arches ($p<0.001$), whereas no significant difference was observed for zone 2 (Table 2.2). Within each landing zone, maximum diameter (Table 2.3) remained comparable ($p>0.089$) between standard and CILCA aortic arches. However, both standard and CILCA configurations exhibited progressive maximum diameter reduction from zone 0 to zone 3 ($p<0.0001$). In particular, CILCA arches showed a slightly greater variation of maximum diameter between adjacent zones (zone 0 – zone 2) with respect to standard arches (zone 0 – zone 1). In standard arches, the mean value of the maximum diameters ranged between 34.7 mm in zone 0 and 28.43 mm in zone 1. In CILCA arches, the mean value of the maximum diameters varied between 34.0 mm in zone 0 and 26.30 mm in zone 2. To better analyze this difference in vessel tapering, a 120 mm arch tract was extracted from every arch centerline starting from the most proximal point of zone 0 (Figure 2.6), and pointwise vessel diameters were plotted as means with 95% confidence intervals (Figure 2.6). In CILCA arches, a steeper curve was observed between 30 and 80 mm, indicating a more markedly tapered geometry of the vessel lumen. Significant differences were found in β angles (Table 2.4) between standard and CILCA aortic arch configurations in both zone

Variable	Standard	CILCA	<i>p</i> value
Centerline curvature radius [mm]	50.26[27.94, 71.31]	48.19[26.47, 68.45]	0.001
Outer curvature radius [mm]	65.94 ± 8.37	63.68 ± 8.19	0.004
Outer/centerline curvature ratio	1.14[0.97, 2.68]	1.13[0.95, 2.28]	0.153
Centerline tortuosity	3.10[1.82, 6.38]	2.96[1.63, 9.15]	0.021

Table 2.1: Whole arch-based geometric parameters calculated by the automatic tool for standard and CILCA arch cases. Median [min, max] or mean ± standard deviations are reported together with the *p* values obtained by the statistical analysis

Variable	Standard	CILCA	<i>p</i> -value ^a
Arc length [mm]			
Zone 0	48.63[21.46, 112.15]	53.02[22.49, 103.67]	< 0.001
Zone 1	8.08[2.25, 20.82]	–	–
Zone 2	16.66[3.36, 80.53]	16.87[0.64, 36.17]	0.142
<i>p</i> -value ^b	< 0.0001	< 0.0001	

Table 2.2: Zone arc lengths in mm calculated by the developed automatic tool for standard and CILCA arch cases. Median [min, max] are reported together with *p* values obtained from the statistical analyses.

^a Standard vs. CILCA. ^b Anova among landing zones.

0 (*p*=0.001) and zone 3 (*p*=0.015), with CILCA arches reporting larger median values. For both arches, β angles significantly differed (*p*<0.0001) among the arch zones, with zone 3 reaching the highest values (Table 2.4). Similarly, a statistically significant difference in tortuosity angle between the two configurations was found in landing zones 0 (*p*=0.001) and 3 (*p*=0.048), with larger angles observed in CILCA arches (Table 2.5). Tortuosity angles were observed to be significantly different among landing zones for both standard and CILCA arches (*p*<0.0001). In both arch types, tortuosity angle reached the highest mean values in zone 3.

Measurements on TAA cases In order to be considered suitable, the proximal landing zone for TEVAR must be at least 20 mm long [13]. For the 9 TAA cases, β and tortuosity angles of zones fulfilling this criterion are reported in Figure 2.5. Among these anatomies, only 3 cases had a sufficiently long zone 2 (Figure 2.5 c, e, i). Severe angulation ($\beta < 60^\circ$) was observed in zone 0 (Fig. 5 b), in zone 2 (Fig.5 e), and in zone 3 of one arch (Figure 2.5, a). Tortuosity angles greater than 30° were observed in zone 0 for two arches (Figure 2.5b, g) and in zone 3 for one arch (Figure 2.5a).

Comparison vs. Semi-Automated Measurements

The differences between the measurements yielded by the proposed automated tool and those obtained by an experienced operator using the semi-automated software

Variable	Standard	CILCA	p-value ^a
Maximum diameter [mm]			
Zone 0	34.70[21.03, 52.78]	34.00[21.36, 53.32]	0.247
Zone 1	28.43[17.09, 43.47]	–	–
Zone 2	26.49[15.25, 42.66]	26.30[14.23, 52.75]	0.089
Zone 3	25.22[7.13, 36.04]	24.76[12.19, 50.35]	0.152
p-value ^b	< 0.0001	< 0.0001	

Table 2.3: Zone maximum diameters in mm calculated by the developed automatic tool for standard and CILCA arch cases. Median [min, max] are reported together with p values obtained from the statistical analyses.

^a Standard vs. CILCA. ^b Anova among landing zones.

Variable	Standard	CILCA	p-value ^a
β angle [°]			
Zone 0	50.17[13.22, 102.66]	54.13[21.90, 77.38]	0.001
Zone 1	46.07[9.68, 104.49]	–	–
Zone 2	54.55 ± 14.07	53.15 ± 14.92	0.297
Zone 3	57.09[13.69, 132.41]	59.61[15.45, 98.88]	0.015
p-value ^b	< 0.0001	< 0.0001	

Table 2.4: Zone β angles in degrees calculated by the developed automatic tool for standard and CILCA arch cases. Mean ± standard deviations are reported together with p values obtained from the statistical analyses.

^a Standard vs. CILCA. ^b Anova among landing zones.

Variable	Standard	CILCA	p-value ^a
Tortuosity angle [°]			
Zone 0	24.81[6.06, 53.70]	27.18[11.63, 39.94]	0.001
Zone 1	24.05[5.58, 60.25]	–	–
Zone 2	27.14 ± 7.53	26.60 ± 8.14	0.462
Zone 3	28.68[6.74, 76.38]	29.85[10.75, 54.19]	0.048
p-value ^b	< 0.0001	< 0.0001	

Table 2.5: Zone tortuosity angles in degrees calculated by the developed automatic tool for standard and CILCA arch cases. Mean ± standard deviations are reported together with p values obtained from the statistical analyses.

^a Standard vs. CILCA. ^b Anova among landing zones.

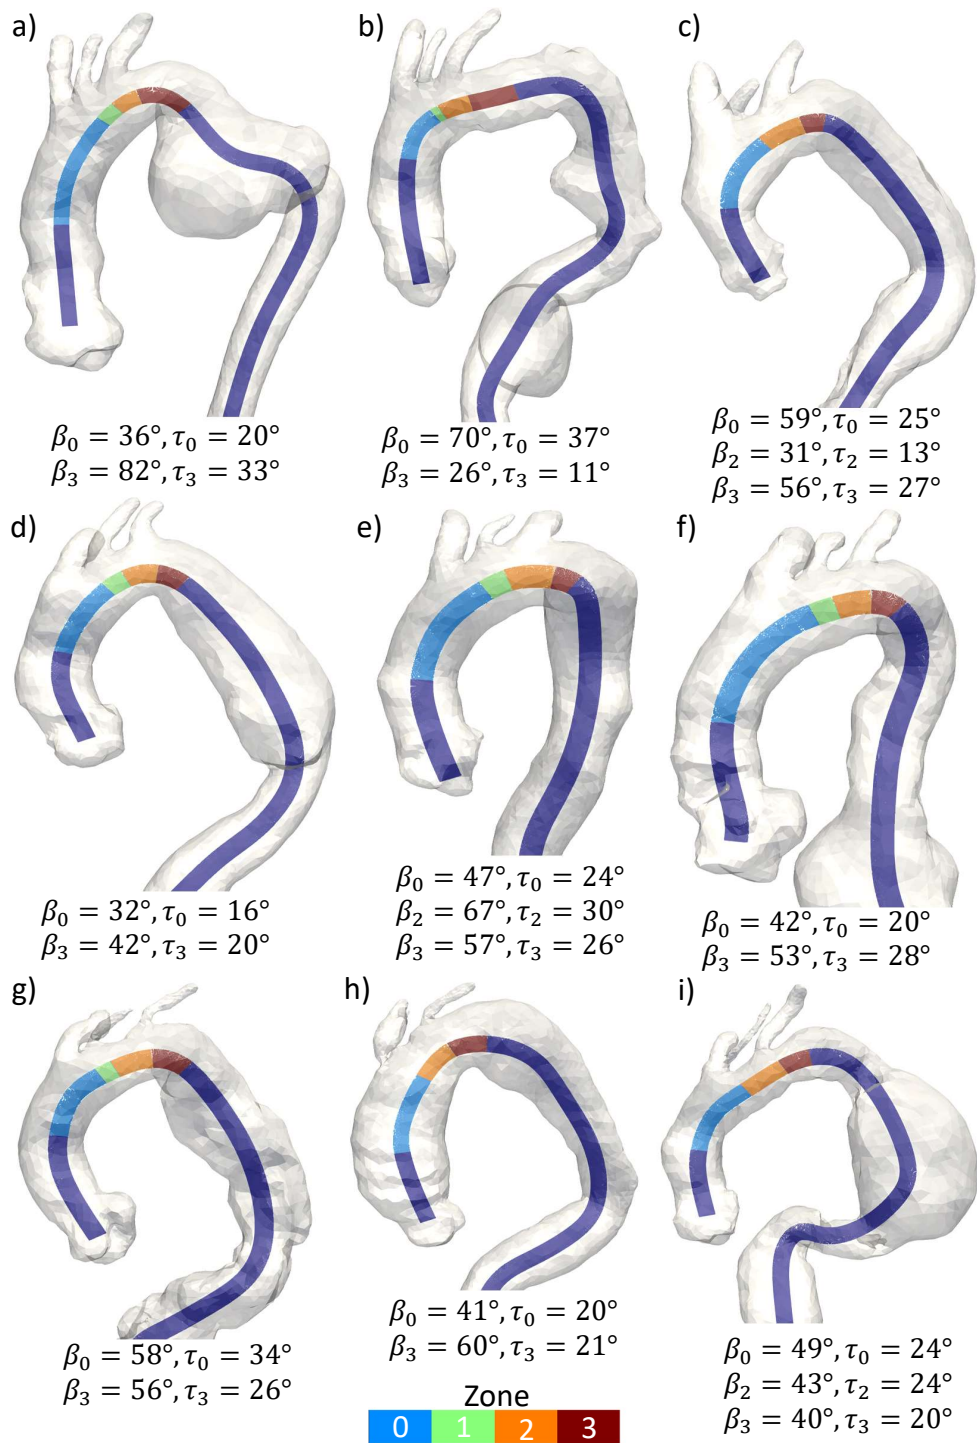


Figure 2.5: Aortic arches with thoracic aortic aneurysm from the test set with center-lines colored by landing zone. Six arches present a standard configuration with the presence of zone 1; all other arches present a CILCA arch configuration without zone 1. β and tortuosity (τ) angles for zones longer than 20 mm are reported, with subscripts indicating the zone they refer to.

package 3Mensio are detailed in Table 2.6. Notably, 3Mensio does not allow to compute a best-fitting circle to estimate curvature radii; hence, these were calculated by dividing the centerline (or outer curvature) tract by π , thus considering it equal to a semicircle. Furthermore, the commercial software now calculates the tortuosity angle using an arm length of 15 mm, rather than 20 mm (used by our tool), and it does not allow for changing this parameter. Overall, good agreement was found between the two measurement approaches, with linear measurements (lengths and diameters) differing by less than 5.4 mm, and angular measurements (β and tortuosity angles) differing by less than 6.5° .

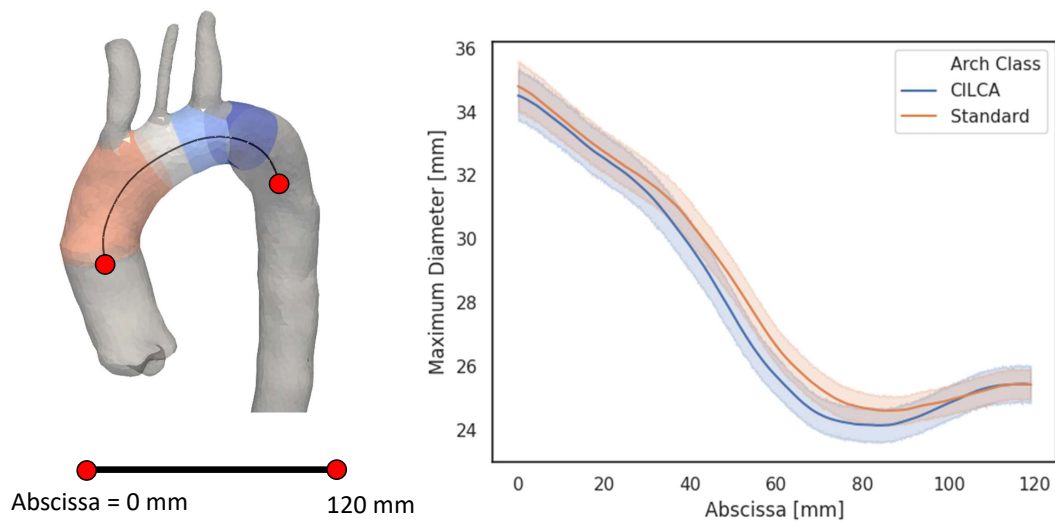


Figure 2.6: Representation of a 60 mm centerline tract covering the landing zones (left), and maximum diameter plotted against the centerline abscissa (right). Values for standard (orange) and CILCA (blue) cases are plotted as mean (solid lines) with 95% confidence intervals (transparent bands)

			Bias	95% limits of agreement
	3Mensio	Our tool	(3Mensio – our tool)	
Centerline curvature radius [mm]	52.31	52.27	0.05	[-2.1, 2.19]
Outer curvature radius [mm]	67.31	68.41	-1.09	[-4.03, 1.84]
Outer/centerline curvatures ratio	1.29	1.31	-0.02	[-0.03, 0.0]
Centerline tortuosity	1.66	3.05	-1.39	[-1.63, -1.15]
Arc length [mm]				
Zone 0	50.67	50.24	0.43	[-2.7, 3.56]
Zone 1	10.06	7.52	2.55	[0.89, 4.2]
Zone 2	23.28	17.89	5.39	[-1.43, 12.22]
Maximum diameter [mm]				
Zone 0	37.82	36.43	1.39	[0.37, 2.42]
Zone 1	30.09	28.73	1.36	[0.22, 2.5]
Zone 2	27.41	26.95	0.47	[-0.04, 0.97]
Zone 3	25.74	25.54	0.19	[-0.43, 0.82]
Beta angle [°]				
Zone 0	56.4	49.95	6.46	[2.76, 10.15]
Zone 1	45.18	46.82	-1.64	[-6.21, 2.92]
Zone 2	49.37	51.47	-2.11	[-5.85, 1.63]
Zone 3	57.07	56.68	0.39	[-2.3, 3.08]
Tortuosity angle [°]				
Zone 0	22.4	24.72	-2.31	[-4.75, 0.12]
Zone 1	15.41	25.1	-9.69	[-13.0, -6.37]
Zone 2	17.03	26.2	-9.17	[-11.61, 6.72]
Zone 3	24.2	28.6	-4.4	[-6.46, -2.33]

Table 2.6: Comparison between geometric measurements obtained with the commercial software 3Mensio vs. our tool. Mean of differences and 95% limits of agreement are reported. Zone 1 metrics refer to standard arches only, while all other metrics refer to both standard and CILCA arches

2.4 Comparison with state of the art and solutions

In the present work, we developed, applied and evaluated a fully automated pipeline for TEVAR pre-procedural planning support and aortic arch geometric analysis from CT. Our pipeline hinges on a CNN with a 3D U-Net architecture to perform automatic thoracic aorta segmentation, and on a series of geometrical computations to quantify several aortic metrics of interest for TEVAR pre-procedural planning. The adopted 3D CNN is at the base of the developed application, providing multi-structure end-to-end segmentation that does not require any manual input nor parameter fine-tuning, thus ruling out any dependency on the operator. This proved sufficient to correctly identify supra-aortic related landmarks in all 70 test cases in our dataset. For the thoracic aorta, average Dice score, Hausdorff distance and intersection-over-union (IoU) coefficient were 0.954, 11.97 mm and 0.91, respectively. This level of accuracy is higher than the results reported by other studies focusing on CNNs for automatic aorta segmentation, including Noothout et al. [45] who reported Dice scores of 0.88, and [46] who obtained Dice scores and Hausdorff distances up to 0.93 and 50.0 mm, respectively. Nonetheless,

results from [45] refer to a different test set of low-dose, non-contrast enhanced CT data, hence a direct comparison cannot be made. Our IoU coefficients are greater than the scores obtained by other studies focusing on lung nodule segmentation [47]. Despite being trained only with CT scans of healthy subjects, our network was also able to perform an accurate segmentation of a small subset of diseased aortas. In fact, our training strategy, including non-rigid image deformation and scaling during data augmentation, allowed the trained U-Net to generalize to 9 pathological TAA cases, yielding an average Dice score of 0.94. After automated segmentation, a series of algorithms for detection of anatomical landmarks were implemented through differential geometry solutions. Computation of the LBO and its eigenvalues, was essential to achieve a robust workflow that enabled fully automated mapping of proximal landing zones for all the subjects in our cohort. To our knowledge, the developed methodologies represent a unique application of complex digital surface processing algorithms to subject-specific anatomies in a real and relevant clinical context. From a clinical standpoint, our work addresses three relevant issues in pre-procedural TEVAR planning. First, extraction of clinically relevant geometric parameters [10, 16, 17], namely angulation and tortuosity, is technically difficult and time consuming, and may require ad hoc commercial software. Our approach does not require any specific skill from the user, as it allows quantifying automatically β and tortuosity angles of proximal landing zones, thus proving clinicians with more comprehensive information on the anatomy to treat by TEVAR. Second, current commercial solutions always require manual input, which introduces inter-operator variability. This is not the case in our tool, which delivers an objective, fully repeatable and systematic framework. Third, our solution shortens the processing time, making it compatible with the clinical routine, and applicable to large CT series for research purposes. On average, the time required by the user to process one case using commercially available software is around 30 minutes, whereas less than 7 minutes are required by our tool, i.e., about 4 and 3 minutes for segmentation and geometric analysis, respectively. The reliability of our automated tool was confirmed by the comparison with geometric measurements obtained with the commercial 3Mensio Vascular 8.0 software (3Mensio Medical Imaging B.V., Bilthoven, The Netherlands). Good agreement was found between arch curvature radii, zone lengths and maximum diameters. Given the average voxel spacings of our dataset ($0.75 \times 0.75 \times 1.25 \text{ mm}^3$), the obtained difference biases never exceeded a length equal to twice the voxel diagonal. Larger discrepancies were observed between zone angles and centerline tortuosity. The reason behind these differences lies in the way the aortic centerlines are calculated by the two methodologies. On the one hand, in our tool, centerlines are defined as the geometric locus of the points that are centers of the maximum spheres inscribed in the vessel surface. On the other hand, 3Mensio does not require a segmented aortic surface, but allows for manual placement of centerline points on 2D views. In complex anatomies, this may lead to an oversimplified centerline characterized by remarkably lower values of tortuosity and tortuosity angles, thus requiring additional and manual correction. Our objective analysis framework allowed investigating in a systematic fashion the differences between standard and CILCA arches in terms of metrics that

are relevant for TEVAR planning. The extracted geometric parameters were compared with aortic features generally quantified using dedicated commercial software as proposed in previous works [21, 23]. With our approach, we found relevant differences in whole-arch metrics between standard and CILCA arches, with the latter showing significantly lower centerline curvature radius ($p=0.001$), outer curvature radius ($p=0.004$) and centerline tortuosity ($p=0.021$). The average radii of curvature were 48.6 mm and 53.0 mm for standard and CILCA arches, respectively. Marrocco-Trischitta et al. [23], reported radii of curvature between 22.7 mm and 37.1 mm for a cohort of 60 standard arches. Such a considerable difference is due to the different methodologies adopted for the geometric analyses. In particular, in the previously referred studies the centerline radius of curvature was approximated as half the distance between two points, on the same axial plane, one in the ascending and one in the descending aorta. Implicitly, that approach assumes that the arch centerline delimited by these two points is exactly a half circumference, which is inherently planar. Our approach improves upon this, by computing the least squares best fitting circle to approximate aortic arch curvature radius, thus accounting for the actual deviations from the idealized profile assumed in the previous work [23]. Our analysis involving zone-based metrics revealed statistically significant geometrical differences in zones 0 and 3 between standard and CILCA arches. Zone 0 of CILCA arches had greater arc length ($p<0.001$), β angle ($p=0.001$) and tortuosity angle ($p=0.001$). Furthermore, greater β and tortuosity angles were found in zone 3 of CILCA arches. For standard arches, zone 3 β and tortuosity angles were equal to 57.09° [$13.69^\circ, 132.41^\circ$] and 28.68° [$6.74^\circ, 76.38^\circ$], respectively. Similar ranges of values were found by Marrocco-Trischitta et al. [23], who obtained β angles going from 52.3 ± 14.9 to 71.1 ± 14.1 and tortuosity angles between 25.9 ± 8.1 and 39.5 ± 8.2 on a cohort of 60 subjects. In a previous study from our group [11], zone 3 was found to be consistently associated with hostile geometric and hemodynamic features in terms of angulation, tortuosity and hemodynamic displacement forces [12]. The present analysis confirms those conclusions, highlighting the role of the β and tortuosity angles for CILCA arches. Furthermore, in the present study, the unfavorable geometric pattern displayed by CILCA arches is also suggested by the more noticeably tapered configuration followed by lumen enlargement (Figure 2.6). This feature may explain the tendency of CILCA arches to develop mild aortic dilation proximal to the aortic isthmus [10, 17, 18, 48].

Limitations. The obtained segmentation results were computed with respect to ground truth labels that were obtained using semi-automated tools, e.g. level-sets, region growing and thresholding. These tools also carry some inherent bias as opposed to a purely manual voxel-by-voxel labeling.

The CNN training, which requires a large number of CT scans, was performed on images of healthy subjects, and therefore our work must be conservatively regarded as a proof-of-concept study. Nevertheless, the CNN was successfully tested also for pathological CT scans, even though in a small group of aneurismatic patients. Regarding the applicability of our finding in a clinical scenario, however, previous confirmatory stud-

ies [49] showed that the aortic arch proximal landing zones present the same geometric pattern in both healthy and diseased aorta in aneurismatic patients. This is related to the fact that our analyses are based on centerline measurements, which are less likely to be affected by the modification of the aortic wall induced by the onset of pathological derangements. Notably, this applies only to patients affected with aneurysm and penetrating aortic ulcer [49], while the applicability of CNN training in aortic dissections remains to be proven, considering the peculiar anatomical features of this pathology.

2.5 Conclusions

In this work, we developed a new tool specifically tailored to automatically compute geometric parameters relevant to TEVAR planning. Our CNN successfully segmented the thoracic aorta and pulmonary vasculature of 70 CT scans unseen during training, including 9 pathological cases of TAA. Our fully automated and systematic analysis pinpointed significant differences in geometric aortic features that may explain the less satisfactory clinical results obtained by TEVAR in CILCA arches compared to standard ones. The reported tool is automated, quick and reliable, and appears to be an innovative solution to improve the decision-making process before TEVAR. Future efforts will be aimed at including more and diverse pathological anatomies in the training set, including namely aortic dissections, thus broadening the clinical scope of this methodology.

CHAPTER 3

A CT-based deep learning system for automatic assessment of aortic root morphology for TAVI planning

Based on:

Saitta, S., Sturla, F., Gorla, R., Oliva, O.A., Votta, E., Bedogni, F. & Redaelli, A. (2023). A CT-based deep learning system for automatic assessment of aortic root morphology for TAVI planning. arXiv preprint arXiv:2302.05378

Chapter Summary

An accurate planning of TAVI is important to minimize complications, and it requires anatomic evaluation of the aortic root (AR), commonly done through 3D CT image analysis. Currently, there is no standard automated solution for this process. Two CNNs with 3D U-Net architectures (model 1 and model 2) were trained on 310 CT scans for AR analysis. Model 1 performed AR segmentation and model 2 identified the aortic annulus and sino-tubular junction (STJ) contours. After training, the two models were integrated into a fully automated pipeline for geometric analysis of the AR. Results were validated against manual measurements of 178 TAVI candidates. The trained CNNs effectively segmented the AR, annulus and STJ, resulting in mean Dice scores of 0.93 for the AR, and mean surface distances of 0.73 mm and 0.99 mm for the annulus and STJ, respectively. Automatic measurements were in good agreement with manual annotations, yielding annulus diameters that differed by 0.52 [-2.96, 4.00] mm (bias and 95% limits of agreement for manual minus algorithm). Evaluating the area-derived diameter, bias and limits of agreement were 0.07 [-0.25, 0.39] mm. STJ and sinuses diameters computed by the automatic method yielded differences of 0.16 [-2.03, 2.34] and 0.1 [-2.93, 3.13] mm, respectively. The proposed tool is a fully automatic solution to quantify morphological biomarkers for pre-TAVI planning. The method was validated against manual annotation from clinical experts and showed to be quick and effective in assessing AR anatomy, with potential for time and cost savings.

3.1 TAVI planning and state of the art

TAVI has emerged as an alternative to traditional open-heart surgery to treat severe aortic stenosis, proving effective in reducing morbidity and mortality in high-risk patients [50, 51]. Despite its benefits, TAVI still carries risk of post-operative complications, including paravalvular leakage, device migration, annulus rupture and conductive disturbances [52, 53, 54]. A meticulous preprocedural planning is thus crucial for minimizing the risk of complications and should envisage an accurate anatomic assessment of the AR apparatus, which is essential for selecting the optimal prosthetic device size [55]. Accurate preoperative assessment is centered around the aortic annulus, and includes quantification of its diameters [56], angulation [57] and perimeter [58]. Furthermore, a comprehensive analysis of the whole AR can provide anatomical measurements of the STJ and sinuses of Valsalva [57, 58, 59]. To this purpose, three-dimensional (3D) CT angiography is the preferred imaging modality to quantify AR anatomy before a TAVI [60]. A comprehensive CT-based TAVI planning involves three main operations: segmentation of the anatomy, landmark detection, and measurement extraction [55, 59, 61, 62, 63]. However, there is currently no standardized fully automated solution, and taking the necessary measurements can be a time-consuming process which often involves several manual operations that introduce operator-dependency and may limit reproducibility. Automatic and semi-automatic systems have already been proposed to identify and quantitatively assess AR features from CT images. Lalys et al. [61] exploited semi-automatic segmentation tools relying on atlas-based methods to segment the AR, localize a wide range of anatomical landmarks (e.g., leaflet and coronary ostium positions) and obtain accurate quantification of annulus diameter. More recent approaches have taken advantage of CNNs to fully automate the detection of AR landmarks from 3D CT [55, 62]. Given the black-box nature of deep neural networks, when employing these tools extensive validation against manual landmark tracing by experts should be performed. To date, the largest patient validation set has been reported by Astudillo et al. [55], who proved the feasibility of building fast and accurate CNN-based systems for detection of the three aortic cusp nadirs and coronary ostia, training their model on 444 CT scans and validating their landmark detection accuracy on 100 patients. Nonetheless, the proposed approach could only be applied to contrast-enhanced CT images and did not include automatic measurement extraction. Hence, full automation of the entire process, including segmentation, landmark detection and extraction of aortic features relevant for TAVI, is still lacking. Given the increasing adoption of TAVI also for intermediate and low-risk patients [50], a reliable, fast and efficient method for assessing AR anatomy and determining the appropriate device size could have an increasingly broad impact by making TAVI planning faster, reliable and fully repeatable.

The main contribution of the present work was the development of an automatic workflow that combines deep learning techniques and tools from differential geometry to segment the AR from 3D CT data, extract AR-specific anatomical landmarks and compute clinically relevant measurements for TAVI planning. The effectiveness of the pro-

posed method was extensively evaluated on a group of 178 patients.

3.2 Development of an automatic method for aortic root characterization

3.2.1 Data collection and manual annotation

CT scans of 512 subjects acquired between 2010 and 2022 were retrospectively collected. Pixel spacing ranged from $0.26 \times 0.26 \text{ mm}^2$ to $0.87 \times 0.87 \text{ mm}^2$, while slice thickness ranged between 0.25 and 1 mm. 24 acquisitions were excluded because of reconstruction artifacts, presence of metal devices, or with slice thickness greater than 1 mm. Accordingly, the dataset was split in two subsets: dataset A, for which no AR measurement was available (N=310), and dataset B (N=178), for which an expert operator manually took measurements of the aortic annulus, STJ and Valsalva sinuses plane using the commercial software 3mensio Structural Heart (v8.2, Pie Medical Imaging, Maastricht, Netherlands) (Figure 3.1). After manual positioning, 3mensio allowed to compute maximum and minimum diameters, area, and perimeter of each plane. All patients in dataset B were TAVI candidates with severe aortic stenosis. For all CT scans in dataset A, an initial segmentation of the ascending aorta was obtained using a neural network previously trained by our group [64]. Segmentations that were erroneously inferred (e.g. incomplete filling) by the pre-trained model were adjusted by an experienced operator using a semi-automatic region growing algorithm [65]. All the resulting aortic segmentations include the left ventricle outflow tract (LVOT), the whole AR and the proximal portion of the ascending aorta according to the scan-specific field-of-view. Both aortic annulus and STJ were manually segmented for all CT scans in dataset A. To expedite the annotation process, a graphic-user interface (GUI) was appositely developed using VTK [66]. Through the designed GUI, the user could easily position a plane in the 3D image space and, exploiting the previously obtained segmentation of the aorta, the annulus or STJ labels were automatically assigned to the image points within an Euclidean distance of 3 mm the chosen plane and inside the aorta. All image data used in the present study were collected from IRCCS Policlinico San Donato (Milan, Italy). The study was approved by the local ethics committee and informed consent was waived because of the retrospective nature of the study and the analysis of anonymized data.

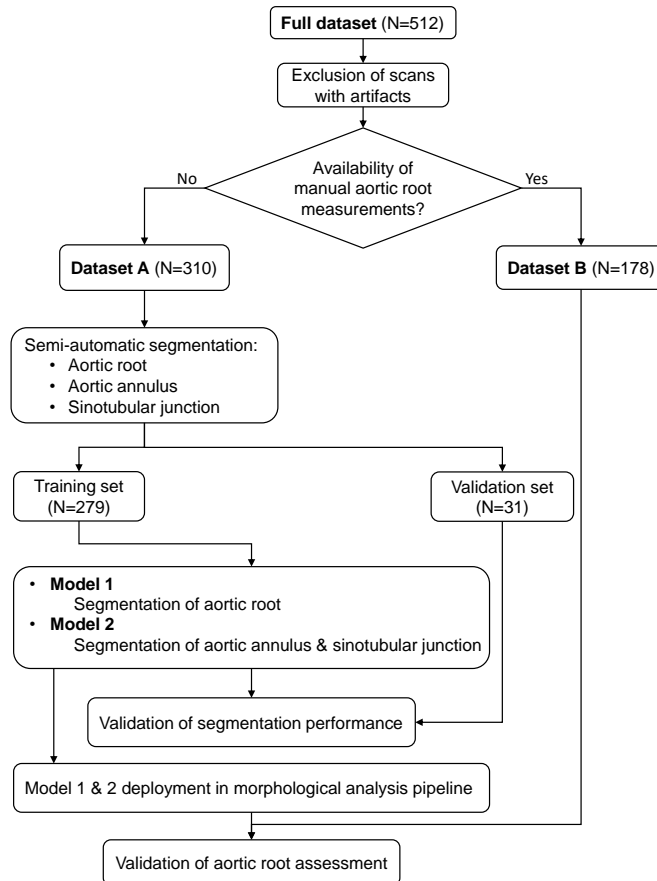


Figure 3.1: Schematic representation of the adopted workflow. Dataset A includes manual segmentations of the aortic root, annulus and STJ; it was used to train and validate two neural networks (model 1 and model 2) for automatic segmentation. Dataset B included manual annotations of aortic root measurements, used to validate the developed morphological analysis pipeline.

3.2.2 Analysis of inter-operator variability

To obtain an unbiased assessment of the inter-operator variability of aortic segmentation, a randomly selected subset of 20 CT scans belonging to dataset B was segmented by three experienced, completely independent operators (Op1, Op2 and Op3). Op1 and Op3 used the open-source segmentation software ITK-snap [65], exploiting a combination of region growing and paintbrush tools to annotate the images. Op2 was a certified user and used the CE-marked commercial software Mimics Medical v21.0 (Materialise, Leuven).

3.2.3 Image preprocessing: aortic root region detection

For each 3D CT scan in our datasets, a region of interest (ROI) encompassing the aortic root was cropped out (Figure 3.2). The detection of the ROI to crop is fully automatic and was achieved through a 3D template-matching approach [67]. An image X with its ground truth aortic root segmentation S was previously selected from dataset A. When processing a new case (target image, Y), the image X is roughly aligned to the target image using a rigid transformation (\mathbf{R}) followed by an affine transformation (\mathbf{A}) able to account for scaling and shear deformations. Defining the generic final transformation:

$$\mathbf{T}_{\alpha^*}(X) = (\mathbf{A} \circ \mathbf{R})(X), \quad (3.1)$$

parameterized by α^* , the optimal mapping is found as:

$$\alpha^* = \arg \max_{\alpha} I(\mathbf{T}_{\alpha^*}(X), Y), \quad (3.2)$$

where $I(\mathbf{T}_{\alpha^*}(X), Y)$ is the mutual information between the two images, which can be considered a nonlinear generalization of cross-correlation [68]. α^* is found using a gradient descent algorithm. Once found, \mathbf{T}_{α^*} is used to map the known ground truth aortic root segmentation into the target image domain, obtaining $\mathbf{T}_{\alpha^*}(S)$. The bounding box of $\mathbf{T}_{\alpha^*}(S)$ is dilated outward in each direction by 20 pixels and the resulting ROI is cropped out from the original image.




Figure 3.2: Template-matching approach for automatic identification of the ROI from 3D CT. A template image (X) for which aortic segmentation (S) is known, is affinely registered to the processed image (Y). The transformation mapping X to Y (T_{α^*}) is applied to S and the bounding box of $T_{\alpha^*}(S)$ is used to crop the ROI from Y .

3.2.4 Automatic segmentation

Dataset A was divided into training and test sets: 279 (90%) scans with their corresponding ground truth segmentations of the aorta, annulus and STJ were randomly selected and used to train two different neural networks, respectively dedicated to the segmentation of the aorta (model 1) and to the multi-class segmentation of the aortic annulus and STJ (model 2). The remaining 31 (10%) scans were used for validation (Figure 3.1). Both models were based on the 3D U-Net architecture proposed in [33], with encoding and decoding branches of 5 resolution levels each, defined using residual units as introduced by [69]. The number of filters used for the encoding branch was 32, 64, 128, 256, 512. Each encoding block consisted of a convolution layer with kernel size of $3 \times 3 \times 3$ and stride of 2, followed by parametric rectified linear units (PReLU) [70] and instance normalization. The encode and decode paths were connected with skip connections via concatenation. Model 1 and model 2 only differed by the number of output channels; given the different tasks, model 1 outputs a tensor with one channel, whereas model 2 returns a tensor with three channels (background, annulus and STJ). The number of trainable parameters for each of the two models was approximately 20M.

For comparison, we considered two additional automatic approaches for semantic segmentation of the aorta and of the aortic annulus and STJ.

Atlas-based

As a baseline, an atlas-based segmentation method was implemented. Similar to the ROI detection solution previously described, this approach exploits the registration of a reference image with its associated labels of the aorta, annulus and STJ, with the

image to segment. With respect to the approach described in the previous section, a free form deformation \mathbf{D} , based on cubic b-splines is added to the final transformation \mathbf{F}_{β^*} parameterized by β^* such that:

$$\mathbf{F}_{\beta^*}(X) = (\mathbf{D} \circ \mathbf{A} \circ \mathbf{R})(X). \quad (3.3)$$

Similar to the formalization presented in equation 3.2, the β^* that maximizes $I(\mathbf{F}_{\beta^*}(X), Y)$ is found using a gradient descent algorithm. More details of the algorithm can be found in [71]. Once found, \mathbf{F}_{β^*} is used to map the known labels of the aorta, annulus and STJ into the target image domain.

3D U-Net Transformer

Similarly to what was described above, two additional separate models with a 3D U-Net transformer architecture were implemented; model 3 for segmentation of the aorta and model 4 for segmentation of the annulus and STJ. Models 3 and 4 were based on the transformer architecture proposed in [72]. During a forward pass of a U-Net transformer, the 3D tensor is divided into a series of non-overlapping patches of size $16 \times 16 \times 16$. These patches are projected into an embedding space using a linear layer with 128 hidden channels. A position embedding is added to the resulting sequence, which is then used as input to a transformer model with 4 attention heads. The transformer model's various encoded representations are extracted from different layers, and they are combined with a decoder using skip connections. Finally, this process is used to predict the final segmentation. During training of all neural networks, an on-line data augmentation routine including random Gaussian noise, cropping, mirroring and rotation was implemented in the MONAI framework [73]. A Dice loss was used to train model 1 and 3, while a combination of Dice and Focal loss [74] with equal weights was used to train model 2 and 4. Training was carried out on an NVIDIA A100 GPU over 1000 epochs, using an Adam optimizer with learning rate of 0.0001.

3.2.5 Pipeline implementation

After training, models 1 and 2 were embedded in a fully automated pipeline for AR analysis and TAVI pre-procedural planning. The different sequential steps of the implemented pipeline are described below and exemplified in Figure 3.3.

Aortic surface processing

After ROI cropping, the CT scan is processed by model 1, generating a binary mask of the ascending aorta including the LVOT. A marching cubes algorithm extracts the corresponding contour as a triangulated surface; surface smoothing is applied with a windowed sinc function interpolation kernel with passband of 0.01, and adaptive remeshing is performed following the approach described in [75]. For each processed surface, the LBO [39] is computed. LBO eigenvectors form a set of bases for the definition of continuous functions on the surface manifold. For tubular structures with a

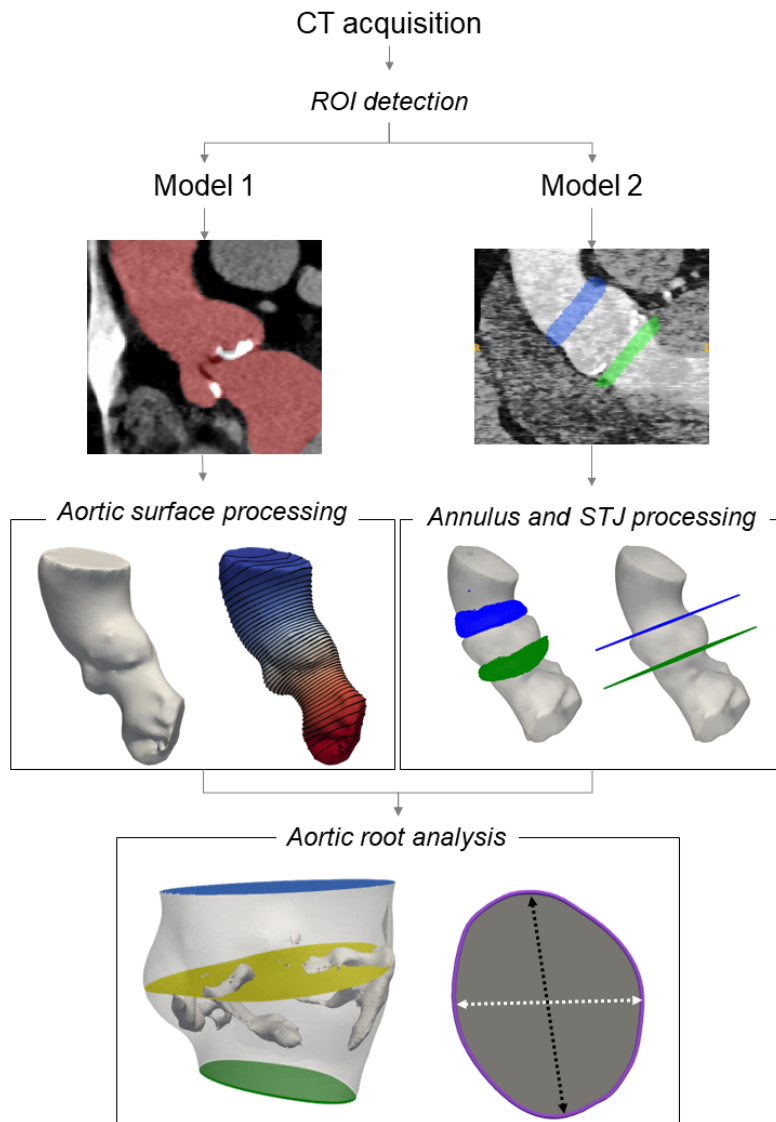


Figure 3.3: Schematic representation of the implemented automatic pipeline. The ROI is detected from the input CT scan. Model 1 infers the segmentation of the AR with the left ventricle outflow tract (LVOT), shown as the red label (left side). Model 2 infers the segmentation of the aortic annulus and STJ (in green and blue, respectively) (right side). The aortic surface processing step computes the second eigenvector of the LBO and its contours (left side). The annulus and STJ processing step performs refinement of model 2 segmentation. In the AR analysis step (bottom), the AR is isolated, and the Sinuses of Valsalva are detected. Anatomical measurements are computed together with calcium volume (shown in white, bottom panel).

dominant longitudinal direction, the second LBO eigenvector (E1) represents a scalar field approximating the curvilinear abscissa of the surface centerline; E1 isocontours on the aortic surface are shown in Figure 3.3.

Annulus and STJ processing

The cropped CT scan is processed by model 2, which infers the masks and the corresponding set of points X_{ann} and X_{STJ} for the annulus and STJ, respectively (Figure 3.4). Using a random sample consensus (RANSAC) iterative algorithm with inlier threshold of 1.5 mm, two best-fitting planes are identified, namely Π_{ann} and Π_{STJ} from X_{ann} and X_{STJ} , respectively. Specifically, a generic plane Π is represented by a bounded rectangular mesh grid identified by its normal vector \mathbf{n} and center of mass \mathbf{c} : $\Pi = \Pi(\mathbf{n}, \mathbf{c})$. The normal vector and center of mass of both Π_{ann} and Π_{STJ} , are refined (rotated and shifted) following a constrained optimization procedure, initialized using \mathbf{n} and \mathbf{c} computed by RANSAC. The optimal plane normal vector \mathbf{n}^* and center of mass $\mathbf{c} + \boldsymbol{\delta}$ are defined as the ones that minimize the area of intersection with the aortic surface:

$$\begin{aligned} \Pi^* = \min_{\mathbf{n}^*, \boldsymbol{\delta}} \text{area}(\Pi(\mathbf{n}^*, \mathbf{c} + \boldsymbol{\delta}) \cap \Gamma), \\ \text{subject to } \|\boldsymbol{\delta}\| < K, \end{aligned} \quad (3.4)$$

where Γ represents the AR and LVOT surface and $K = 5mm$. The described procedure is applied to both Π_{ann} and Π_{STJ} and the yielded planes are used to clip and isolate the AR (Figure 3.4.d). Within the extracted AR surface, the E1 isocontour with maximum area is used to identify the plane of the sinuses of Valsalva (Figure 3.4e).

Aortic root analysis

For the annulus, the STJ and the plane of Valsalva sinuses, the following metrics are automatically computed: i) area; ii) perimeter; iii) the maximum (D_{max}), iv) minimum (D_{min}) and v) mean ($D_{mean} = (D_{max} + D_{min})/2$) diameters. Maximum and minimum diameters are defined as the lengths of the largest and shortest segments connecting two opposite points of the perimeter while passing through the center of mass. In addition, for the aortic annulus the annulus angle, i.e., the angle formed by the plane normal to the foot-head image axis [57], is extracted. The AR calcium score is automatically quantified as the volume enclosing pixels with Hounsfield units (HUs) greater than 800 by a simple thresholding of the corresponding segmented AR [76].

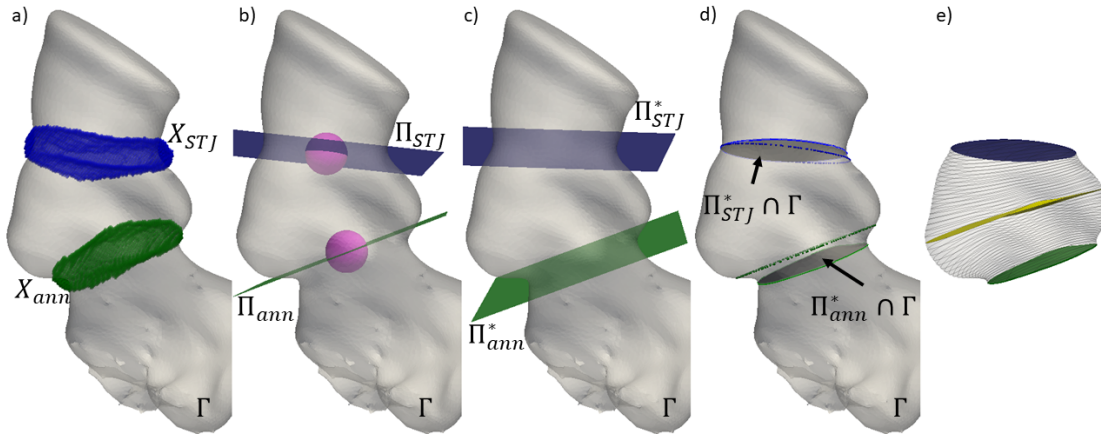


Figure 3.4: steps implemented for aortic annulus and STJ plane detection refinement. a) the point sets yielded by model 2 for the annulus (green) and STJ (blue), visualized with the aortic segmentation produced by model 1 (transparent gray). b) the best fitting planes for the two structures (Π_{ann} and Π_{STJ}), together with the region inside which their center of mass is allowed to move during the refinement procedure (magenta spheres). c) the resulting planes (Π_{ann}^* and Π_{STJ}^*) after the optimization procedure. d) visualization of the refined planes in dark gray and the initialization planes obtained by model 2. e) the isolated AR with the contours of the second eigenvalue of the LBO (in black) and the plane identifying the sinuses of Valsalva (yellow).

3.2.6 Statistical analysis

Measurements based on expert manual annotations were used as reference values in assessing the proposed method's accuracy. Accuracy of areas, perimeters and diameters was evaluated using Bland-Altman analysis. Analyses were performed using python scipy 1.5.1 statistics library.

3.3 Results

3.3.1 Evaluation of segmentation performance

The generalization performance of the segmentation models was evaluated on the validation set ($n=31$), for which manual ground truth labels of the aorta, annulus and STJ were available (dataset A). For all subjects in the validation set, both models 1 and 2 proved able to accurately trace the region encompassing the AR, annulus and STJ. Segmentation performance was quantitatively assessed, in terms of Dice coefficient and MSD with respect to manual ground truth segmentations; for both indices, mean value and [min, max] range were computed. Model 1 and 2 outputs were compared against model 3 and 4, respectively, and on the atlas-based method on the same validation set. Results are reported in Table 3.1. For the AR, the U-Net architecture (model 1) gave

0.02 higher Dice scores and 0.08 mm lower MSDs than the U-Net transformer (model 3). For the aortic annulus and STJ, model 1 slightly outperformed model 3, yielding 0.03 higher Dice scores and 0.1 mm lower MSDs. In general, the atlas-based method gave the worst segmentation performance, reaching the lowest Dice score of 0.16 for the STJ.

	Aorta			Annulus			STJ		
	Model 1 (U-Net)	Model 3 (UNETR)	Atlas based	Model 2 (U-Net)	Model 4 (UNETR)	Atlas based	Model 2 (U-Net)	Model 4 (UNETR)	Atlas based
Dice	0.93 [0.80, 0.97]	0.91 [0.77, 0.97]	0.53 [0.0, 0.91]	0.57 [0.34, 0.87]	0.54 [0.18, 0.71]	0.25 [0.0, 0.78]	0.66 [0.29, 0.90]	0.54 [0.18, 0.73]	0.16 [0.0, 0.66]
MSD [mm]	0.39 [0.0, 4.7]	0.31 [0.0, 6.4]	3.6 [0.1, 18]	0.73 [0.0, 3.9]	0.63 [0.0, 4.3]	12 [0.2, 76]	0.99 [0.0, 1.5]	0.85 [0.0, 4.3]	14 [0.2, 91]

Table 3.1: Comparison between automatic segmentation approaches and ground truth manual segmentations for the AR, annulus and STJ. Mean value and [min, max] ranges of Dice score and mean surface distance (MSD) are reported.

The test cases corresponding to the best and worst performance of models 1 and 2 are shown in Figure 3.5.

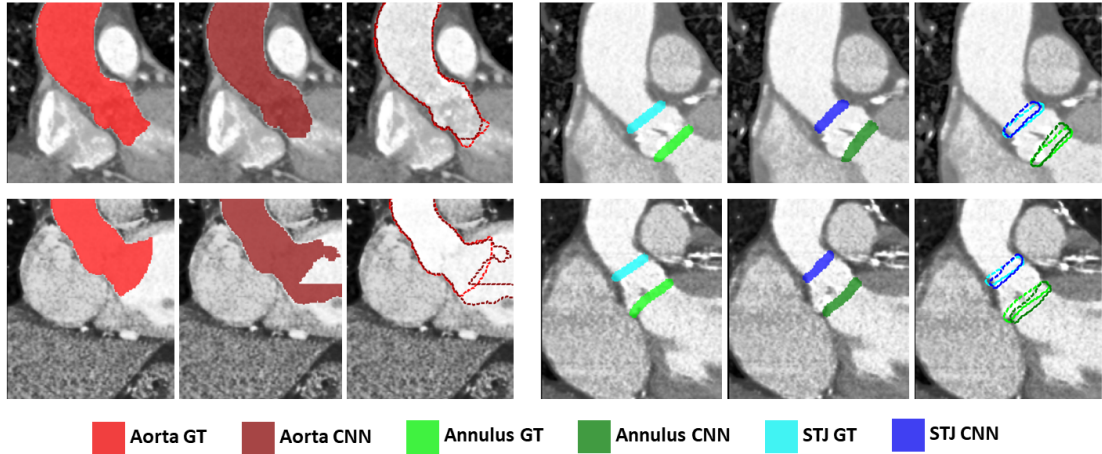


Figure 3.5: Comparison of manual (ground truth, GT) segmentation in bright red (aorta), bright green (annulus) and light blue (STJ) vs. automatic segmentation obtained by the trained neural networks (CNN) in dark red (aorta), dark green (annulus) and dark blue (STJ) on four test cases: the best (top row) and the worst (bottom row) for the two models. Coronal slices with color fill and with dashed contours are shown.

Inter-operator variability

For the 20 randomly selected subjects, the aortic segmentations obtained by the three independent operators were compared pairwise, in terms of Dice scores, MSDs and physical volume differences (ΔV). For a generic comparison OpX vs. OpY, ΔV is computed by subtracting the label mask created by OpY from the one of OpX. Comparison

	Op1 vs. Op2	Op1 vs. Op3	Op2 vs. Op3
Dice	0.96 [0.93, 0.99]	0.89 [0.78, 0.95]	0.91 [0.78, 0.96]
MSD	0.11 [-0.58, 0.67]	1.0 [-0.02, 5.3]	1.2 [-0.21, 6.1]
ΔV	1.6 [-10, 4.9]	7.7 [-4.1, 37]	7.7 [-7.7, 38]

Table 3.2: Inter-operator variability analysis results. Dice scores, mean surface distance (MSD) and physical volume differences (ΔV) are computed for each pairwise comparison. Median [minimum, maximum] values are reported.

results are reported in Table 3.2. In general, good Op1 and Op2 produced almost identical segmentations (mean Dice of 0.96), whereas Op1 and Op3 had the lowest agreement (0.89 mean Dice score).

3.3.2 Comparison vs. manual measurements

Bland-Altman plots (Figure 3.6) allowed to analyze the agreement between the automatic pipeline output and manual measurements. Herein, we report comparisons as biases and 95% limits of agreement, i.e., bias [lower limit, upper limit], where differences were computed as manual – algorithm. In general, automatic and manual anatomical measurements were in good agreement. A tendency of the automatic system to underestimate annulus diameters vs. manual measurements was observed, as evidenced by positive biases (solid horizontal lines in Figure 3.6). For D_{max} , measurement differences were 0.51 [-2.79, 3.81] mm. A similar trend was found for D_{min} , resulting in 0.89 [-2.8, 4.62] mm. Significantly smaller discrepancies between the two measurement techniques were obtained for the annulus area. Evaluating the area-derived diameter, bias and limits of agreement were 0.07 [-0.24, 0.38] mm. For the annulus angle, an average difference $<3^\circ$ was found between measurements, while limits of agreement were $[-17^\circ, 11^\circ]$. D_{mean} computed at the STJ and Sinuses by the automatic method slightly underestimated manual measurements, yielding differences of 0.05 [-1.98, 2.07] and 0.17 [-2.63, 2.97] mm, respectively. As compared to diameter measurements, perimeter measurements showed larger differences: -1.8 [-8.06, 11.74] mm for the annulus perimeter, and 1.09 [-6.18, 8.37] mm for the STJ perimeter.

Critical cases

For few cases, a large discrepancy was observed between the automatic and expert measurements. The largest differences in annulus diameters were observed for patient 057 and were equal to $\Delta D_{max} = -4.9mm$ and $\Delta D_{min} = -5.7mm$, where Δ is defined as (expert minus algorithm). However, a detailed analysis of case 057 showed that AR segmentation was precise; thus, discrepancies are not due to errors in automated segmentation (Figure 3.7, top row). Concerning STJ measurements, the largest discrepancies were obtained for patient 029 ($\Delta D_{mean} = 4.4mm$); these were likely due to the presence of calcifications around the aortic bulb (Figure 3.7, mid row), possibly introducing larger uncertainty in STJ measurements. The largest difference in D_{mean} for the




Figure 3.6: Bland–Altman plots of the proposed algorithm versus expert manual measurements. Mean differences are shown as continuous horizontal lines, while 95% limits of agreement are shown as dashed horizontal lines. Critical cases are highlighted with a red arrow and identification number.

sinuses of Valsalva was obtained for patient 140 ($\Delta D_{mean} = -4.3mm$). In this case, the high curvature of the aortic bulb surface could cause larger discrepancies between the two methods (Figure 3.7, bottom row).




Figure 3.7: Results of the implemented automatic pipeline for 3 cases for which the discrepancies between expert and automatic measurements were largest. The first column shows the segmentation inferred by model 1 and bounded to the AR region (in red). The second column shows the 3D anatomy reconstructions together with either the annulus (green), STJ (blue) or sinuses plane (yellow). The third column shows an interpolation of the image onto the plane identified by either the annulus, STJ or sinuses planes. Segmentation contours are show in green (annulus), blue (STJ) and yellow (sinuses), together with D_{max} (dashed arrows, black) and D_{min} (dashed arrows, white).

3.4 Discussion and comparison with existing approaches

In this work, we presented a fully automatic pipeline for AR morphological analysis and TAVI pre-procedural planning support from 3D CT. Our system provides 3D segmentations of the anatomical structures of interest without requiring any human supervision and extracts quantitative morphological parameters of the AR with good accuracy as compared to manual annotations performed by a clinical expert. To obtain a quantitative morphometric assessment of the AR, our approach relies on two CNNs trained to perform automatic 3D segmentation of the ascending aorta including the LVOT (model 1), and of the aortic annulus and STJ (model 2). The CNN trained to segment the aorta and LVOT (model 1) proved able to accurately delineate the anatomical region of interest in all 31 validation cases from dataset A and all 178 test cases in dataset B. For the validation cases, the average Dice score and MSDs were comparable to state-of-the-art approaches focusing on similar anatomical districts, with values of 0.93 [0.80, 0.97] and 0.39 [0.0, 4.7] mm, respectively. In a previous paper by our group dealing with automatic segmentation of the thoracic aorta from CT, a mean Dice score of 0.954 was reported [64]. Elattar et al. [77] adopted an AR segmentation approach based on normalized cuts and achieved Dice scores of 0.95 [0.85, 0.98] and MSDs of 0.74 ± 0.39 mm. In our study, model 1 was trained on ground truth segmentations encompassing different portions of the left ventricle due to CT scan field of view or inter-operator variability. This led to label map predictions with variable extents (as it is visible in Figure 5), entailing slightly worse accuracy with respect to other similar studies. This inconsistency was irrelevant to the downstream tasks of our pipeline and did not invalidate the overall effectiveness of our approach, which only requires the LVOT to be segmented correctly. In addition, when considering only the AR region mean values of Dice score and MSD of 0.96 and 0.14 mm were found. These errors are in the range of twice the average pixel spacing of 0.56 mm of our image dataset. In our experiments, models with 3D U-Net architectures (models 1 and 2) performed slightly better than transformer-based architectures (models 3 and 4), and significantly better than the implemented atlas-based approach. In some cases, the b-spline registration method underlying the atlas-based approach proved to be unsuitable for the task at hand, yielding segmentations with no overlap with the ground truth. The realm of neural networks for image segmentation is vast, and further studies are necessary to investigate whether different architectures such as UNet++ [78] or models based on atrous convolutions [79] may provide better performance. Concerning the identification of the aortic annulus and STJ, our approach (model 2) differs from previously published techniques since it derives anatomical landmarks from image segmentation and differential geometry. Elattar et al. [59] based their methodology on extracting the AR centerline, computing the Gaussian curvature and an harmonic decomposition of the aortic shape. Their automatic measurements were in excellent agreement with manual ones, with differences in annulus radii of 0.24 ± 0.70 mm and 0.37 ± 0.82 mm for two observers, respectively, and differences in annulus angles of $6.86 \pm 5.39^\circ$ and $6.34 \pm 4.00^\circ$. We obtained slightly larger differences for annulus mean radii (0.34 ± 0.78 mm), lower

differences for annulus area-derived radii (0.006 ± 0.16 mm), and similar discrepancies for annulus angles (-3.07 ± 7.15 °). However, [59] still represents a proof of concept, lacking generality, since automatic measurements are validated against a limited set of 40 patients. In our experience, Gaussian curvature-based criteria for landmark detection would perform suboptimally or fail for anatomies with less pronounced bulb curvature. Atlas-based methods could be an alternative, but could still be limited in their capability to generalize to very heterogeneous anatomies unless the reference atlas is heterogeneous enough [61]. The significantly larger number of test patients included in our study allowed us to devise robust solutions that proved able to deal with great anatomical variability among subjects. To date, with 178 patients, our study presents the most comprehensive validation of an automatic AR assessment system. Another type of landmark detection approach implies training a neural network to directly localize key anatomical points. Astudillo et al. [55] adopt this kind of strategy using a DenseVNet architecture to detect coronary ostia position and height with respect to the annulus, validating their accuracy against manual annotations on 100 patients. Compared to manual annotations, they report mean differences of 0.54 mm and -0.16 mm in left and right coronary heights, and of 1.4 mm in 3D Euclidean distances of coronary ostia. However, when utilizing a landmark-based method for anatomical measurement quantification, an error in identifying even one landmark can result in significant errors in more complex anatomical metrics. In contrast, in our approach we did not rely on direct neural network-based landmark detection, but rather on a series of geometric computations for robust identification of the aortic annulus and STJ. In particular, the annulus and STJ plane refinement step, enabled our system to cope with segmentation errors, requiring only a rough localization of the two anatomical regions. This makes the results of our pipeline more dependent on the accuracy of AR and LVOT segmentation, which are more clearly defined and visible from CT scans, and thus pose an easier challenge to deep learning segmentation systems. Given the large number of cases on which we tested our automatic measurement method, it is reasonable to expect some large discrepancies with respect to manual measurements. In the critical cases reported in Figure 3.7, the automatic segmentation algorithms were able to segment the anatomical structures with good precision. However, some anatomical features such as high curvature of the aortic bulb, and presence of abundant calcium deposits could cause differences in how diameters are computed by the two approaches. Being able to automatically segment calcium, our system potentially enables the end-user to choose whether to include or exclude calcifications in the computed measurements. Overall, our approach is significantly faster than commercially available semi-automated tools [60], requiring less than 45 seconds to run from CT to measurements on a GPU-accelerated workstation, therefore it shortens the processing time, making it compatible with the clinical routine, and applicable to large population studies. On average, the time required by the user for a full aortic assessment using commercially available software is around 30 minutes.

3.5 Final remarks

We have presented a pipeline for automating quantification of complex morphological biomarkers that are relevant for preprocedural planning of TAVI from CT images, providing an unprecedented validation of our approach against 178 patients, the largest to date. Our proposed method demonstrates a quick, accurate, and consistent assessment of AR anatomy from CT data, yielding 3D segmentations of the AR (mean Dice score 0.96). Incorporating this deep learning-based tool into the preoperative planning routine in TAVI environments could potentially lead to time and cost savings, as well as improved accuracy.

CHAPTER 4

Mixed Reality and Deep Learning for External Ventricular Drainage Placement: A Fast and Automatic Workflow for Emergency Treatments

Based on:

Palumbo, M. C., Saitta, S., Schiariti, M., Sbarra, M. C., Turconi, E., Raccuia, G., Fu, J., Dal-lolio, V., Ferrolì, P., Votta, E., De Momi, E. & Redaelli, A. Mixed Reality and Deep Learning for External Ventricular Drainage Placement: A Fast and Automatic Workflow for Emergency Treatments. In Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VII (pp. 147-156).

Summary

The treatment of hydrocephalus is based on anatomical landmarks to guide the insertion of an EVD. This procedure can benefit from the adoption of mixed reality (MR) technology. In this study, we assess the feasibility of a fully automatic MR and deep learning-based workflow to support emergency EVD placement, for which CT images are available and a fast and automatic workflow is needed. The proposed study provides a tool to automatically i) segment the skull, face skin, ventricles and FoM from CT scans; ii) import the segmented model in the MR application; iii) register holograms on the patient's head via a marker-less approach. An ad-hoc evaluation approach including 3D-printed anatomical structures was developed to quantitatively assess the accuracy and usability of the registration workflow.

4.1 Clinical problem and current solutions

EVD represents the optimal surgical treatment for emergency cases of hydrocephalus [80]. This procedure consists in draining cerebrospinal fluid by inserting a catheter into the brain ventricles, positioning the tip on the FoM. EVD placement is usually performed blindly, relying on paradigmatic anatomical landmarks localized freehand on the patient's head, thus carrying the risk of catheter malpositioning that can lead to hemorrhage, infections, brain tissue damage, multiple passes and reinsertion, entailing unnecessary brain trauma. Despite being a relatively simple procedure, a nearly 50% inaccuracy rate has been reported [80]. To address these issues, pre-operative images can provide helpful guidance during EVD placement and insertion [80, 81], for instance, when combined with neuronavigation systems [82]. However, the high cost and significant encumbrance of these systems makes them hard to be adopted in the operating room. MR-based solutions can represent a low-cost, easily portable guidance system for EVD placement. Previous studies have shown the feasibility of adopting MR based solutions for ventriculostomy procedures, providing visualization of internal brain structures [83], or to register the manually segmented holographic content on the patient's anatomy using marker based methods [84, 85, 86]. Nonetheless, these approaches do not represent adequate solutions for emergency cases where both automatic segmentation and quick registration need to coexist in a fast multistep routine. A first attempt to superimpose holograms on real world object using HoloLens2 (H2) head mounted display (HMD) (Microsoft, Redmond, WA) and its depth camera via a marker-less approach, has been reported [87], resulting in low accuracy. The goal of this study was to assess the feasibility and accuracy of a MR and deep learning-based workflow to support EVD placement, in emergency procedures for which preoperative CT images are only available and a fast and automatic workflow is critical. The proposed study provides a tool to automatically i) segment the skull, skin, ventricles and FoM from CT scans; ii) import the segmented model in the MR application; iii) register holograms on the patient's head via a marker-less approach.

4.2 Approach

The envisioned surgical workflow graphically depicted in Figure 4.1 is divided into consecutive steps resulting in a final MR application specifically tailored for H2 device. The following sections provide a detailed description on the development of the automatic segmentation (Section 4.2.1), automatic registration (Section 4.2.2) and registration accuracy assessment (Section 4.2.3) methods.

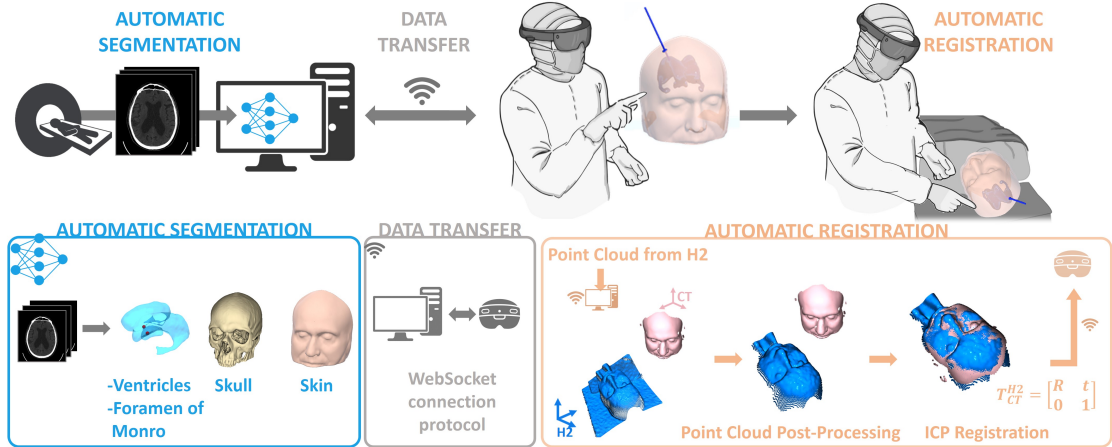


Figure 4.1: Schematic representation of the developed workflow: the 3D CT images are automatically segmented to obtain the brain structures that will be sent to H2 device through an internet connection protocol. The user is then able to visualize the holographic models and set a path between the FoM and an entry point on the skin layer. The PtC of the skin surface of the patient is acquired using H2 depth camera and this data is used to estimate the transformation matrix T_{CT}^{H2} .

4.2.1 Automatic Segmentation

3D CT scans of 200 subjects (including 41 cases of hydrocephalus) acquired between 2018 and 2022 were used as benchmark (pixel spacing ranging from 0.543×0.543 to 0.625×0.625 mm², slice thickness between 0.625 and 1.25 mm). The study was approved by the local ethical committee and informed consent was waived because of the retrospective nature of the study and the analysis of anonymized data.

Skin and Skull Segmentation was performed through i) automatic elimination of the mouth region in case of dental implant inducing artifacts; ii) thresholding (bone tissue: high-pass filter, threshold of 500 Hounsfield units (HU) [88]; skin: band-pass filter, lower and upper thresholds of -100 HU and 50 HU, respectively); iii) hole filling and smoothing.

Ventricle Segmentation was performed by a neural network (NN) based on a 3D UNET architecture, composed of an input layer of 16 filters, encoder and decoder branches of 5 resolution levels each, a bottleneck block and skip connections via concatenation [32]. The NN was trained and tested using all the CT volumes: the two lateral ventricles and the third ventricle were manually segmented by two experienced operators, guaranteeing that each acquisition was segmented by at least one user. The whole dataset was divided into training and test sets: 180 (90%) scans with their corresponding ground truth segmentations were randomly selected and used for NN training. A data augmentation routine was implemented to increase dataset diversity. Random rotations, mirroring and affine transformations were applied to the processed patches. A combination of Lovasz-Softmax loss and focal loss [64] was used to trained

the NN to perform simultaneous segmentation of the lateral and third ventricles. The remaining 20 scans (10%), including 5 cases of hydrocephalus, were used for testing.

Detection of FoM, which is the target for the tip of the catheter during EVD, was performed by a recursive thresholding approach [64].

4.2.2 Automatic Registration

The registration process was conceived to allow the end user to reliably superimpose the MR representation of the segmented skin model (\mathcal{S}_{CT}) on the head of the real patient, whose face is acquired as a 3D PtC using the H2 depth sensing camera. This step was conceived to be performed before any surgical draping, once the position of the head is already fixed for the entire duration of the procedure. The holographic interface that permits the surgeon to acquire the 3D PtC and then run the registration algorithm, was designed in Unity3D (Unity Technologies, version 2020.3.1LTS) using the MR toolkit library MRTK v2.7.0. The algorithm consisted of the following steps (Figure 4.2):

- Acquire the PtC of the facing surface $\{p_{i,H2}\}$, using H2 research mode [89], and the position of the H2 device itself ($p_{camera,H2}$), both referred to the H2 global coordinate reference system.
- Initialize the position of \mathcal{S}_{CT} with respect to the target $\{p_{i,H2}\}$ based on $p_{camera,H2}$.
- Apply an Hidden Point Removal algorithm [90] to \mathcal{S}_{CT} to filter out points in the back of the head to avoid unnecessary computations.
- Extract $\{p_{i,H2}\}$ points belonging to the face through a DBSCAN density based clustering algorithm [91].
- Apply a Fast Global Registration algorithm [92] between the simplified \mathcal{S}_{CT} and the cleaned $\{p_{i,H2}\}$, obtaining a first alignment, which is then refined via Local Refined Registration method based on a Point-to-Plane iterative closest point (ICP) algorithm [93] to obtain T_{CT}^{H2} . All the described steps related to the PtC tweaking and registration were implemented using the open3D library [94].
- Send back T_{CT}^{H2} to H2 and apply it to the holographic model of \mathcal{S}_{CT} , thus allowing the surgeon to visualize the segmented model aligned with the patient face $\mathcal{S}_{H2} = T_{CT}^{H2}\mathcal{S}_{CT}$.

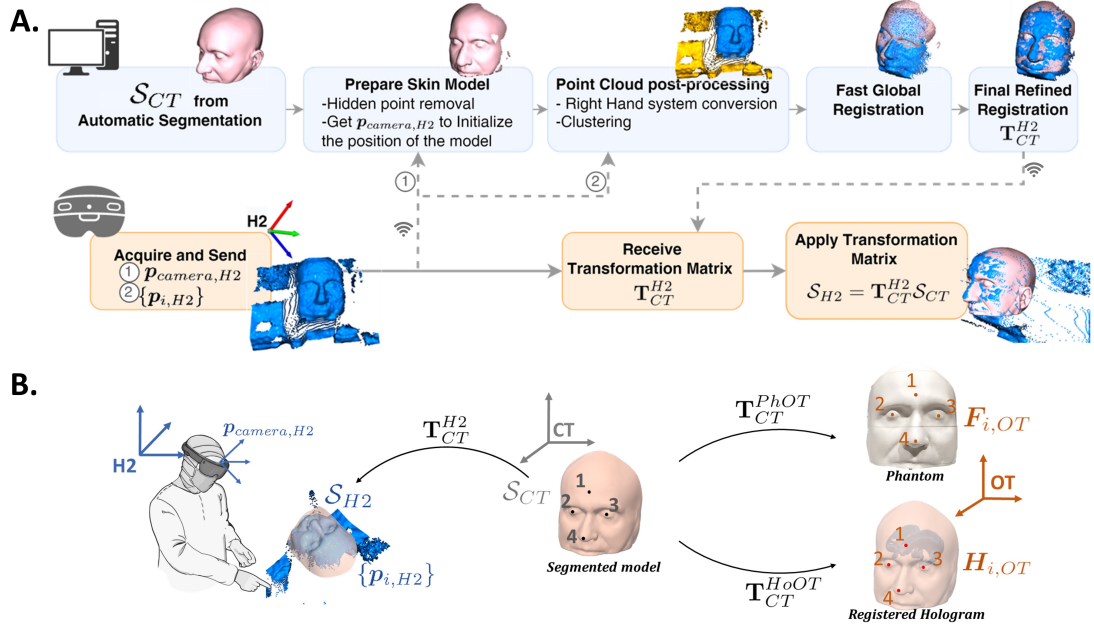


Figure 4.2: A. Detailed representation of the registration workflow for the MR environment with B. the considered coordinate systems and transformation matrices.

4.2.3 Registration Accuracy Assessment

An experimental setup to evaluate the automatic registration accuracy was developed (Figure 4.3A). The evaluation was divided into two different procedures to evaluate: 1. the accuracy of the hologram-to-phantom registration; 2. the targeting accuracy in a ventriculostomy-like procedure. For both tests, a 3D model of a patient with hydrocephalus was reconstructed from automated CT scan segmentation. The extracted skin model was modified to add four markers consisting in conical cavities on the skin surface at the left eye, right eye, nose and forehead, respectively. The modified model was 3D-printed and used to simulate a real patient's head. An optical tracker (OT) system (NDI Polaris Vicra) equipped with a trackable probe was used to acquire the position of the conical markers on the phantom, designed to accommodate the tip in a unique manner. The position of the markers was also reported in the holographic model where four red spheres of 1mm of diameter were added to highlight each marker position. To evaluate the accuracy of the hologram-to-phantom registration, an H2 expert user was asked to repeat the following procedure 10 times:

- Acquire the 4 markers position on the printed phantom in the OT space ($F_{i, OT}$, with $i = 1, ..4$) (Figure 4.3B).
- Wearing the H2 device execute the automatic registration and then, moving away the phantom, place the tip of the probe in correspondence of the 4 holographic

marker to obtain their positions in the OT space ($\mathbf{H}_{i,OT}$, with $i = 1, \dots, 4$). For this phase the probe is attached to a tripod which stabilizes the probe movement not limiting its orientation in the space (Figure 4.3C).

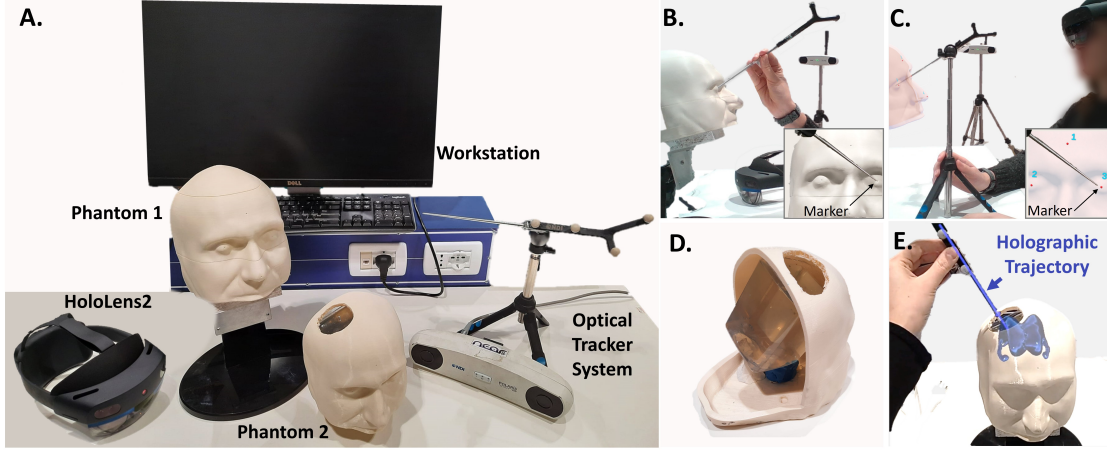


Figure 4.3: A. Experimental setup showing the hardware and physical components needed to test the registration accuracy; user executing the first test is evaluating the position of the 4 marker on the printed phantom B., and on the registered holographic model C.; D. phantom used by neurosurgeons for the second test while E. trying to insert the OT probe to match the holographic trajectory.

The distance between the phantom and holographic model was assessed both by computing the Euclidean distance $e_{ij} = \|\mathbf{F}_{ij,OT} - \mathbf{H}_{ij,OT}\|$ of the i -th marker and j -th coordinate (x, y, z) , but also providing an estimate of the error for each point of the skin surface and ventricle. To do this, a first transformation matrix \mathbf{T}_{CT}^{PhOT} was computed to map \mathcal{S}_{CT} with the 4 markers of the printed phantom $\mathbf{F}_{i,OT}$, and a second one \mathbf{T}_{CT}^{HoOT} to map \mathcal{S}_{CT} with the 4 markers of the registered hologram $\mathbf{H}_{i,OT}$ (Figure 4.2B). A least-square solution based on the singular value decomposition algorithm [95] was applied to compute the transformation matrices. Finally, the point-to-point distance d_j with $j = 1, \dots, N$ vertices, was computed as:

$$d_j = \|\mathbf{T}_{CT}^{PhOT} \mathbf{m}_{j,CT} - \mathbf{T}_{CT}^{HoOT} \mathbf{m}_{j,CT}\| \quad (4.1)$$

with $\mathbf{m}_{j,CT}$ being the vertices of the segmented model of the skin and ventricles in the CT space. The root mean square error (RMSE) was computed as:

$$RMSE = \sqrt{\frac{1}{N} \sum_{j=1}^N d_j^2} \quad (4.2)$$

To evaluate targeting accuracy, another phantom was 3D-printed from the same segmented model, provided with a large hole in proximity of a common entry point

for a ventriculostomy procedure and with an internal box full of gelatin mimicking brain tissue (6% gelatin in water) (Figure 4.3D). A membrane was placed on top of the hole entrance to reduce the bias experienced by the users when choosing the point of entrance. For this test, 7 experienced neurosurgeons were asked to:

- acquire the positions of the 4 conical markers on the phantom with the OT probe $F_{i,OT}$.
- do a blind (i.e., without wearing the H2) ventriculostomy procedure using the OT probe as a catheter, trying to reach the FoM based on anatomical references. Once deemed arrived to the target, the tip position of the probe was acquired in the OT space ($p_{blind,OT}$)
- wear the H2 in which the segmented model is visualized to perform a registration procedure on the phantom annotating T_{CT}^{H2} .
- guided by the visualization of an holographic path that goes from the skin to the FoM, insert the probe and try matching the holographic trajectory (Figure 4.3E). Again, once the FoM was deemed reached, the position of the probe in the OT space was acquired ($p_{Hguided,OT}$).

The transformation matrix T_{CT}^{PhOT} was computed to refer the position of the inserted probe to the internal structures of the brain, thus evaluating the tip position in the blind and hologram guided procedure. To quantify the improvement brought by adopting the proposed methodology, 7 neurosurgeons independently simulated EVD procedures on the gelatin phantom. Of note, none of the participant was familiar with the mixed reality device. The targeting accuracy improvement was computed as: $\Delta\% = \left(\frac{|p_{blind,OT} - p_{Hguided,OT}|}{|p_{blind,OT} - p_{t,OT}|} \times 100 \right) \%$, where $p_{t,OT}$ is the target position in OT space.

4.3 Result assessment

4.3.1 Automatic segmentation

For each 3D image, all the processing steps of automatic segmentation, from reading raw data to post-processing 3D surfaces, required less than 2 minutes, on average.

Skin and skull 20 subjects were randomly selected and for both the face skin and skull, the distance from every point in the automatically obtained surface mesh to its closest point in the ground truth manually reconstructed mesh was computed. For the skin, mean and maximum distances of 1.75 mm and 17.6 mm were found, whereas for the skull, mean and maximum distances of 0.44 mm and 3.50 mm were found, respectively.

Brain ventricles test set results The trained NN was applied to run inference on the test set made of 20 CT scans, 5 of which presented hydrocephalus. Our automatic segmentation algorithm was able to generalize to new cases, extracting the 3D ventricle regions with accuracy (Figure 4.4), requiring approximately 130 seconds per scan. For the lateral ventricles, mean Dice score and 95% Hausdorff distance were equal to of 0.88

and 14.4 mm, respectively. For the third ventricle, mean Dice score and 95% Hausdorff distance were equal to 0.62 and 19.4 mm. Considering the five test cases with hydrocephalus, a mean Dice score of 0.94 was obtained for the lateral ventricles and a mean value of 0.81 for the third ventricles, while 95% Hausdorff distances were equal to 18.3 mm and 7.48 mm, respectively.

For all test cases with hydrocephalus, the ventricle surface reconstruction accuracy was sufficient to automatically detect the FoM.




Figure 4.4: Comparison of manual (ground truth, GT) segmentation in bright red (lateral ventricles, LV) and bright green (third ventricle, TV), vs. automatic segmentation obtained by the trained neural network (NN) in dark red (lateral ventricles, LV) and dark green (third ventricle, TV) on a test case.

4.3.2 Registration Accuracy Assessment

Hologram-to-Phantom registration The global time required by the algorithm to compute the transformation matrix \mathbf{T}_{CT}^{H2} and send it back to H2 was 0.26 ± 0.02 s, with a fitness of $92.93 \pm 3.91\%$ which measures the overlapping area between the cleaned PtC and the processed skin model, averaged over the 10 measurements. The median distance with interquartile range between $F_{i,OT}$ and $H_{i,OT}$ for the 4 markers, in the 10 repeated procedures, resulted in $1.25(0.65 - 1.72)$ mm, $1.75(0.98 - 2.42)$ mm, $1.35(0.41 - 1.83)$ mm, for the x, y, z coordinates in the OT space respectively (corresponding to the axial, sagittal and coronal plane of the phantom). The distances referred to each of the 4 markers in the 3 coordinates are reported in Figure 4.5A. The average 3D Euclidean distance for all the markers over the 10 procedures resulted in 2.72 ± 0.67 mm. The point-to-point accuracy over the whole model as obtained from Eq. 4.2 showed a $RMSE$ of 3.19 ± 0.69 mm and 3.61 ± 1.13 mm for the face and ventricle surfaces respectively. Point-to-point distance maps between the surfaces are represented in Figure 4.5C.

Ventriculostomy procedure To quantify the improvement brought by adopting the proposed methodology, 7 neurosurgeons independently simulated EVD procedures on the gelatin phantom, aiming at the FoM blindly, i.e. relying on anatomical landmarks, and subsequently with the holographic guidance displaying the FoM target and the trajectory to follow. To note, none of the participant was familiar with the MR device. The targeting accuracy improvement was computed as: $\Delta\% = \frac{|d_{blind} - d_{holo}|}{d_{blind}} \cdot 100$, where d_{blind} and d_{holo} are the target distances of the probe tip without and with holographic support, respectively. Guided by the holographic system, a 42% higher accuracy was

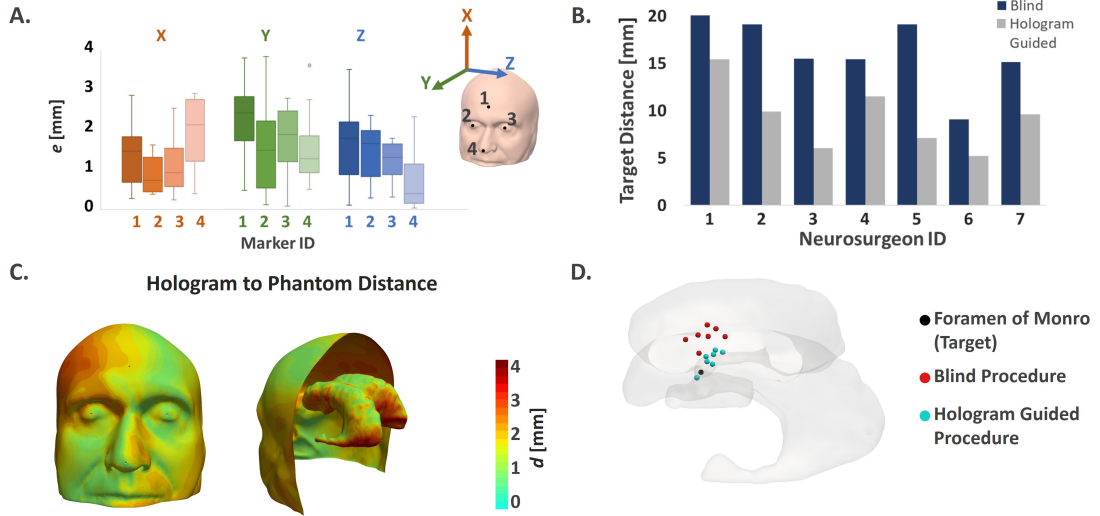


Figure 4.5: Experimental results from evaluation tests. A: hologram-to-phantom accuracy evaluated on the 4 markers; B: quantitative targeting accuracy from the tests done by 7 neurosurgeons; C: hologram to phantom accuracy evaluated over the whole model surface; D: visualization of catheter tip positions and target from the EVD simulations.

shown, on average. Euclidean target distances of blind vs. holographically-guided simulations are reported in Figure 4.5B. In the blind procedure, the mean distance to the target was equal to 16.2 ± 3.48 mm, compared to a mean distance of 9.26 ± 3.27 mm for the hologram guided procedure, resulting in a statistically significant improvement on a paired t -test ($p=0.001$). Figure 4.5D graphically shows the final tip position in all the procedures done by the neurosurgeons compared to the FoM target in the ventricle.

4.4 Conclusions

The present study proposed a novel solution for supporting neurosurgeons during EVD procedures. Our system leverages a deep NN for automatic image segmentation and a markerless registration method based on a HMD depth sensor to align internal brain structures with the physical position of the patient. We quantitatively assessed the accuracy of our workflow at each of its steps, proving the added value to ventriculostomy procedures under experimental settings; adopting our technology, neurosurgeons targeting accuracy improved by 42%, on average. Of note, all the described steps require little or no manual intervention, and computations required globally less than 3 minutes, thus making our system suitable for emergency procedures where timing can be crucial. Our future efforts will focus on including an holographic navigation of the probe using electromagnetic sensors and evaluate the accuracy of adding a further, independent, Azure Kinect (Microsoft, Redmond, WA) depth camera to achieve a more stable and robust point cloud acquisition. Moreover, a phantom reproducing the dif-

ferent resistance of the ventricles with respect to the brain will be adopted to simulate the surgical procedure in a more realistic way.

Part II

Unsupervised approaches for hemodynamic assessment

CHAPTER 5

Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta

Based on:

Saitta, S., Maga, L., Armour, C., Votta, E., O'Regan, D. P., Salmasi, M. Y., Athanasiou, T., Weinsaft, J.W., Xu, Y.X., Pirola, S. & Redaelli, A. (2022). Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta. Computer Methods and Programs in Biomedicine

Chapter Summary

Background and Objective: Numerical simulations of blood flow are a valuable tool to investigate the pathophysiology of ascending thoracic aortic aneurysms (ATAA). To accurately reproduce *in vivo* hemodynamics, computational fluid dynamics (CFD) models must employ realistic inflow boundary conditions (BCs). However, the limited availability of *in vivo* velocity measurements, still makes researchers resort to idealized BCs. The aim of this study was to generate and thoroughly characterize a large dataset of synthetic 4D aortic velocity profiles sampled on a 2D cross-section along the ascending aorta with features similar to clinical cohorts of patients with ATAA.

Methods: Time-resolved 3D phase contrast magnetic resonance (4D flow MRI) scans of 30 subjects with ATAA were processed through in-house code to extract anatomically consistent cross-sectional planes along the ascending aorta, ensuring spatial alignment among all planes and interpolating all velocity fields to a reference configuration. Velocity profiles of the clinical cohort were extensively characterized by computing flow morphology descriptors of both spatial and temporal features. By exploiting PCA, a statistical shape model (SSM) of 4D aortic velocity profiles was built and a dataset of 437 synthetic cases with realistic properties was generated.

Results: Comparison between clinical and synthetic datasets showed that the synthetic data presented similar characteristics as the clinical population in terms of key morphological parameters. The average velocity profile qualitatively resembled a parabolic-shaped profile, but was quantitatively characterized by more complex flow patterns which an idealized profile would not replicate. Statistically significant correlations were found between PCA principal modes of variation and flow descriptors.

Conclusions: We built a data-driven generative model of 4D aortic inlet velocity profiles, suitable to be used in computational studies of blood flow. The proposed software system also allows to map any of the generated velocity profiles to the inlet plane of any virtual subject given its coordinate set.

5.1 Significance and state of the art

Thoracic aortic aneurysms (TAA) is a life-threatening condition involving an abnormal dilatation of the aortic wall [96]. An accurate assessment of blood flow plays an essential role in clinical diagnosis, risk stratification and treatment planning of TAA [97, 98, 99]. CFD is a well established tool to quantify hemodynamics [100, 101] through *in silico* trials [102, 103]. To achieve a high level of fidelity, CFD models need to account for patient-specific BCs. When choosing inflow BCs, prescribing patient-specific data in the form of 3-directional velocity profiles allows to obtain significantly more accurate results compared to using idealized profiles, as amply shown by several recent studies that make use of velocity information extracted from phase-contrast magnetic resonance imaging (PC-MRI) [104, 105, 106, 107]. Nonetheless, the limited availability of *in vivo* velocity measurements, still makes researchers resort to idealized BCs. Moreover, such lack of clinical data represents an obstacle for setting up population-based *in silico* trials and for building datasets for training ML models. Generative models can be used to overcome this limitation by creating larger data-driven synthetic datasets [108]. In particular, SSMs have been adopted in the cardiovascular field [109, 110]. SSMs are data-driven approaches for assessing shape variability and creating large virtual cohorts from clinical ones. An SSM is typically based on PCA and describes the shape probability distribution of the input data by a mean shape and modes of shape variations [111]. Several studies have effectively applied SSMs to study TAA geometry [112, 113, 114]. Nonetheless, aortic hemodynamics, which have been shown to play a key role in pathophysiology of this disease [101, 115, 116], have not received the same attention. An exception is the work of Catalano *et al.*, who exploited SSMs to build an atlas of aortic hemodynamics in subjects with bicuspid aortic valve (BAV) and tricuspid aortic valve (TAV) [97]. However, the authors imposed an idealized parabolic velocity profile as inlet BC for their CFD models. Despite revealing important insights on BAV vs. TAV biomarkers, the study is hampered by the use of such simplified inlet BCs, which significantly affect the computed aortic blood flow, especially in regions that are close to the inlet, namely the ascending aorta [106, 107, 117].

Motivated by the need for boosting the impact and the fidelity of numerical studies involving ATAA hemodynamics, the present work leverages a SSM to pursue three specific aims. We provide: i) a quantitative and detailed characterization of a representative 4D ATAA inlet velocity profile as a valid alternative to idealized inlet BCs for numerical simulations; ii) a synthetic virtual cohort of 4D ATAA inlet velocity profiles with features that are consistent with those of real ATAA inlet profiles and potentially large enough to allow for ML approaches to be used; iii) insights into both spatial and temporal hemodynamic features of ATAA velocity fields in the ascending aorta.

5.2 Statistical shape model of aortic velocity profile

5.2.1 Image data

Thoracic 4D flow MRI scans of 30 subjects with ATAA acquired between 2017 and 2019 were retrospectively retrieved. Our dataset included fully deidentified images provided by Weill Cornell Medicine, (NY, USA) and Hammersmith Hospital (London, United Kingdom). None of the subjects in our cohort were BAV-affected. Respiratory compensated 4D flow acquisitions were performed with the following settings: spatial resolution (voxel size) 1.4 – 2.0 mm (range), field of view = 360 mm, flip angle = 15°, VENC = 150 – 200 cm/s (range), time resolution 20 – 28 frames/cardiac cycle (range). Data usage was approved by the Weill Cornell Medicine Institutional Review Board (New York, NY, USA) and by the Health Research Authority (HRA) (17/NI/0160) in the UK and was sponsored by the Imperial College London Joint Research and Compliance Office, as defined under the sponsorship requirements of the Research Governance Framework (2005). The participants provided their written informed consent to participate in this study.

5.2.2 Data preprocessing

4D flow MRI data were preprocessed using in-house Python code and following the workflow presented in Figure 5.1. No correction for eddy currents was applied, based on the high VENC values set in the acquisitions and on the dimensions and position of the ascending aorta within the scanner [118]. Visual assessment of the measured velocity data was performed to verify that no aliasing was present. For each patient, a 3D binary mask of the aorta was extracted from phase-contrast magnetic resonance angiography (PC-MRA) images using semi-automatic tools available in the open source software ITK-SNAP [30]. Aortic centerlines were extracted using the VMTK [37]. To guarantee consistency of inlet plane location among all ATAA subjects, inlet planes were defined with respect to a commonly used anatomical landmark represented by the bifurcation of the pulmonary artery (PA) [64]. Briefly, the PA bifurcation point (P_B) was identified from the PC-MRA images (in green in Figure 5.2). Then, the point (P_{BAo}) on the aortic centerline with the same axial (foot-head) coordinate of P_B was identified (pink sphere in Figure 5.2). An inlet point (P_{in}) was extracted 30 mm upstream of P_{BAo} along the centerline curvilinear abscissa; the inlet plane was defined as the plane passing through P_{in} and locally perpendicular to the centerline. Being aware of the inter-variability among patient-specific geometries, this methodology was preferred rather than extracting the inlet plane at the STJ. In fact, the limited space-resolution of 4D flow MRI does not allow to locate the STJ as clearly as in, e.g., CT images. A triangulated mesh of the selected plane within the aortic lumen was generated; 4D flow velocity data were then probed at inlet plane nodal locations. For the generic subject indexed by j , we defined the inlet plane nodal coordinates as $\tilde{\Xi}^{(j)} = [\tilde{\xi}_1^{(j)}, \dots, \tilde{\xi}_\tau^{(j)}, \dots, \tilde{\xi}_{\mathcal{T}^{(j)}}^{(j)}]^\top$, and the corresponding measured velocity vector field as $\tilde{\mathcal{V}}^{(j)} = [\tilde{v}_1^{(j)}, \dots, \tilde{v}_\tau^{(j)}, \dots, \tilde{v}_{\mathcal{T}^{(j)}}^{(j)}]^\top$, with

$\tilde{\xi}_\tau^{(j)}, \tilde{v}_\tau^{(j)} \in \mathbb{R}^{\mathcal{N}^{(j)} \times 3}$ and where $\mathcal{T}^{(j)}$ and $\mathcal{N}^{(j)}$ are the number of frames in the cardiac cycle of subject j and the number of probed nodal locations on the inlet plane, respectively; therefore, in general, the dimensions of $\tilde{\Xi}^{(j)}$ and $\tilde{\mathcal{V}}^{(j)}$ vary among subjects.

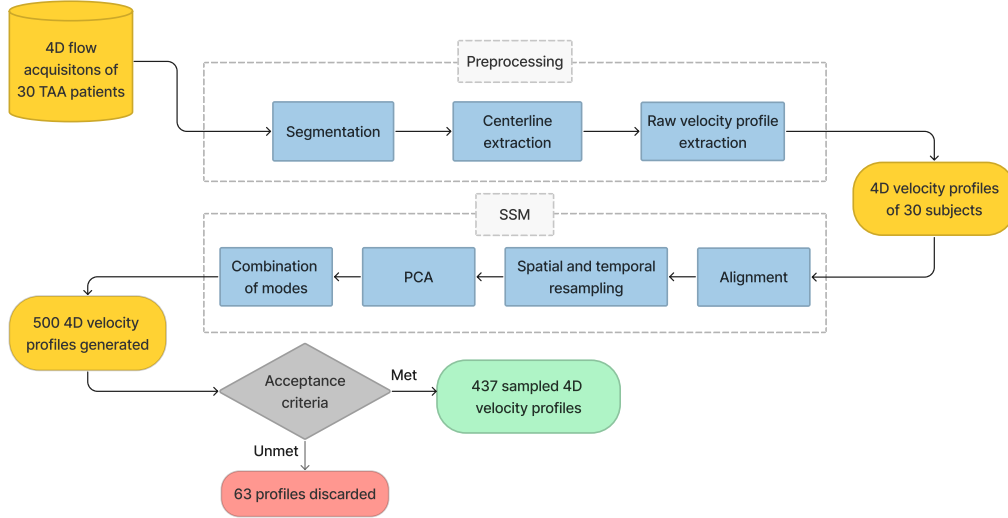


Figure 5.1: Schematic representation of the adopted workflow. All 4D flow acquisitions go through a preprocessing pipeline for the extraction of velocity profiles. The SSM process consists in a common alignment and spatiotemporal resampling of the profiles and then a combination of PCA modes to generate new ones. Only the profiles that meet specific acceptance criteria are included in the final dataset.




Figure 5.2: Method for ascending aortic plane selection. The pulmonary artery bifurcation point (P_B , green) is selected from the PC-MRA images. The corresponding point on the aortic centerline (P_{BAo}) is identified (pink, bottom right), and the selected plane (red, bottom right) is chosen at a 30 mm arc length upstream along the centerline and locally normal to the centerline.

5.2.3 Statistical shape modeling of inlet velocity profiles

Alignment Consistent spatial orientation among the extracted inlet velocity profiles was ensured through two steps: first, each inlet plane was centered at the origin by applying the translation $\mathbf{T}^{(j)}$ to transform the nodal coordinates $\tilde{\xi}_\tau^{(j)}$ into $\hat{x}_\tau^{(j)} = \tilde{\xi}_\tau^{(j)} + \mathbf{T}^{(j)}$. Second, two consecutive rigid rotations were applied to the translated coordinates $\hat{x}_\tau^{(j)}$ and to the corresponding velocities $\tilde{v}_\tau^{(j)}$. The first rotation ($\mathbf{R}_1 \in \mathbb{R}^{3 \times 3}$) transformed $\hat{x}_\tau^{(j)}$ and the corresponding velocity profile $\tilde{v}_\tau^{(j)}$ to $\hat{x}_\tau^{(j)} = \mathbf{R}_1^{(j)} \hat{x}_\tau^{(j)}$ and $\hat{v}_\tau^{(j)} = \mathbf{R}_1^{(j)} \tilde{v}_\tau^{(j)}$, respectively, so to make the inlet plane containing $\hat{x}_\tau^{(j)}$ normal to the z-axis. The second rigid rotation ($\mathbf{R}_2 \in \mathbb{R}^{3 \times 3}$) transformed $\hat{x}_\tau^{(j)}$ and $\hat{v}_\tau^{(j)}$ to $\mathbf{x}_\tau^{(j)} = \mathbf{R}_2^{(j)} \hat{x}_\tau^{(j)}$ and $\mathbf{v}_\tau^{(j)} = \mathbf{R}_2^{(j)} \hat{v}_\tau^{(j)}$, and it ensured that the x-axis was aligned with the right-to-left direction of the subject.

Resampling After alignment, each velocity profile $\mathbf{v}_\tau^{(j)}$ was mapped onto a reference disk with unit radius using linear radial basis functions (RBFs), effectively enabling the resampling of each velocity profile at $N = 1071$ fixed spatial locations uniformly distributed over the reference disk (Figure 5.3a and b). For points of the reference disk that did not fall within the original plane (as it would happen for non-circular cross-sections) zero velocity vectors were set, limiting alterations from the original flow rate. Each velocity profile time sequence was temporally interpolated to a reference temporal interval $t \in [0, 1]$ discretized in $T = 20$ frames, using cubic polynomials. Finally, for the generic subject j , the spatiotemporally aligned and resampled velocity profiles are defined as: $\mathbf{V}^{(j)} = [\mathbf{v}_1^{(j)}, \dots, \mathbf{v}_t^{(j)}, \dots, \mathbf{v}_T^{(j)}]^\top$, with $\mathbf{v}_t^{(j)} \in \mathbb{R}^{N \times 3}$.

Principal component analysis The 30 aligned 4D velocity profiles were rearranged into column vectors and assembled into a matrix $\mathbf{V} = [\mathbf{V}^{(1)}, \dots, \mathbf{V}^{(j)}, \dots, \mathbf{V}^{(30)}]$, with $\mathbf{V} \in \mathbb{R}^{P \times J}$ where J is the number of subjects (30) and $P = 3 \times N \times T$. Matrix \mathbf{V} was used as

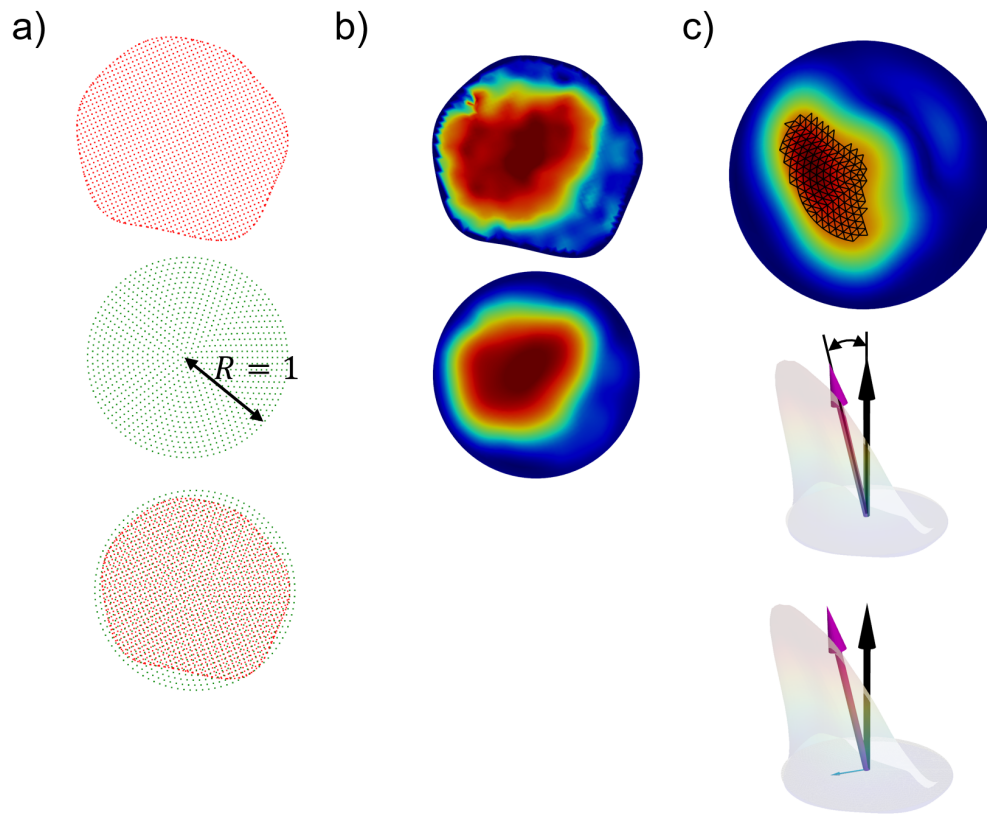


Figure 5.3: a) Representation of an inlet plane coordinate set from a representative subject in the clinical cohort (red), as obtained upon rigid roto-translation, and of the corresponding point set in the fixed disk domain (green). b) A velocity profile from the clinical cohort (top), resampled to the fixed reference disk using RBFs (bottom). c) Exemplification of the computed flow descriptors. FDI (top): $area^{top15\%}$ is highlighted in black; FJA (middle): angle between plane normal (black) and mean velocity direction (magenta); SFD (bottom): ratio between the normal component (black) and in-plane component (cyan) of the mean velocity vector (magenta).

input for a PCA. Standard PCA begins by computing the mean velocity profile defined as:

$$\bar{\mathbf{V}} = \frac{1}{J} \sum_{j=1}^J \mathbf{v}^{(j)}, \quad (5.1)$$

and by assembling the covariance matrix \mathbf{C} , given by:

$$\mathbf{C} = \frac{1}{J} \sum_{j=1}^J (\mathbf{v}^{(j)} - \bar{\mathbf{V}})(\mathbf{v}^{(j)} - \bar{\mathbf{V}})^T. \quad (5.2)$$

The eigenvalues and eigenvectors of \mathbf{C} were sorted in descending order to obtain a sequence of eigenvectors that progressively maximize the explained variance. The first 18 eigenvectors were considered since their cumulative variance was $\geq 90\%$. Each mode represents a shape direction of variation from the mean velocity profile that is representative of data variability.

Shape sampling Starting from the mean shape $\bar{\mathbf{V}}$, the SSM can be built, and a synthetic 4D profile \mathbf{U} can be generated by adding to $\bar{\mathbf{V}}$ a shape variation, i.e., a linear combination of the selected modes [112] as:

$$\mathbf{U} = \bar{\mathbf{V}} + \sum_{m=1}^M b^{(m)} \sqrt{\lambda^{(m)}} \mathbf{a}^{(m)} \quad (5.3)$$

where M denotes the number of selected modes, $\mathbf{a}^{(m)}$ is the eigenvector of \mathbf{C} associated to the m -th selected mode, and $\lambda^{(m)}$ is the corresponding eigenvalue. New profile shapes can be sampled from the SSM by using a set of coefficients, or weights, $\mathbf{b} = [b^{(1)}, \dots, b^{(m)}, \dots, b^{(M)}]$. In particular, to study the shape variations captured by a specific mode, one can sample shapes by considering only the selected mode and varying the coefficient $b^{(m)}$.

A uniform sampling technique for the first 18 modes of variation of the PCA was used to generate the virtual dataset of inlet velocity profiles:

$$b^{(m)} \sim \mathcal{U}(-1.5, 1.5). \quad (5.4)$$

5.2.4 Flow morphology descriptors

To quantitatively characterize inlet velocity profiles belonging to the clinical cohort and to the synthetic one, several descriptors of flow morphology were computed, both at systolic peak and as time-averaged quantities: positive peak velocity (PPV), flow dispersion index (FDI), flow jet angle (FJA), secondary flow degree (SFD), and retrograde flow index (RFI). Representations of FDI, FJA and SFD is included in Figure 5.3c. For a generic profile $\mathbf{v}_t^{(j)}$ consisting of nodal velocity vectors $\mathbf{v}_{t,n}^{(j)}$, with $n = 1, \dots, N$, flow descriptors were computed as follows.

Flow dispersion (FDI) was calculated to determine whether flow displayed a broad or peaked in-plane distribution. FDI was computed as the ratio between the area of the region characterized by the top 15% of peak velocity magnitudes ($area^{top15\%}$) and the inlet area [106]:

$$FDI = \frac{area^{top15\%}}{area} \times 100\%, \quad (5.5)$$

Accordingly, the higher the FDI value the more homogeneous the velocity profile; the lower the FDI value the sharper the velocity profile.

Flow jet angle (FJA) represents the angle formed by the mean velocity direction (jet direction v_{mean}) and the unit vector orthogonal to the inlet surface \mathbf{n} :

$$FJA = \arccos(v_{mean} \cdot \mathbf{n}), \quad (5.6)$$

FJA quantifies the skewness of the inlet flow towards aortic walls. An FJA value of 0° represents a mean jet direction perpendicular to the inlet plane.

Secondary flow degree (SFD) is computed as the ratio between the mean in-plane (radial) velocity magnitude v_{\parallel} and the mean axial velocity magnitude (through-plane velocity) v_{\perp} as:

$$SFD = \frac{\|v_{\parallel}\|}{\|v_{\perp}\|}, \quad (5.7)$$

Retrograde flow index (RFI) was calculated as the fraction of negative area under the curve of the flow rate time-course over the whole area under the curve [119]:

$$RFI = \frac{\left| \int_0^T Q_r dt \right|}{\left| \int_0^T Q_a dt \right| + \left| \int_0^T Q_r dt \right|} \times 100\%, \quad (5.8)$$

where Q_r and Q_a are the total retrograde and antegrade flow rate, respectively. A higher RFI value implies an increasing flow direction inversion during the cardiac cycle.

5.2.5 Acceptance criteria

To avoid generation of unrealistic velocity profiles and restrict the synthetic population only to plausible cases acceptance criteria were introduced. Such criteria were based on the flow features of the clinical profiles. Specifically, intervals of acceptance I_d were defined based on the statistical distributions of the flow descriptors defined in section 5.2.4 and computed as:

$$I_d = \left[\mu_d - 2\sqrt{\lambda_d} \right] \cup \left[\mu_d + 2\sqrt{\lambda_d} \right] \quad (5.9)$$

where μ_d denotes the mean value of the considered flow descriptor d :

$$\mu_d = \frac{1}{J} \sum_j \frac{1}{T} \sum_t d_t^{(j)}, \quad (5.10)$$

and $\sqrt{\lambda_d}$ its standard deviation. Flow descriptors were extracted from the synthetic velocity profiles; those synthetic profiles characterized by at least one parameter falling outside the acceptance intervals were automatically rejected.

5.2.6 Statistical analysis

To assess statistical differences between the clinical and synthetic sets, comparisons were made using unpaired t-tests for normally distributed variables and Mann–Whitney U tests for non-normally distributed data. Data normality was determined using the Shapiro–Wilk test.

Pearson’s correlation coefficients were calculated to assess correlations between shape modes and flow descriptors. p -values <0.05 were considered statistically significant.

5.3 Results

5.3.1 Statistical shape model analysis

Characterization of the clinical cohort For the original cohort, the mean 4D profile $\bar{\mathbf{V}}$ was obtained and characterized in terms of flow descriptors. Time-averaged descriptors for the mean profile of the clinical cohort were: $PPV = 0.42 \text{ m/s}$, $FDI = 13.34\%$, $FJA = 0.42^\circ$ and $SFD = 7.57$. Values at peak systole (PS) were: $PPV = 1.47 \text{ m/s}$, $FDI = 12.8\%$, $FJA = 13^\circ$, $SFD = 0.23$; while $RFI = 2\%$. $\bar{\mathbf{V}}$ (orientated as in Figure 5.4b) can be visualized at three key time points through the cardiac cycle: early systole (ES), PS and late systole (LS) (Figure 5.4c); colormaps visualizations were generated using the *warp by vector* filter available in the open-source Python library *pyvista*.

Generation of the synthetic inlet velocity profiles The sampling of the SSM through the process described in section 5.2.3 led to generate 500 synthetic 4D profiles. Out of these, the acceptance criteria led to accepting 437 velocity profiles that constituted the final synthetic cohort.

Comparison of synthetic vs. clinical velocity profiles Synthetic and clinical inlet velocity profiles were compared based on time-averaged flow descriptors shown in Table 5.1 and Figure 5.5. No statistically significant differences were found between the two cohorts, except for PPV ($p = 0.040$). Nonetheless, PPV mean values only differed by 0.03 m/s and similar standard deviations were observed (0.12 and 0.08 for clinical and synthetic cohorts, respectively).

Descriptor	Clinical cohort	Synthetic cohort	<i>p</i> value
PPV [m/s]	0.59 ± 0.12	0.56 ± 0.08	0.040 [†]
FDI [%]	10.36[5.78, 18.53]	9.72[6.30, 16.11]	0.187
FJA [°]	66.77[27.94, 106.29]	68.74[30.09, 101.60]	0.929
SFD [-]	3.84[0.62, 20.45]	2.85[0.65, 15.72]	0.190
RFI [%]	4.11[0.00, 32.10]	5.86[0.00, 27.74]	0.728

Table 5.1: Comparison of time-averaged flow descriptors between real and synthetic cohorts. Normally distributed variables are expressed as mean ± standard deviation; non-normally distributed variables are expressed as median [min, max]. [†] indicates statistical significance $p \leq 0.05$.

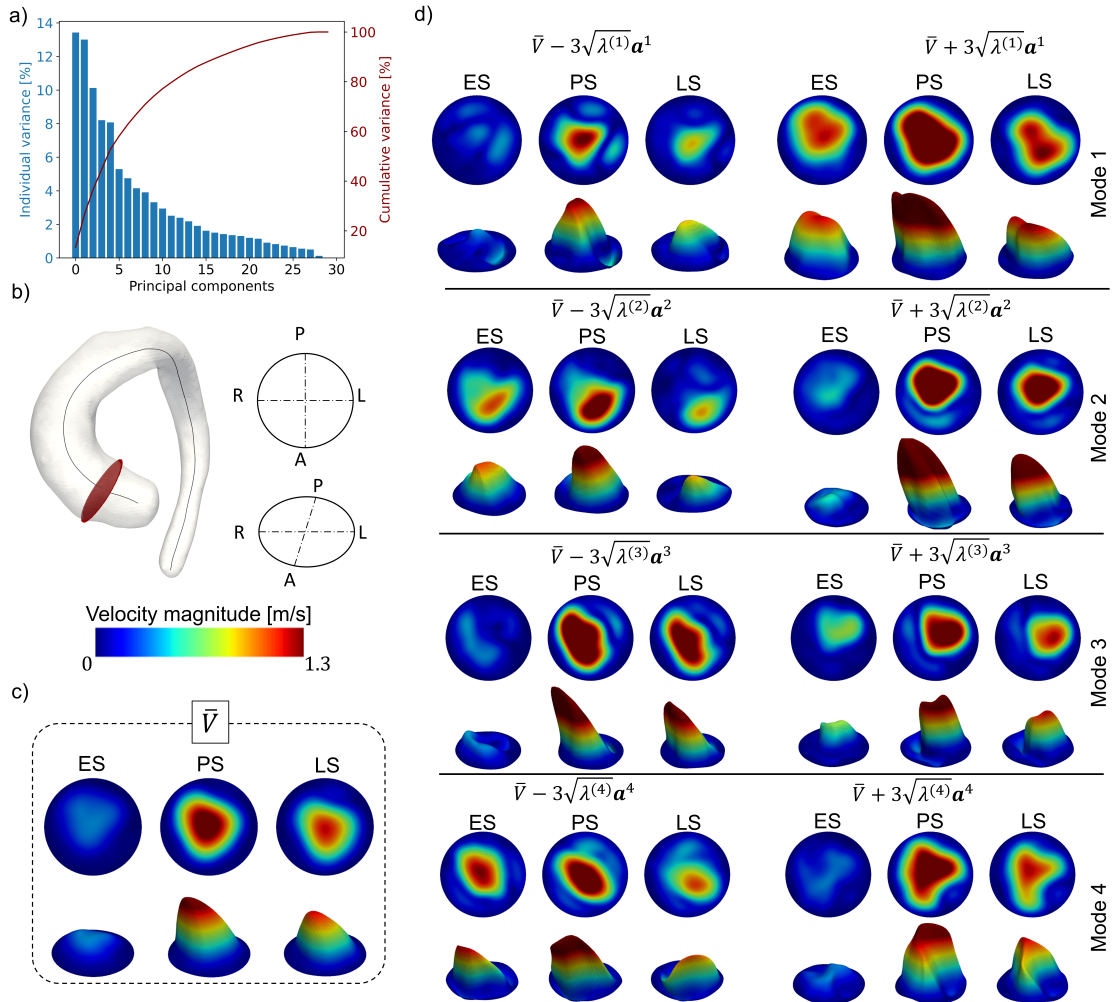


Figure 5.4: a. Individual and cumulative variance associated with each principal component. b. Orientation of the displayed planes and profiles in with respect to the subject in 2D and 3D. c. Mean velocity profile (\bar{V}) colored by velocity magnitude in 2D (top row) and 3D (bottom row) at early systole (ES), peak systole (PS) and late systole (LS). d. 2D and 3D visualizations of velocity profiles deformed towards minimum and maximum for the first 4 modes and colored by velocity magnitude; color scale shown in panel b.




Figure 5.5: Box plots showing distributions of time-averaged flow descriptors for real and synthetic cohorts. Similar medians and ranges were observed. No significant differences ($p \geq 0.05$) were found except for PPV ($p = 0.040$). Whiskers indicate 1.5 interquartile ranges; diamonds highlight outliers.

5.3.2 Associations of velocity profile modes with flow morphology descriptors

The devised SSM was exploited to analyze the modes of variation and assess their correlation with flow morphological features. By containing the majority of statistical information, the first four modes were responsible for approximately 45% of the total dataset variability (Figure 5.4a). For the first four modes, the obtained extreme shape variations are visualized in Figure 5.4d. Although it is not straightforward to associate each mode of variation with a specific flow feature, it can be hypothesized that the first PCA modes are related to some physically meaningful characteristic of the profile. In practice, the potential physical meaning of each mode was assessed by evaluating equation 5.3 taking into account only one mode at a time and choosing a set of 10 evenly spaced coefficients $b^{(m)} \in [-3, 3]$ as described in [97]. The flow morphology descriptors introduced in section 5.2.4, were computed for each generated profile as $b^{(m)}$ was gradually increased. For simplicity, correlations with spatial and temporal flow features were analyzed separately. Correlations with spatial features were computed at PS.

Modes correlation with spatial features at peak systole Correlation results at PS are reported in Table 5.2. Mode 1 seemed related to the spatial heterogeneity of velocity magnitude and FDI, with larger high velocity regions observed when $b^{(1)} = 3$. This was confirmed by the statistically significant positive correlations of mode 1 with PPV ($r = 0.94, p < 0.0001$) and FDI ($r = 0.99, p < 0.0001$). Overall, with increasing $b^{(1)}$, PPV, FDI, SFD, and FJA increased, indicating a tendency of the profile to be less aligned with the plane normal. Similar trends were obtained for velocity profiles generated by varying $b^{(2)}$, which appeared related to overall flow rate together with the size of the region with intermediate velocity as it can be visualized by the less pronounced jet for $b^{(2)} = -3$ in Figure 5.4. With increasing $b^{(2)}$, PPV, SFD, and FJA significantly increased ($r = 0.86, p = 0.001$; $r > 0.99, p < 0.0001$; $r > 0.99, p < 0.0001$, respectively). No statistically significant correlation was found between mode 2 and FDI ($r = -0.14, p = 0.702$). Profiles corresponding to $b^{(3)} = -3$ showed a more spatially concentrated velocity jet than profiles sampled at $b^{(3)} = 3$. Quantitatively, this was reflected in increasing SFD ($r = 0.88, p = 0.001$) and FJA ($r = 0.88, p = 0.001$) and in decreasing PPV ($r = -0.67, p = 0.032$) as $b^{(3)}$ increased. No correlation was found between mode 3. Finally, profiles generated by varying $b^{(4)}$ showed differences in the shape of the high velocity region and in jet orientation. Specifically, with increasing $b^{(4)}$, FJA ($r = -0.91, p < 0.0001$) and SFD ($r = -0.90, p < 0.0001$) decreased while FDI increased ($r = 0.96, p < 0.0001$). No correlation was found between mode 4 and PPV ($r = -0.27, p = 0.458$).

Modes correlation with temporal features For each of the first four modes, the flow rates computed on the 10 evenly spaced coefficients $b^{(m)} \in [-3, 3]$ were analyzed (figure 5.6). Profiles corresponding to mode 1 showed the largest variation in peak flow rate, ranging from 0.62 for $b^{(1)} = -3$ to 2.37 for $b^{(1)} = 3$, and the largest temporal shift

	$b^{(m)}$										r	p
	-3.00	-2.33	-1.67	-1.00	-0.33	0.33	1.00	1.67	2.33	3.00		
Mode 1												
PPV [m/s]	1.34	1.35	1.37	1.39	1.42	1.52	1.62	1.73	1.87	2.03	0.94	< 0.0001
FDI [%]	7.42	8.09	8.86	9.83	11.78	14.01	15.92	17.53	18.70	19.59	0.99	< 0.0001
FJA [°]	9.48	11.54	13.08	14.25	12.61	13.44	14.09	14.61	15.02	15.36	0.88	0.001
SFD [-]	0.16	0.20	0.23	0.25	0.22	0.23	0.25	0.26	0.26	0.27	0.88	0.001
Mode 2												
PPV [m/s]	1.53	1.40	1.31	1.28	1.38	1.56	1.76	1.96	2.16	2.37	0.86	0.001
FDI [%]	8.93	9.86	11.63	14.66	13.76	12.02	11.01	10.30	9.78	9.54	-0.14	0.702
FJA [°]	7.56	8.82	10.06	11.27	12.47	13.64	14.78	15.91	17.01	18.09	>0.99	< 0.0001
SFD [-]	0.13	0.15	0.17	0.20	0.22	0.24	0.26	0.28	0.30	0.32	>0.99	< 0.0001
Mode 3												
PPV [m/s]	2.03	1.83	1.65	1.53	1.49	1.46	1.43	1.43	1.51	1.64	-0.67	0.032
FDI [%]	12.40	14.16	15.33	15.25	13.38	12.40	12.75	13.45	12.94	11.29	-0.52	0.122
FJA [°]	12.86	12.79	12.78	12.82	12.95	13.19	13.56	14.09	14.82	15.78	0.88	0.001
SFD [-]	0.22	0.22	0.22	0.22	0.23	0.23	0.24	0.25	0.26	0.28	0.88	0.001
Mode 4												
PPV [m/s]	1.66	1.58	1.51	1.49	1.47	1.47	1.50	1.54	1.60		-0.27	0.458
FDI [%]	11.60	11.50	11.43	12.21	12.31	13.14	14.12	15.09	16.08	16.52	0.96	< 0.0001
FJA [°]	41.68	35.70	28.65	21.83	15.74	10.66	7.13	6.26	7.90	10.40	-0.91	< 0.0001
SFD [-]	0.89	0.71	0.54	0.40	0.28	0.18	0.12	0.10	0.13	0.18	-0.90	< 0.0001

Table 5.2: PCA modes correlations with velocity profile flow descriptors at peak systole. Results of Pearson correlation analyses are reported. $p < 0.05$ indicates statistical significance.

($\Delta t = 0.096$). Furthermore, a negative linear correlation was found between $b^{(1)}$ and RFI ($r = -0.87$, $p = 0.001$), with larger portions of retrograde flow corresponding to lower values of $b^{(1)}$. A similar trend was observed for mode 4, whose associated flow rate peaks increased with increasing $b^{(4)}$, but within a narrower range ([0.90, 1.76]) and with a smaller temporal shift ($\Delta t = 0.04$). On the other hand, an opposite trend was found for mode 3, with higher flow rates for decreasing $b^{(3)}$, and with negligible temporal shifts. Flow rates associated to profiles generated by varying $b^{(2)}$ all showed similar flow rate curves over time.

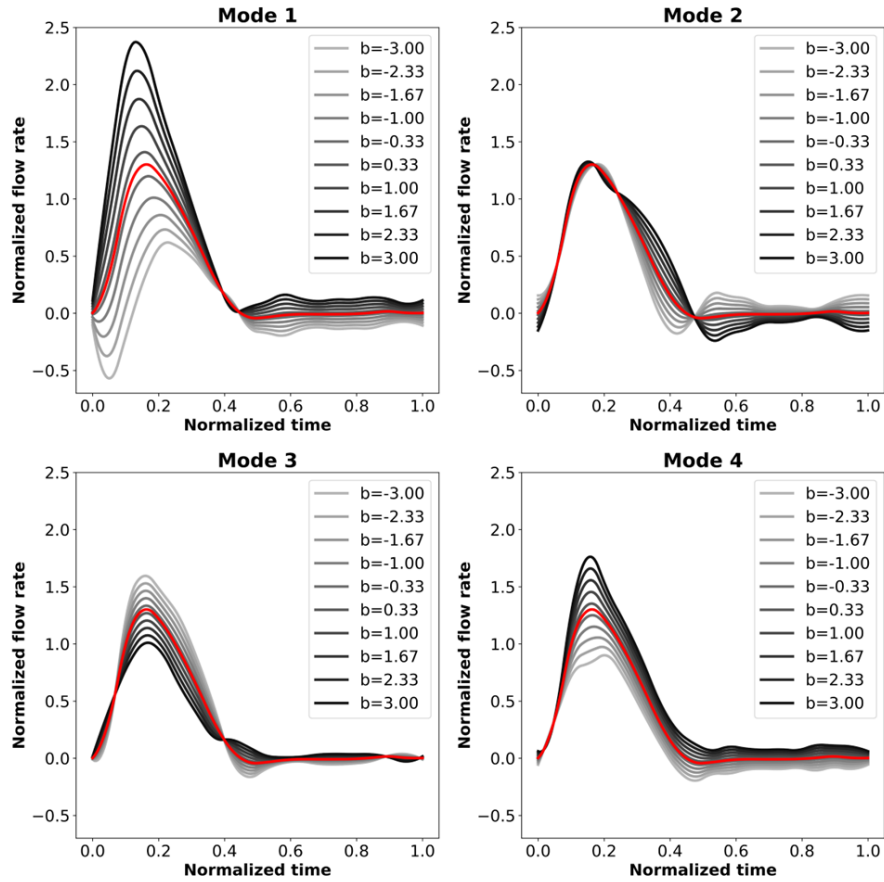


Figure 5.6: Plots of normalized flow rates over time computed for the first four modes by varying their corresponding coefficient $b^{(m)}$. Mean flow rate over time corresponding to \bar{V} is shown in red for all subplots.

5.4 Results impact and comparison with previous work

The lack of patient-specific information on *in vivo* flow features has led to the widespread use of simplified inlet velocity profile BCs when setting up CFD models. However, this assumption can significantly impact results, particularly when modeling the ascending aorta [104, 105, 106]. The present work addressed the issue of scarcity of patient-specific hemodynamic data by proposing a valid alternative to idealized inlet velocity profiles. In particular, we focused on cases of ATAA. The first key achievement of the present work was exploiting SSMs to generate a dataset of 437 synthetic velocity profiles starting from a clinical cohort of 30 ATAA subjects. The proposed methodology allowed to create aortic inlet velocity fields that presented similar spatiotemporal characteristics as compared to real ATAA patients, and that are suitable to be prescribed as inlet BC for CFD simulations. We found an average velocity profile \bar{V} that qualitatively resembled

a 3D paraboloid-like shape (Figure 5.4c). This finding suggests that imposing an idealized parabolic profile as inflow BC of an ATAA CFD model is the best choice in the absence of patient-specific flow data. Similar findings were reported by [106], who also suggested that a parabolic profile is a reasonable choice when patient-specific data are missing and in case of TAV. Nonetheless, our $\bar{\mathbf{V}}$ was characterized by $FJA = 13^\circ$ and $SFD = 0.23$ at PS, whereas a perfectly symmetric and centered paraboloid would have null FJA and SFD values. In particular, the velocity jet was directed toward the right side of the subject, i.e., the extrados of the ascending aorta. As recognized by [106] and confirmed by additional studies ([105, 107]), including in-plane velocity components significantly affects flow dynamics predictions in the ascending aorta. Therefore, for an accurate assessment of ATAA hemodynamics, the mean profile proposed in this study represents a better baseline choice for inlet BC specification in CFD simulations.

Our SSM was also sampled to generate new realistic 4D profiles; the individual variance associated to the first mode was equal to $\approx 13\%$, whereas the first four modes accounted for $\approx 45\%$ of cumulative variance (Figure 5.4a). Thamsen *et al.* [113] built an SSM of aortic coarctation anatomy and found that the first mode covered 43.7% of the total variation, while the first three modes were responsible for its 70%. The authors also built an SSM of inlet vector fields based on 4D flow MRI data, but did not provide full insights into profile variability and only considered PS, limiting their analysis to steady conditions. Our results indicate substantially lower individual mode explained variance. Such discrepancy can be attributed mainly to the fact that we dealt with 3D velocity vectors that change over time, which represent considerably more complex, higher dimensional features as compared to 3D positions describing static geometries.

The second key achievement of this study was providing insights into the hemodynamic features that are responsible for the most significant variations of ascending aorta spatiotemporal velocity profiles. Although it was not possible to find a one to one correspondence between shape modes and flow descriptors, we were able to demonstrate that the first four modes can be related to a unique combinations of spatial flow morphology descriptors (Table 5.2 and Figure 5.4). A more straightforward interpretation of PCA modes has been reported in studies dealing with aortic anatomy [110, 114]. In our results, an overlapping of feature contributions to PCA modes is likely due to the complexity and high dimensionality of our 4D velocity vector fields that makes our data significantly more heterogeneous. Nonetheless, it is interesting to note that mode 1 was clearly linked to a temporal feature: flow rate over time (Figure 5.6). Spatially, this resulted in a positive correlation between $b^{(1)}$ and systolic PPV ($r = 0.94$, $p < 0.0001$) and FDI ($r = 0.99$, $p < 0.0001$), which both contribute to increase net flow rate. Therefore, when exploiting the proposed SSMs, users that wish to investigate the impact of specific flow features will be able to adjust mode weights to generate profiles with the desired characteristics.

Limitations It is important to address the limitations of the present study. First, a clinical cohort size of 30 subjects limits the statistical power of our conclusions. This downside is related to the retrospective nature of the study and to the choice of con-

sidering only subjects with TAV and ATAA. Future studies could focus on analyzing aortic velocity profiles in BAV patients, which are known to display considerably different flow features [106]. Second, our choice of plane positioning may result in slightly different anatomical locations among subjects. Having an additional imaging modality, such as CT or MRA, aligned with 4D flow images, one could select a plane at the level of the STJ. In general, the low quality of 4D flow images does not allow to detect the STJ with certainty, whereas the chosen anatomical landmark (bifurcation of the pulmonary artery) was always clearly visible in our flow-encoded images. Finally, PCA is fundamentally a linear transformation. Non-linear approaches based on deep learning could potentially increase the effectiveness of data-driven hemodynamic features extraction [120].

5.5 Conclusions

In this work we built the first data-driven generative model of time-dependent 3D aortic inlet velocity profiles, suitable to be used in numerical simulations of blood flow. With the aim of expediting the development of future *in silico* analyses, the proposed software system also allows to map any of the generated velocity profiles to the inlet plane of any virtual subject given its coordinate set. The present work thus sets a new standard for the computational bioengineering community, allowing to replace the common practice of prescribing idealized inflow BC in numerical simulations of blood flow with more realistic conditions.

Data availability and usability

We provide the generated synthetic cohort of 4D velocity profiles ready to be used for time-dependent CFD simulations with a Github repository also containing the necessary scripts to replicate our study. The code can be found at: <https://github.com/saitta-s/flow4D> and all synthetic velocity profiles can be downloaded from <https://doi.org/10.5281/zenodo.7251987>.

CHAPTER 6

Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI

Based on:

Saitta, S., Carioni, M., Mukherjee, S., Schönlieb, C.B. & Redaelli, A. (2022). Implicit neural representations for unsupervised super-resolution and denoising of 4D flow MRI. arXiv:2302.12835

Summary

Context. 4D flow magnetic resonance imaging (4D flow MRI) is the only non-invasive imaging method that can provide time-resolved measurements of blood flow velocities. However, velocity fields detected by this technique have significant limitations that prevent accurate quantification of blood flow markers. These limitations are mainly related to low spatio-temporal resolution and measurement noise. Several lines of research have been pursued to overcome the main limitations. Among these, coordinate-based neural networks have shown to be able to represent complex signals as continuous functions, making them suitable for super-resolution tasks. In this work we investigate sinusoidal representation networks (SIRENS) for time-varying 3-directional velocity fields measured in the aorta by 4D flow MRI, achieving denoising and super-resolution.

Method. We trained our method on 4D voxel coordinates and enforce the no-slip condition at the vessel wall. First, we benchmarked our approach using synthetic measurements generated from a CFD simulation for which ground truth velocity fields are available. Three different levels of noise were simulated: *mild*, *medium* and *extreme*. Then, we test our method on a real 4D flow MRI scan of a patient with aneurysm of the ascending aorta. We assessed both velocity and WSS fields obtained with our method.

Results. The performance of different SIREN architectures was evaluated on synthetic measurements. The best configuration was chosen as the one that minimized the sum of velocity normalized root mean square error (vNRMSE), magnitude normalized root mean square error (mNRSME) and direction error (DE). A SIREN with 20 layers and 300 neurons per layer gave the lowest error for all levels of noise. The results of the SIREN with the chosen configuration were compared against state-of-the-art methods for 4D flow denoising and super-resolution. Our approach outperformed all existing techniques, giving up to 50% lower vNRMSE, 42% lower mNRSME and 15% lower DE on velocity and WSS fields. Additionally, the developed approach produced denoised and super-resolved velocity fields from clinical data, while maintaining low discrepancies in macroscopic flow measurements.

Conclusions. We showed the feasibility of SIRENS to representing complex and multi-dimensional blood flow velocity fields obtained from 4D flow MRI. Our approach is both quick to execute and straightforward to implement for novel cases. By meticulously optimizing our SIREN architecture, we leverage the spectral bias to generate a functional representation of our data with minimized noise, surpassing current solutions. Our method produces continuous velocity fields that can be queried at any spatio-temporal location, effectively achieving 4D super-resolution.

6.1 Background and significance

Accurate hemodynamic assessment is essential for having a deeper understanding of cardiovascular pathophysiology [121, 122]. For certain cardiovascular conditions such as aortic coarctation, valvular dysfunction and vascular aneurysm, disease diagnosis and management are based on *in vivo* hemodynamic biomarkers [123, 124] that can be obtained by either catheter insertion or non-invasive blood flow imaging methodologies. To date, 4-dimensional flow encoded magnetic resonance imaging (4D flow MRI or 4D flow) is the only existing non-invasive imaging technique that provides true time-resolved 3-dimensional (3D) and 3-directional blood flow velocity measurements [125]. 4D flow is based on the PC-MRI principle, which makes use of bipolar magnetic gradients to calculate the phase shift of moving protons. PC-MRI encodes tissue velocity $\mathbf{v}(\mathbf{x}, t) \in \mathbb{R}^3$ at spatial location \mathbf{x} during cardiac phase $t (1 \leq t \leq N_t)$ according to:

$$\rho_i(\mathbf{x}, t) = \rho_0(\mathbf{x}, t) \exp \left(j\pi \frac{(\Phi \mathbf{v}(\mathbf{x}, t))_i}{VENC} \right), \quad (6.1)$$

where $VENC$ is a manually set parameter determining the maximum velocity that can be recorded, $i = 0, \dots, 3$ are the encoded velocity components, and adopting a four-point velocity encoding, Φ is defined as:

$$\Phi = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (6.2)$$

Hence, the measured tissue velocity component i is proportional to the phase shift of the reconstructed images ρ_i [126].

Considering $\rho_{it} \in \mathbb{R}^{N_r \times N_c \times N_s}$ a discretized complex PC image on a Cartesian $N_r \times N_c \times N_s$ grid corresponding to cardiac phase t and velocity component i , the reconstructed image $\mathbf{H}(\rho_{it}) \in \mathbb{R}^{N_r \times N_c \times N_s}$ can be modeled as:

$$\mathbf{H}(\rho_{it}) = \mathbf{F}^{-1}(\mathbf{M}(\mathbf{F}(\rho_{it}) + \epsilon)), \quad (6.3)$$

where \mathbf{F} is the Fourier transform, $\mathbf{M} \in \{0, 1\}^{N_r \times N_c \times N_s}$ defines the undersampling mask in k -space, and $\epsilon \in \mathbb{C}^{N_r \times N_c \times N_s}$ is the additive complex noise. Herein, we have neglected coil sensitivity maps for simplicity, a more rigorous description of the PC-MRI measurement operator can be found in [126].

Velocity fields measured by 4D flow MRI can be processed after reconstruction to quantify more complex and clinically relevant hemodynamic biomarkers such as WSS [127, 128, 129], relative pressure [130, 131] and vortex structure [132]. Nonetheless, the large amount of raw data collected during a 4D flow MRI scan limits this imaging technique, entailing longer image reconstruction time and more difficult image analysis with respect to standard magnetic resonance angiography (MRA). Modern compressed sensing has enabled substantial shortening of 4D flow MRI acquisition time [133, 134] by

undersampling k -space data and exploiting priors about data regularity during reconstruction [135, 136]. Although modern techniques could achieve 4D flow MRI reconstructions in under 5 minutes [137], the resulting flow images still suffer from important limitations that make them inadequate for accurate quantification of blood flow markers. These limitations mainly concern spatio-temporal resolution, velocity encoding (VENC) related signal to noise ratio (SNR) and k -space noise [138]. For cardiothoracic acquisitions, image spacing is usually isotropic and in the range of 1.5-3 mm³, whilst typical temporal resolution is 30-50 ms [139]. Furthermore, 4D flow MR images are corrupted by noise, which is commonly assumed to be zero-mean Gaussian in frequency space [140]. Overcoming 4D flow limitations by reducing noise levels and enhancing its spatio-temporal resolution, would lead to more accurate hemodynamic assessment, boosting its widespread adoption and increasing its clinical usefulness. In this work, we present a novel application of implicit neural representations (INRs) [141] to achieve super-resolution (SR) and denoising of 4D flow MRI velocity fields. Our approach provides a continuous reconstruction of blood flow in both space and time, showing superior performance with respect to state-of-the-art methods for 4D flow denoising.

6.1.1 Related Work

Over the past two decades, several lines of research have been pursued to overcome the main limitations of 4D flow MRI velocity measurements and provide more accurate non-invasive hemodynamic assessment. Herein, we categorize these efforts broadly into two groups: model-based and data-based approaches. Model-based methods are based on CFD informed using MR images, whereas data-based methods directly operate on image data and apply regularized interpolation to enhance flow-encoded images [142, 143, 144, 145, 146, 147, 148].

Model-based approaches

As an alternative to 4D flow MRI, CFD has been widely applied to study cardiovascular flows [100, 115, 149, 150]. In contrast to 4D flow, CFD simulations can provide noise-free blood flow velocity fields at arbitrarily high spatio-temporal resolutions. By solving the Navier-Stokes (N-S) equations, numerical hemodynamic solutions inherently satisfy mass and linear momentum conservation laws. Nevertheless, when modeling subject specific cardiovascular flows, the accuracy of CFD simulations heavily relies on the choice of boundary conditions [107, 117] and blood constitutive model [151, 152]. When comparing CFD solutions to 4D flow velocity measurements, these assumptions inevitably lead to substantial discrepancies between hemodynamic markers computed via the two approaches [153, 154, 155]. Model-based approaches typically formulate their approach as an inverse N-S problem, in which one or more unknown parameters in the governing partial differential equations (PDEs) are optimized to minimize an objective functional representing the discrepancy between the N-S solution and the measured data. Model-based methods effectively address the limitations of noise and low spatio-temporal resolution of flow measurements, yielding velocity fields defined

on fine computational grids and at arbitrarily small time steps (<10 ms). In the seminal works of Funke et al. [142] and Koltukluoğlu et al. [143], this approach came down to solving a N-S boundary value problem formulated in a variational data assimilation framework, in which one or more boundary conditions are optimized in either 3D [143] or 4D [142]. Recently, a more comprehensive approach was proposed by Kontogiannis et al. [144] to jointly recover both boundary conditions, domain boundary and kinematic viscosity by assimilating noisy flow image data into the N-S solution. The same well-designed settings were later extended by the same authors to cope with undersampled synthetic phase contrast data [145], resulting in the first N-S informed compressed sensing reconstruction method. However, their method was only applied to 2D cases with steady flow conditions. In fact, despite the considerable progress made, all the mentioned techniques are computationally intensive, and their cost can enormously increase in higher dimensions, undermining their feasibility for real medical case-scenarios. For this reason, there is still lack of model-based approaches that can efficiently assimilate time-dependent 3D velocity fields from real medical flow images for concrete human data applications.

Data-based approaches

In contrast to model-based approaches, data-based methods do not rely on solving the governing PDEs. Instead, they employ interpolation techniques to approximate flow data, incorporating some form of regularization to impart desirable characteristics to the resulting velocity fields. Within this category, we discern between conventional denoising approaches and recent neural network-based methods.

Conventional denoising approaches adopt signal processing tools to enhance the acquired flow MRI data after reconstruction. To incorporate prior physical knowledge in their processing stage, several works have exploited the assumption of blood incompressibility. This physical condition is usually enforced through a divergence-free constraint on the reconstructed velocity field [146, 147]. Busch et al. [147], achieved effective denoising by projecting noisy flow measurements onto a 3D space of divergence-free RBFs. Moreover, their approach allows for incorporation of boundary conditions in the reconstructed flow field. Nonetheless, strict enforcement of the divergence-free condition may result in over-regularized velocity fields, especially in flow regions near edges of flow (boundaries). An approach to *softly* enforce the divergence-free condition was proposed by Ong et al. [146]. By constructing divergence-free wavelets (DFWs), the authors were able to decompose measured velocity fields into divergence-free and non-divergence-free wavelet coefficients, promoting the former and penalizing the latter according to suitably tuned thresholds. The described process is essentially analogous to wavelet denoising, but the different choice of wavelets enables correction of divergent flow components, providing flow fields with more coherent streamlines when compared to unfiltered medical data. By directly processing images, data-based methods are in general more computationally efficient than model-based techniques.

However, none of the proposed approaches so far has taken into account the time-dependent nature of 4D flow data. Additionally, working on structured grids (images) prevents a precise quantification of near-wall blood flow markers that are often sought-after in medical applications, such as WSS fields. In a recent study [156], 4D RBFs were used to reconstruct velocity fields measured with MRI. To make their approach computationally light, the authors combined multi-quadric RBF interpolation with a partition of unity scheme. Since it operates on scattered data, this approach results efficient in high dimensions, and well-suited for enforcing Dirichlet boundary conditions on the reconstructed velocity fields and improving quantification of near-wall features.

Neural network approaches are a relatively new class of data-based methods that can be thought of as approximations of complex nonlinear functions. Within the scientific computing community, a popular paradigm for exploiting the high expressive power of NNs is represented by physics informed neural networks (PINNs). In the formulation proposed in [157], PINNs are parametrized multi-layer perceptrons (MLPs) that learn a function mapping coordinates to outputs. PINNs can be seen as continuous functions fitting sparse observations and simultaneously minimizing the residual of a PDE (in differential form) that is identified as the mathematical model that generated the data. This is done by constructing the loss function as a weighted sum of two terms: a data fidelity term and a PDE residual. By leveraging automatic differentiation, evaluation of the PDE residual does not introduce severe numerical errors, albeit entailing considerable increase in computational cost. The relative ease of implementation of PINNs and their potential to seamlessly combine measurements with PDEs has made this approach appealing for incorporating the N-S equations in fluid flow reconstruction applications [158, 159, 160, 161]. Nonetheless, when training PINNs, choosing the correct hyperparameters is often the result of a trial and error procedure. When dealing with high dimensional problems, hyperparameter tuning can become unfeasible. In particular: *i*) a sufficient number of collocation points (coordinates where the PDE residual is evaluated) needs to be considered to achieve good accuracy [157], and *ii*) the relative weights of the PDE residual in the loss function needs to be carefully calibrated [162]. These issues make PINNs still difficult to apply to real-world fluid mechanics applications where data is often noisy and high dimensional. To date, the only application of PINNs to 4D flow MRI data is the work of Fathi et al. [148]. The authors tested their approach on synthetic 4D flow data from a reference CFD simulation for which ground truth velocity and pressure fields were known, achieving SR and superior denoising with respect to DFWs [146]. Successively, they evaluated their method on *in vitro* 4D flow velocity measurements, but not on *in vivo* medical data.

PINNs can be thought of as a subset of coordinate-based MLPs [141, 163, 164]. These networks can represent complex signals by taking low-dimensional coordinates as input and returning the value of the signal at the input locations. Such signal representations are often referred to as INRs [141]. Signals represented by coordinate-based MLPs can be orders of magnitude more compact than their grid-based counterparts [165]. Coordinate-based MLPs are known to suffer from severe spectral bias, namely they

struggle to learn the high frequency signal components. To overcome this limitation, recent studies have proposed the introduction of a sinusoidal mapping of input coordinates [165, 166, 167]. In practice, this consists in a Fourier feature encoding [168] of input coordinates \mathbf{x} to $\gamma(\mathbf{x}) = [a_1 \cos(2\pi \mathbf{b}_1^\top \mathbf{v}), a_1 \sin(2\pi \mathbf{b}_1^\top \mathbf{v}), \dots, a_m \cos(2\pi \mathbf{b}_m^\top \mathbf{v}), a_m \sin(2\pi \mathbf{b}_m^\top \mathbf{v})]$, where, in most cases, $a_j = 1$ and \mathbf{b}_j is sampled from an isotropic distribution [165]. Successful biomedical applications of coordinate-based networks include 3D vascular surface reconstruction [169] and non-rigid medical image registration [170]. A further improvement of coordinate-based MLPs to capture high order derivatives in the output signal was achieved by employing sinusoidal activation functions for every hidden neural network layer, as introduced by Sitzmann et al. [141]. Sinusoidal representation networks (SIRENs) have been shown to be suited for representing complex natural signals, including images, solutions to Poisson equations and 3D shapes [141].

In this work, we employ SIRENs for learning time-varying velocity fields measured by 4D flow MRI. We leverage the MLP’s architecture to introduce an implicit prior to constrain the space of solutions and investigate such implicit regularization bias towards lower frequencies, which simultaneously prevents overfitting and reduces noise in flow-encoded MR images.

6.2 Implicit neural representations of 4D flow MRI

6.2.1 Problem setting

Partially borrowing the notation from [145], from here on we will use the superscripts $(\cdot)^*$ to denote measured quantities and $(\cdot)^\bullet$ to denote ground truth quantities (when they exist). Additionally, quantities defined on unstructured point sets (or meshes) will be denoted with lowercase letters and parentheses, while quantities defined on structured grids (images) will be denoted with uppercase letter and square brackets.

Let $V^* \in \mathbb{R}^{N_r \times N_c \times N_s \times N_t \times 3}$ be a time-resolved flow image volume sequence obtained from a reconstructed 4D flow acquisition, with N_r , N_c , N_s and N_t being the number of rows, columns, slices and cardiac frames, respectively. Let $\Delta \mathbf{x} \in \mathbb{R}^3$ and $\Delta t \in \mathbb{R}$ be the physical spacings between adjacent voxels along the spatial and temporal dimensions. At a voxel center with 4D coordinates $[\mathbf{x}_i, t_j]$, where $\mathbf{x}_i \in \mathbb{R}^3$, we denote the measured velocity vector as: $V^*[\mathbf{x}_i, t_j] \in \mathbb{R}^3$, with $i \in [1, N_r \times N_c \times N_s]$, and $j \in [1, N_t]$. Herein, the index i refers to the flattened spatial coordinates of the image.

We are interested in representing V^* with a continuous function $f : \mathbb{R}^4 \rightarrow \mathbb{R}^3$, where $f(\mathbf{x}_i, t_j) = V^*[\mathbf{x}_i, t_j]$. To approximate f , we use an MLP f_Θ with weights Θ and with sinusoidal activation functions (SIREN). The l -th layer of a SIREN receiving a generic input tensor $\mathbf{x}_l \in \mathbb{R}^{Q_l}$ performs the following operation:

$$\mathbf{x}_{l+1} = \sin(\Theta_l \mathbf{x}_l + \mathbf{b}_l), \quad (6.4)$$

where $\Theta_l \in \mathbb{R}^{P_l \times Q_l}$ and $\mathbf{b}_l \in \mathbb{R}^{P_l}$ are the weight matrix and biases of the l -th layer, respectively. Following [141], each weight θ is initialized so that $\theta \sim \mathcal{U}(-\sqrt{6/c}, \sqrt{6/c})$, where c is the generic input feature size. Furthermore, as proposed by [141], the first

layer of the SIREN is modified as: $\sin(\omega_0 \cdot \Theta \mathbf{x} + \mathbf{b})$. Following [141], we set $\omega_0 = 30$. For our purposes, a key advantage of this formulation lies in representing a high dimensional image as a continuous function that can be queried at arbitrary spatio-temporal resolutions.

6.2.2 Training a SIREN in 4D

In most practical cases dealing with blood vessels, one is only interested in reconstructing blood flow within a bounded region $\Omega \subset \mathbb{R}^3$ with inflow boundary Γ_i , outflow boundary Γ_o and wall boundary Γ_w , and within a time interval $[t_a, t_b]$. The inner fluid region is denoted by $\Omega_f = \Omega \setminus \{\Gamma_i \cup \Gamma_o \cup \Gamma_w\}$. In our approach, velocity field reconstruction is achieved by sampling N_f spatial voxel coordinates $\mathbf{x}^{(f)}$ from Ω_f repeated over the time interval $[t_a, t_b]$, i.e., $(\mathbf{x}_i^{(f)}, t_j)$, with $i = 1, 2, \dots, N_f; j = 1, 2, \dots, N_t$. Additionally, we enforce the no-slip condition on the vessel wall by sampling N_w spatial coordinates from Γ_w repeated over the time interval $[t_a, t_b]$, i.e., $(\mathbf{x}_p^{(w)}, t_q)$, with $p = 1, 2, \dots, N_w; q = 1, 2, \dots, N_t$. Of note, we oversample spatial coordinates from Γ_w by setting $N_f \approx N_w$. We denote the total number of spatial coordinates used for SIREN training as $N = N_f + N_w$. Unlike the original SIREN formulation, this approach allows us to only use a relatively small set of coordinates compared to the total number of image points, greatly reducing training time and memory cost.

For each generic input coordinate pair (\mathbf{x}, t) , the following non-dimensionalization is performed:

$$\hat{\mathbf{x}} = \frac{\mathbf{x} - \mathbf{x}_{min}}{\Delta \mathbf{x}} D, \quad \hat{t} = \frac{t - t_{min}}{\Delta t} D, \quad (6.5)$$

where \mathbf{x}_{min} and t_{min} are the minimum spatial and temporal coordinates, and we set $D = 0.01$.

Training f_Θ implies solving the following minimization problem:

$$\min_{\Theta} \mathcal{L}(\Theta), \quad (6.6)$$

where the loss function \mathcal{L} is given by the misfit between the MLP prediction and the measured data plus a boundary condition term:

$$\mathcal{L}(\Theta) = \sum_{i=1, j=1}^{N_f, N_t} \|f_\Theta(\hat{\mathbf{x}}_i^{(f)}, \hat{t}_j) - V^*[\mathbf{x}_i^{(f)}, t_j]\|_2^2 + \sum_{p=1, q=1}^{N_w, N_t} \|f_\Theta(\hat{\mathbf{x}}_p^{(w)}, \hat{t}_q)\|_2^2. \quad (6.7)$$

Hence, the only supervision comes from the image values and fixed Dirichlet boundary conditions, making the approach fully unsupervised. Training is carried out using a limited memory limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm with learning rate of 1 until $\nabla_{\Theta} \mathcal{L} = 0$.

6.2.3 SIREN evaluation

Once trained, f_Θ can be queried at arbitrary spatio-temporal collocation points denoted as (\mathbf{x}', t') , sampled from the spatio-temporal domain $\Omega \times [t_a, t_b]$. To evaluate the gen-

eralization capabilities of a trained SIREN, collocation points corresponding to N' spatial coordinates and N'_t temporal coordinates, i.e., (\mathbf{x}'_n, t'_n) , with $n = 1, 2, \dots, N'$; $m = 1, 2, \dots, M'$, are defined at $\approx \times 20$ higher spatial resolution and $\times 10$ higher temporal resolution than image coordinates. Non-dimensionalization of evaluation coordinates is operated consistently with the one used for the training set:

$$\hat{\mathbf{x}}' = \frac{\mathbf{x}' - \mathbf{x}_{min}}{\Delta \mathbf{x}} D, \quad \hat{t}' = \frac{t' - t_{min}}{\Delta t} D. \quad (6.8)$$

Therefore, we evaluate f_{Θ} on: $(\hat{\mathbf{x}}'_n, \hat{t}'_m), n = 1, \dots, N'; m = 1, \dots, M'$.

6.2.4 Error quantification

To quantify errors obtained in experiments, three different metrics were used. Differences between a reference vector field \mathbf{u}_{ref} and another generic vector field \mathbf{u} , were evaluated by computing vNRMSE, mNRSME, and the DE as:

$$mNRMSE = \frac{1}{\max |\mathbf{u}_{ref}|} \sqrt{\frac{1}{K} \sum_{k=1}^K (|\mathbf{u}| - |\mathbf{u}_{ref}|)_k^2}, \quad (6.9)$$

$$vNRMSE = \frac{1}{\max |\mathbf{u}_{ref}|} \sqrt{\frac{1}{K} \sum_{k=1}^K (\mathbf{u} - \mathbf{u}_{ref})_k^2}, \quad (6.10)$$

$$DE = \frac{1}{K} \sum_{k=1}^K \left(1 - \frac{|\mathbf{u}_{ref,k} \cdot \mathbf{u}_k|}{|\mathbf{u}_{ref,k}| |\mathbf{u}_k|} \right), \quad (6.11)$$

where K is the generic number of 4D points where the two velocity fields are evaluated.

6.2.5 Wall shear stress analysis

From a generic velocity field, the WSS field was calculated following the approach described in [171]. From the definition of WSS for a Newtonian fluid:

$$WSS = \mu \left(\frac{\partial v}{\partial y} \right)_{y=0}, \quad (6.12)$$

where μ is the dynamic viscosity, v is the component of the velocity vector that is locally parallel to the wall, and y is the Euclidean distance from the wall, for each spatial 3D point on the vessel wall (Γ_w), the implemented WSS calculation method requires interpolation of the velocity fields at 2 points evenly spaced by a distance δ_n along the inward normal. The so obtained local velocity profile is interpolated with a quadratic function, whose analytical derivative is used to approximate $\partial v / \partial y$. In our experiments, we set $\delta_n = 0.5mm$

6.2.6 Case 1: synthetic 4D flow MRI

CFD simulation. To have a benchmark for evaluating the proposed method, a synthetic 4D flow MRI acquisition was created from a reference CFD simulation. First, the ascending aorta of a subject with TAA was segmented from 3D MRA images using open-source software [65]. The segmented domain Ω , was divided into 3 subdomains: inlet Γ_i , outlet Γ_o and wall Γ_w (Figure 6.1a). A 3D tetrahedral mesh with a base size of 0.6 mm was generated using the VMTK library [38]. The final volumetric mesh consisted of $\approx 800k$ nodes. Time-varying 3-directional velocity profiles (Figure 6.1b) were prescribed as inlet boundary conditions, mapping a realistic TAA inlet velocity profile to the Γ_i following the approach described in [172] and producing the flow waveform represented in Figure 6.1c. A zero-pressure condition was enforced on Γ_o and a homogeneous Dirichlet boundary condition (no-slip) was assumed on Γ_w . Blood was modeled as a Newtonian fluid with constant density $\rho = 1060 \text{ kg/m}^3$ and dynamic viscosity $\mu = 0.004 \text{ Pa}\cdot\text{s}$. A finite volume simulation was run at a fixed time step of 0.001 s. An implicit scheme with splitting of operators (PISO) was used to solve the governing equations of blood flow in Star-CCM+. Results were exported at every timestep within the time interval [0.2 - 0.399] s (peak to late systole), yielding a sequence of $M' = 200$ velocity fields \mathbf{u}^\bullet defined on the computational nodes $\mathbf{x}'_n, n = [1, \dots, N']$ over the simulation time coordinates $t'_m, m = [1, \dots, M']$.




Figure 6.1: a) Computational domain. b) Time-varying 3-directional velocity profiles prescribed at as inlet boundary condition. c) Flow waveform imposed at Γ_i

Synthetic data creation. Noise-free, high resolution CFD velocity fields were processed to obtain low resolution velocity fields corrupted with noise typical of 4D flow MRI measurements. To achieve this, the following steps were implemented partially following [173]:

1. the sequence of CFD solution snapshots denoted by \mathbf{u}^\bullet was temporally down-sampled to a sequence of $M = \frac{M'}{h}$ velocity fields $\bar{\mathbf{u}}$, using a moving average such that: $\bar{\mathbf{u}}(\mathbf{x}'_n, t'_j) = \frac{1}{h} \sum_{k=1}^{m+h-1} \mathbf{u}^\bullet(\mathbf{x}'_n, t'_k)$;

2. each time frame of \bar{u} was converted to a uniform Cartesian grid with voxel size of 1 mm^3 using a linear interpolation scheme to assign velocity vector values to grid cells, yielding a sequence of Cartesian grids $\tilde{U}[\tilde{x}_i, t_j] \in \mathbb{R}^{N_r \times N_c \times N_s \times N_t \times 3}$, with $N_r = 76$, $N_c = 112$ and $N_s = 292$;
3. each velocity grid in \tilde{U} was converted to a complex tensor containing magnitude and phase images using suitable VENC values as formalized in equation 6.1;
4. the fast Fourier transform was applied to obtain the corresponding k -space data;
5. a truncation of the 3D k -space data (high frequencies) was performed, effectively decreasing the spatial resolution by a factor of 2;
6. a zero-mean Gaussian noise with standard deviation σ corresponding to the desired SNR (calculated according to [173]) was added to the k -space data;
7. a randomized sampling mask drawn from a normal distribution and covering $S\%$ percent of the k -space was applied to further undersample the frequency content, keeping a fully sampled calibration region of $5 \times 5 \times 5$ in the center of k -space;
8. the inverse Fourier transform was applied to the undersampled, noise-corrupted k -space, yielding a complex tensor of magnitude and phase images;
9. complex images were converted back to real images of velocity fields using VENC values consisted with step 3, obtaining a sequence of noisy synthetic velocity measurements U^* defined at voxel coordinates $x_i, i = 1, \dots, N$ with isotropic voxel size of $2 \times 2 \times 2 \text{ mm}^3$ and at times $t_j, j = 1, \dots, M$.

For each velocity direction, the VENC value was chosen 10% larger than the maximum velocity, so that the phase wrapping/unwrapping would not introduce aliasing artifacts. Figure 6.2 shows the effect of the different implemented steps to create our synthetic 4D flow data.

The level of degradation of our synthetic measurements is regulated by three parameters: h , SNR and S . We chose $h = 40$ to obtain $\Delta t = 0.04\text{s}$, which is compatible with real medical measurements. We tested our approach on three different levels of degradation, *mild*, *medium* or *extreme* reported in Table 6.1 and shown in Figure 6.6.

Noise level	SNR	S
<i>mild</i>	20	99
<i>medium</i>	5	95
<i>extreme</i>	2	68

Table 6.1: Degradation parameters for generating synthetic measurements from the reference velocity field.

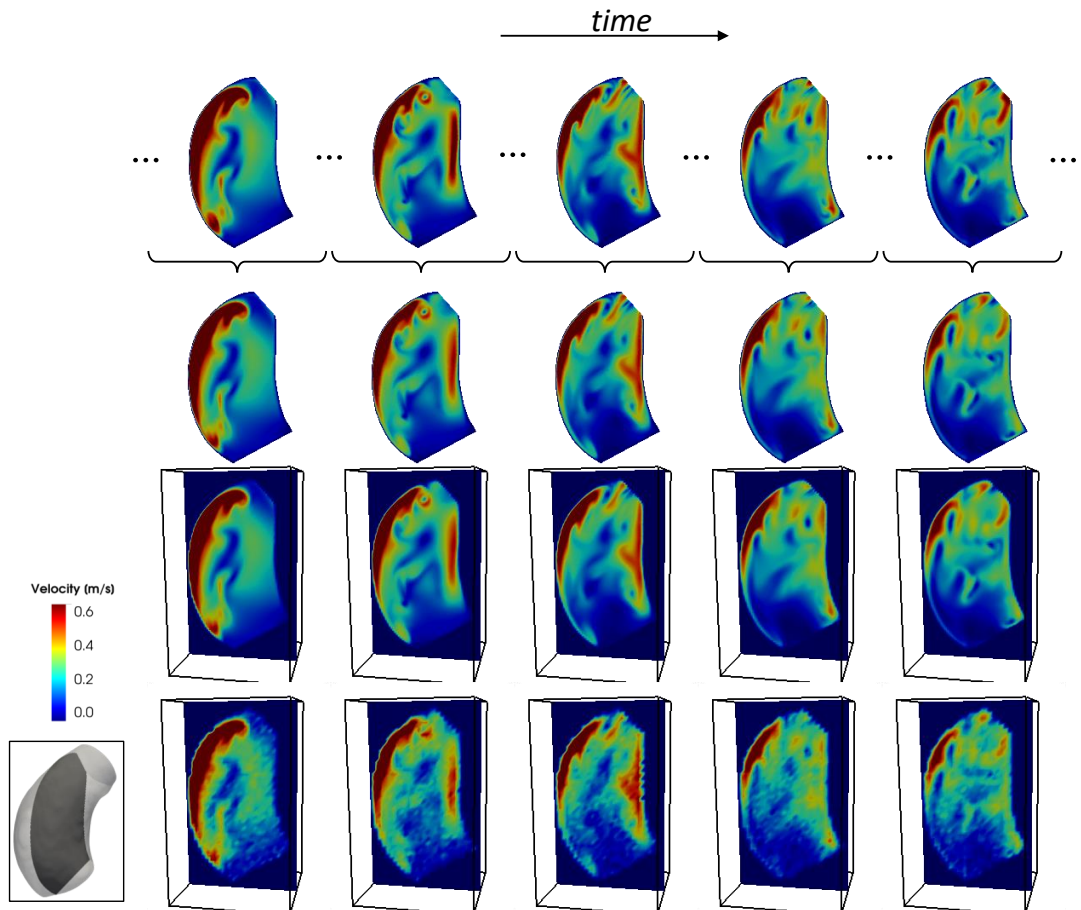


Figure 6.2: Visualization of velocity magnitude colormaps on a sagittally oriented 2D slice (bottom left) corresponding to the different steps to create synthetic 4D flow data from CFD results. CFD results (top row) are temporally averaged (second row). Temporally average velocity fields defined on unstructured meshes are resampled to a fine 3D Cartesian grid (third row). Fine cartesian grids are downsampled in k -space and complex Gaussian noise is added to produce synthetic flow measurements (bottom row).

Comparison with existing methods

The denoising and SR performances of our approach were compared against existing methods. Among them, only 4D RBFs lend themselves to the joint task of denoising and SR. Hence, pure denoising approaches were combined with simple interpolation schemes to achieve SR. The following methods were tested:

1. linear interpolation (LITP) as a baseline for SR;
2. DFW with automatic threshold selection based on *SureShrink* [146] for denoising and LITP for SR. Velocity fields generated by this approach will be denoted as DFW;
3. 3D divergence-free radial basis functions (DF-RBFs) as described in [147] for denoising and LITP for SR. Velocity fields generated by this approach will be denoted as DF-RBF;
4. an approach based on 4D RBFs [156], but for which no official implementation was available. Hence, we implemented our version of 4D RBFs with multi-quadric kernel and local support for denoising and SR. We set 10 as the number of nearest neighboring points for local kernel support. A *soft* enforcement of the no-slip condition on the vessel wall was applied by setting null velocity values at point belonging to the wall region. Velocity fields generated by this approach will be denoted as 4D-RBF.

6.2.7 Case 2: *in vivo* 4D flow MRI

A thoracic 4D flow MRI scan of a subject with ascending thoracic aortic aneurysm was retrospectively retrieved. Images were fully deidentified and provided by Weill Cornell Medicine, (NY, USA). A respiratory compensated technique was adopted with the following settings: spatial resolution (voxel size) $1.14 \text{ mm} \times 1.14 \text{ mm} \times 0.9 \text{ mm}$, field of view = 360 mm, flip angle = 15° , VENC = 200 cm/s in all 3 directions, time between consecutive frames = 30 ms, for a total of 20 frames per cardiac cycle. DICOM images were processed using open-source code [172] to compute the PC-MRA image and extract the segmentation of the enlarged ascending aortic tract (Figure 6.3).

The velocity components measured through flow MRI can be visualized in Figures 6.4 and 6.5 for a sagittally and an axially oriented 2D slice, respectively.




Figure 6.3: Left: segmented 3D geometry (red) superimposed on a slice representation of PC-MRA images. Right: fine 3D mesh for spatial super-resolution




Figure 6.4: Sagittally oriented 2D slice within the aortic aneurysm (bottom left), together with colormaps of left-to-right (first and fourth rows), posterior-to-anterior (second and fifth rows) and foot-to-head (third and sixth rows) velocity components. Rows 4 to 6 show velocities sampled on a 3D mesh by linear interpolation. Columns from left to right correspond to increasing time points.




Figure 6.5: Axially oriented 2D slice within the aortic aneurysm (bottom left), together with colormaps of left-to-right (first and fourth rows), posterior-to-anterior (second and fifth rows) and foot-to-head (third and sixth rows) velocity components. Rows 4 to 6 show velocities sampled on a 3D mesh by linear interpolation. Columns from left to right correspond to increasing time points.

6.3 Experiment results

6.3.1 Case 1: quantification of image degradation

The implemented degradation process to transform CFD velocity fields into synthetic flow MR measurements consisted of a sequence of steps that gradually worsened the quality of the data. The contribution of each step described in Section 6.2.6 is reported in Table 6.2. To make these comparisons, the CFD velocity fields were evaluated at coarse spatio-temporal discretizations as denoted in Section 6.2.6. Qualitatively, the addition of *mild* complex noise produced velocity fields that well kept the low frequency features present in the original fields (Figures 6.6 and 6.7, rows 1 and 2). Synthetic images with *medium* noise level produced more degraded images, with clearly visible artifacts and loss of high frequency details (Figures 6.6 and 6.7, rows 1 and 3). The *extreme* noise level yielded severely worsened velocity field with respect to ground truth, with a visible complete loss of fine flow details and generally higher velocities (Figures 6.6 and 6.7, rows 1 and 4).

	$\mathbf{u}_{ref} = \mathbf{u}^\bullet(\mathbf{x}', t)$ $\mathbf{u} = \bar{\mathbf{u}}(\mathbf{x}', t)$	$\mathbf{u}_{ref} = \mathbf{u}^\bullet(\tilde{\mathbf{x}}, t)$ $\mathbf{u} = \tilde{\mathbf{U}}[\tilde{\mathbf{x}}, t]$	$\mathbf{u}_{ref} = \mathbf{u}^\bullet(\mathbf{x}, t)$ $\mathbf{u} = \mathbf{U}^*[\mathbf{x}, t]$
mNRMSE [%]	1.79	1.75	3.63 5.24 8.23
vNMRSE[%]	2.06	2.01	3.22 5.60 9.64
DE [%]	4.95	0.59	1.41 7.96 15.9

Table 6.2: Errors introduced by the implemented degradation steps.

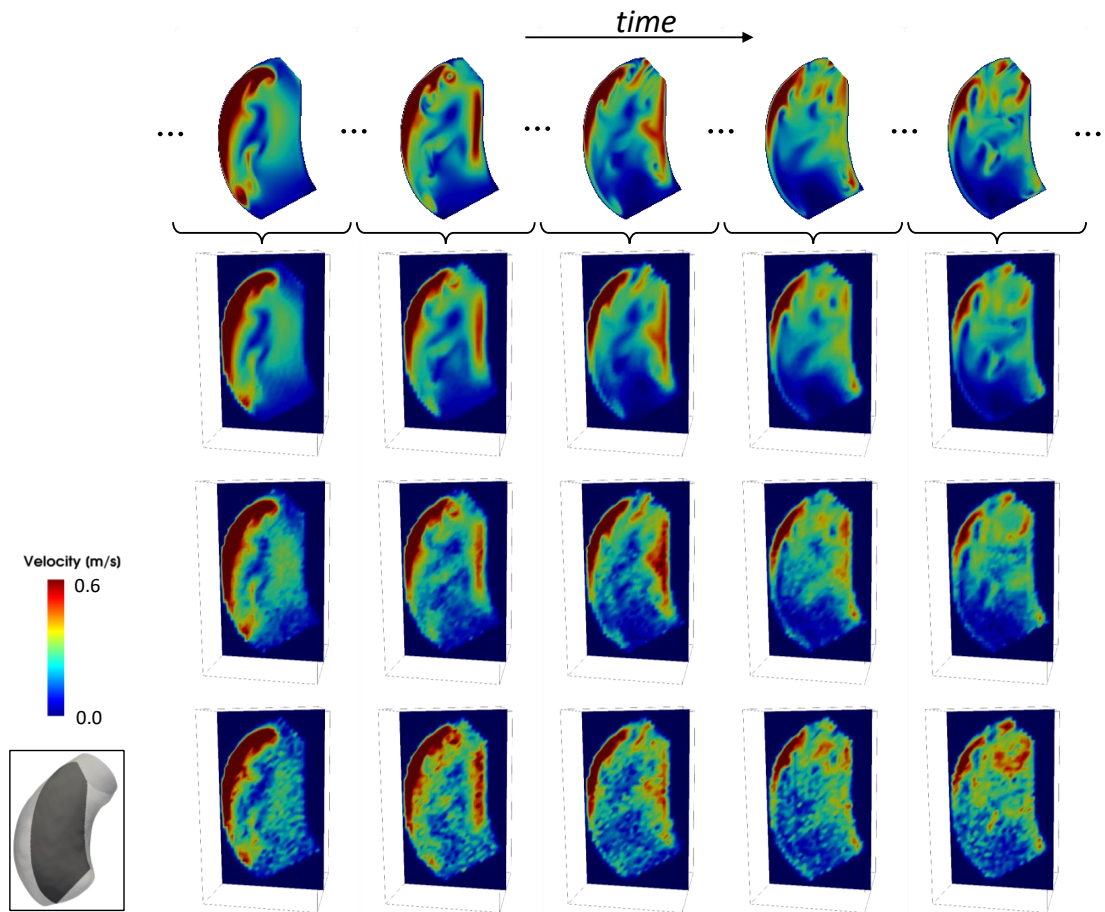


Figure 6.6: Velocity magnitude colormaps on a sagittal slice (bottom left corner) corresponding to: CFD solution (top row), images with *mild* added noise (second row), images with *medium* added noise (third row), images with *extreme* added noise (bottom row). Columns from left to right correspond to increasing time points.

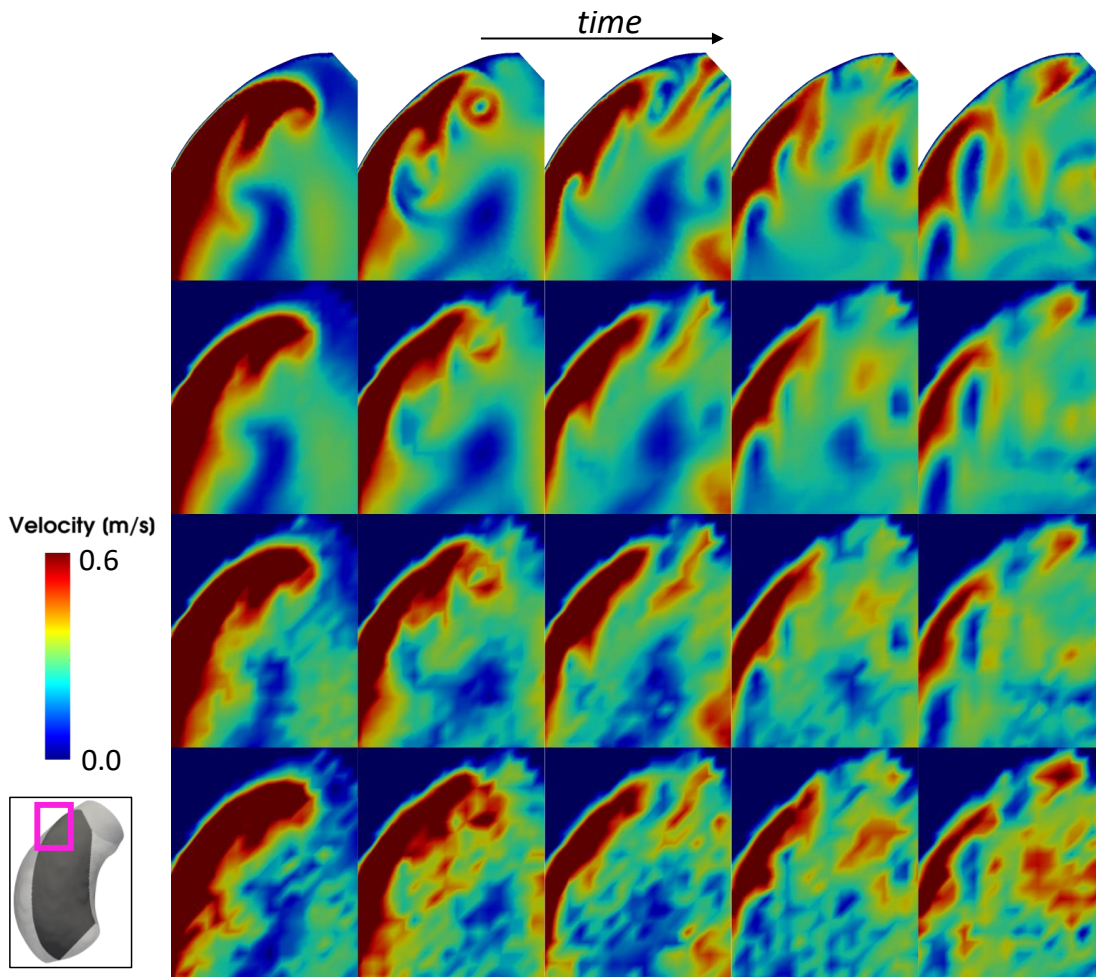


Figure 6.7: Detailed view of velocity magnitude colormaps on a sagittal slice (magenta rectangle in bottom left corner) corresponding to: CFD solution (top row), images with *mild* added noise (second row), images with *medium* added noise (third row), images with *extreme* added noise (bottom row).

6.3.2 Case 1: hyperparameter tuning

The effects of SIREN’s depth (number of layers) and width (number of neurons per layer) on the denoising and SR performances were assessed by training different configurations of f_{Θ} on synthetic images with *mild*, *medium* and *extreme* noise level. Trained models were evaluated on CFD nodal mesh coordinates x' and at time points t' evenly spaced by 0.004 s, effectively oversampling Ω by $\approx \times 20$ and the time interval [0.2 - 0.399] by $\times 10$. Results are reported in Tables 6.3, 6.4 and 6.5, for *mild*, *medium* and *extreme* noise levels, respectively. For *mild* noise levels, all the tested SIREN configurations gave low errors with respect to ground truth velocity fields. Networks with greater width and depth resulted in only slightly lower mNRMSE, vNRMSE and DE. For *medium* noise levels, better results were obtained by wider models, but not necessarily by deeper ones. For this noise settings, the best performing SIREN consisted of 12 layers, each with 500 neurons. In the case of *extreme* noise, wider networks gave worse results than narrower ones, while deeper architectures generally produced more accurate velocity fields compared to ground truth. Wider models showed a tendency to overfit high frequency noise, as shown in Figures 6.8 and 6.9.

We chose the single best configuration as the one that minimized the sum of vNRMSE, mNRMSE and DE for all three noise levels, hence a SIREN 20 layer deep and with 300 neurons per layer was selected for all downstream comparisons against other methods and on the medical dataset.

		<i>Depth</i>				
		4	8	12	16	20
<i>Width</i>	100	3.52	3.56	3.53	3.42	3.46
		3.43	3.36	3.32	3.22	3.27
		6.8	6.4	6.3	6.22	6.31
	200	3.48	3.46	3.52	3.52	3.41
		3.36	3.27	3.27	3.27	3.21
		6.43	6.23	6.16	6.16	6.14
	300	3.51	3.53	3.62	3.50	3.43
		3.40	3.30	3.33	3.24	3.21
		6.47	6.24	6.19	6.14	6.13
	400	3.49	3.47	3.60	3.49	3.53
		3.34	3.25	3.31	3.25	3.29
		6.37	6.16	6.17	6.12	6.15
	500	3.59	3.47	3.55	3.54	3.45
		3.45	3.26	3.30	3.29	3.23
		6.42	6.19	6.17	6.19	6.14

Table 6.3: Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *mild* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

		<i>Depth</i>				
		4	8	12	16	20
<i>Width</i>	100	4.48	4.63	4.65	5.86	5.29
		4.73	4.92	4.94	6.13	5.61
		10.5	10.8	10.9	13.6	12.3
	200	3.98	3.98	4.42	4.49	4.85
		4.24	4.24	4.72	4.81	5.15
		9.66	9.61	10.5	10.7	11.5
	300	3.94	3.94	3.97	4.02	4.20
		4.20	4.20	4.21	4.30	4.51
		9.59	9.56	9.5	9.67	10.21
	400	4.03	3.81	3.89	3.87	4.03
		4.35	4.09	4.16	4.12	4.28
		9.87	9.39	9.48	9.44	9.69
	500	3.95	3.82	3.83	3.85	3.84
		4.24	4.11	4.08	4.10	4.10
		9.62	9.41	9.38	9.33	9.36

Table 6.4: Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *medium* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

		<i>Depth</i>				
		4	8	12	16	20
<i>Width</i>	100	5.31	5.37	5.26	5.30	5.87
		5.73	5.80	5.73	5.75	6.29
		13.01	13.33	12.93	13.02	14.17
	200	5.84	5.99	5.16	5.15	5.21
		6.62	6.91	5.68	5.67	5.64
		14.83	15.11	12.87	12.78	12.84
	300	7.30	7.74	5.40	5.21	5.16
		8.50	9.05	6.09	5.77	5.63
		18.86	19.91	13.41	12.95	12.7
	400	8.86	8.14	7.61	5.56	5.44
		10.3	9.52	8.89	6.33	6.10
		22.34	21.02	19.68	13.98	13.51
	500	9.41	8.07	7.95	7.46	6.50
		10.9	9.45	9.28	8.65	7.56
		23.52	20.88	20.5	19.16	16.66

Table 6.5: Effect of SIREN number of layers (depth) and number of neurons per layer (width) on denoising and super-resolution using data with *extreme* noise level. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.




Figure 6.8: Sagittal view of the effect of different SIREN architectures on synthetic images with *extreme* noise level. Rows 1 and 5: images with *extreme* noise added; rows 2 and 6: CFD solution; rows 3 and 7: results of a SIREN with 20 layers and 500 neurons per layer; rows 4 and 8: results of a SIREN with 20 layers and 300 neurons per layer.

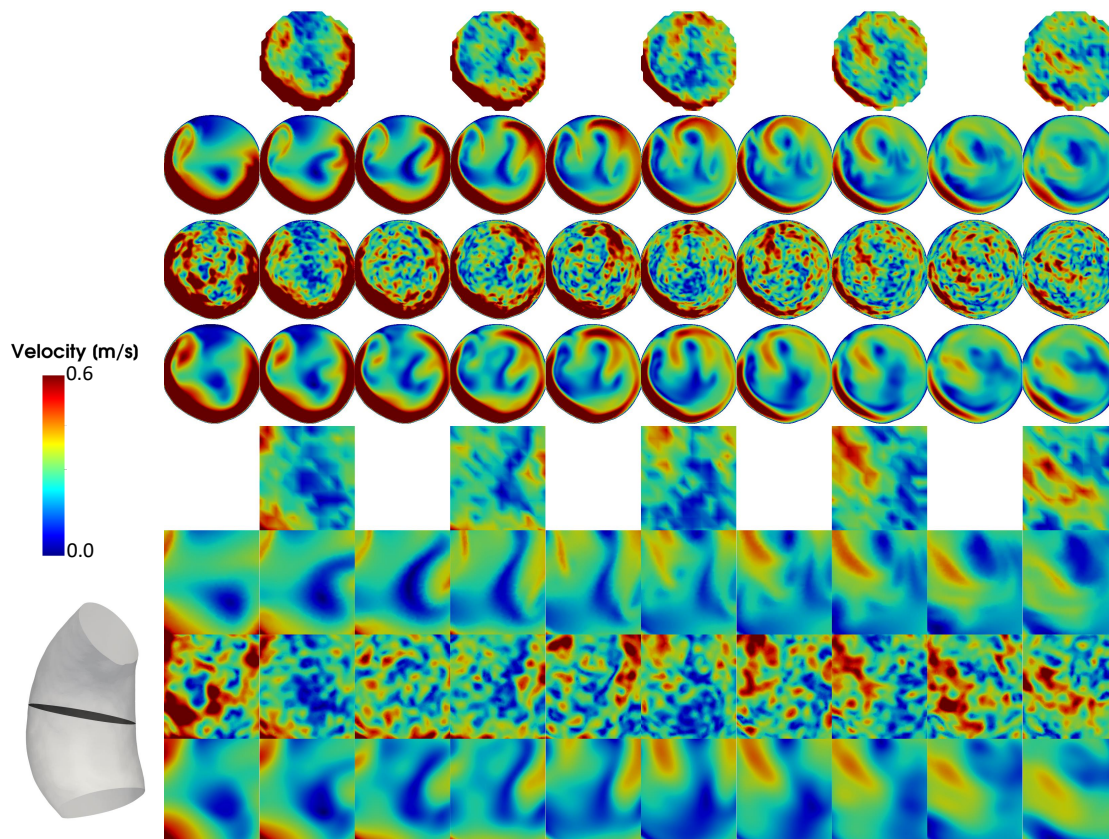


Figure 6.9: Axial view of the effect of different SIREN architectures on synthetic images with *extreme* noise level. Rows 1 and 5: images with *extreme* noise added; rows 2 and 6: CFD solution; rows 3 and 7: results of a SIREN with 20 layers and 500 neurons per layer; rows 4 and 8: results of a SIREN with 20 layers and 300 neurons per layer.

6.3.3 Case 1: comparison with existing methods

Velocity fields

The SIREN configuration that showed the best results was evaluated on the synthetic images with three different levels of noise. Velocity fields obtained with our approach were compared against the existing methods listed in Section 6.2.6. The proposed SIREN gave lower mNRMSE, vNRMSE and DE for all noise levels (Table 6.6). Among the LITP-based methods, DF-RBF [147] performed better than LITP alone, while DFW provided the worst performance. Qualitatively, all existing methods produced very similar velocity fields for the *mild* noise level case, showing suboptimal capabilities of reconstructing finer flow details (Figure 6.10). For the same level of noise, our method was able to fit flow details more accurately. Similar results were observed for the *medium* and *extreme* noise levels. As shown in Figures 6.11 and 6.12, all baseline approaches tended to perform denoising by oversmoothing the data. On the other hand, our method was able to filter out noise but better preserving finer flow structures.

	LITP	DFW +	DF-RBF +	4DRBF	SIREN
		LITP	LITP		
<i>mild</i>	5.88	6.47	5.90	6.91	3.50
	4.96	5.40	4.99	5.68	3.24
	7.13	7.96	7.2	8.15	6.14
<i>medium</i>	6.11	7.26	6.05	6.79	4.02
	5.50	6.18	5.42	6.01	4.30
	10.1	10.16	9.78	10.53	9.67
<i>extreme</i>	7.06	8.53	6.85	7.57	5.21
	6.96	7.43	6.62	7.04	5.77
	14.16	13.01	13.26	13.9	12.95

Table 6.6: Velocity field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.

Wall shear stress fields

The same methodology described in Section 6.2.5 was applied to the velocity fields obtained by CFD (ground truth), DFW, DF-RBF, 4D-RBF and our best SIREN (20 layers, 300 neurons per layer). Results obtained by the different approach are reported in Table 6.7. Our method gave the lowest mNRMSE, vNRMSE and DE, outperforming the others for all the tested noise levels. WSS fields computed from the denoised and super-resolved velocity field obtained by the different approaches can be visualized as 3D colormaps in Figures 6.13, 6.15 and 6.17 for *mild*, *medium* and *extreme* noise levels, respectively, and as unwrapped 2D surfaces in figures 6.14, 6.16 and 6.18. More visible




Figure 6.10: Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *mild* noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.




Figure 6.11: Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *medium* noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.




Figure 6.12: Velocity colormaps on a 2D sagittal slice (bottom left). Top row: synthetic measurements with *extreme* noise level at time points evenly spaced by $\Delta t = 40ms$. Second row: ground truth velocity fields from CFD at time points evenly spaced by $\Delta t = 20ms$. Third row: LITP results. Fourth row: DFW results. Fifth row: DF-RBF results. Sixth row: 4D-RBF results. Seventh row: our method.

differences produced by the different methods can be appreciated for the *extreme* noise level (Figures 6.17 and 6.18). In this case, all methods, including ours, underestimated WSS magnitudes. Our SIREN showed superior performance, giving WSS fields in good agreement with CFD-derived results, with high WSS magnitudes on the outer curve of the vessel wall, as expected.

	LITP	DFW +	DF-RBF +	4DRBF	SIREN
		LITP	LITP		
<i>mild</i>	8.76	8.65	8.68	18.6	5.62
	6.03	6.10	6.00	12.31	4.02
	3.43	4.64	3.52	3.84	2.1
<i>medium</i>	8.43	7.93	8.29	18.6	6.53
	6.05	5.99	5.96	12.6	4.76
	4.94	5.83	4.86	6.85	3.98
<i>extreme</i>	10.32	9.06	10.1	19.5	7.82
	7.59	7.01	7.41	13.4	6.04
	7.74	8.41	7.5	10.34	7.0

Table 6.7: WSS field comparison with existing methods. In each cell values of mNRMSE (top), vNRMSE (middle) and DE (bottom) are reported.




Figure 6.13: Results corresponding to *mild* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).




Figure 6.14: Results corresponding to *mild* noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).




Figure 6.15: Results corresponding to *medium* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).




Figure 6.16: Results corresponding to *medium* noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).




Figure 6.17: Results corresponding to *extreme* noise levels. WSS colormaps on the aortic wall at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).




Figure 6.18: Results corresponding to *extreme* noise levels. WSS colormaps on the aortic wall represented as a 2D surface at time points evenly spaced by $\Delta t = 20ms$. Top row: WSS computed on ground truth velocity fields from CFD. LITP (second row), DFW (third row), DF-RBF (fourth row), 4D-RBF (fifth row), our method (sixth row).

6.3.4 Case 2: velocity field assessment

Table 6.8 shows quantitative measurements from raw MRI data and reconstructed velocity fields with the proposed SIREN. Overall, our method gave high resolution velocity fields that maintained low discrepancies in macroscopic quantitative measurements that are considered to be reasonably accurate when assessed from unprocessed MRI measurements [174], giving differences in mean and maximum flow rate of <5% and underestimating reverse flow index (RFI) by 9.4%, as computed in [119]. Qualitatively, our approach produced cleaner velocity fields, but maintaining the high velocity regions observed at the extrados of the ascending aorta as in the measured data (Figures 6.19 and 6.20). From the 3D vector visualization (third and fourth rows in Figures 6.19 and 6.20), it can be appreciated how the reconstructed velocity vector field kept the swirling patterns formed in the middle ascending aorta, but filtered out the spurious vector components in the near-wall regions.

	Measured data	SIREN	$\Delta\%$
Mean flow rate [L/min]	4.66	4.88	4.9%
Max flow rate [L/min]	27.7	27.2	1.8%
RFI [%]	21	19	9.4%

Table 6.8: Quantitative measurements for Case 2 on raw flow MR measurements and on velocity fields reconstructed with our method.




Figure 6.19: Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on a sagittally oriented 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.




Figure 6.20: Velocity magnitude colormaps (rows 1 and 2) and velocity vectors (rows 3 and 4) for Case 2 on an axially oriented 2D slice (bottom left). Rows 1 and 3: 4D flow measurements, rows 2 and 4: SIREN results.

6.3.5 Case 2: wall shear stress field assessment

As observed for the synthetic case, the high velocity region near the wall of the outer curve of the aneurysm caused higher WSS in this region (Figure 6.21), with a maximum WSS value of 20 Pa and mean time-averaged WSS (TAWSS) of 7.5 mPa.




Figure 6.21: WSS magnitude colormaps (row 1) and WSS vectors (rows 2) computed from SIREN results.

6.4 Discussion of SIRENs for 4D flow MRI

The use of 4D flow MRI in analyzing blood flow in major vessels has been widely studied, but the limitations of this imaging technique, such as noise and low spatio-temporal resolution have hindered its use in more advance velocity-based hemodynamic analysis. For instance, some studies have attempted to quantify WSS directly from flow-encoded MR images [175, 176], evaluating their accuracy at different levels of image noise or spatial resolution. Although these approaches enable fast assessment of near-wall quantities, they often tend to underestimate the true WSS values with respect to numerical simulation. In pursuit of more accurate quantification of WSS fields, researchers have turned to increasingly complex numerical simulations of blood flow [142, 177]. Although CFD studies are a powerful tool for estimating WSS, they require significant computational resources and can take several hours, or even days, to run due to their computational cost.

In this study, we proposed an unsupervised learning method based on INRs for denoising and SR of velocity fields measured by 4D flow MRI. We evaluated our approach on realistic synthetic data with various levels of noise and showed the superior performance of our method with respect to state-of-the-art methods in terms of both denois-

ing and spatio-temporal SR. The proposed approach was able to denoise and super-resolve 4D velocity data while maintaining the integrity of the dominant flow features.

Among neural network approaches, it is worth mentioning the studies of Ferdian et al. [173], Rutkowski et al. [178] and [179]. In these works, CFD simulations are used to create synthetic 4D flow MRI datasets, so to have measurements - ground truth pairs to train convolutional neural networks for denoising and SR. Hypothetically, if large realistic training data is generated, these approaches would be able to learn the denoising and SR tasks, and could generalize to new domains without the need of re-training. Nonetheless, by operating convolutions in the image domain, these models are usually engineered to super-resolve $\times 2$ or $\times 4$ in space alone, and they do not allow for precise evaluation of near-wall quantities. Additionally, these methods require full ground truth supervision. For these reasons, we did not include them in our comparison against existing methods.

The success of our method relies on two main properties of MLPs. First, we leverage SIREN’s spectral bias [180, 181] to achieve velocity field denoising. This property of dense fully-connected networks prevents them from learning high frequency functions. Assuming noise in MR to be Gaussian in k -space results in high frequency artifacts in image space that are effectively removed by an INR. Our experiments on simulated 4D flow data (Case 1) allowed us to gain insights into the capabilities of SIRENS to denoise velocity data. When exploiting the spectral bias for signal denoising, one should carefully choose model’s width and depth such that high frequency noise is filtered out, while desirable fine signal details are maintained. Overall, with *mild* and *medium* noise levels, all the tested combinations of depths and widths were able to fit the data, with a tendency of larger models to give slightly lower errors (Tables 6.3, 6.4). For *extreme* noise levels, SIRENS with more than 300 neurons per layer overcame the spectral bias and overfitted high frequency noise. Once the best architecture was identified, our approach outperformed all other tested denoising methods. The second key strength of the proposed approach relies on the fact that the trained MLP fits the velocity vector field as a continuous function of space-time coordinates. In practice, this was achieved by building upon the work of Sitzmann et al. [141], who showed how SIRENS are better suited to fit complicated, feature-rich signals, such as natural images and solutions to simple PDEs. The present work is the first to adopt SIRENS for fitting 4D velocity measurements. By acting pointwise on 4D coordinates, a trained SIREN can be queried at continuous spatio-temporal locations, theoretically providing SR at arbitrarily fine spatial and temporal scales. Additionally, by incorporating time as an input feature, SIREN’s pointwise outputs are implicitly affected by temporal neighboring point features.

In principle, our approach is closely related to the formulation introduced by PINNs [157], with the main difference lying in the definition of the loss function. We train our models only using a data fidelity term, neglecting physical priors. The introduction of N-S residuals in differential form in the loss function could potentially lead to more physically consistent velocity fields, but would significantly slow down the training process. In the present study, we were interested in demonstrating the feasibility of

a method that could potentially be applied to real medical scenarios where speed of execution is required. For the TAA patient (Case 2), using an NVIDIA A100 graphics card, our method took approximately 4 minutes to train and less than 2 seconds to evaluate at a fine spatio-temporal resolution.

In contrast to most data-based approaches, the devised method requires the definition of a bounded domain $\Omega \times [t_a, t_b]$ (Section 6.2.2). On one hand, this choice represents a limitation of our workflow, entailing longer processing times. On the other hand, precise definition of a smooth vessel wall surface enables robust computation of WSS, a clinically important hemodynamic biomarker. Results on synthetic data showed good quantitative and qualitative agreement between predicted and reference CFD data (Table 6.7 and Figures 6.13, 6.15, 6.17). Results on real medical data revealed a high WSS region on the outer curvature of the ascending aorta (Figure 6.21), with most values ranging from 0 to 15 Pa. These results are in good agreement with a recent study on TAA biomechanics [116], who reported maximum WSS values of 10.18 ± 4.14 Pa for a cohort of 10 patients. Even though more detailed analyses are needed on the use of neural networks for *in vivo* WSS estimation, our results indicate that velocity fields produced by our method are suitable for extraction of derived biomarkers, otherwise difficult to assess.

6.5 Final considerations

In this work we showed the feasibility of SIRENS to represent complex, high dimensional blood flow velocity fields measured by 4D flow MRI. By training on low resolution coordinates, our method is quick to execute for new cases and easy to implement. By carefully tuning our SIREN architecture, we exploit the spectral bias to obtain a functional representation of our data with reduced noise, outperforming state-of-the-art solutions. Our method provides continuous velocity fields that can be queried at arbitrary spatio-temporal locations, effectively achieving 4D super-resolution.

CHAPTER 7

Final remarks

This chapter offers an overview and analysis of the research presented throughout this thesis.

The work presented in **Chapter 2**, involved developing a tool that can automatically compute geometric parameters relevant to TEVAR planning using CNNs and computational geometry. The tool successfully segmented the thoracic aorta and pulmonary vasculature of 70 CT scans, including 9 pathological cases of TAA. The study found significant differences in geometric aortic features that may explain the less satisfactory clinical results obtained by TEVAR in CILCA arches compared to standard ones. The tool is automated, quick, reliable, and considered an innovative solution to improve the decision-making process before TEVAR.

In **Chapter 3**, we built a pipeline that automates the quantification of complex morphological biomarkers for preprocedural planning of TAVI using CT images. To validate our approach, we tested it on the largest patient dataset to date, comprising 178 patients. Our proposed method delivers quick, accurate, and consistent assessment of aortic root anatomy, yielding 3D segmentations of the aortic root with a mean Dice score of 0.96. By incorporating this deep learning-based tool into the preoperative planning routine in TAVI environments, we anticipate potential time and cost savings, as well as improved accuracy.

Chapter 4 proposed a system that supports neurosurgeons during EVD procedures. It utilizes a CNN for automatic segmentation of the brain ventricles, detection of the procedural target and a markerless registration method based on a depth sensor to align internal brain structures with the physical position of the patient. The system was tested and proven to improve the accuracy of neurosurgeons targeting by 42% on average, with little to no manual intervention and computations required in less than 3 minutes.

Chapter 5 presented a system that generates time-dependent 3D aortic velocity profiles using a data-driven statistical shape model. The generated velocity profiles can be used in numerical simulations of blood flow, and can be mapped to the inlet plane of any virtual subject. The proposed software system replaces the common practice of prescribing idealized inflow boundary conditions with more realistic conditions, setting a new standard for the computational bioengineering community.

Chapter 6 demonstrated the use of multilayer perceptrons with sinusoidal activation functions (SIRENs) to represent high-dimensional blood flow velocity fields measured by 4D flow MRI. By training the network on low-resolution coordinates, the method is computationally efficient and easy to implement for new cases. We showed how choosing the proper SIREN architecture is crucial to effectively take advantage of its spectral bias, resulting in a functional representation of the data with reduced noise and superior performance compared to existing methods. The continuous velocity fields generated by this approach can be queried at any spatio-temporal location, achieving 4D super-resolution.

Future directions and perspectives

This thesis presents various applications of machine learning in biomedical sciences, highlighting its significant impact in this field.

From a methodological standpoint, one common limitation of the approaches presented in Part I is the separation between semantic segmentation and geometric feature extraction. On one hand, this choice enabled a step-by-step control over both the segmentation phase and the advanced post-processing phase. On the other hand, it did not involve end-to-end training of machine learning models that, in principle, could perform both segmentation and feature extraction steps in a single operation. For instance, in Chapter 2, the computed morphological biomarkers are based on the vessel centerlines. One could train a model to extract centerlines directly from 3D image data; such a model could make use of a convolutional encoder followed by an MLP, to predict the set of point coordinates belonging to the vessel's skeleton. In a supervised setting, this kind of task could be achieved by, e.g., minimizing the Chamfer distance of the prediction with respect to the ground truth centerlines. Similarly, in Chapter 3, the anatomical measurements of the aortic root are automatically extracted through a disjoint segmentation – post-processing approach. In this case, morphological measurement of interest for pre-procedural planning rely on specific landmarks that could be detected by a neural network in an end-to-end fashion. A successful example of how this could be achieved is presented by [45], who adopt a global-to-local supervised landmark detection approach that can be applied to multiple imaging modalities. Within our work, instead of focusing on methodological improvements, current efforts are being devoted to train a classifier that analyzes pre-operative CT scans of subjects before they underwent TAVI and predict the risk of post-operative complications. Early attempts have shown a tendency of common classifier architectures (e.g. 3D ResNet and DenseNet) to overfit the training distribution. In the future, multiple regularization strategies will be investigated, including the adoption of an image reconstruction convolutional block that will encourage the encoder to extract features that are also representative of the true image content.

In general, the post-processing approaches described in Chapters 2, 3 and 4 dealt with analyzing unstructured data that can be represented as a mesh, or graph. In the future, the use of graph neural networks (GNNs) should be investigated. GNNs constitute a relatively new class of machine learning methods suitable for working with high-dimensional, unstructured data. By performing aggregations and pooling operations directly on graphs, GNNs can learn functions and representations on non-Cartesian data structures. A generic landmark detection on a graph could be formulated as a node classification problem; hence, GNNs could represent a unifying solution for many different anatomical feature quantification problems. In the future, besides well-established paradigms based on graph convolutions, it will be interesting to leverage unsupervised approaches such as Deep Graph Infomax [182], since a meticulous annotation of key anatomical landmarks on complex medical data can often be daunting.

A possible future extension of the work proposed in Chapter 5 could consist of replacing the SSM with a deep autoencoder. Compared to PCA, which is a linear trans-

formation, a deeper non-linear model could, in principle, learn a more representative latent space of the complex blood flow velocity data. Following this approach, one could reduce the number of components that are combined to reconstruct the training data and generate new velocity profiles.

Finally, current efforts are being focused on improving the formulation presented in Chapter 6. In particular, adding a physically meaningful constrain to enhance blood flow fields measured by 4D flow MRI would be of great impact for many applications. Besides adding a PDE residual term in the loss function (PINN approach), novel and more efficient ways of enforcing a *soft* physics constraint will be investigated. For instance, one could achieve super-resolution with increased temporal coherency by exploiting either advection-based reconstruction techniques [183] or a Lagrangian approach to incorporate the fluid dynamics equations [184]. These methodologies have been successfully employed in different fluid dynamics scenarios, including 3D particle image velocimetry and smoke simulation. Nonetheless, their applicability to blood flow velocity data remains unexplored.

Each chapter of this thesis focused on a specific imaging modality; however, future research efforts should prioritize the integration of multi-modality data to improve our understanding of physiology and pathophysiology. Building models capable of handling high-dimensional data from various imaging modalities to capture complementary information is essential. However, constructing datasets that include multi-modality imaging for a large number of subjects is a significant challenge. Despite the increase in medical data digitization, there is relatively less data available for analysis and modeling, particularly for rare and specific pathologies. Thus, promoting data sharing among clinical institutions to build more extensive and better medical datasets that researchers worldwide can access is crucial.

As machine learning becomes increasingly integrated into medical applications, model explainability and interpretability are critical challenges that need to be addressed. Lack of transparency in these models can make it difficult to understand their decision-making process, hindering their adoption in clinical practice. Deep learning models, in particular, can be highly complex and challenging to interpret. To overcome these challenges, future research in machine learning for medical applications must develop models that are both highly accurate and interpretable. This could involve new paradigms that share a common ground with more mathematically consolidated approaches. For example, Bayesian networks offer a flexible and powerful framework for uncertainty quantification that can be used in a broad range of applications, including decision-making under uncertainty.

Bibliography

- [1] P. Lambin, R. T. Leijenaar, T. M. Deist, J. Peerlings, E. E. De Jong, J. Van Timmeren, S. Sanduleanu, R. T. Larue, A. J. Even, A. Jochems, *et al.*, “Radiomics: the bridge between medical imaging and personalized medicine,” *Nature reviews Clinical oncology*, vol. 14, no. 12, pp. 749–762, 2017.
- [2] P. Schoenhagen, U. Numburi, S. S. Halliburton, P. Aulbach, M. Von Roden, M. Y. Desai, L. L. Rodriguez, S. R. Kapadia, E. M. Tuzcu, and B. W. Lytle, “Three-dimensional imaging in the context of minimally invasive and transcatheter cardiovascular interventions using multi-detector computed tomography: from pre-operative planning to intra-operative guidance,” *European Heart Journal*, vol. 31, no. 22, pp. 2727–2740, 2010.
- [3] F. Alam, S. U. Rahman, S. Ullah, and K. Gulati, “Medical image registration in image guided surgery: Issues, challenges and research opportunities,” *Biocybernetics and Biomedical Engineering*, vol. 38, no. 1, pp. 71–89, 2018.
- [4] G. Coorey, G. A. Figtree, D. F. Fletcher, V. J. Snelson, S. T. Vernon, D. Winlaw, S. M. Grieve, A. McEwan, J. Y. H. Yang, P. Qian, *et al.*, “The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field,” *npj Digital Medicine*, vol. 5, no. 1, p. 126, 2022.
- [5] A. Badano, A. Badal, S. Glick, C. G. Graff, F. Samuelson, D. Sharma, and R. Zeng, “In silico imaging clinical trials for regulatory evaluation: initial considerations for VICTRE, a demonstration study,” in *Medical Imaging 2017: Physics of Medical Imaging*, vol. 10132, pp. 494–499, SPIE, 2017.
- [6] E. Abadi, W. P. Segars, B. M. Tsui, P. E. Kinahan, N. Bottenus, A. F. Frangi, A. Maidment, J. Lo, and E. Samei, “Virtual clinical trials in medical imaging: a review,” *Journal of Medical Imaging*, vol. 7, no. 4, pp. 042805–042805, 2020.

- [7] Y. S. Chatzizisis, A. U. Coskun, M. Jonas, E. R. Edelman, P. H. Stone, and C. L. Feldman, "Risk stratification of individual coronary lesions using local endothelial shear stress: a new paradigm for managing coronary artery disease," *Current Opinion in Cardiology*, vol. 22, no. 6, pp. 552–564, 2007.
- [8] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.
- [9] V. Rimbau, D. Böckler, J. Brunkwall, *et al.*, "Management of descending thoracic aorta diseases. clinical practice guidelines of the european society for vascular surgery (esvs)," *Eur J Vasc Endovasc Surg*, vol. 53, pp. 4–52, 2017.
- [10] K. S. Bellamkonda, S. Yousef, N. Nassiri, A. Dardik, R. J. Guzman, A. Geirsson, and C. I. O. Chaar, "Trends and outcomes of thoracic endovascular aortic repair with open concomitant cervical debranching," *Journal of Vascular Surgery*, vol. 73, no. 4, pp. 1205–1212, 2021.
- [11] M. M. Marrocco-Trischitta, H. W. de Beaufort, G. Piffaretti, S. Bonardelli, M. Gargiulo, M. Antonello, J. A. van Herwaarden, S. Boveri, R. Bellosta, S. Trimarchi, *et al.*, "The modified arch landing areas nomenclature predicts proximal endograft failure after thoracic endovascular aortic repair," *European Journal of Cardio-Thoracic Surgery*, vol. 58, no. 2, pp. 309–318, 2020.
- [12] C. A. Figueroa, C. A. Taylor, A. J. Chiou, V. Yeh, and C. K. Zarins, "Magnitude and direction of pulsatile displacement forces acting on thoracic aortic endografts," *Journal of Endovascular Therapy*, vol. 16, no. 3, pp. 350–358, 2009.
- [13] A. F. members, R. Erbel, V. Aboyans, C. Boileau, E. Bossone, R. D. Bartolomeo, H. Eggebrecht, A. Evangelista, V. Falk, H. Frank, *et al.*, "2014 esc guidelines on the diagnosis and treatment of aortic diseases: document covering acute and chronic aortic diseases of the thoracic and abdominal aorta of the adult the task force for the diagnosis and treatment of aortic diseases of the european society of cardiology (esc)," *European heart journal*, vol. 35, no. 41, pp. 2873–2926, 2014.
- [14] M. Grabenwöger, F. Alfonso, J. Bachet, R. Bonser, M. Czerny, H. Eggebrecht, A. Evangelista, R. Fattori, H. Jakob, L. Lönn, *et al.*, "Thoracic endovascular aortic repair (tevar) for the treatment of aortic diseases: a position statement from the european association for cardio-thoracic surgery (eacts) and the european society of cardiology (esc), in collaboration with the european association of percutaneous cardiovascular interventions (eapci)," *European heart journal*, vol. 33, no. 13, pp. 1558–1563, 2012.
- [15] J. S. Matsumura, R. P. Cambria, M. D. Dake, R. D. Moore, L. G. Svensson, S. Snyder, T. C. T. Investigators, *et al.*, "International controlled clinical trial of thoracic endovascular aneurysm repair with the zenith tx2 endovascular graft: 1-year results," *Journal of vascular surgery*, vol. 47, no. 2, pp. 247–257, 2008.

- [16] S. Ishimaru, "Endografting of the aortic arch," *Journal of endovascular therapy*, vol. 11, no. 6_suppl, pp. II-62, 2004.
- [17] H.-E. Altnji, B. Bou-Saïd, and H. Walter-Le Berre, "Morphological and stent design risk factors to prevent migration phenomena for a thoracic aneurysm: a numerical analysis," *Medical engineering & physics*, vol. 37, no. 1, pp. 23-33, 2015.
- [18] T. Ueda, H. Takaoka, B. Raman, J. Rosenberg, and G. D. Rubin, "Impact of quantitatively determined native thoracic aortic tortuosity on endoleak development after thoracic endovascular aortic repair," *AJR-American Journal of Roentgenology*, vol. 197, no. 6, p. W1140, 2011.
- [19] M. M. Marrocco-Trischitta, M. Alaidroos, R. M. Romarowski, V. Milani, F. Ambrogi, F. Secchi, M. Glauber, and G. Nano, "Aortic arch variant with a common origin of the innominate and left carotid artery as a determinant of thoracic aortic disease: a systematic review and meta-analysis," *European Journal of Cardio-Thoracic Surgery*, vol. 57, no. 3, pp. 422-427, 2020.
- [20] S. N. Mylonas, A. Barkans, M. Ante, J. Wippermann, D. Böckler, and J. S. Brunkwall, "Prevalence of bovine aortic arch variant in patients with aortic dissection and its implications in the outcome of patients with acute type b aortic dissection," *European Journal of Vascular and Endovascular Surgery*, vol. 55, no. 3, pp. 385-391, 2018.
- [21] M. M. Marrocco-Trischitta, M. Alaidroos, R. M. Romarowski, F. Secchi, P. Righini, M. Glauber, and G. Nano, "Geometric pattern of proximal landing zones for thoracic endovascular aortic repair in the bovine arch variant," *European Journal of Vascular and Endovascular Surgery*, vol. 59, no. 5, pp. 808-816, 2020.
- [22] M. M. Marrocco-Trischitta, R. M. Romarowski, M. Alaidroos, F. Sturla, M. Glauber, and G. Nano, "Computational fluid dynamics modeling of proximal landing zones for thoracic endovascular aortic repair in the bovine arch variant," *Annals of Vascular Surgery*, vol. 69, pp. 413-417, 2020.
- [23] M. M. Marrocco-Trischitta, H. W. de Beaufort, F. Secchi, T. M. van Bakel, M. Ranucci, J. A. Van Herwaarden, F. L. Moll, and S. Trimarchi, "A geometric reappraisal of proximal landing zones for thoracic endovascular aortic repair according to aortic arch types," *Journal of vascular surgery*, vol. 65, no. 6, pp. 1584-1590, 2017.
- [24] J. Parodi, R. Berguer, P. Carrascosa, K. Khanafer, C. Capunay, and E. Wizauer, "Sources of error in the measurement of aortic diameter in computed tomography scans," *Journal of vascular surgery*, vol. 59, no. 1, pp. 74-79, 2014.
- [25] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghahfarokan, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, vol. 42, pp. 60-88, 2017.

- [26] S. Bakas, M. Reyes, A. Jakab, S. Bauer, M. Rempfler, A. Crimi, R. T. Shinohara, C. Berger, S. M. Ha, M. Rozycki, *et al.*, "Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge," *arXiv preprint arXiv:1811.02629*, 2018.
- [27] C. Tobon-Gomez, A. J. Geers, J. Peters, J. Weese, K. Pinto, R. Karim, M. Ammar, A. Daoudi, J. Margeta, Z. Sandoval, *et al.*, "Benchmark for algorithms segmenting the left atrium from 3d ct and mri datasets," *IEEE transactions on medical imaging*, vol. 34, no. 7, pp. 1460–1473, 2015.
- [28] P. Bilic, P. Christ, H. B. Li, E. Vorontsov, A. Ben-Cohen, G. Kaissis, A. Szeskin, C. Jacobs, G. E. H. Mamani, G. Chartrand, *et al.*, "The liver tumor segmentation benchmark (lits)," *Medical Image Analysis*, vol. 84, p. 102680, 2023.
- [29] A. L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze, *et al.*, "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *arXiv preprint arXiv:1902.09063*, 2019.
- [30] P. A. Yushkevich and G. Gerig, "Itk-snap: an intractive medical image segmentation tool to meet the need for expert-guided segmentation of complex medical images," *IEEE pulse*, vol. 8, no. 4, pp. 54–57, 2017.
- [31] R. Kikinis, S. D. Pieper, and K. G. Vosburgh, "3d slicer: a platform for subject-specific image analysis, visualization, and clinical support," in *Intraoperative imaging and image-guided therapy*, pp. 277–289, Springer, 2014.
- [32] D. Shen, T. Liu, T. M. Peters, L. H. Staib, C. Essert, S. Zhou, P.-T. Yap, and A. Khan, *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II*, vol. 11765. Springer Nature, 2019.
- [33] F. Isensee, J. Petersen, A. Klein, D. Zimmerer, P. F. Jaeger, S. Kohl, J. Wasserthal, G. Koehler, T. Norajitra, S. Wirkert, *et al.*, "nnu-net: Self-adapting framework for u-net-based medical image segmentation," *arXiv preprint arXiv:1809.10486*, 2018.
- [34] M. Berman, A. R. Triki, and M. B. Blaschko, "The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4413–4421, 2018.
- [35] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.
- [36] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *ACM siggraph computer graphics*, vol. 21, no. 4, pp. 163–169, 1987.

- [37] L. Antiga and D. A. Steinman, "Robust and objective decomposition and mapping of bifurcating vessels," *IEEE transactions on medical imaging*, vol. 23, no. 6, pp. 704–713, 2004.
- [38] R. Izzo, D. Steinman, S. Manini, and L. Antiga, "The vascular modeling toolkit: a python library for the analysis of tubular structures in medical images," *Journal of Open Source Software*, vol. 3, no. 25, p. 745, 2018.
- [39] M. Reuter, F.-E. Wolter, and N. Peinecke, "Laplace–beltrami spectra as 'shape-dna' of surfaces and solids," *Computer-Aided Design*, vol. 38, no. 4, pp. 342–366, 2006.
- [40] H. Zhang, O. Van Kaick, and R. Dyer, "Spectral mesh processing," in *Computer graphics forum*, pp. 1865–1894, Wiley Online Library, 2010.
- [41] A. Bône, M. Louis, B. Martin, and S. Durrleman, "Deformetrica 4: an open-source software for statistical shape analysis," in *International Workshop on Shape in Medical Imaging*, pp. 3–13, Springer, 2018.
- [42] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," in *Computer graphics forum*, pp. 1383–1392, Wiley Online Library, 2009.
- [43] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*. MIT press, 2022.
- [44] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, "Comparing images using the hausdorff distance," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 15, no. 9, pp. 850–863, 1993.
- [45] J. M. Noothout, B. D. De Vos, J. M. Wolterink, and I. Išgum, "Automatic segmentation of thoracic aorta segments in low-dose chest ct," in *Medical Imaging 2018: Image Processing*, vol. 10574, pp. 446–451, SPIE, 2018.
- [46] A. Fantazzini, M. Esposito, A. Finotello, F. Auricchio, B. Pane, C. Basso, G. Spinella, and M. Conti, "3d automatic segmentation of aortic computed tomography angiography combining multi-view 2d convolutional neural networks," *Cardiovascular engineering and technology*, vol. 11, no. 5, pp. 576–586, 2020.
- [47] S. Pramod Kumar, M. V. Latte, and S. K. Siri, "Volumetric lung nodule segmentation in thoracic ct scan using freehand sketch," *IET Image Processing*, vol. 14, no. 14, pp. 3456–3462, 2020.
- [48] C. Malone, T. Urbania, S. Crook, and M. Hope, "Bovine aortic arch: a novel association with thoracic aortic dilation," *Clinical radiology*, vol. 67, no. 1, pp. 28–31, 2012.

- [49] M. M. Marrocco-Trischitta, R. M. Romarowski, H. W. De Beaufort, M. Conti, R. Vitale, F. Secchi, F. Auricchio, and S. Trimarchi, "The modified arch landing areas nomenclature identifies hostile zones for endograft deployment: a confirmatory biomechanical study in patients treated by thoracic endovascular aortic repair," *European Journal of Cardio-Thoracic Surgery*, vol. 55, no. 5, pp. 990–997, 2019.
- [50] M. B. Leon, C. R. Smith, M. J. Mack, R. R. Makkar, L. G. Svensson, S. K. Kodali, V. H. Thourani, E. M. Tuzcu, D. C. Miller, H. C. Herrmann, *et al.*, "Transcatheter or surgical aortic-valve replacement in intermediate-risk patients," *New England Journal of Medicine*, vol. 374, no. 17, pp. 1609–1620, 2016.
- [51] M. J. Mack, M. B. Leon, C. R. Smith, D. C. Miller, J. W. Moses, E. M. Tuzcu, J. G. Webb, P. S. Douglas, W. N. Anderson, E. H. Blackstone, *et al.*, "5-year outcomes of transcatheter aortic valve replacement or surgical aortic valve replacement for high surgical risk patients with aortic stenosis (partner 1): a randomised controlled trial," *The Lancet*, vol. 385, no. 9986, pp. 2477–2484, 2015.
- [52] A. Khosravi and O. Wendler, "Tavi 2018: from guidelines to practice," *Eur J Cardiol Pract*, vol. 15, p. 29, 2018.
- [53] P. Breitbart, J. Minners, M. Hein, H. Schröfel, F.-J. Neumann, and P. Ruile, "Implantation depth and its influence on complications after tavi with self-expanding valves," *The International Journal of Cardiovascular Imaging*, vol. 37, no. 10, pp. 3081–3092, 2021.
- [54] S. Bhushan, X. Huang, Y. Li, S. He, L. Mao, W. Hong, and Z. Xiao, "Paravalvular leak after transcatheter aortic valve implantation its incidence, diagnosis, clinical implications, prevention, management, and future perspectives: a review article," *Current Problems in Cardiology*, vol. 47, no. 10, p. 100957, 2022.
- [55] P. Astudillo, P. Mortier, J. Bosmans, O. De Backer, P. de Jaegere, F. Iannaccone, M. De Beule, and J. Dambre, "Automatic detection of the aortic annular plane and coronary ostia from multidetector computed tomography," *Journal of Interventional Cardiology*, vol. 2020, 2020.
- [56] A. G. Cerillo, M. Mariani, S. Berti, and M. Glauber, "Sizing the aortic annulus," *Annals of cardiothoracic surgery*, vol. 1, no. 2, p. 245, 2012.
- [57] R. Gorla, F. De Marco, A. Garatti, G. Bianchi, A. Popolo Rubbio, E. Acerbi, M. Casenghi, P. Spagnolo, N. Brambilla, L. Testa, *et al.*, "Impact of aortic angle on transcatheter aortic valve implantation outcome with evolut-r, portico, and acurate-neo," *Catheterization and Cardiovascular Interventions*, vol. 97, no. 1, pp. E135–E145, 2021.
- [58] C. J. Schultz, A. Moelker, N. Piazza, A. Tzikas, A. Otten, R. J. Nuis, L. A. Neefjes, R. J. Van Geuns, P. De Feyter, G. Krestin, *et al.*, "Three dimensional evaluation of the aortic annulus using multislice computer tomography: are manufacturer's

- guidelines for sizing for percutaneous aortic valve replacement helpful?," *European heart journal*, vol. 31, no. 7, pp. 849–856, 2010.
- [59] M. Elattar, E. Wiegerinck, F. van Kesteren, L. Dubois, N. Planken, E. Vanbavel, J. Baan, and H. Marquering, "Automatic aortic root landmark detection in cta images for preprocedural planning of transcatheter aortic valve implantation," *The international journal of cardiovascular imaging*, vol. 32, pp. 501–511, 2016.
- [60] V. Delgado, A. C. Ng, J. D. Schuijf, F. Van Der Kley, M. Shanks, L. F. Tops, N. R. van de Veire, A. De Roos, L. J. Kroft, M. J. Schalij, *et al.*, "Automated assessment of the aortic root dimensions with multidetector row computed tomography," *The Annals of thoracic surgery*, vol. 91, no. 3, pp. 716–723, 2011.
- [61] F. Lalys, S. Esneault, M. Castro, L. Royer, P. Haigron, V. Auffret, and J. Tomasi, "Automatic aortic root segmentation and anatomical landmarks detection for tavi procedure planning," *Minimally invasive therapy & allied technologies*, vol. 28, no. 3, pp. 157–164, 2019.
- [62] P. G. Tahoces, R. Varela, and J. M. Carreira, "Deep learning method for aortic root detection," *Computers in Biology and Medicine*, vol. 135, p. 104533, 2021.
- [63] Y. Zheng, M. John, R. Liao, J. Boese, U. Kirschstein, B. Georgescu, S. K. Zhou, J. Kempfert, T. Walther, G. Brockmann, *et al.*, "Automatic aorta segmentation and valve landmark detection in c-arm ct: application to aortic valve implantation," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010: 13th International Conference, Beijing, China, September 20–24, 2010, Proceedings, Part I 13*, pp. 476–483, Springer, 2010.
- [64] S. Saitta, F. Sturla, A. Caimi, A. Riva, M. C. Palumbo, G. Nano, E. Votta, A. D. Corte, M. Glauber, D. Chiappino, *et al.*, "A deep learning-based and fully automated pipeline for thoracic aorta geometric analysis and planning for endovascular repair from computed tomography," *Journal of Digital Imaging*, vol. 35, no. 2, pp. 226–239, 2022.
- [65] P. A. Yushkevich, Y. Gao, and G. Gerig, "Itk-snap: An interactive tool for semi-automatic segmentation of multi-modality biomedical images," in *2016 38th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pp. 3342–3345, IEEE, 2016.
- [66] W. J. Schroeder, L. S. Avila, and W. Hoffman, "Visualizing with vtk: a tutorial," *IEEE Computer graphics and applications*, vol. 20, no. 5, pp. 20–27, 2000.
- [67] K. Marstal, F. Berendsen, M. Staring, and S. Klein, "Simpleelastix: A user-friendly, multi-lingual library for medical image registration," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 134–142, 2016.
- [68] F. Maes, D. Vandermeulen, and P. Suetens, "Medical image registration using mutual information," *Proceedings of the IEEE*, vol. 91, no. 10, pp. 1699–1722, 2003.

- [69] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual u-net," *IEEE Geoscience and Remote Sensing Letters*, vol. 15, no. 5, pp. 749–753, 2018.
- [70] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- [71] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, "Elastix: a toolbox for intensity-based medical image registration," *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 196–205, 2009.
- [72] A. Hatamizadeh, Y. Tang, V. Nath, D. Yang, A. Myronenko, B. Landman, H. R. Roth, and D. Xu, "Unetr: Transformers for 3d medical image segmentation," in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 574–584, 2022.
- [73] M. J. Cardoso, W. Li, R. Brown, N. Ma, E. Kerfoot, Y. Wang, B. Murrey, A. Myronenko, C. Zhao, D. Yang, *et al.*, "Monai: An open-source framework for deep learning in healthcare," *arXiv preprint arXiv:2211.02701*, 2022.
- [74] M. Yeung, E. Sala, C.-B. Schoenlieb, and L. Rundo, "Unified focal loss: Generalising dice and cross entropy-based losses to handle class imbalanced medical image segmentation," *Computerized Medical Imaging and Graphics*, vol. 95, p. 102026, 2022.
- [75] S. Valette, J. M. Chassery, and R. Prost, "Generic remeshing of 3d triangular meshes with metric-dependent discrete voronoi diagrams," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 2, pp. 369–381, 2008.
- [76] N. Bettinger, O. K. Khalique, J. M. Krepp, N. B. Hamid, D. J. Bae, T. C. Pulerwitz, M. Liao, R. T. Hahn, T. P. Vahl, T. M. Nazif, *et al.*, "Practical determination of aortic valve calcium volume score on contrast-enhanced computed tomography prior to transcatheter aortic valve replacement and impact on paravalvular regurgitation: elucidating optimal threshold cutoffs," *Journal of Cardiovascular Computed Tomography*, vol. 11, no. 4, pp. 302–308, 2017.
- [77] M. A. Elattar, E. Wiegerinck, R. Planken, E. VanBavel, H. C. van Assen, J. Baan, and H. A. Marquering, "Automatic segmentation of the aortic root in ct angiography of candidate patients for transcatheter aortic valve implantation," *Medical & biological engineering & computing*, vol. 52, pp. 611–618, 2014.
- [78] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested u-net architecture for medical image segmentation," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pp. 3–11, Springer, 2018.

- [79] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2017.
- [80] A. AlAzri, K. Mok, J. Chankowsky, M. Mullah, and J. Marcoux, "Placement accuracy of external ventricular drain when comparing freehand insertion to neuronavigation guidance in severe traumatic brain injury," *Acta neurochirurgica*, vol. 159, no. 8, pp. 1399–1411, 2017.
- [81] T. J. Wilson, W. R. Stetler, W. N. Al-Holou, and S. E. Sullivan, "Comparison of the accuracy of ventricular catheter placement using freehand placement, ultrasonic guidance, and stereotactic neuronavigation," *Journal of neurosurgery*, vol. 119, no. 1, pp. 66–70, 2013.
- [82] M. Mahan, R. F. Spetzler, and P. Nakaji, "Electromagnetic stereotactic navigation for external ventricular drain placement in the intensive care unit," *Journal of Clinical Neuroscience*, vol. 20, no. 12, pp. 1718–1722, 2013.
- [83] J. A. van Doormaal, T. Fick, M. Ali, M. Köllen, V. van der Kuijp, and T. P. van Doormaal, "Fully automatic adaptive meshing based segmentation of the ventricular system for augmented reality visualization and navigation," *World Neurosurgery*, vol. 156, pp. e9–e24, 2021.
- [84] E. Azimi, Z. Niu, M. Stiber, N. Greene, R. Liu, C. Molina, J. Huang, C.-M. Huang, and P. Kazanzides, "An interactive mixed reality platform for bedside surgical procedures," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 65–75, Springer, 2020.
- [85] Y. Li, X. Chen, N. Wang, W. Zhang, D. Li, L. Zhang, X. Qu, W. Cheng, Y. Xu, W. Chen, *et al.*, "A wearable mixed-reality holographic computer for guiding external ventricular drain insertion at the bedside," *Journal of neurosurgery*, vol. 131, no. 5, pp. 1599–1606, 2018.
- [86] M. Schneider, C. Kunz, A. Pal'a, C. R. Wirtz, F. Mathis-Ullrich, and M. Hlaváč, "Augmented reality–assisted ventriculostomy," *Neurosurgical focus*, vol. 50, no. 1, p. E16, 2021.
- [87] F. von Haxthausen, Y. Chen, and F. Ernst, "Superimposing holograms on real world objects using hololens 2 and its depth camera," *Current Directions in Biomedical Engineering*, vol. 7, no. 1, pp. 111–115, 2021.
- [88] S. Patrick, N. P. Birur, K. Gurushanth, A. S. Raghavan, S. Gurudath, *et al.*, "Comparison of gray values of cone-beam computed tomography with hounsfield units of multislice computed tomography: an in vitro study," *Indian Journal of Dental Research*, vol. 28, no. 1, p. 66, 2017.

- [89] D. Ungureanu, F. Bogo, S. Galliani, P. Sama, X. Duan, C. Meekhof, J. Stühmer, T. J. Cashman, B. Tekin, J. L. Schönberger, *et al.*, “Hololens 2 research mode as a tool for computer vision research,” *arXiv preprint arXiv:2008.11239*, 2020.
- [90] S. Katz, A. Tal, and R. Basri, “Direct visibility of point sets,” in *ACM SIGGRAPH 2007 papers*, pp. 24–es, 2007.
- [91] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *kdd*, vol. 96, pp. 226–231, 1996.
- [92] Q.-Y. Zhou, J. Park, and V. Koltun, “Fast global registration,” in *European conference on computer vision*, pp. 766–782, Springer, 2016.
- [93] Y. Chen and G. Medioni, “Object modelling by registration of multiple range images,” *Image and vision computing*, vol. 10, no. 3, pp. 145–155, 1992.
- [94] Q.-Y. Zhou, J. Park, and V. Koltun, “Open3d: A modern library for 3d data processing,” *arXiv preprint arXiv:1801.09847*, 2018.
- [95] K. S. Arun, T. S. Huang, and S. D. Blostein, “Least-squares fitting of two 3-d point sets,” *IEEE Transactions on pattern analysis and machine intelligence*, no. 5, pp. 698–700, 1987.
- [96] J. A. Elefteriades, “Natural history of thoracic aortic aneurysms: indications for surgery, and surgical versus nonsurgical risks,” *The Annals of thoracic surgery*, vol. 74, no. 5, pp. S1877–S1880, 2002.
- [97] C. Catalano, V. Agnese, G. Gentile, G. M. Raffa, M. Pilato, and S. Pasta, “Atlas-based evaluation of hemodynamic in ascending thoracic aortic aneurysms,” *Applied Sciences*, vol. 12, no. 1, p. 394, 2021.
- [98] J. J. Yeung, H. J. Kim, T. A. Abbruzzese, I. E. Vignon-Clementel, M. T. Draney-Blomme, K. K. Yeung, I. Perakash, R. J. Herfkens, C. A. Taylor, and R. L. Dalman, “Aortoiliac hemodynamic and morphologic adaptation to chronic spinal cord injury,” *Journal of vascular surgery*, vol. 44, no. 6, pp. 1254–1265, 2006.
- [99] S. Chien, S. Li, and J. Y. Shyy, “Effects of mechanical forces on signal transduction and gene expression in endothelial cells,” *Hypertension*, vol. 31, no. 1, pp. 162–169, 1998.
- [100] S. Pirola, B. Guo, C. Menichini, S. Saitta, W. Fu, Z. Dong, and X. Y. Xu, “4-D flow MRI-based computational analysis of blood flow in patient-specific aortic dissection,” *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 12, pp. 3411–3419, 2019.

- [101] V. Mendez, M. Di Giuseppe, and S. Pasta, "Comparison of hemodynamic and structural indices of ascending thoracic aortic aneurysm as predicted by 2-way fsi, cfd rigid wall simulation and patient-specific displacement-based fea," *Computers in biology and medicine*, vol. 100, pp. 221–229, 2018.
- [102] Z. Li and C. Kleinstreuer, "Blood flow and structure interactions in a stented abdominal aortic aneurysm model," *Medical engineering & physics*, vol. 27, no. 5, pp. 369–382, 2005.
- [103] M. Peirlinck, F. S. Costabal, J. Yao, J. Guccione, S. Tripathy, Y. Wang, D. Ozturk, P. Segars, T. Morrison, S. Levine, *et al.*, "Precision medicine in human heart modeling," *Biomechanics and modeling in mechanobiology*, vol. 20, no. 3, pp. 803–831, 2021.
- [104] U. Morbiducci, R. Ponzini, D. Gallo, C. Bignardi, and G. Rizzo, "Inflow boundary conditions for image-based computational hemodynamics: impact of idealized versus measured velocity profiles in the human aorta," *Journal of biomechanics*, vol. 46, no. 1, pp. 102–109, 2013.
- [105] S. Pirola, O. Jarral, D. O'Regan, G. Asimakopoulos, J. Anderson, J. Pepper, T. Athanasiou, and X. Xu, "Computational study of aortic hemodynamics for patients with an abnormal aortic valve: The importance of secondary flow at the ascending aorta inlet," *APL bioengineering*, vol. 2, no. 2, p. 026101, 2018.
- [106] P. Youssefi, A. Gomez, C. Arthurs, R. Sharma, M. Jahangiri, and C. Alberto Figueroa, "Impact of patient-specific inflow velocity profile on hemodynamics of the thoracic aorta," *Journal of biomechanical engineering*, vol. 140, no. 1, 2018.
- [107] C. H. Armour, B. Guo, S. Pirola, S. Saitta, Y. Liu, Z. Dong, and X. Y. Xu, "The influence of inlet velocity profile on predicted flow in type b aortic dissection," *Biomechanics and modeling in mechanobiology*, vol. 20, no. 2, pp. 481–490, 2021.
- [108] P. Romero, M. Lozano, F. Martínez-Gil, D. Serra, R. Sebastián, P. Lamata, and I. García-Fernández, "Clinically-driven virtual patient cohorts generation: An application to aorta," *Frontiers in Physiology*, p. 1375, 2021.
- [109] A. A. Young and A. F. Frangi, "Computational cardiac atlases: from patient to population and back," *Experimental physiology*, vol. 94, no. 5, pp. 578–596, 2009.
- [110] M. E. Casciaro, D. Craiem, G. Chironi, S. Graf, L. Macron, E. Mousseaux, A. Simon, and R. L. Armentano, "Identifying the principal modes of variation in human thoracic aorta morphology," *Journal of Thoracic Imaging*, vol. 29, no. 4, pp. 224–232, 2014.
- [111] I. T. Jolliffe and J. Cadima, "Principal component analysis: A review and recent developments," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 374, no. 2065, p. 20150202, 2016.

- [112] L. Liang, M. Liu, C. Martin, J. A. Elefteriades, and W. Sun, "A machine learning approach to investigate the relationship between shape features and numerically predicted risk of ascending aortic aneurysm," *Biomechanics and modeling in mechanobiology*, vol. 16, no. 5, pp. 1519–1533, 2017.
- [113] B. Thamsen, P. Yevtushenko, L. Gundelwein, A. A. Setio, H. Lamecker, M. Kelm, M. Schafstedde, T. Heimann, T. Kuehne, and L. Goubergrits, "Synthetic database of aortic morphometry and hemodynamics: overcoming medical imaging data availability," *IEEE Transactions on Medical Imaging*, vol. 40, no. 5, pp. 1438–1449, 2021.
- [114] F. Cosentino, G. M. Raffa, G. Gentile, V. Agnese, D. Bellavia, M. Pilato, and S. Pasta, "Statistical shape analysis of ascending thoracic aortic aneurysm: correlation between shape and biomechanical descriptors," *Journal of Personalized Medicine*, vol. 10, no. 2, p. 28, 2020.
- [115] G. Nannini, A. Caimi, M. C. Palumbo, S. Saitta, L. N. Girardi, M. Gaudino, M. J. Roman, J. W. Weinsaft, and A. Redaelli, "Aortic hemodynamics assessment prior and after valve sparing reconstruction: A patient-specific 4d flow-based fsi model," *Computers in Biology and Medicine*, vol. 135, p. 104581, 2021.
- [116] M. Y. Salmasi, S. Pirola, S. Sasidharan, S. M. Fisichella, A. Redaelli, O. A. Jarral, D. P. O'Regan, A. Y. Oo, J. E. Moore Jr, X. Y. Xu, *et al.*, "High wall shear stress can predict wall degradation in ascending aortic aneurysms: an integrated biomechanics study," *Frontiers in Bioengineering and Biotechnology*, p. 935, 2021.
- [117] S. Pirola, Z. Cheng, O. Jarral, D. O'Regan, J. Pepper, T. Athanasiou, and X. Xu, "On the choice of outlet boundary conditions for patient-specific analysis of aortic flow using computational fluid dynamics," *Journal of biomechanics*, vol. 60, pp. 15–21, 2017.
- [118] R. Lorenz, J. Bock, J. Snyder, J. G. Korvink, B. A. Jung, and M. Markl, "Influence of eddy current, Maxwell and gradient field corrections on 3D flow visualization of 3D CINE PC-MRI data," *Magnetic Resonance in Medicine*, vol. 72, no. 1, pp. 33–40, 2014.
- [119] S. Saitta, B. Guo, S. Pirola, C. Menichini, D. Guo, Y. Shan, Z. Dong, X. Y. Xu, and W. Fu, "Qualitative and quantitative assessments of blood flow on tears in type b aortic dissection with different morphologies," *Frontiers in Bioengineering and Biotechnology*, vol. 9, 2021.
- [120] L. Liang, W. Mao, and W. Sun, "A feasibility study of deep learning for predicting hemodynamics of human thoracic aorta," *Journal of Biomechanics*, vol. 99, p. 109544, 2020.
- [121] P. F. Davies, "Hemodynamic shear stress and the endothelium in cardiovascular pathophysiology," *Nature clinical practice Cardiovascular medicine*, vol. 6, no. 1, pp. 16–26, 2009.

- [122] P. Bagan, R. Vidal, E. Martinod, M.-D. Destable, B. Tremblay, J. L. Dumas, and J. F. Azorin, “Cerebral ischemia during carotid artery cross-clamping: predictive value of phase-contrast magnetic resonance imaging,” *Annals of vascular surgery*, vol. 20, no. 6, pp. 747–752, 2006.
- [123] W. C. Members, E. M. Isselbacher, O. Preventza, J. Hamilton Black III, J. G. Augoustides, A. W. Beck, M. A. Bolen, A. C. Braverman, B. E. Bray, M. M. Brown-Zimmerman, *et al.*, “2022 ACC/AHA Guideline for the diagnosis and management of aortic disease: a report of the American Heart Association/American College of Cardiology Joint Committee on Clinical Practice Guidelines,” *Journal of the American College of Cardiology*, vol. 80, no. 24, pp. e223–e393, 2022.
- [124] C. Bertoglio, R. Núñez, F. Galarce, D. Nordsletten, and A. Osses, “Relative pressure estimation from velocity measurements in blood flows: State-of-the-art and new approaches,” *International journal for numerical methods in biomedical engineering*, vol. 34, no. 2, p. e2925, 2018.
- [125] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben, “4D flow MRI,” *Journal of Magnetic Resonance Imaging*, vol. 36, no. 5, pp. 1015–1036, 2012.
- [126] V. Vishnevskiy, J. Walheim, and S. Kozerke, “Deep variational network for rapid 4D flow MRI reconstruction,” *Nature Machine Intelligence*, vol. 2, no. 4, pp. 228–235, 2020.
- [127] A. J. Barker, M. Markl, J. Bürk, R. Lorenz, J. Bock, S. Bauer, J. Schulz-Menger, and F. von Knobelsdorff-Brenkenhoff, “Bicuspid aortic valve is associated with altered wall shear stress in the ascending aorta,” *Circulation: Cardiovascular Imaging*, vol. 5, no. 4, pp. 457–466, 2012.
- [128] C. Trenti, M. Ziegler, N. Bjarnegård, T. Ebbers, M. Lindenberg, and P. Dyverfeldt, “Wall shear stress and relative residence time as potential risk factors for abdominal aortic aneurysms in males: a 4D flow cardiovascular magnetic resonance case–control study,” *Journal of Cardiovascular Magnetic Resonance*, vol. 24, no. 1, pp. 1–12, 2022.
- [129] M. M. Bissell, A. T. Hess, L. Biasioli, S. J. Glaze, M. Loudon, A. Pitcher, A. Davis, B. Prendergast, M. Markl, A. J. Barker, *et al.*, “Aortic dilation in bicuspid aortic valve disease: flow pattern is a major contributor and differs with valve fusion type,” *Circulation: Cardiovascular Imaging*, vol. 6, no. 4, pp. 499–507, 2013.
- [130] J. Bock, A. Frydrychowicz, R. Lorenz, D. Hirtler, A. J. Barker, K. M. Johnson, R. Arnold, H. Burkhardt, J. Hennig, and M. Markl, “In vivo noninvasive 4D pressure difference mapping in the human aorta: phantom comparison and application in healthy volunteers and patients,” *Magnetic resonance in medicine*, vol. 66, no. 4, pp. 1079–1088, 2011.

- [131] S. Saitta, S. Pirola, F. Piatti, E. Votta, F. Lucherini, F. Pluchinotta, M. Carminati, M. Lombardi, C. Geppert, F. Cuomo, *et al.*, “Evaluation of 4d flow mri-based non-invasive pressure assessment in aortic coarctations,” *Journal of biomechanics*, vol. 94, pp. 13–21, 2019.
- [132] G. Reiter, U. Reiter, G. Kovacs, H. Olschewski, and M. Fuchsjäger, “Blood flow vortices along the main pulmonary artery measured with MR imaging for diagnosis of pulmonary hypertension,” *Radiology*, vol. 275, no. 1, pp. 71–79, 2015.
- [133] E. Bollache, A. J. Barker, R. S. Dolan, J. C. Carr, P. van Ooij, R. Ahmadian, A. Powell, J. D. Collins, J. Geiger, and M. Markl, “k-t accelerated aortic 4D flow MRI in under two minutes: feasibility and impact of resolution, k-space sampling patterns, and respiratory navigator gating on hemodynamic measurements,” *Magnetic resonance in medicine*, vol. 79, no. 1, pp. 195–207, 2018.
- [134] A. Rich, L. C. Potter, N. Jin, Y. Liu, O. P. Simonetti, and R. Ahmad, “A Bayesian approach for 4D flow imaging of aortic valve in a single breath-hold,” *Magnetic resonance in medicine*, vol. 81, no. 2, pp. 811–824, 2019.
- [135] D. Kim, H. A. Dyvorne, R. Otazo, L. Feng, D. K. Sodickson, and V. S. Lee, “Accelerated phase-contrast cine MRI using k-t SPARSE-SENSE,” *Magnetic resonance in medicine*, vol. 67, no. 4, pp. 1054–1064, 2012.
- [136] G. Valvano, N. Martini, A. Huber, C. Santelli, C. Binter, D. Chiappino, L. Landini, and S. Kozerke, “Accelerating 4 D flow MRI by exploiting low-rank matrix structure and hadamard sparsity,” *Magnetic resonance in medicine*, vol. 78, no. 4, pp. 1330–1341, 2017.
- [137] L. E. Ma, M. Markl, K. Chow, H. Huh, C. Forman, A. Vali, A. Greiser, J. Carr, S. Schnell, A. J. Barker, *et al.*, “Aortic 4D flow MRI in 2 minutes using compressed sensing, respiratory controlled adaptive k-space reordering, and inline reconstruction,” *Magnetic resonance in medicine*, vol. 81, no. 6, pp. 3675–3690, 2019.
- [138] J. Jiang, P. Kokeny, W. Ying, C. Magnano, R. Zivadinov, and E. M. Haacke, “Quantifying errors in flow measurement using phase contrast magnetic resonance imaging: comparison of several boundary detection methods,” *Magnetic resonance imaging*, vol. 33, no. 2, pp. 185–193, 2015.
- [139] P. Dyverfeldt, M. Bissell, A. J. Barker, A. F. Bolger, C.-J. Carlhäll, T. Ebbers, C. J. Francios, A. Frydrychowicz, J. Geiger, D. Giese, *et al.*, “4D flow cardiovascular magnetic resonance consensus statement,” *Journal of Cardiovascular Magnetic Resonance*, vol. 17, no. 1, pp. 1–19, 2015.
- [140] A. Cárdenas-Blanco, C. Tejos, P. Irarrazaval, and I. Cameron, “Noise in magnitude magnetic resonance images,” *Concepts in Magnetic Resonance Part A: An Educational Journal*, vol. 32, no. 6, pp. 409–416, 2008.

- [141] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7462–7473, 2020.
- [142] S. W. Funke, M. Nordaas, Ø. Evju, M. S. Alnæs, and K. A. Mardal, "Variational data assimilation for transient blood flow simulations: Cerebral aneurysms as an illustrative example," *International journal for numerical methods in biomedical engineering*, vol. 35, no. 1, p. e3152, 2019.
- [143] T. S. Koltukluoğlu and P. J. Blanco, "Boundary control in computational haemodynamics," *Journal of Fluid Mechanics*, vol. 847, pp. 329–364, 2018.
- [144] A. Kontogiannis, S. V. Elgersma, A. J. Sederman, and M. P. Juniper, "Joint reconstruction and segmentation of noisy velocity images as an inverse Navier–Stokes problem," *Journal of Fluid Mechanics*, vol. 944, p. A40, 2022.
- [145] A. Kontogiannis and M. P. Juniper, "Physics-informed compressed sensing for PC-MRI: an inverse Navier-Stokes problem," *IEEE Transactions on Image Processing*, vol. 32, pp. 281–294, 2022.
- [146] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanaawala, and M. Lustig, "Robust 4D flow denoising using divergence-free wavelet transform," *Magnetic resonance in medicine*, vol. 73, no. 2, pp. 828–842, 2015.
- [147] J. Busch, D. Giese, L. Wissmann, and S. Kozerke, "Reconstruction of divergence-free velocity fields from cine 3D phase-contrast flow measurements," *Magnetic resonance in medicine*, vol. 69, no. 1, pp. 200–210, 2013.
- [148] M. F. Fathi, I. Perez-Raya, A. Baghaie, P. Berg, G. Janiga, A. Arzani, and R. M. D'Souza, "Super-resolution and denoising of 4D-flow MRI using physics-informed deep neural nets," *Computer Methods and Programs in Biomedicine*, vol. 197, p. 105729, 2020.
- [149] A. L. Marsden, A. J. Bernstein, V. M. Reddy, S. C. Shadden, R. L. Spilker, F. P. Chan, C. A. Taylor, and J. A. Feinstein, "Evaluation of a novel Y-shaped extracardiac Fontan baffle using computational fluid dynamics," *The Journal of Thoracic and Cardiovascular Surgery*, vol. 137, no. 2, pp. 394–403, 2009.
- [150] P. D. Morris, A. Narracott, H. von Tengg-Kobligk, D. A. S. Soto, S. Hsiao, A. Lungu, P. Evans, N. W. Bressloff, P. V. Lawford, D. R. Hose, *et al.*, "Computational fluid dynamics modelling in cardiovascular medicine," *Heart*, vol. 102, no. 1, pp. 18–28, 2016.
- [151] J. Xiang, M. Tremmel, J. Kolega, E. I. Levy, S. K. Natarajan, and H. Meng, "Newtonian viscosity model could overestimate wall shear stress in intracranial aneurysm domes and underestimate rupture risk," *Journal of neurointerventional surgery*, vol. 4, no. 5, pp. 351–357, 2012.

- [152] S. N. Doost, L. Zhong, B. Su, and Y. S. Morsi, "The numerical analysis of non-Newtonian blood flow in human patient-specific left ventricle," *Computer methods and programs in biomedicine*, vol. 127, pp. 232–247, 2016.
- [153] T. Naito, S. Miyachi, N. Matsubara, H. Isoda, T. Izumi, K. Haraguchi, I. Takahashi, K. Ishii, and T. Wakabayashi, "Magnetic resonance fluid dynamics for intracranial aneurysms—comparison with computed fluid dynamics," *Acta neurochirurgica*, vol. 154, no. 6, pp. 993–1001, 2012.
- [154] K. Jain, J. Jiang, C. Strother, and K.-A. Mardal, "Transitional hemodynamics in intracranial aneurysms—comparative velocity investigations with high resolution lattice Boltzmann simulations, normal resolution ANSYS simulations, and MR imaging," *Medical physics*, vol. 43, no. 11, pp. 6186–6198, 2016.
- [155] J. Jiang, K. Johnson, K. Valen-Sendstad, K.-A. Mardal, O. Wieben, and C. Strother, "Flow characteristics in a canine aneurysm model: A comparison of 4D accelerated phase-contrast MR measurements and computational fluid dynamics simulations," *Medical physics*, vol. 38, no. 11, pp. 6300–6312, 2011.
- [156] J. Thewlis, D. Stevens, H. Power, D. Giddings, P. Gowland, and M. Vloeberghs, "4-dimensional local radial basis function interpolation of large, uniformly spaced datasets," *Computer Methods and Programs in Biomedicine*, vol. 228, p. 107235, 2023.
- [157] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [158] X. Jin, S. Cai, H. Li, and G. E. Karniadakis, "NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations," *Journal of Computational Physics*, vol. 426, p. 109951, 2021.
- [159] M. Raissi, A. Yazdani, and G. E. Karniadakis, "Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations," *Science*, vol. 367, no. 6481, pp. 1026–1030, 2020.
- [160] X. Jin, S. Laima, W.-L. Chen, and H. Li, "Time-resolved reconstruction of flow field around a circular cylinder by recurrent neural networks based on non-time-resolved particle image velocimetry measurements," *Experiments in Fluids*, vol. 61, pp. 1–23, 2020.
- [161] J.-X. Wang, J.-L. Wu, and H. Xiao, "Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data," *Physical Review Fluids*, vol. 2, no. 3, p. 034603, 2017.

- [162] S. Wang, Y. Teng, and P. Perdikaris, “Understanding and mitigating gradient flow pathologies in physics-informed neural networks,” *SIAM Journal on Scientific Computing*, vol. 43, no. 5, pp. A3055–A3081, 2021.
- [163] M. Tancik, B. Mildenhall, T. Wang, D. Schmidt, P. P. Srinivasan, J. T. Barron, and R. Ng, “Learned initializations for optimizing coordinate-based neural representations,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2846–2855, 2021.
- [164] J. N. Martel, D. B. Lindell, C. Z. Lin, E. R. Chan, M. Monteiro, and G. Wetzstein, “Acorn: Adaptive coordinate networks for neural scene representation,” *arXiv preprint arXiv:2105.02788*, 2021.
- [165] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, “Fourier features let networks learn high frequency functions in low dimensional domains,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 7537–7547, 2020.
- [166] E. D. Zhong, T. Bepler, J. H. Davis, and B. Berger, “Reconstructing continuous distributions of 3D protein structure from cryo-EM images,” *arXiv preprint arXiv:1909.05215*, 2019.
- [167] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [168] A. Rahimi and B. Recht, “Random features for large-scale kernel machines,” *Advances in neural information processing systems*, vol. 20, 2007.
- [169] D. Alblas, C. Brune, K. K. Yeung, and J. M. Wolterink, “Going Off-Grid: Continuous Implicit Neural Representations for 3D Vascular Modeling,” in *Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers: 13th International Workshop, STACOM 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Revised Selected Papers*, pp. 79–90, Springer, 2023.
- [170] J. M. Wolterink, J. C. Zwienenberg, and C. Brune, “Implicit neural representations for deformable image registration,” in *International Conference on Medical Imaging with Deep Learning*, pp. 1349–1359, PMLR, 2022.
- [171] S. Petersson, P. Dyverfeldt, and T. Ebbens, “Assessment of the accuracy of mri wall shear stress estimation using numerical simulations,” *Journal of Magnetic Resonance Imaging*, vol. 36, no. 1, pp. 128–138, 2012.
- [172] S. Saitta, L. Maga, C. Armour, E. Votta, D. P. O’Regan, M. Y. Salmasi, T. Athanasiou, J. W. Weinsaft, X. Y. Xu, S. Pirola, *et al.*, “Data-driven generation of 4D velocity profiles in the aneurysmal ascending aorta,” *arXiv preprint arXiv:2211.00551*, 2022.

- [173] E. Ferdian, A. Suinesiaputra, D. J. Dubowitz, D. Zhao, A. Wang, B. Cowan, and A. A. Young, “4DFlowNet: super-resolution 4D flow MRI using deep learning and computational fluid dynamics,” *Frontiers in Physics*, p. 138, 2020.
- [174] H. Ha, G. B. Kim, J. Kweon, S. J. Lee, Y.-H. Kim, D. H. Lee, D. H. Yang, and N. Kim, “Hemodynamic measurement using four-dimensional phase-contrast MRI: quantification of hemodynamic parameters and clinical applications,” *Korean journal of radiology*, vol. 17, no. 4, pp. 445–462, 2016.
- [175] J. Sotelo, L. Dux-Santoy, A. Guala, J. Rodríguez-Palomares, A. Evangelista, C. Sing-Long, J. Urbina, J. Mura, D. E. Hurtado, and S. Uribe, “3D axial and circumferential wall shear stress from 4D flow MRI data using a finite element method and a laplacian approach,” *Magnetic Resonance in Medicine*, vol. 79, no. 5, pp. 2816–2823, 2018.
- [176] F. Piatti, F. Sturla, M. M. Bissell, S. Pirola, M. Lombardi, I. Nesteruk, A. Della Corte, A. C. Redaelli, and E. Votta, “4D flow analysis of BAV-related fluid-dynamic alterations: evidences of wall shear stress alterations in absence of clinically-relevant aortic anatomical remodeling,” *Frontiers in physiology*, vol. 8, p. 441, 2017.
- [177] F. Mut, R. Löhner, A. Chien, S. Tateshima, F. Viñuela, C. Putman, and J. R. Cebal, “Computational hemodynamics framework for the analysis of cerebral aneurysms,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 27, no. 6, pp. 822–839, 2011.
- [178] D. R. Rutkowski, A. Roldán-Alzate, and K. M. Johnson, “Enhancement of cerebrovascular 4D flow MRI velocity fields using machine learning and computational fluid dynamics simulation data,” *Scientific reports*, vol. 11, no. 1, pp. 1–11, 2021.
- [179] S. Shit, J. Zimmermann, I. Ezhov, J. C. Paetzold, A. F. Sanches, C. Pirkl, and B. Menze, “SRflow: Deep learning based super-resolution of 4D-flow MRI data,” *Frontiers in Artificial Intelligence*, p. 171, 2022.
- [180] R. Basri, M. Galun, A. Geifman, D. Jacobs, Y. Kasten, and S. Kritchman, “Frequency bias in neural networks for input of non-uniform density,” in *International Conference on Machine Learning*, pp. 685–694, PMLR, 2020.
- [181] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville, “On the spectral bias of neural networks,” in *International Conference on Machine Learning*, pp. 5301–5310, PMLR, 2019.
- [182] P. Velickovic, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, “Deep graph infomax,” *ICLR (Poster)*, vol. 2, no. 3, p. 4, 2019.

- [183] J. F. Schneiders, R. P. Dwight, and F. Scarano, "Time-supersampling of 3d-piv measurements with vortex-in-cell simulation," *Experiments in Fluids*, vol. 55, pp. 1–15, 2014.
- [184] B. Ummenhofer, L. Prantl, N. Thuerey, and V. Koltun, "Lagrangian fluid simulation with continuous convolutions," in *International Conference on Learning Representations*, 2020.