

Patient Specific Computational Fluid Dynamic in the Left Atrium.
Atrial Fibrillation vs Sinus Rhythm

Savier Yogita Sharda

A thesis

submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Aeronautics & Astronautics

University of Washington

2025

Committee:

Dana Dabiri

Juan C. del Alamo

Program Authorized to Offer Degree:

Aeronautics & Astronautics

© Copyright 2025

Savier Yogita Sharda

University of Washington

Abstract

Patient Specific Computational Fluid Dynamic in the Left Atrium.

Atrial Fibrillation vs Sinus Rhythm

Savier Yogita Sharda

Chair of the Supervisory Committee:

Dana Dabiri

William E. Boeing Department of Aeronautics & Astronautics

Atrial fibrillation (AF) is the most common sustained arrhythmia and a leading cause of stroke, yet the direct hemodynamic consequences of AF remain difficult to isolate in clinical practice. Imaging captures anatomy, and ECGs record rhythm, but neither alone can resolve the full, three-dimensional, time-resolved blood flow patterns that link AF to clot formation.

This study develops a patient-specific computational framework to address this limitation. Critically, we analyze the same patients under both sinus rhythm (SR) and AF, enabling direct comparison of how rhythm alone alters atrial flow. Using imaging and ECG data, we reconstruct left atrial (LA) geometries, incorporate motion across the cardiac cycle, and simulate blood flow to evaluate changes in transport dynamics and stasis between the two conditions.

The comparisons reveal clear hemodynamic signatures of AF: reduced atrial emptying, irregular contraction, and disrupted flow organization, all of which contribute to conditions favorable for thrombus formation. At the same time, differences in atrial appendage morphology and rhythm-specific electrical patterns highlight the patient-to-patient variability that shapes risk.

By directly contrasting SR and AF within the same individuals, this framework bridges a critical gap between clinical observation and mechanistic understanding. It shows how computational modeling can uncover hidden flow dynamics inaccessible to imaging alone and offers a foundation for personalized assessment of AF-related stroke risk.

TABLE OF CONTENTS

List of Figures	iii
List of Tables	vii
Chapter 1. Introduction	10
1.1 Clinical Background and Motivation.....	10
1.2 Cardiac Electrical Activity and Its Representation on the Electrocardiogram	12
1.3 Hemodynamic Alterations in Atrial Fibrillation (AF)	16
1.3.1 Blood Flow Differences in Atrial Fibrillation (AF) and Sinus Rhythm (SR).....	18
1.3.2 Shear Stress Regulation of Endothelial Function	20
1.3.3 Molecular and Cellular changes in Response to Altered Hemodynamics	21
1.3.4 Virchow’s Triad	22
1.4 CFD in understanding hemodynamics.....	24
1.5 Risk Stratification	26
1.6 Limitations	28
1.6.1 Project Contributions	29
Chapter 2. Methodology	31
2.1 Geometry Creation.....	32
2.1.1 Generation of patient-specific models	32
2.1.2 Manual Refinement and alignment of anatomical features	35
2.2 Model preparation for simulation	36
2.2.1 Smoothing of imaging data	38

2.2.2	Meshing of Anatomical Geometry.....	40
2.2.3	Motion Registration	41
2.2.4	Anatomical Boundary Conditions.....	42
2.3	CFD Simulation	44
2.3.1	Immersed Boundary Flow Framework	44
2.3.2	TUCAN Solver and Numerical Discretization	45
2.3.3	Residence Time Transport Equation.....	46
2.3.4	Simulation Setup and Flow Convergence.....	47
Chapter 3.	Results	49
3.1	Geometry Creation.....	53
3.2	Model preparation for simulation	58
3.3	CFD Simulation	70
3.4	Post Processing	71
Chapter 4.	Conclusion.....	85
Bibliography		87
Appendix A		92

LIST OF FIGURES

Figure 1.1 Illustration of the R–R Interval on an ECG, Representing One Complete Cardiac Cycle.	12
Figure 1.2 The Cardiac Conduction System, showing the pathway of electrical impulses.	14
Figure 1.3 Normal vs Abnormal Electrical Pathways. Normal conduction follows a clear path from the SA to AV node, while atrial fibrillation involves disorganized electrical signals within the atria.	14
Figure 1.4 ECG Comparison of Normal Sinus Rhythm and Atrial Fibrillation [5].	15
Figure 1.5 Anatomic Variants of LAA Morphology	18
Figure 1.6 Schematic representation of the LA during AF, illustrating the three components of Virchow’s triad [15].	23
Figure 2.1 Methodology flowchart outlining the steps from geometry creation and model preparation to CFD simulation and post-processing.....	32
Figure 2.2 3D nnU-Net Architecture for Volumetric Image Segmentation	34
Figure 2.3 Pulmonary vein (PV) in LA. Shaded areas show deviations from the typical four-vein pattern. A: Typical (4 PVs). B: Short common left trunk. C: Long common left trunk. D: Right middle lobe PV. E: Two right middle PVs. F: Right middle plus extra right “upper” PV [32].	38
Figure 2.4 Pulmonary vein inlets and mitral valve outlet with cork structures for boundary condition setup.	42
Figure 2.5 Non-Physiological Inlet/Outlet Placement extending beyond geometry boundaries.	43
Figure 3.1 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 1.....	50
Figure 3.2 DICOM Images for Atrial Fibrillation (AF) – Case 1.....	50
Figure 3.3 DICOM Images for Sinus Rhythm (SR) – Case 1	51
Figure 3.4 ECG Comparison Between AF and SR – Case 1	51

Figure 3.5 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF – Case 1	54
Figure 3.6 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 1	54
Figure 3.7 nnU-Net-based 3D segmentation from stacked multiplanar CT slices	55
Figure 3.8 Segmentations using nnU-Net AF and SR – Case 1	56
Figure 3.9 Automated nnU-Net segmentation vs. post manual refinement, showing improved anatomical accuracy AF – Case 1	57
Figure 3.10 Segmentations post Manual Refinement AF and SR – Case 1	58
Figure 3.11 . Maximum and minimum left atrial (LA) volumes during the cardiac cycle for each case in sinus rhythm (SR, green) and atrial fibrillation (AF, red).	60
Figure 3.12 Left atrial (LA) emptying fractions for each case in sinus rhythm (SR, green) and atrial fibrillation (AF, red)	60
Figure 3.13 Left atrial (LA) volume over time from 3 seconds onward. SR shows greater variation than AF, indicating stronger contractile function.	62
Figure 3.14 Final cardiac cycle showing reduced and abnormal volume changes in AF, reflecting impaired atrial contraction.	63
Figure 3.15 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.	64
Figure 3.16 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.	65
Figure 3.17 Mesh AF and SR – Case 1	66
Figure 3.18 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 1	68
Figure 3.19 Inlets / outlet which are for the CFD simulation for AF and SR – Case 1....	69
Figure 3.20 Velocity contours of the LA for Case 1 AF	72
Figure 3.21 Velocity contours of the LA for Case 1 SR.....	73
Figure 3.22 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.	75

Figure 3.23 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.	76
Figure 3.24 Residence Time in the LA for Case 1 AF	77
Figure 3.25 Residence time in the LA for Case 1 SR.....	78
Figure 3.26 LAA volume AF and SR – Case 1	79
Figure 3.27 Average Residence Time compared in different regions of the LAA for	80
Figure 3.28 Average Residence Time in LAA for AF and SR – Case 1.	81
Figure 3.29 Average Residence Time compared in different regions of the LAA in AF and SR – Case 1.....	82
Figure 3.30 Violin plots of residence time distributions for AF and SR cases.....	83
Figure 3.31 Maximum residence time trends for AF and SR cases across all frames.....	84
Figure A.1 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 2.....	92
Figure A.2 DICOM Images for Atrial Fibrillation (AF) – Case 2.....	92
Figure A.3 DICOM Images for Sinus Rhythm (SR) – Case 2.....	92
Figure A.4 ECG Comparison Between AF and SR – Case 2	93
Figure A.5 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 3.....	93
Figure A.6 DICOM Images for Atrial Fibrillation (AF) – Case 3.....	93
Figure A.7 DICOM Images for Sinus Rhythm (SR) – Case 3.....	94
Figure A.8 ECG Comparison Between AF and SR – Case 3	94
Figure A.9 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF – Case 2	94
Figure A.10 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 2.....	95
Figure A.11 Segmentations using nnU-Net AF and SR – Case 2	95
Figure A.12 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF –Case 3.....	96
Figure A.13 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 3.....	96

Figure A.14 Segmentations using nnU-Net AF and SR – Case 3	96
Figure A.15 Mesh AF and SR – Case 2.....	97
Figure A.16 Mesh AF and SR – Case 3.....	97
Figure A.17 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 2	98
Figure A.18 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 3	98
Figure A.19 Inlets / outlet which are for the CFD simulation for AF and SR – Case 2..	99
Figure A.20 Inlets / outlet which are for the CFD simulation for AF and SR – Case 3.	100
Figure A.21 Velocity contours of the LA for Case 2 AF.....	100
Figure A.22 Velocity contours of the LA for Case 2 SR.....	101
Figure A.23 Velocity contours of the LA for Case 3 AF.....	102
Figure A.24 Velocity contours of the LA for Case 3 SR.....	103
Figure A.25 Residence Time in the LA for Case 2 AF	104
Figure A.26 Residence Time in the LA for Case 2 SR.....	105
Figure A.27 Residence Time in the LA for Case 3 AF	106
Figure A.28 Residence Time in the LA for Case 3 SR.....	107

LIST OF TABLES

Table 1.1: Electromechanical Correlates of the Cardiac Cycle as Represented on the ECG	13
Table 1.2: Quantitative assessment of Virchow’s triad in AF-related stroke risk.	26
Table 2.1: Recommended Mesh Keep Ratios for Different Spatial Resolutions.	40
Table 3.1: Case-wise hematocrit values.....	53
Table 3.2: LAA Morphology	57
Table 3.3: Segmentation Time Frames for AF and SR Across Different Cases.....	59
Table 3.4: Used Segmentation Time Points in AF Cases with Time Zero as Reference .	59
Table 3.5: Mesh simplification parameters (mesh_keepratio and faces_cyl_thresh0) assigned for each case in the surface preprocessing stage.	67
Table 3.6: Total Number of vertices and faces	67
Table 3.7: Eularian Grid Resolution for each case	67
Table 3.8: Assigned colors for key LA structures used for visualization.	69
Table 3.9: Simulation Runtime and GPU Used for Each Case and Rhythm Type.....	70

ACKNOWLEDGEMENTS

This thesis, *Patient Specific Computational Fluid Dynamics in the Left Atrium: Atrial Fibrillation vs. Sinus Rhythm*, would not have been possible without the support and guidance of many individuals. I would like to express my deepest gratitude to my advisors, Dr. Dana Dabiri and Dr. Juan Carlos del Álamo, for their invaluable expertise, mentorship, and encouragement throughout the course of this research. Their insights into fluid dynamics, medical imaging, and scientific rigor have been instrumental in shaping the direction and quality of this work. I am grateful to Alejandro Gonzalo Grande for teaching me the fundamentals of the project and subject area, and for supporting me throughout my learning process. My sincere thanks go to Akshaykumar Gupta for being a supportive colleague and collaborator, and to the University of Washington Hyak team for providing computational resources that made large-scale simulations possible.

I would also like to thank my family for their unwavering support, patience, and belief in me throughout my graduate journey. Their encouragement has kept me motivated during challenging times.

DEDICATION

To my mother, who believed in me even more than I believed in myself

Chapter 1. INTRODUCTION

1.1 CLINICAL BACKGROUND AND MOTIVATION

Atrial fibrillation (AF) and sinus rhythm (SR) represent two distinct cardiac rhythm states with markedly different hemodynamic characteristics, particularly in the left atrium (LA). AF is associated with irregular atrial contractions and impaired blood transport, creating conditions conducive to thrombus formation, while SR maintains coordinated contractions that promote efficient flow. While clinical imaging can capture valuable snapshots of atrial anatomy and limited flow information, it cannot fully resolve the time-resolved, three-dimensional hemodynamics that drive these differences. Computational modeling bridges this gap by reconstructing and analyzing patient-specific flow fields over the cardiac cycle, offering insights into physiological dynamics that are otherwise inaccessible in vivo.

Atrial fibrillation (AF) is a progressive cardiac arrhythmia characterized by rapid and disorganized electrical impulses in the atria, resulting in an irregular and often elevated heart rate. This electrical disruption interferes with the normal conduction pathway governed by the sinoatrial (SA) node, which typically maintains a steady and coordinated rhythm known as SR. Under SR, the atria contract effectively, allowing for efficient ventricular filling and optimal cardiac output. In contrast, AF overrides the sinoatrial node's activity with erratic impulses, leading to ineffective atrial contractions and irregular ventricular response.

AF is the most common sustained arrhythmia encountered in clinical practice. According to the Global Burden of Disease Study, over 33 million individuals worldwide were living with AF as of 2019 [1]. This number is projected to double by 2050, driven largely by an aging population and the increasing prevalence of comorbid conditions such as hypertension, diabetes, obesity, and

coronary artery disease [2]. The highest incidence is observed in individuals over the age of 65, although prevalence varies across regions due to differences in healthcare infrastructure and diagnostic practices.

The clinical implications of AF are significant and wide-ranging. One of the most severe complications is thromboembolic stroke, which arises when ineffective atrial contraction causes blood stasis and subsequent clot formation. Patients with AF are approximately five times more likely to experience a stroke than those in SR [3], and such strokes are typically more severe and debilitating. Furthermore, AF is linked to a twofold increase in both cardiovascular and all-cause mortality [4]. This elevated risk reflects not only thromboembolic events but also the progression of heart failure, impaired cardiac output, and structural remodeling of the atria and ventricles over time.

In addition to its clinical burden, AF presents a major challenge for healthcare systems globally. It is associated with frequent hospitalizations, repeat emergency department visits, and substantial long-term costs due to the need for ongoing medical therapy like anticoagulants, rate or rhythm control drugs and interventional procedures like catheter ablation. Despite advancements in AF management, many patients suffer from recurrent episodes, diminished quality of life, and increased psychological stress.

The objective of this study is to develop a patient-specific computational framework capable of simulating blood flow through the LA under both SR and AF. By integrating anatomically accurate geometries and physiologically relevant boundary conditions into computational fluid dynamics

(CFD) models, the framework enables the detailed analysis of rhythm-specific hemodynamics. This approach allows us to investigate whether there are fundamental differences in flow dynamics between SR and AF, with particular emphasis on identifying patterns related to blood stasis and thrombogenic risk.

1.2 CARDIAC ELECTRICAL ACTIVITY AND ITS REPRESENTATION ON THE ELECTROCARDIOGRAM

The electrocardiogram (ECG) is a critical tool for understanding the electromechanical differences between SR and AF, as it is a non-invasive procedure that reveals electrical patterns closely linked to cardiac mechanical function.




Figure 1.1 Illustration of the R–R Interval on an ECG, Representing One Complete Cardiac Cycle.

Figure 1.1 illustrates a normal ECG waveform, with its major components described in Table 1. Among these, the P wave, which reflects atrial depolarization, holds particular significance when comparing AF and SR, as it directly relates to the coordinated contraction of the atria and is markedly altered in AF.

Table 1.1: Electromechanical Correlates of the Cardiac Cycle as Represented on the ECG

ECG Component	Electrical Event	Mechanical Phase	Heart Action
P Wave	Atrial depolarization	Diastole	Atria contract, ventricles fill
PR Interval	AV node delay	Diastole	Ventricles fill
QRS Complex	Ventricular depolarization	Systole	Ventricles contract (start pumping)
ST Segment	Ventricles fully depolarized	Systole	Ventricles still pumping
T Wave	Ventricular repolarization	Systole → Diastole	Ventricles relax
Post-T	Electrical rest	Diastole	Ventricles fill again

Atrial Depolarization starts with the signal sent from the sinoatrial node (SA) which is present in the right atria (RA) which relays the signal towards the Atrioventricular node (AV) during the electrical signal is spread around the LA with the help of the Bachmann's bundle which spreads the electrical signal thought the Myocardium (Cardiac Muscle Tissue) which triggers the compression of the LA leading to the emptying of the blood into the left ventricle (LV). This is the process which happens during the P wave. The PR segment which is the AV node delay which just the duration in which the blood keeps on filling in the left-over blood into the LV. The P wave and the PR segment together is responsible for the atrial function.

A combination of the all the different wave sections shown in Figure 1.1 from the R peak to another R peak is considered as 1 cardiac cycle which consist of a Systole phase (Ventricular contraction leading to the ejection of blood into the arteries) and a Diastole phase (ventricular filling).




Figure 1.2 The Cardiac Conduction System, showing the pathway of electrical impulses.




Figure 1.3 Normal vs Abnormal Electrical Pathways. Normal conduction follows a clear path from the SA to AV node, while atrial fibrillation involves disorganized electrical signals within the atria.




Figure 1.4 ECG Comparison of Normal Sinus Rhythm and Atrial Fibrillation [5].

Now when it comes to the AF the greatest problem is that the Cardiomyocytes which are the cells which make up the myocardium fail and the conduction of electrical signals and lead to irregular and rapid disorganization of the electrical impulse. This can also be clearly seen when we compare the ECG's of a patient in SR vs a patient in AF as shown in Figure 1.4, it becomes very evident with the irregularities which can be seen in the pattern like different phase durations, narrow QRS wave, erratic baseline which has a section which mimics the P wave and there are sections where there is an absence of the P wave. With major wave phases compromised this leads to improper cardiac activity.

Now looking at the PR segment which is responsible for the proper functioning of the atria it is visible and in some of the phases the PR sections that the P wave is absent or there is a baseline activity which is trying to mimic the function of the P wave [6].

By understanding the function of the P wave, we can assess whether a patient is experiencing atrial fibrillation (AF), as its absence or distortion is a key indicator of disorganized atrial activity. Once AF is identified, it can be further classified into different types based on its presentation and

duration. A first episode refers to the initial detection of AF, regardless of symptoms or how long it has lasted. When a patient experiences more than two episodes, it is termed recurrent AF. Paroxysmal AF describes episodes that terminate spontaneously and last less than seven days, while persistent AF refers to episodes that last more than seven days and do not resolve on their own. When AF continues for over a year, it is classified as long-standing persistent AF. In cases where AF has persisted for more than a year and rhythm control interventions are either not attempted or have been unsuccessful, it is termed permanent (or accepted) AF. It is important to note that paroxysmal AF lasting longer than 48 hours is unlikely to spontaneously revert to sinus rhythm, and in such cases, anticoagulation should be considered to reduce the risk of thromboembolism.

1.3 HEMODYNAMIC ALTERATIONS IN ATRIAL FIBRILLATION (AF)

The hemodynamic changes associated with atrial fibrillation (AF) play a central role in promoting thrombus formation by disrupting the heart's normal anti-thrombotic state and creating a pro-thrombotic environment.

A key consequence of these changes is blood stasis and increased residence time within the atria, particularly in the left atrial appendage (LAA), which is a common site for thrombus development. The LAA is a small, pouch-like extension of the LA and is known for its anatomical variability across individuals. Understanding the morphology of the LAA is essential, as its shape influences the likelihood of thrombus formation. Broadly, LAA morphologies are classified into four types: chicken wing, windsock, cauliflower, and cactus, as illustrated in Figure 1.5.

Building on this classification, a recent study using multidetector computed tomography (MDCT) and cardiac magnetic resonance (CMR) in patients with drug-refractory AF identified the

prevalence of these four LAA shapes. The chicken wing configuration was the most common, observed in 48 percent of cases, followed by cactus at 30 percent, windsock at 19 percent, and cauliflower at just 3 percent [7].

Among the different morphologies, the cauliflower type has been most frequently associated with embolic events. It is described as having a short overall length, an irregular orifice, and a complex internal structure with multiple lobes, none of which are dominant. In contrast, the cactus shape has a central dominant lobe, from which smaller lobes extend upward and downward. The windsock morphology features a prominent primary lobe, with variability in the number and positioning of secondary or even tertiary lobes. The chicken wing type includes a dominant lobe with a distinct bend in its proximal or middle section that folds back on itself, often accompanied by smaller secondary lobes [8].

These anatomical differences are clinically significant because the structure of the LAA contributes to blood flow dynamics within the heart. Under normal conditions, the LAA acts as a reservoir, helping regulate blood flow. However, during AF, the loss of effective atrial contraction causes this anatomical feature to become problematic. Blood velocity within the LAA is consistently lower than in the main body of the LA, and the intricate structure of the LAA promotes the formation of recirculation zones and areas of stagnation, both of which significantly increase the risk of thrombus development.




Figure 1.5 Anatomic Variants of LAA Morphology

Sample images taken from explanted hearts demonstrate different LAA morphologies (top). (A) Chicken wing. (B) Windsock. (C) Cauliflower. (D) Cactus [9].

1.3.1 *Blood Flow Differences in Atrial Fibrillation (AF) and Sinus Rhythm (SR)*

During normal SR, the LA contracts in a highly coordinated sequence with the rest of the cardiac cycle. This atrial contraction, sometimes referred to as the "atrial kick," actively propels blood from the LA into the left ventricle (LV), complementing passive filling and ensuring efficient ventricular preload. Within this process, the left atrial appendage (LAA) also empties effectively, with brisk flow velocities that minimize the risk of blood stagnation.

In AF, the electrical activity of the atria becomes disorganized, causing the LA walls and LAA to fibrillate rather than contract. This chaotic motion replaces the synchronized atrial kick with small, rapid, and ineffective contractions. As a result:

- LAA flow velocity decreases which is a well-recognized hemodynamic marker of blood stasis.
- Blood residence time increases in both the LA and LAA due to incomplete and inconsistent emptying.
- Regions of near-zero flow develop, with one of the most common sites being the LAA apex, where the anatomy naturally predisposes to stagnation.

Computational fluid dynamics (CFD) studies have quantified these changes, demonstrating that AF patients experience significantly prolonged residence time compared to control subjects, even during periods of restored sinus rhythm in paroxysmal AF [10]. This suggests that structural remodeling and flow pattern alterations persist beyond episodes of arrhythmia. The prolonged residence time enhances the likelihood of blood constituents interacting long enough to activate the coagulation cascade, creating favorable conditions for thrombus formation.

Clinical imaging studies reinforce these findings: patients with an existing LAA thrombus have been observed to exhibit markedly lower peak LAA velocities than those without thrombus (0.17 vs. 0.26 m/s) [11]. This evidence establishes a clear, direct relationship between reduced flow velocity and elevated thrombotic risk, emphasizing the importance of hemodynamic assessment in AF management.

1.3.2 *Shear Stress Regulation of Endothelial Function*

In SR the blood flow is Laminar in nature that creates shear stress on the endothelium which is crucial for maintaining the endothelial health and function. In AF the laminar flow is disrupted, creating turbulent flow patterns; this leads to reduced and oscillatory shear stress alter the endothelial gene expression.

The gene expression is responsible for the production of Nitric oxide (NO) which is produced by the endothelial nitric oxide synthase (eNOS). The synthesis of NO is significantly reduced during AF [12]. This reduction in NO production contributes to endothelial dysfunction and creates a pro-thrombotic environment.

In normal SR, the blood in the LA has a laminar flow. Laminar flow creates a steady rubbing force on the inner wall of the atrium, known as shear stress. This force doesn't damage the surface and it's essential for keeping the endothelial cells (the thin layer of cells lining the heart and blood vessels) healthy.

One of the ways shear stress maintains endothelial health is by controlling certain genes. These genes include those that make endothelial nitric oxide synthase (eNOS), an enzyme that produces nitric oxide (NO). NO is important because it relaxes blood vessels (vasodilation), reduces clot formation (by preventing platelets from sticking together), and limits inflammation along the vessel wall.

In AF, the atria stop contracting in a coordinated way, which disrupts laminar flow and replaces it with swirling, irregular patterns creating a turbulent flow pattern. Turbulent and irregular flows reduce the average shear stress and cause it to change direction frequently (oscillatory shear stress). These changes confuse the endothelial cells, altering their gene activity and reducing eNOS production. With less eNOS, the amount of NO in the blood drops. Lower NO levels mean the

vessel walls lose some of their natural protection making them more prone to dysfunction, inflammation, and clot formation [12].

1.3.3 *Molecular and Cellular changes in Response to Altered Hemodynamics*

In AF, changes in blood flow, known as hemodynamic alterations, do more than just slow circulation; they set off a chain of molecular and cellular events that make the blood more likely to clot. This combination of effects creates a prothrombotic environment, meaning the body is in a state where clot formation is favored. Clinical studies have shown that patients with a LAA thrombus, a blood clot in the LAA, display higher levels of several blood markers compared to those without thrombus. Platelet factor 4 (PF4), a protein released when platelets are activated, is significantly higher in patients with thrombus (6.1 vs. 3.5 IU/ml). Beta-thromboglobulin, another marker of platelet activation, is also elevated (56.8 vs. 30.4 IU/ml). Thrombin and antithrombin complexes (TAT), which indicate active clot formation, are likewise higher (5.59 vs. 3.06 µg/ml [11]. Together, these findings show that AF is associated with increased platelet activation and coagulation activity, both of which raise the risk of thrombosis.

Another important marker is the von Willebrand factor (vWF), a protein involved in platelet adhesion. vWF levels rise when the endothelium, the thin layer of cells lining the heart and blood vessels, is damaged or not functioning properly. Remarkably, vWF levels remain elevated in AF patients even when they have no other cardiovascular risk factors [12]. This strongly suggests that AF itself, through its abnormal flow patterns, directly damages the endothelium. Researchers have consistently identified three major predictors of LAA thrombus across different patient groups: reduced LAA flow velocity, representing blood stasis; elevated beta-thromboglobulin, representing platelet activation; and increased vWF, representing endothelial damage [13].

Other biomarkers further highlight the clot-promoting effects of AF. Fibrinogen, a glycoprotein produced in the liver, serves as the raw material for fibrin, the mesh-like structure that stabilizes a blood clot. Elevated fibrinogen levels create a denser fibrin network that is more resistant to breakdown by plasmin, resulting in clots that are more stable and persistent. D-dimer, a fibrin degradation product, is another important marker [14]. Elevated D-dimer levels indicate ongoing clot formation and breakdown, providing direct evidence that the clotting process is actively taking place in the body. Together, these changes demonstrate that AF does not simply passively increase the risk of clot formation by slowing blood flow. Instead, it actively alters blood composition and vessel wall biology in ways that promote clot formation, increase clot stability, and cause chronic endothelial injury.

1.3.4 *Virchow's Triad*

Virchow's triad provides a unifying framework for understanding why AF carries such a high risk of thromboembolism. By recognizing that thrombus formation arises from the interplay of blood stasis, endothelial injury, and hypercoagulability, it becomes possible to integrate a wide range of clinical and mechanistic findings into a single explanatory model. In AF, each arm of the triad is not only present but often reinforces the others. Stasis promotes endothelial injury, endothelial injury amplifies coagulability, and hypercoagulability increases the consequences of stasis [15]. This interconnectedness means that interventions targeting only one component may reduce, but rarely eliminate, thrombotic risk [16].

The importance of this model lies in its clinical utility. Virchow's triad does not simply describe pathophysiology; it informs diagnosis, risk assessment, and treatment planning. In AF, biomarkers and imaging findings corresponding to each component such as reduced LAA velocity for stasis, elevated von Willebrand factor for endothelial injury, and raised fibrinogen or D-dimer for

hypercoagulability can help stratify patients by their thromboembolic risk. This layered understanding allows clinicians to move beyond generic risk scores and instead adopt a more tailored approach, identifying which patients might benefit from intensified anticoagulation, antiplatelet therapy, or procedural interventions.

Ultimately, the value of Virchow's triad is that it connects molecular biology, hemodynamics, and clinical medicine into one coherent picture. In the context of AF, it underscores that preventing thrombus formation requires addressing all three pathways. These include restoring more physiological flow patterns where possible, protecting endothelial function, and controlling the coagulation cascade [16]. By applying this triad-based perspective, both researchers and clinicians can better target their efforts to reduce the devastating consequences of AF-related thromboembolism.




Figure 1.6 Schematic representation of the LA during AF, illustrating the three components of Virchow's triad [15]

1.4 CFD IN UNDERSTANDING HEMODYNAMICS

Computational Fluid Dynamics (CFD) has emerged as a powerful in-silico approach for investigating complex blood flow patterns in the cardiovascular system, particularly within the LA. By generating detailed numerical simulations of cardiac hemodynamics, CFD provides insights that are often beyond the reach of conventional clinical imaging. It enables precise quantification of parameters such as blood velocity, wall shear stress, and regions of flow stagnation factors that play a critical role in thrombus formation. This ability to visualize and measure otherwise inaccessible flow features makes CFD a valuable tool for improving thromboembolic risk assessment and guiding personalized treatment strategies in patients with atrial fibrillation.

Traditional clinical methods for assessing hemodynamics, such as transthoracic echocardiography (TTE) and transesophageal echocardiography (TEE), provide only limited information, mainly average flow velocity and direction within a specific sampling area. In contrast, CFD offers detailed three-dimensional insights into local blood flow patterns throughout the heart. Although 4D flow MRI can also capture complex flow information, its spatial resolution is currently too low to accurately assess small structures like the LAA [17].

The computational approach enables the simultaneous calculation of pressure and velocity, along with the ability to visualize blood flow dynamics in patient-specific 3D models. Importantly, CFD allows for the calculation of wall shear stress (WSS) and energy loss—parameters that are difficult to measure directly but carry significant clinical implications for understanding endothelial function, which is critical for maintaining optimal circulation [18].

Patient-specific LA models typically begin with high-resolution cardiac imaging, such as computed tomography (CT) or magnetic resonance imaging (MRI). These images are processed

using segmentation techniques to extract the 3D anatomical geometry of the LA and its appendages. More advanced models integrate high-resolution CT-based LA reconstructions with motion data and pulmonary inflow measurements from 4D flow MRI, resulting in multimodal CFD approaches that better capture the dynamic nature of the LA [19]

CFD studies have revealed marked differences in LA hemodynamics between normal SR and AF. One of the most relevant parameters in this context is residence time (RT), which measures how effectively blood is cleared from a given region. The left atrial appendage (LAA) is a particular focus since approximately 90% of LA thrombi responsible for thromboembolic events in AF originate there [13]. Particle-tracking studies have shown that in SR, only about 26% of particles remain in the LAA after three cardiac cycles, whereas in AF, around 45.6% remain, demonstrating a greater tendency for blood stasis in AF [20]

Over the years, CFD has also enabled the development of novel analytical tools, such as Age Stasis, which uses neural networks (NN) and CFD data to predict the effects of aging on the cardiac system [21].

Studies comparing AF patients with and without a history of stroke have shown clear differences in the blood residual ratio and particle residual ratio within the left atrial appendage (LAA). In patients with prior stroke, both ratios were significantly lower, suggesting that these measures could be useful in identifying individuals at higher risk for thromboembolic events [22].

Researchers have also used reduced-order models to investigate wall shear stress patterns and endothelial cell activation in the LAA. These models highlight how atrial remodeling disrupts flow and promotes stasis. By mapping the results onto a standardized LAA coordinate system, they make it possible to compare flow characteristics across patients with different anatomical shapes and functional conditions [23].

Table 1.2: Quantitative assessment of Virchow’s triad in AF-related stroke risk.

Triad Component	Assessment Method	Clinical Impact
Blood Stasis	CFD-derived residence time maps [24], [25]	Identifies >90% of thrombus locations [26]
Endothelial Damage	LA strain imaging and von Willebrand factor (vWF) levels [27]	Predicts 3.2× higher stroke risk when impaired [26]
Hypercoagulability	Fibrinogen, D-dimer, and thrombin kinetics [24]	Explains 40% of residual stroke risk in anticoagulated patients [25]

Table 2 summarizes how each component of Virchow’s triad: blood stasis, endothelial damage, and hypercoagulability can be quantitatively assessed using modern diagnostic tools and how these measures translate into clinical relevance. Blood stasis can be characterized through CFD derived residence time maps, which have been shown to identify more than 90 percent of thrombus formation sites [24], [25], [26]. Endothelial damage can be evaluated using LA strain imaging and plasma von Willebrand factor (vWF) levels, with impaired measures predicting a 3.2-fold increase in stroke risk [27]. Hypercoagulability can be assessed using biomarkers such as fibrinogen, D-dimer, and thrombin generation kinetics, which together explain up to 40 percent of residual stroke risk in anticoagulated patients. By aligning each component of the triad with objective diagnostic modalities, this framework supports a mechanistic and individualized approach to thromboembolic risk stratification.

1.5 RISK STRATIFICATION

Current stroke risk stratification in AF primarily relies on the CHA₂DS₂-VASc score, a simple clinical tool that assigns weighted points to factors such as heart failure, hypertension, diabetes, vascular disease, age, and prior stroke or transient ischemic attack. This approach has become a cornerstone in guiding anticoagulation therapy, offering a straightforward method to categorize

patients into low-, intermediate-, and high-risk groups. While effective in broad clinical decision-making, it is based solely on the presence of comorbidities and does not directly evaluate the underlying hemodynamic and structural mechanisms that contribute to thromboembolism. As a result, the CHA₂DS₂-VASc score, despite its widespread use, demonstrates notable limitations in fully capturing patient-specific stroke risk.

One important shortcoming lies in its insensitivity to hemodynamic risk factors, as it does not account for LA flow abnormalities. Studies have shown that 35% of AF patients with a CHA₂DS₂-VASc score of 1 or lower can still develop LA thrombi, and 11.4% of those deemed low risk have experienced thromboembolic events [27]. This underscores the score's limited capacity to detect elevated stroke risk in individuals with minimal clinical burden. Another limitation is its incomplete assessment of the atrial substrate, as it does not incorporate factors such as LA fibrosis, endothelial dysfunction, or atrial hypocontractility, all of which may persist even in sinus rhythm and independently contribute to thromboembolic risk. Moreover, its predictive performance remains modest, with a C-statistic between 0.57 and 0.67, meaning that a substantial proportion of high-risk patients have no LA thrombi while 15–20% of low-risk patients still develop thrombotic complications.

To address these gaps, emerging approaches integrate the principles of Virchow's triad—blood stasis, endothelial injury, and hypercoagulability, through the use of advanced imaging, biomarkers, and computational fluid dynamics (CFD). This enables a mechanistic and individualized assessment of thromboembolic risk. A proposed integrated model begins with a baseline assessment that combines clinical scoring via CHA₂DS₂-VASc, structural evaluation

through LA volumetric imaging by cardiac MRI, and molecular markers such as plasma fibrinogen and D-dimer levels. Patient-specific CFD simulations can then be performed by reconstructing LA geometry from MRI data, mapping wall shear stress, and quantifying blood residence time and stagnation zones. This integrated strategy has demonstrated improved predictive accuracy, achieving a C-statistic of 0.87 compared to 0.67 for CHA₂DS₂-VASc alone, and by capturing all three components of Virchow's triad, it offers a more comprehensive framework for identifying patients at elevated thromboembolic risk, particularly those who might be misclassified by conventional scoring systems.

1.6 LIMITATIONS

While CFD has greatly advanced our understanding of blood flow in the LA, current models still have limitations that can affect both accuracy and real-world usefulness. A key issue is the rigid wall assumption; many simulations treat the LA walls as fixed, unmoving surfaces. The atrium moves and deforms throughout the heartbeat, and this motion has a strong impact on blood flow patterns. Research shows that ignoring wall motion can lead to underestimating important measures like time averaged wall shear stress (TAWSS) and oscillatory shear index (OSI), both of which are linked to clot formation risk [19].

Another simplification is treating blood as a Newtonian fluid, assuming its viscosity stays constant. In truth, blood is shear thinning (it flows more easily at higher shear rates). This property is especially important in low shear regions where clots tend to form. Using the Newtonian assumption can overpredict stagnant zones and miss subtle flow changes that influence clot risk.

Finally, when building 3D models from Cine MRI scans, limitations in imaging such as slice thickness can reduce how precisely the LAA shape and motion are captured. These small inaccuracies in geometry can carry over into the simulation and affect the results [17].

1.6.1 *Project Contributions*

This project takes important steps toward overcoming several of the limitations in current CFD-based atrial flow analysis.

First, it incorporates dynamic wall motion by modeling atrial wall deformation in sync with the cardiac cycle. By capturing the distinct phases of systole and diastole, the simulations better reflect realistic blood flow patterns in both SR and AF. This leads to more reliable estimates of residence time (RT) and more physiologically accurate flow dynamics.

Second, the project advances anatomical reconstruction by using Cine CT rather than Cine MRI. Cine CT offers higher resolution and improved contrast, enabling more accurate segmentation of the LAA and complex pulmonary vein (PV) inlet geometries. This also improves wall motion tracking, enhancing model fidelity.

However, several challenges remain. Limited AF-specific patient data especially for paroxysmal AF which makes it difficult to capture AF episodes during CT scans. Segmentation inaccuracies can occur, particularly around the PV inlets and LAA, due to contrast loss and highly complex geometries. The high computational demands of running these models on high-performance computing (HPC) clusters also pose a barrier for real-time or near-clinical-time applications. Additionally, the workflow assumes periodic cardiac cycles stitched together from image sequences, but in AF, beat-to-beat variability means that true periodicity cannot be guaranteed.

By addressing some of these limitations while acknowledging those that persist this project contributes toward more accurate, patient-specific CFD analysis, moving the field closer to tools that could eventually support clinical decision-making.

Chapter 2. METHODOLOGY

This chapter outlines the step-by-step workflow used to develop patient-specific CFD simulations of LA blood flow under AF and SR. The goal is to characterize and compare hemodynamic differences between the two rhythms.

The methodology combines medical imaging and computational modeling. This pipe is organized in the following stages:

1. Geometry Creation

- Generation of patient-specific models: Constructing a 3D point cloud which is extracted from the imaging data. Using the data through a nn-UNET processing to get a 3D segmentation of the left side of the heart.
- Manual refinement and alignment of anatomical features: Using the 3D point cloud generated from the nn-UNET segmentation, manual adjustments are performed to ensure anatomical accuracy.

2. Model Preparation for Simulation

- Temporal smoothing of imaging data: All frames from the CT dataset are combined and smoothed in both spatial and temporal dimensions to eliminate sudden disturbances and reduce noise.
- Meshing of anatomical geometry: The 3D point cloud is discretized into a mesh, defining the volumetric domain and surface boundaries required for simulation.
- Motion registration: Frame-to-frame point correspondence is established to align cardiac motion and accurately capture contraction and expansion dynamics.
- Anatomical boundary conditions: Inlet and outlet regions are identified and classified to support boundary condition setup in subsequent modeling.

3. CFD Simulation

- The prepared geometry and motion data are used to run CFD simulations, enabling the analysis of blood flow dynamics.

4. Post-Processing

- The CFD output is processed to generate readable and interpretable data formats, enabling clear visualization and quantitative analysis of flow dynamics.

This pipeline ensured consistency between the AF and SR datasets of the same patient, enabling valid comparison. Each step is described in detail in the sections below.




Figure 2.1 Methodology flowchart outlining the steps from geometry creation and model preparation to CFD simulation and post-processing

2.1 GEOMETRY CREATION

2.1.1 Generation of patient-specific models

The creation of a patient-specific model of the LA is a foundational step in the computational pipeline. This process begins with the acquisition of cardiac imaging data in DICOM (Digital Imaging and Communications in Medicine) [28] format, which stores anatomical information as a sequence of 2D slices in the XY-plane. To reconstruct a full three-dimensional volume, these slices

are stacked along the Z-axis using known inter-slice spacing. An in-house Python script was developed to automate this conversion process, generating a 3D volume in NIfTI (Neuroimaging Informatics Technology Initiative) format. The NIfTI format offers improved compatibility with segmentation tools and machine learning frameworks. At this stage, the volumetric data can be visualized using software such as ITK-SNAP, allowing for initial inspection and verification of anatomical structures.

To automate the segmentation of the LA and surrounding regions, the NIfTI data is processed using nnU-Net (no-new-U-Net), a self-configuring deep learning framework based on the U-Net architecture. Unlike conventional convolutional neural networks, nnU-Net automatically adjusts its preprocessing steps, network architecture, and training parameters based on the properties of the input dataset. This adaptability is particularly valuable in cardiac imaging, where anatomical features such as the LA often exhibit large variations in shape, size, and contrast.

In this study, nnU-Net was trained in-house using a curated dataset of patient-specific segmentations and matching imaging volumes. The dataset included manually segmented LA geometries from past cases, allowing the model to learn features closely aligned with the study population. This custom training approach improved segmentation accuracy, reduced the need for post-processing, and ensured that the extracted geometries reflected anatomical realism.

nnU-Net supports several configuration modes, each tailored to specific segmentation needs:

- 2D Configuration: Processes data slice-by-slice; useful when slice thickness is large relative to in-plane resolution [29].
- 3D Full Resolution: Processes the entire volume at original resolution to preserve fine anatomical detail [30].

- 3D Low Resolution: Uses downsampled input to reduce computational load while maintaining acceptable accuracy [30].
- 3D Cascade: Combines low- and high-resolution passes by first detecting the region of interest with a coarse model, then refining it with a full-resolution network.
- Ensemble Models: Aggregates outputs from multiple configurations to improve performance, though with mixed success in cardiac applications [28].




Figure 2.2 3D nnU-Net Architecture for Volumetric Image Segmentation

The typical architecture of nnU-Net includes five encoding and decoding layers, each composed of convolutional and activation operations. For cardiac segmentation, the encoder commonly begins with two $3 \times 3 \times 3$ convolutions and Leaky ReLU activations, followed by downsampling and deeper feature extraction layers [31].

The output of the nnU-Net pipeline is a segmented NIfTI volume in which anatomical regions such as the LA, LV and the PVs are labeled. This 3D segmented volume provides a structured and spatially resolved dataset that forms the basis for subsequent modeling steps.

2.1.2 *Manual Refinement and alignment of anatomical features*

While nnU-Net provides an effective baseline segmentation of the LA, additional manual refinement is necessary to ensure anatomical fidelity particularly in regions where automated methods often struggle. This step is critical to transforming a generic model into a patient-specific geometry that can support high-quality simulations.

The LAA is a critical focus in stroke risk assessment, as it is the primary site for thrombus formation in atrial fibrillation. Its complex and highly variable morphology, differing in shape, size, and orientation between individuals, makes it particularly challenging to model with precision. These difficulties are amplified in cases with low image contrast, where the LAA's proximity to the pulmonary veins creates overlapping anatomical features and dynamic motion that hinder clear boundary detection. While the nnU-Net model performs well for most atrial structures, targeted refinement of the LAA is essential to capture patient specific geometry and preserve clinically significant anatomical differences.

Manual correction is equally important near the pulmonary vein inlets and the mitral valve outlet, where inlet and outlet planes are adjusted to align orthogonally with the physiological direction of blood flow. Ensuring this alignment is crucial for setting up appropriate boundary conditions and for avoiding non-physiological flow patterns in later simulations.

These refinements also play a key role in ensuring consistency across cardiac temporal frames. Particularly the inlet and outlet regions must remain spatially aligned across all frames. This consistency directly supports downstream operations such as meshing, registration, and cork fitting, which rely on geometric continuity from frame to frame.

The final output of the geometry creation process is an anatomically accurate and temporally coherent 3D volume of the LA. This patient-specific model captures detailed morphological features, including the LAA, and provides a consistent geometric representation across all cardiac phases. Enabling reliable flow analysis under SR and AF rhythm.

2.2 MODEL PREPARATION FOR SIMULATION

Before the patient-specific geometry can be used for computational simulations, several preprocessing steps are necessary to ensure that the anatomical and temporal data are consistent, clean, and physiologically meaningful. This step acts as a bridge between raw segmented data and simulation-ready input, enabling accurate flow modeling across the cardiac cycle. To perform this preparation correctly, the following input variables are required:

- Segmentation Time:

This variable specifies the temporal position of each frame within the cardiac cycle as a percentage of the R–R interval, ranging from 0 to 100. A value of 0 corresponds to the peak of the first R-wave, and 100 marks the next R-wave, completing one full cardiac cycle. In SR, the ECG signal is regular, enabling consistent and reliable assignment of (seg_time) values. In AF, the rhythm is highly irregular, making R-peak detection more challenging. As a result, some frames may be assigned timestamps that occur before a confirmed R-


peak, producing negative (seg_time) values. This requires a screening process to retain only the valid frames that accurately represent a physiologically meaningful cycle


- Compressed Segmentation Frame:

The segmented volumetric data is stored and loaded as a sequence of NIFTI files in compressed format. Each frame represents the geometry of the LA at a specific point in the cardiac cycle. These files are the primary input for both smoothing and motion registration tasks that follow.

- Number of Pulmonary Veins (PVs):

The number of PVs connected to the LA must be defined as an input variable. Although four PVs (two anterior and two posterior) is the most common configuration, anatomical variations exist, including patients with three or six PVs. Accurately specifying this number is essential for correctly identifying and labeling the PV inlet regions, which in turn affects boundary condition setup in CFD simulations.





*Figure 2.3 Pulmonary vein (PV) in LA. Shaded areas show deviations from the typical four-vein pattern. **A:** Typical (4 PVs). **B:** Short common left trunk. **C:** Long common left trunk. **D:** Right middle lobe PV. **E:** Two right middle PVs. **F:** Right middle plus extra right “upper” PV [32]*

- Time per Cardiac Cycle (s):

This input represents the cardiac period in seconds (s) and is used to assign realistic timing to the motion data. Adjusting this value allows the simulation to reflect the actual duration of a patient’s heartbeat, supporting more physiologically accurate flow calculations, especially when comparing different rhythm conditions such as AF and SR.

Together, these inputs allow the preprocessing pipeline to interpret the temporal structure of the imaging data, account for anatomical variations, and align the cardiac cycle with physical time. The following sections describe how these inputs are used to perform spatial and temporal smoothing, mesh preparation, and motion registration.

2.2.1 Smoothing of imaging data

Once the full set of segmented frames across the cardiac cycle is obtained, a temporal smoothing operation is applied to improve the continuity and quality of the dynamic geometry. This process

addresses minor inconsistencies and noise artifacts that can arise due to limitations in CT acquisition or segmentation errors, particularly in spatiotemporal datasets.

The smoothing algorithm operates across both spatial and temporal domains and is governed by a parameter called *smoothing_stencil*. This parameter is specified in the format [X Y Z t], where the first three values represent the number of neighboring points considered in the X, Y, and Z directions, and the fourth value indicates the number of consecutive frames included in the temporal smoothing. Spatial smoothing works by averaging values over local neighborhoods of points in the mesh, which helps reduce small surface irregularities and micro-abnormalities. These minor variations can otherwise introduce non-physical fluctuations in wall shape or curvature, potentially affecting the accuracy and stability of flow simulations.

In the temporal dimension, smoothing is performed by averaging the geometry at a given frame with its preceding and succeeding frames, as determined by the stencil range. Since the input data spans a full cardiac cycle from one R-peak to the next, the geometry is treated as cyclic in time, meaning the first and last frames are considered adjacent for the purposes of smoothing. This wraparound behavior ensures consistency at the boundaries of the temporal domain and preserves physiological periodicity

Importantly, the smoothing function does not rely on the actual percent completion of the cardiac cycle but rather on the number of available frames for averaging. This allows the method to be generalizable across datasets with varying frame counts or non-uniform sampling rates. The result is a set of geometries that exhibit smooth, continuous motion across the cardiac cycle minimizing abrupt deformations that could interfere with motion tracking, registration, or CFD simulations.

By applying both spatial and temporal smoothing, the dataset is better prepared for downstream tasks which require consistency in both anatomical structure and temporal evolution.

2.2.2 *Meshing of Anatomical Geometry*

After segmentation and refinement, the anatomical structure of the LA is discretized into a finite difference mesh suitable for computational simulation. The meshing process is performed independently for each temporal frame, without accounting for inter-frame motion or temporal correlation at this stage. The objective is to generate smooth, clean, and watertight meshes that accurately reflect the anatomical geometry at every captured time point.

Meshing is performed using *iso2mesh*, a MATLAB-based open-source toolbox capable of generating high-quality 3D tetrahedral meshes from medical imaging data, such as segmented CT or MRI volumes.

A key parameter in this process is the *mesh_keepratio*, which determines what fraction of the initial surface points are retained in the final mesh. This directly influences mesh resolution. The recommended values vary by resolution as shown:

Table 2.1: *Recommended Mesh Keep Ratios for Different Spatial Resolutions.*

Spatial Resolution	Points per Box Length (nx)	mesh_keepratio
Low	144	0.1 – 0.3
Nominal	256	0.3 – 0.5
High	384	0.9 – 1.0

To maintain mesh quality, a sanity check is implemented by counting the number of surface triangles exceeding a calculated *MaxArea*, based on the x-direction spacing. If more than 30 such triangles are detected in any frame, the mesh is considered suboptimal, and the *mesh_keepratio* is adjusted accordingly to improve mesh fidelity.

Another essential parameter is *faces_cyl_thresh0*, which plays a key role in the classification of inlet and outlet surfaces, especially around the PVs and the MV. Adjusting *faces_cyl_thresh0* helps in correctly defining inlet and outlet planes for each frame, ensuring that boundary surfaces remain physiologically meaningful even as the geometry changes through the cardiac cycle.

2.2.3 *Motion Registration*

To analyze the heart's motion and simulate flow accurately, it is important to understand how the shape of the LA changes smoothly over time. Since only a limited number of 3D frames (usually 8–12) are available from imaging, we use two key techniques to align and expand these frames into a consistent, high-resolution sequence: Coherent Point Drift (CPD) and cubic spline interpolation.

CPD is a registration algorithm that helps us match the shape of the LA across all available frames by aligning corresponding anatomical features. For example, a point near a PVs in one frame will be tracked and connected to its most likely position in the next frame. Unlike basic matching techniques, CPD does not require the same number of points in every frame. It works by treating one frame as a reference (fixed) and the others as moving targets that need to be smoothly deformed to match the reference.

CPD aligns shapes by smoothly shifting one set of points to match another, ensuring nearby points move together rather than independently. This is important in modeling heart anatomy, where tissues deform in a coordinated manner.

CPD uses a Gaussian kernel to calculate smooth, realistic transformations for each point over time, producing a sequence of registered point clouds one per frame where anatomical landmarks are consistently aligned. In our case, this results in 8 to 12 frames across the cardiac cycle.

To better capture the heart's continuous motion, we upsample these to 40 frames using cubic spline interpolation. This method fits a smooth curve through known point positions at specific times (given as percentages of the cardiac cycle) and predicts positions in between.

Applying this process to all points in the 3D geometry produces evenly spaced frames, even when the original data was unevenly timed. The resulting 40-frame dataset provides a smooth and consistent representation of the full heartbeat, facilitating simulations and enabling direct comparisons between patients by ensuring each case has the same temporal resolution.

2.2.4 Anatomical Boundary Conditions

To enable accurate simulation of blood flow through the LA, it is essential to define physiologically realistic boundary conditions at the PV inlets and MV outlet. These boundaries guide how flow enters and exits the simulation domain and directly affect the behavior of the CFD solver.




Figure 2.4 Pulmonary vein inlets and mitral valve outlet with cork structures for boundary condition setup.

The anatomical boundaries are derived from the mesh geometry using a surface classification approach, where inlet and outlet planes are fitted at locations corresponding to the PVs and MV. The definition of these surfaces is governed by the *faces_cyl_thresh0* parameter, which determines the cylindrical threshold used to identify boundary regions based on their geometric orientation and curvature. A properly set *faces_cyl_thresh0* value ensures that boundary planes are created at appropriate anatomical locations, avoiding overlaps or misalignment with adjacent walls.

The goal is to align flow direction orthogonal to the local inlet or outlet plane, thereby preserving a physically meaningful representation of inflow and outflow patterns. This is especially important in patient-specific models where PVs may be oriented in various directions due to anatomical variability.




Figure 2.5 Non-Physiological Inlet/Outlet Placement extending beyond geometry boundaries.

Incorrect setting of these boundaries such as placing them too close to the atrial wall or at points where the surface curves sharply can lead to non-physiological flow entry shown in Figure 2.3. For example, boundary faces may extend beyond the intended anatomical region, leading to artificial flow entry through non-physiological areas.

By tuning *faces_cyl_thresh0* for each frame and patient case, we ensure that the PV inlets and MV outlet are cleanly and consistently segmented across the cardiac cycle. This step is essential to avoid mismatches to maintain consistent boundary behavior in the presence of temporally varying geometries.

The output of this step is a set of inlet and outlet surface planes aligned and labeled for each frame that are ready for the assignment of boundary conditions in the CFD solver.

2.3 CFD SIMULATION

This section describes the simulation framework used to evaluate blood flow in the deforming LA. We use a combination of image-derived anatomical motion, non-Newtonian flow modeling, and high-resolution scalar transport to simulate and analyze patient-specific cardiac hemodynamics. The simulations are executed using TUCAN, a GPU-based solver designed for immersed boundary flows in cardiovascular geometries. Our objective is to compute residence time (RT) a measure of flow stagnation and potential thrombotic risk under AF and SR conditions.

2.3.1 *Immersed Boundary Flow Framework*

We adopt the Immersed Boundary Method (IBM) to simulate blood flow in the dynamically deforming LA without requiring remeshing or body-fitted grids. The Lagrangian mesh: extracted from time-resolved CT segmentation which captures the surface motion of the atrial wall across the cardiac cycle. This mesh is overlaid onto a stationary Eulerian mesh that discretizes the fluid domain using a structured Cartesian grid.

The Lagrangian mesh is overlaid onto the Eulerian mesh, which creates a volumetric body that captures a fluid inside where no-slip boundary conditions are applied. To model fluid–structure

interaction, we use one-way coupling: wall motion drives the fluid, but fluid forces do not influence the wall. This simplification is physiologically grounded. The wall motion obtained from imaging inherently includes the effect of intraventricular fluid pressure during scanning. Therefore, enforcing an additional force-based interaction would introduce deformation that was not present in the physiological state and result in unrealistic wall behavior.

2.3.2 *TUCAN Solver and Numerical Discretization*

The simulations are performed using TUCAN (TUBular Cardiovascular Adaptive Navier–Stokes), an open-source computational solver designed for modeling blood flow in deforming cardiovascular structures. This solver integrates seamlessly with our workflow because it can directly incorporate frame-by-frame wall motion data extracted from CT segmentations into the flow field. TUCAN employs a fractional-step method to handle pressure–velocity coupling, uses second-order central finite difference schemes to compute spatial derivatives, and applies semi-implicit Runge–Kutta schemes for time integration. Its architecture is optimized for high performance on GPU platforms, specifically the Nvidia A40 and A100, enabling efficient handling of simulations with high frame counts over extended cardiac cycles. During execution, the solver sequentially reads and registers each frame’s geometry into the Eulerian mesh, allowing accurate simulation of flow in a continuously evolving domain. To ensure computational robustness, TUCAN incorporates a restart and sanity-check mechanism that enables simulations interrupted by high-performance computing (HPC) time constraints to resume from the last successfully completed frame without loss of progress.

2.3.3 *Residence Time Transport Equation*

In modeling blood residence time (RT) within the LA, it is important to account for the movement of fluid particles over time. This requires retaining the advection term in the transport equation, as it captures how blood elements travel through the domain and how long they remain in specific regions. Such tracking is particularly relevant in assessing areas prone to stasis, where blood flow is slow and clot formation risk is elevated.

The balance between advection and diffusion in scalar transport is described by the Schmidt number, which compares the rate at which momentum diffuses to the rate at which the scalar quantity spreads. A high Schmidt number indicates that momentum diffuses much faster than the scalar field, resulting in sharp gradients in RT that are highly sensitive to numerical accuracy. In blood, viscosity and therefore the Schmidt number increases with hematocrit, leading to a stronger dominance of advection over diffusion.

To further characterize the transport regime, we also consider the Péclet number, which relates the relative importance of advective to diffusive transport over a characteristic length and velocity scale. In the LA, and particularly within complex anatomical regions like the LAA, these values indicate that scalar transport is strongly advection-dominated. This means that the movement of blood is governed far more by bulk flow patterns than by molecular diffusion.

Advection-dominated transport poses numerical challenges. Conventional discretization schemes often introduce artificial smoothing, known as numerical diffusion, which can blur sharp features in RT distributions. Such smoothing can mask subtle but clinically relevant flow retention patterns, especially near the pulmonary veins or in the LAA. To address this, we use a Weighted Essentially Non-Oscillatory (WENO) scheme for solving the advection term. WENO adaptively adjusts its stencils based on local flow gradients, allowing accurate resolution of steep transitions in RT

without introducing spurious oscillations or excessive smoothing. This approach ensures that physiologically important flow features are preserved, enabling more reliable analysis of stasis and thrombus risk.

2.3.4 *Simulation Setup and Flow Convergence*

The simulations are organized into two distinct phases to ensure both numerical stability and physiological accuracy. The first stage, the stabilization phase, involves simulating three cardiac cycles at a coarse temporal resolution of 30 frames in total, corresponding to roughly six frames per cycle. This reduced frame rate is computationally efficient and is primarily intended to bring the system to a stable operating state. During this phase, the velocity and pressure fields are monitored to confirm convergence, while residence time values are tracked to ensure they remain physically bounded. By the end of this stage, transient start-up effects are minimized, and the flow is in a physiologically consistent regime.

The final flow and scalar field distributions from the stabilization phase serve as the initial conditions for the second stage, the production phase. In this phase, twenty full cardiac cycles are simulated at a high temporal resolution of 800 frames, or approximately 40 frames per cycle. These frames are generated using spline-interpolated motion fields obtained during preprocessing, enabling smooth wall kinematics and improved accuracy in fluid–structure interaction modeling. This fine resolution allows for the continuous evolution of both flow and scalar quantities, capturing subtle fluctuations in hemodynamic parameters that would be lost at coarser resolutions.

This two-phase simulation strategy offers several advantages. The initial stabilization prevents the propagation of non-physical transients into the main dataset, while the production phase provides a physiologically detailed representation of the cardiac cycle suitable for quantitative analysis.

Furthermore, the high-resolution data from the production phase are optimized for downstream tasks such as residence time computation, vorticity analysis, and visualization. By separating the simulation into these two stages, the methodology achieves a balance between computational efficiency and the generation of high-quality, clinically relevant hemodynamic datasets.

Chapter 3. RESULTS

As discussed in the previous chapter, the computational framework developed for this study involves a sequence of interconnected processes where each step is designed to extract, refine, and simulate patient-specific data that captures the dynamics of blood flow in the LA under SR and AF. Every step, from image segmentation to CFD simulation, was guided by the goal of preserving physiological realism while enabling precise quantitative assessment. The inputs were carefully examined for anatomical accuracy, temporal coherence, and computational consistency to ensure that the simulated outcomes reflect meaningful physiological behavior.

In this section, we present and interpret the results generated from this pipeline. We focus on comparing SR and AF conditions across key hemodynamic metrics, visual flow structures, and region-specific outcomes, particularly in the LAA. Through this layered analysis, we aim to evaluate how rhythm-specific flow dynamics contribute to thrombogenic risk and to demonstrate the capability of this model to serve as a tool for individualized stroke risk assessment. The results will be presented by following the same order and structure used in the previous sections of model development.

This study utilized three patient datasets with case IDs 44243, CVC, and 46483, which will be referred to as Case 1, Case 2, and Case 3 throughout the remainder of this work. Each patient dataset included paired cardiac CT acquisitions one under SR and another under AF enabling a consistent and direct comparison of hemodynamic behavior across cardiac rhythms while holding anatomical features constant. This paired structure ensured that rhythm-induced changes in blood flow dynamics could be evaluated independently of patient-specific anatomical variation.

For each dataset, a set of core physiological inputs was collected and used throughout the modeling process. This included the DICOM-formatted CT image sequences, hematocrit (Hct) values, and ECG signals. Together, these inputs formed the foundation of the entire computational pipeline and ensured that the simulations maintained both physical and clinical relevance.

Case1 | 44243




Figure 3.1 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 1.




Figure 3.2 DICOM Images for Atrial Fibrillation (AF) – Case 1





Figure 3.3 DICOM Images for Sinus Rhythm (SR) – Case 1



(A) AF (B) SR
Figure 3.4 ECG Comparison Between AF and SR – Case 1

To provide context for the simulations, we first examined the ECG traces and patient characteristics for the three cases studied. Case 1 involves a female patient, while Cases 2 and 3 involve male patients. Each case was analyzed in both AF and SR to understand the electrical activity patterns that influence blood flow simulations.

In the AF ECG for Case 1, shown in Figure 3.4(A), the heartbeat pattern is noticeably irregular. Within a 1-second window, there are three R-peaks in AF compared to only two R-peaks in the SR trace, meaning the heart is beating faster and less predictably. Another striking difference is the absence of the P-wave in AF, which signals that the atria are no longer contracting in a

coordinated way. In SR, the P-wave is clearly visible before each QRS complex, reflecting normal atrial depolarization.

For the simulations, a cardiac cycle is defined from one R-peak to the next, with timing normalized as a percentage of the cycle. In SR, this is straightforward because the rhythm is regular. In AF, the RR intervals vary so much that the system cannot predict the next R-peak ahead of time. Instead, it records frames continuously and, after detecting an R-peak, retrospectively defines the cycle from roughly -40% to +90% of the RR interval. This irregularity directly affects how the motion data is extracted and applied in our CFD workflow

In Case 2, the AF ECG looks more stable than in Case 1 at first glance, with the RR intervals appearing regular. However, closer inspection shows subtle disruptions in the ST segment compared to SR. These mid-cycle variations may reflect incomplete ventricular repolarization or inconsistent recovery between beats. The P-wave is still absent in AF, confirming the loss of synchronized atrial contraction.

From a modeling perspective, the relatively stable RR intervals in AF for this case make frame selection less challenging than in Case 1. However, the irregularities in the ST segment could still influence downstream hemodynamics, especially during the diastolic filling phase in our simulations. (The ECG plots for this case are provided in the Appendix.)

Case 3 shows the most chaotic AF ECG pattern of all. The trace is filled with erratic electrical spikes, making it difficult to identify clear phases of the cardiac cycle. The RR intervals vary widely, and the absence of P-waves confirms disorganized atrial activity. This high level of electrical noise not only reflects severe rhythm disturbance but also makes it harder to define precise segment timings for modeling.

In the simulations, this meant relying on only the clearest RR interval sequences for constructing motion input. Because the AF signal here was so irregular, fewer usable segmentation points were available, which limited the temporal resolution we could achieve in the CFD setup. (The ECG plots for this case are provided in the Appendix.)

Table 3.1: Case-wise hematocrit values

Case ID	Hematocrit
44243	48.3
CVC	42
46483	40.9

In the Table 3.1, we report the hematocrit (Hct) values for each patient, which serves as a critical parameter in modeling the non-Newtonian behavior of blood. Hematocrit reflects the proportion of red blood cells in the blood volume and directly influences its viscosity and flow resistance. Incorporating patient-specific hematocrit values allows for more physiologically accurate simulation. These values were integrated into the solver as part of the viscosity model, ensuring that each simulation reflects the unique hemodynamic environment of the corresponding patient.

3.1 GEOMETRY CREATION

The process of geometric creation is a critical stage in this study, as it forms the structural foundation upon which all subsequent simulation steps are built. While many parts of the processing pipeline are automated, this stage is where the researcher’s anatomical knowledge and physiological understanding have the greatest influence. The segmented geometry serves as the direct input for meshing, registration, and CFD simulations, meaning that errors introduced here will propagate downstream. Complex regions such as the PVs and the LAA require special attention to ensure that the reconstructed anatomy reflects physiological reality. This section

describes the two main stages of segmentation: deep learning–based automated segmentation and manual refinement.

3.1.1 Deep Learning–Based Automated Segmentation

Initial segmentation of the LA and its associated structures was performed using nnU-Net, a state-of-the-art deep learning–based medical image segmentation framework. nnU-Net was selected for its robust performance across a range of cardiac imaging modalities and its ability to adapt to domain-specific datasets without extensive manual tuning.




Figure 3.5 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF – Case 1




Figure 3.6 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 1




Figure 3.7 nnU-Net-based 3D segmentation from stacked multiplanar CT slices

Once the deep learning solver processes the cardiac CT images, it generates surface labels for all anatomical structures relevant to the simulation domain, including the LA body, PVs, MV annulus, and LAA. These segmented surfaces are represented in the three standard spatial planes: axial, coronal, and sagittal. They are then combined to form a single volumetric segmentation through a process known as stacking. The stacking operation merges segmentation slices from all three orientations into a unified 3D volume, which can then be exported as a frame-by-frame series representing the entire cardiac cycle.

Figure 3.7 illustrates the stacking process, which combines the post-nnU-Net DICOM images shown in Figures 3.5 and 3.6 into a unified 3D volume. Figure 3.8 presents example nnU-Net segmentations for both AF and SR cases. As shown, the model accurately captures the overall geometry of the LA, including the PV inlets, LAA morphology, and MV region. However, despite its strong baseline performance, automated segmentation can struggle in low-contrast areas, particularly near the PV inlets and LAA base, which often requires targeted manual refinement.




Figure 3.8 Segmentations using nnU-Net AF and SR – Case 1

3.1.2 Manual Refinement

Following automated segmentation, a manual refinement step was performed to improve anatomical accuracy and ensure compatibility with downstream CFD simulations. This step is essential in regions where the automated output contains non-physiological geometries, artifacts, or poorly aligned boundaries that could introduce non-physical flow patterns during simulation.

As illustrated in Figure 3.9, certain areas in the nnU-Net output, most often around the PV inlets, MV outlet, and LAA required manual correction. Using ITK-SNAP [33] and other 3D visualization tools, local adjustments were made to redefine inlet and outlet cut planes so that they were orthogonal to the expected direction of blood flow. Maintaining this orthogonality is critical for applying physiologically accurate boundary conditions and ensuring stable numerical performance in CFD simulations.





Figure 3.9 Automated nnU-Net segmentation vs. post manual refinement, showing improved anatomical accuracy AF – Case 1

Additionally, manual refinement was used to restore anatomical integrity in areas where the deep learning model produced ambiguous boundaries due to poor image contrast or anatomical variability. This was particularly relevant for the LAA, which exhibits significant inter-patient variation in shape and size.

Table 3.2: LAA Morphology

Case ID	LAA Morphology
Case 1 44243	Chicken Wing
Case 2 CVC	Cactus
Case 3 46483	Windsock

Table 3.2 summarizes the morphological characteristics of the LAA for each of the three cases analyzed in this study. The observed diversity in LAA morphology provides a valuable opportunity to assess how anatomical variation influences blood flow patterns and thrombogenic risk.



(A) AF (B) SR
Figure 3.10 Segmentations post Manual Refinement AF and SR – Case 1

The final manually refined geometries for all cases are presented in Figure 3.10. These geometries represent the finalized simulation-ready models that will be used in the meshing and CFD pipeline. By combining automated segmentation with expert-guided refinement, we ensured that the final geometries not only align with physiological expectations but are also computationally robust for high-fidelity simulations under both SR and AF conditions.

3.2 MODEL PREPARATION FOR SIMULATION

3.2.1 Volumetric study

With the completion of anatomical segmentation and manual refinement, the resulting geometries serve as the foundation for the meshing process. In this stage, the segmented volumes are transformed into computational domains through surface and volumetric discretization. The

primary objective is to generate a mesh that preserves anatomical accuracy while ensuring numerical stability and compatibility with the solver requirements. These meshes form the structural backbone of the simulation pipeline, enabling the application of boundary conditions and the accurate resolution of flow behavior within the LA chamber.

Table 3.3: Segmentation Time Frames for AF and SR Across Different Cases

Case ID	Original Segmentation Times
44243 SR	[0 10 30 35 40 50 60 70 80 90]
44243 AF	[-27 -18 -9 0 15 30 44 59 73 88]
CVC SR	[0 6 12 18 24 30 36 42 48 54 60 66 72 78 84 90 95]
CVC AF	[-18 -11 -6 0 7 14 21 28 35 42 49 56 63 70 77 84 89]
46843 SR	[0 10 20 30 35 40 50 60 70 80 90 95]
46483 AF	[1 22 43 64 85 104 119 135 150 165 180 195]

Table 3.4: Used Segmentation Time Points in AF Cases with Time Zero as Reference

Case ID	Used Segmentation Times
44243 AF	[0 15 30 44 59 73 88]
CVC AF	[0 7 14 21 28 35 42 49 56 63 70 77 84 89]
46483 AF	[4 19 35 50 65 80 95]

Tables 3.3 and 3.4 present the original and final segmentation time points used for meshing in each AF case. As explained in the methodology, the segmentation time percentage (*seg_time*) indicates the phase within an RR interval when each CT frame was acquired.

In Table 3.3, Case 1 AF and Case 2 AF include irregular or negative *seg_time* values. Negative values occur when frames are captured before the reference R-wave peak, meaning they belong to the previous cardiac cycle. Because the simulations focus on a single, physiologically consistent RR interval, these frames were excluded from meshing. Table 3.4 lists the filtered time points used for the simulations; for both Case 1 AF and Case 2 AF, only positive values within a single cycle were retained to ensure consistency within one cardiac period.

For Case 3 AF, the original dataset captured two complete RR intervals, reflected in segmentation times ranging from 1 to 195. To maintain temporal consistency and adequate frame density, the second RR interval was chosen for simulation, as it offered more complete and evenly spaced coverage. The final segmentation times in Table 3.4 reflect this choice, ensuring that each AF case is modeled using a physiologically representative, single RR cycle, aligned with the SR datasets used in the study.




Figure 3.11 . Maximum and minimum left atrial (LA) volumes during the cardiac cycle for each case in sinus rhythm (SR, green) and atrial fibrillation (AF, red).




Figure 3.12 Left atrial (LA) emptying fractions for each case in sinus rhythm (SR, green) and atrial fibrillation (AF, red)

The trends observed in Figure 3.11 align closely with the ECG patterns discussed earlier. In Case 1 AF, the ECG trace shows a pronounced peak during atrial activity compared to SR, suggesting premature or erratic atrial contraction. This electrical irregularity translates into a higher minimum LA volume during AF (113.18 mL) than in SR (100.89 mL), indicating that the atrium does not contract down to its smallest possible volume.

A similar pattern is seen in Case 2 AF, where AF results in a higher minimum volume and, interestingly, a lower maximum volume compared to SR. This reduction in peak volume can be linked to subtle irregularities in the ST phase of the ECG, which may limit complete atrial relaxation and filling. In Case 3 AF, the disorganized electrical activity observed in the ECG corresponds to a markedly reduced maximum LA volume, further indicating incomplete filling during AF.

Figure 3.12 summarizes these effects across all three cases, showing that AF consistently reduces the emptying fraction (EF) relative to SR. The emptying fraction, a normalized measure of how much blood is ejected relative to the atrium's maximum volume, is lower in AF because of both impaired contraction and incomplete filling.

This relationship is evident in the volume–time and EF plots within Figure 3.12. SR cases display a larger difference between maximum and minimum volumes, reflecting more effective contraction. Additionally, SR cases tend to have higher maximum volumes, indicating better expansion and filling. The EF plots reinforce this trend: in every case, EF is significantly lower for AF, highlighting reduced pumping efficiency and confirming the functional impact of arrhythmia on atrial performance.




Figure 3.13 Left atrial (LA) volume over time from 3 seconds onward. SR shows greater variation than AF, indicating stronger contractile function.

Figure 3.13 illustrates the change in LA volume for case 1 over time for the entire duration of the simulation. The horizontal axis begins at 3 seconds, which corresponds to the end of the initialization phase and the start of the main simulation phase, as discussed earlier. This time window captures the physiological behavior after the flow and structural fields have stabilized.

The plotted curves show the periodic compression and expansion of the LA. The red line represents the AF case, while the SR case is shown for comparison. It is evident that the SR curve consistently exhibits higher peak volumes during expansion and deeper troughs during contraction. This indicates that the LA in SR undergoes larger amplitude changes in volume, reflecting stronger and more coordinated contractile function compared to AF.




Figure 3.14 Final cardiac cycle showing reduced and abnormal volume changes in AF, reflecting impaired atrial contraction.

Figure 3.14 provides a closer view by zooming in on the final cardiac cycle of the simulation, spanning one second. This allows for a clearer comparison between the observed volume changes and expected physiological behavior. In the SR case, the volume decreases toward the end of the cycle, corresponding to atrial contraction, and increases again during the middle of the cycle as the atrium relaxes and fills.

In contrast, the AF case deviates significantly from this physiological pattern. Here, the volume reaches its maximum during the phase when the atria should be contracting, indicating ineffective mechanical activity. This reversal of the expected contraction-filling sequence highlights the loss of synchronized atrial function in AF. Overall, the volumetric analysis confirms that the case is representative of AF, with impaired contraction and reduced volume variation compared to SR.




Figure 3.15 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.

The differences between SR and AF are clearly illustrated in Figure 3.15, which shows the LA volume–time curves for the final cardiac cycle overlaid for direct comparison. Two key observations emerge from this plot. First, the range between peak and minimum volumes is consistently larger in SR cases, indicating a more effective contraction and relaxation cycle of the atria. This reflects the atria’s ability to empty and fill efficiently with each beat. Second, in SR cases, the minimum volume typically occurs just before the end of the cycle, coinciding with atrial contraction and aligning well with expected physiological timing. In contrast, AF cases show no consistent pattern. The timing of minimum volumes varies significantly between beats, with no clear relationship to specific phases of the cardiac cycle. This irregularity highlights the loss of coordinated atrial contraction in AF, leading to less efficient volume changes.

3.2.2 Smoothing of the volume




Figure 3.16 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.

After reconstructing each timepoint of the cardiac cycle, a dedicated smoothing process was applied to improve geometric fidelity and eliminate residual artifacts. This step addressed minor segmentation errors and mesh distortions common in high-resolution, noisy spatiotemporal data. The smoothing combined spatial and temporal components. Spatially, vertex positions were averaged locally to reduce surface irregularities without altering key anatomical features. Temporally, each frame's geometry was adjusted in relation to adjacent timepoints, ensuring smooth transitions and preserving cyclic continuity by linking the first and last frames.

This dual-domain smoothing enhanced anatomical consistency and fluid motion across the cardiac cycle, providing a stable foundation for downstream tasks such as mesh registration, deformation analysis, and CFD simulations.

3.2.3 Meshing




Figure 3.17 Mesh AF and SR – Case 1

Surface preprocessing was carried out for each case to generate simulation-ready meshes shown in figure 3.16 from the segmented LA geometries. The values of *mesh_keepratio* and *faces_cyl_thresh0* were tuned on a case-by-case basis to balance surface detail preservation with mesh smoothness and to ensure physiologically consistent inlet/outlet plane definitions shown in Table 3.6.

In AF geometries, *mesh_keepratio* values ranged from 0.6 to 0.8, reflecting the need to preserve fine-scale structures such as irregular pulmonary vein junctions while still removing surface noise. SR geometries generally used slightly higher ratios (0.7–0.8), consistent with their

smoother surface morphology. Adjustments to `faces_cyl_thresh0` targeted accurate classification of PV and MV boundaries, with lower thresholds applied in more complex AF surfaces to avoid over-classification of irregular regions.

Table 3.5: Mesh simplification parameters (`mesh_keepratio` and `faces_cyl_thresh0`) assigned for each case in the surface preprocessing stage.

Case ID	<code>mesh_keepratio</code>	<code>faces_cyl_thresh0</code>
Case 1 AF	0.6	0.8
Case 1 SR	0.8	0.8
Case 2 AF	0.6	0.5
Case 2 SR	0.7	0.6
Case 3 AF	0.8	0.6
Case 3 SR	0.8	0.6

Table 3.6: Total Number of vertices and faces

Case ID	Number of Vertices	Number of Triangles
Case 1 SR	135545	261822
Case 1 AF	107668	207388
Case 2 SR	154245	308486
Case 2 AF	155717	302329
Case 3 SR	124060	240460
Case 3 AF	159267	309890

Table 3.7: Eulerian Grid Resolution for each case

Case ID	Eulerian Grid Resolution
Case 1 SR	256x256x256
Case 1 AF	256x256x256
Case 2 SR	256x288x270
Case 2 AF	256x274x264
Case 3 SR	256x270x256
Case 3 AF	256x280x256




Figure 3.18 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 1

Figure 3.16 presents representative AF and SR meshes for Case 1, highlighting differences in surface complexity. Figure 3.17 shows the relationship between the number of triangles and the characteristic length scale. The optimal scale (red line) served as a target for parameter tuning. In all cases, adjustments to mesh_keepratio successfully aligned the mesh characteristics closer to this optimal range, reducing both over-simplification and excessive surface refinement, thereby ensuring suitable quality for CFD simulation.

3.2.4 Boundary Condition




Figure 3.19 Inlets / outlet which are for the CFD simulation for AF and SR – Case 1.

Table 3.8: Assigned colors for key LA structures used for visualization.

Yellow	Right Superior PV
Orange	Right Inferior PV
Blue	Left Superior PV
Sky Blue	Left Inferior PV
Red	MV

The boundary classification procedure successfully identified and labeled all PV inlets and the MV outlet for both AF and SR geometries in Case 1. Figure 3.18 illustrates the resulting inlet–outlet surfaces, color-coded according to their anatomical assignment (Table 3.9). For both rhythm types, the right superior PV (yellow) and right inferior PV (orange) are clearly segmented and separated

from adjacent atrial wall structures. Similarly, the left superior PV (blue) and left inferior PV (sky blue) are distinctly resolved, with inlet planes aligned orthogonal to the local vessel axis. The MV outlet (red) appears as a continuous, well-defined surface in both cases, with no evidence of misalignment or overlap with the atrial wall.

3.3 CFD SIMULATION

Once the CPD-registered motion fields and surface meshes are finalized, the CFD simulations are initiated through a fully automated pipeline. This system processes each case by reading the motion and domain files, assigning case-specific parameters such as hematocrit and cardiac period, and launching the immersed boundary solver without manual configuration. The automation also handles post-processing tasks such as exporting flow fields, computing residence time, and organizing output files for visualization.

Table 3.9: Simulation Runtime and GPU Used for Each Case and Rhythm Type

Case ID	Runtime	GPU
Case 1 AF	20h 4m	A100
Case 1 SR	20h 39m	A100
Case 2 AF	3d 22h 56m	A40
Case 2 SR	7d 6h 39m	A40
Case 3 AF	7d 3h 28m	A40
Case 3 SR	4d 14h 51m	A40

Table 3.7 presents the production-phase simulation runtimes for each patient case. These runs include the final 20 cardiac cycles used for data extraction and analysis. It can be observed that simulations executed on Nvidia A100 GPUs consistently outperform those run on A40s.

3.4 POST PROCESSING

Once the CFD simulations are completed, the results undergo a structured post-processing stage to convert raw numerical output into physiologically meaningful data. This step is crucial because the solver produces vast amounts of information that, in its raw form, is not directly interpretable. The post-processing focuses on extracting key hemodynamic metrics, including velocity fields, patterns of kinetic energy distribution, regions of high or low residence time, variations in blood viscosity, and Eulerian residual domains that indicate flow recirculation. These parameters are selected because they provide insight into how blood moves through the geometry, how efficiently it is transported, and where stagnation or abnormal shear conditions may occur.

3.4.1 Velocity

To ensure reliability and consistency, an in-house automated pipeline organizes the simulation output, separates it into meaningful categories, and prepares it for both numerical analysis and visualization. This approach maintains the full spatial and temporal resolution of the original simulations, which is essential for identifying subtle differences between patient cases and heart rhythms. By standardizing the workflow, direct comparisons can be made between cases while minimizing the risk of processing-related bias or data loss.

One full cardiac cycle in these simulations corresponds to 1 second, represented by 40 frames in the velocity field visualizations. As seen in Figure 3.19, the for the AF case the velocity within the LAA remains consistently low throughout the cardiac cycle, with very limited fluid exchange between the LAA and the main atrial cavity. This stagnant environment is particularly evident in Figure 3.19 (B), where a large amount of chaotic, irregular motion can be observed. While this motion appears active, it is primarily localized and disorganized, with minimal net displacement

of blood either into or out of the appendage. Such chaotic fluctuations are characteristic of AF, where the absence of coordinated atrial contraction leads to ineffective washout. This low and erratic flow state is known to increase blood residence time, which in turn heightens the risk of thrombus formation within the LAA.




Figure 3.20 Velocity contours of the LA for Case 1 AF

In contrast, the SR case (Figure 3.20) displays well-organized and periodic patterns of fluid motion within the LAA. The velocity field clearly shows inflow and outflow phases that are synchronized with atrial contraction, enabling more effective blood exchange between the appendage and the

atrial cavity. These structured flow patterns reduce stagnant zones and promote more uniform velocity distributions within the LAA, thereby lowering the likelihood of prolonged blood residence and clot formation. This difference between AF and SR not only highlights the role of atrial contraction in maintaining healthy hemodynamics but also underscores why patients with AF are at significantly higher risk for embolic events originating from the LAA.




Figure 3.21 Velocity contours of the LA for Case 1 SR

In Case 2, the simulations reveal that blood flow from the right pulmonary veins into the LA is noticeably higher during Sinus Rhythm compared to Atrial Fibrillation. This increased velocity is likely a result of the pressure drop generated by coordinated atrial contraction, which effectively pulls blood into the chamber. In contrast, during Atrial Fibrillation, this organized contraction is absent, leading to slower and less efficient inflow.

In Case 3, similar trends are observed within the LA and MV. The Sinus Rhythm case demonstrates a strong, coherent velocity jet through the mitral valve, reflecting effective atrial ejection and coordinated contraction. In the Atrial Fibrillation case, this ejection is significantly diminished, highlighting the reduced contractile function and disrupted flow patterns. Some irregularities in the AF geometry are also apparent, primarily due to noise and distortions in the original imaging data, but despite these limitations, the distinction between coordinated flow in Sinus Rhythm and disorganized, sluggish flow in Atrial Fibrillation is clearly evident. Overall, these comparisons emphasize the functional impact of rhythm on atrial hemodynamics, with Sinus Rhythm supporting efficient filling and ejection while Atrial Fibrillation results in impaired circulation and stasis within the chamber.

3.4.3 Kinetic Energy

For the LA, the velocity contour plots provided a clear visual understanding of the global flow dynamics. However, when focusing specifically on the LAA, these plots were less effective in revealing meaningful qualitative differences in flow patterns. To gain a more quantitative and interpretable measure of activity within the LAA over the cardiac cycle, kinetic energy was analyzed as an alternative metric.

Figure 3.21 shows the kinetic energy in the LAA for all patients and all cardiac cycles, overlaid in a single plot. This representation allows for a direct comparison of overall flow activity in the appendage across different cases and conditions. By combining data from all cycles, it provides a broader perspective on how active or relatively stagnant the LAA flow is under various physiological states.




Figure 3.22 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.

Figure 3.22 focuses on just the final cardiac cycle, making the trends clearer. In every case, the SR patients consistently exhibit higher kinetic energy in the LAA compared to those with AF. This pattern indicates that the LAA experiences greater internal motion and flow during normal sinus rhythm. Conversely, AF cases show reduced kinetic activity, suggesting a more stagnant flow environment that could have clinical implications for thrombus formation.




Figure 3.23 Left atrial (LA) volume time curves for the final cardiac cycle, comparing SR and AF cases to highlight differences in contraction patterns.

3.4.3 Residence Time

In the comparison between the AF and SR cases for the same patient, the differences in LAA flow are particularly striking and provide important insights into the hemodynamic consequences of arrhythmia. As illustrated in Figure 3.23 for AF and Figure 3.24 for SR, the atrial contraction in SR is highly coordinated, producing an effective washout mechanism that displaces older blood from the appendage and replaces it with fresh blood from the atrium. This coordinated motion ensures that blood does not linger in low-flow regions and supports efficient circulation throughout the LA. In contrast, the AF case lacks this synchronized contraction, resulting in much slower blood movement within the appendage. Once blood enters the deeper recesses of the LAA, it tends to remain trapped, forming persistent zones of low velocity and stagnation. These stagnant regions are particularly concerning because they create favorable conditions for thrombus formation and highlight the impaired hemodynamics associated with AF. The visualization of these flow patterns

underscores the critical role of atrial mechanics in maintaining effective blood clearance and demonstrates how disruptions in rhythm can dramatically alter circulation within the appendage, emphasizing the clinical importance of understanding patient-specific LAA flow dynamics.




Figure 3.24 Residence Time in the LA for Case 1 AF




Figure 3.25 Residence time in the LA for Case 1 SR

In Case 2, the pattern is clear AF leads to more regions of high residence time within the LAA, reflecting prolonged blood stagnation. This is most evident in the deeper recesses of the appendage, where the lack of coordinated atrial contraction prevents effective clearance. In SR, however, atrial contraction promotes a steady displacement of older blood, maintaining continuous turnover and healthier flow dynamics.

In Case 3, the difference between rhythms is even more pronounced. Under AF, the blood remains largely static in the appendage throughout the cycle, creating persistent low-flow zones. In SR, by contrast, the flow is more dynamic, with blood actively circulating in and out of the appendage during the cardiac cycle. This enhanced washout in SR points to more efficient circulation and a reduced risk of thrombus formation.




Figure 3.26 LAA volume AF and SR – Case 1

Now turning to a more quantitative assessment for Case 1, the LAA was segmented into three distinct regions: distal, medial, and proximal, as shown in Figure 3.25, to better capture the spatial distribution of blood residence time. The corresponding residence time values are presented in Figure 3.26. Among these regions, the distal segment, located at the farthest end of the appendage, consistently exhibits the highest average residence time in both AF and SR. The contrast between rhythms is particularly pronounced: in AF, the distal region shows substantially elevated residence times, indicating severe stagnation compared to SR.

The time-history plots further highlight these rhythmic differences. In SR, the residence time traces display regular, well-defined oscillations across the cardiac cycle, reflecting the periodic compression and expansion of the LAA that promotes effective washout. In AF, however, these oscillations are markedly dampened or blunted, consistent with the absence of coordinated

contractile activity. This diminished pumping action reduces blood turnover within the appendage, especially in the distal region, and may contribute to the elevated thromboembolic risk associated with AF.




Figure 3.27 Average Residence Time compared in different regions of the LAA for AF and SR – Case 1.

Figure 3.27 presents the consolidated average residence time across the entire LAA for both AF and SR. The most striking feature in both cases is the evolution of the curve over time. At the start, the average residence time increases almost linearly, indicating a steady accumulation. As the simulation progresses, this growth transitions into a square-root-like pattern before gradually flattening into a plateau. This plateau signifies that the residence time is stabilizing, meaning the simulation is approaching a converged state.

In the sinus rhythm case, the curve in Figure 3.27 has nearly plateaued, suggesting that the simulation has either reached or is very close to convergence. In contrast, the AF curve has not fully leveled off, implying that the simulation may require additional cardiac cycles to achieve complete convergence. This difference reinforces the notion that AF involves more prolonged and irregular flow patterns, which take longer to stabilize.




Figure 3.28 Average Residence Time in LAA for AF and SR – Case 1.

Figure 3.28 presents an overlay of the average residence time curves for all simulated cases, directly comparing AF and SR on the same plot. This combined view highlights patterns that are less apparent when cases are examined individually. Across nearly all cases, AF consistently exhibits higher residence times than SR, indicating slower blood clearance and greater stagnation within the LAA. The primary exception is Case 2, where the AF and SR curves lie closer together. Even in this case, however, AF still shows a measurable increase in residence time relative to SR.

This consistent elevation in AF residence time carries important physiological implications. Prolonged blood residence within the LAA can increase the likelihood of clot formation, a major risk factor for thromboembolic events such as stroke. The magnitude of the difference between AF and SR varies across cases, suggesting that anatomical variability, flow patterns, or cycle-to-cycle differences may influence how strongly AF alters residence time. Despite these variations, the overarching trend is clear: AF promotes more stagnant flow within the LAA compared to SR, reinforcing clinical observations of elevated thrombus risk in AF patients.




Figure 3.29 Average Residence Time compared in different regions of the LAA in AF and SR – Case 1.

Figure 3.29 presents the average residence time data in the form of a violin plot. This visualization highlights not only the central tendency of the data but also its overall distribution, illustrating where most values are concentrated and how widely they are spread. the AF cases show a broad distribution, with most residence times clustering between 4 and 7 seconds. The white markers embedded within each violin represent the mean residence times, which generally lie between 4.0 and 4.5 seconds. These findings indicate that during AF, the LAA retains blood for longer

durations, producing conditions that favor stagnation and increase the likelihood of thrombus formation.

In contrast, the figure 3.30 (B), which represents the SR cases, reveals a narrower and lower distribution. Most values fall between 3 and 5 seconds, and the mean residence times are consistently lower, typically between 3.0 and 3.5 seconds. This distribution suggests that during SR, blood is cleared from the LAA more efficiently, reducing the potential for stasis.

The contrast between the two violin plots reinforces the trends observed in earlier figures. AF cases exhibit higher residence times, consistent with slower and more stagnant blood flow, while SR cases show lower residence times, reflecting more effective washout. By simultaneously presenting both averages and distributions, the violin plot strengthens the conclusion that atrial fibrillation, independent of patient-specific variability, systematically results in higher blood residence times in the LAA compared to SR.




Figure 3.30 Violin plots of residence time distributions for AF and SR cases




Figure 3.31 *Maximum residence time trends for AF and SR cases across all frames.*

Figure 3.30 presents the maximum residence time across all frames for each patient, comparing both AF and SR cases. In the AF cases, the maximum residence time continues to rise over the course of the cardiac cycle, indicating that certain regions of the LAA experience persistent stagnation and that the flow has not yet fully converged. This ongoing increase highlights the impaired hemodynamics characteristic of AF, where coordinated atrial contraction is absent, leading to prolonged retention of blood in deeper appendage regions. Conversely, the SR cases show maximum residence time values that plateau relatively quickly, suggesting that the flow stabilizes and becomes more consistent over time. The contrast between AF and SR in these maximum residence time trends underscores a clinically relevant insight: the prolonged and increasing residence times observed in AF may correspond to higher risk for thrombus formation due to blood stasis.

Chapter 4. CONCLUSION

4.1 Conclusions

This work introduces a comprehensive, patient-specific computational framework for simulating left atrial (LA) hemodynamics under sinus rhythm (SR) and atrial fibrillation (AF). By combining ECG-gated cardiac CT imaging, hematocrit-dependent non-Newtonian blood models, and time-resolved wall motion, the workflow provides physiologically realistic CFD simulations that capture both anatomical variability and rhythm-specific flow alterations.

High temporal resolution across multiple cardiac cycles enabled the characterization of intra- and inter-cycle dynamics, particularly within the left atrial appendage (LAA). Results consistently demonstrated that AF leads to elevated residence times and reduced kinetic energy compared to SR, reflecting impaired washout and greater stagnation. In contrast, SR preserved coordinated contraction and more effective clearance. Together, these findings validate the robustness of the workflow and highlight its translational potential for patient-specific assessment of thrombogenic risk.

4.2 Limitations

Several limitations constrain the current study. First, access to high-quality AF-specific imaging remains limited, especially for paroxysmal AF, reducing validation opportunities and case representativeness. Segmentation accuracy presents another challenge, particularly near the pulmonary vein inlets and LAA, where low contrast or noise may introduce geometric uncertainties despite manual refinement.

Computationally, the spatial resolution restricts the capture of small-scale vortical structures, while inlet and outlet boundary conditions are derived from representative cycles, not true beat-to-beat variability. This introduces an assumption of periodicity that may not hold during AF. Moreover, the workflow is computationally intensive, requiring high-performance GPU resources and extended runtimes, which currently limits its suitability for time-sensitive clinical applications.

4.3 Future Opportunities

Future work could expand the framework in several important directions. Extending simulations to incorporate the right atrium (RA) and right ventricles (RV) would enable chamber-to-chamber interaction studies, bridging electromechanical and hemodynamic analyses. Coupling with thrombus-prediction models—such as clot transport equations or shear-based thresholds—could directly quantify thrombogenic potential.

On the technical side, increased automation from segmentation to simulation would improve scalability and reduce operator dependence. Machine learning could also be used to estimate AF-induced wall motion in cases with only SR imaging, broadening applicability to larger patient cohorts. Finally, applying this pipeline to post-intervention studies or patient-specific risk stratification could accelerate its translation into clinical decision-making and personalized treatment planning.

BIBLIOGRAPHY

- [1] M. Jiao, C. Liu, Y. Liu, Y. Wang, Q. Gao, and A. Ma, “Estimates of the global, regional, and national burden of atrial fibrillation in older adults from 1990 to 2019: insights from the Global Burden of Disease study 2019,” *Front. Public Health*, vol. 11, p. 1137230, June 2023, doi: 10.3389/fpubh.2023.1137230.
- [2] G. Lippi, F. Sanchis-Gomar, and G. Cervellin, “Global epidemiology of atrial fibrillation: An increasing epidemic and public health challenge,” *Int. J. Stroke*, vol. 16, no. 2, pp. 217–221, Feb. 2021, doi: 10.1177/1747493019897870.
- [3] J. P. Piccini and G. C. Fonarow, “Preventing Stroke in Patients With Atrial Fibrillation—A Steep Climb Away From Achieving Peak Performance,” *JAMA Cardiol.*, vol. 1, no. 1, p. 63, Apr. 2016, doi: 10.1001/jamacardio.2015.0382.
- [4] A. Odutayo, C. X. Wong, A. J. Hsiao, S. Hopewell, D. G. Altman, and C. A. Emdin, “Atrial fibrillation and risks of cardiovascular disease, renal disease, and death: systematic review and meta-analysis,” *BMJ*, p. i4482, Sept. 2016, doi: 10.1136/bmj.i4482.
- [5] F. Castells, C. Mora, J. Millet, J. J. Rieta, C. Sánchez, and J. M. Sanchís, “Multidimensional ICA for the Separation of Atrial and Ventricular Activities from Single Lead ECGs in Paroxysmal Atrial Fibrillation Episodes,” in *Independent Component Analysis and Blind Signal Separation*, vol. 3195, C. G. Puntonet and A. Prieto, Eds., in Lecture Notes in Computer Science, vol. 3195. , Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 1229–1236. doi: 10.1007/978-3-540-30110-3_155.
- [6] A. Vraká, V. Bertomeu-González, F. Hornero, A. Quesada, R. Alcaraz, and J. J. Rieta, “Splitting the P-Wave: Improved Evaluation of Left Atrial Substrate Modification after

- Pulmonary Vein Isolation of Paroxysmal Atrial Fibrillation,” *Sensors*, vol. 22, no. 1, p. 290, Dec. 2021, doi: 10.3390/s22010290.
- [7] L. Di Biase *et al.*, “Does the Left Atrial Appendage Morphology Correlate With the Risk of Stroke in Patients With Atrial Fibrillation?,” *J. Am. Coll. Cardiol.*, vol. 60, no. 6, pp. 531–538, Aug. 2012, doi: 10.1016/j.jacc.2012.04.032.
- [8] C. Stöcklberger, G. Ernst, E. Bonner, J. Finsterer, and J. Slany, “Left atrial appendage morphology: comparison of transesophageal images and postmortem casts,” *Z. Für Kardiologie*, vol. 92, no. 4, pp. 303–308, Apr. 2003, doi: 10.1007/s00392-003-0903-x.
- [9] R. Beigel, N. C. Wunderlich, S. Y. Ho, R. Arsanjani, and R. J. Siegel, “The Left Atrial Appendage: Anatomy, Function, and Noninvasive Evaluation,” *JACC Cardiovasc. Imaging*, vol. 7, no. 12, pp. 1251–1265, Dec. 2014, doi: 10.1016/j.jcmg.2014.08.009.
- [10] S. Bäck *et al.*, “Elevated atrial blood stasis in paroxysmal atrial fibrillation during sinus rhythm: a patient-specific computational fluid dynamics study,” *Front. Cardiovasc. Med.*, vol. 10, p. 1219021, Aug. 2023, doi: 10.3389/fcvm.2023.1219021.
- [11] R. M. Heppell, K. E. Berkin, J. M. McLenachan, and J. A. Davies, “Haemostatic and haemodynamic abnormalities associated with left atrial thrombosis in non-rheumatic atrial fibrillation,” *Heart*, vol. 77, no. 5, pp. 407–411, May 1997, doi: 10.1136/hrt.77.5.407.
- [12] A. A. Khan, G. N. Thomas, G. Y. H. Lip, and A. Shantsila, “Endothelial function in patients with atrial fibrillation,” *Ann. Med.*, vol. 52, no. 1–2, pp. 1–11, Feb. 2020, doi: 10.1080/07853890.2019.1711158.
- [13] A. N. Ganesan *et al.*, “The impact of atrial fibrillation type on the risk of thromboembolism, mortality, and bleeding: a systematic review and meta-analysis,” *Eur. Heart J.*, vol. 37, no. 20, pp. 1591–1602, May 2016, doi: 10.1093/eurheartj/ehw007.

- [14] Q. Bao *et al.*, “Clinical Significance of Plasma D-Dimer and Fibrinogen in Outcomes after Stroke: A Systematic Review and Meta-Analysis,” *Cerebrovasc. Dis.*, vol. 52, no. 3, pp. 318–343, 2023, doi: 10.1159/000526476.
- [15] W. Y. Ding, D. Gupta, and G. Y. H. Lip, “Atrial fibrillation and the prothrombotic state: revisiting Virchow’s triad in 2020,” *Heart*, vol. 106, no. 19, pp. 1463–1468, Oct. 2020, doi: 10.1136/heartjnl-2020-316977.
- [16] T. Watson, E. Shantsila, and G. Y. Lip, “Mechanisms of thrombogenesis in atrial fibrillation: Virchow’s triad revisited,” *The Lancet*, vol. 373, no. 9658, pp. 155–166, Jan. 2009, doi: 10.1016/S0140-6736(09)60040-4.
- [17] A. Qureshi *et al.*, “MRI-based modelling of left atrial flow and coagulation to predict risk of thrombogenesis in atrial fibrillation,” *Med. Image Anal.*, vol. 101, p. 103475, Apr. 2025, doi: 10.1016/j.media.2025.103475.
- [18] R. Gerrah and S. J. Haller, “Computational fluid dynamics: a primer for congenital heart disease clinicians,” *Asian Cardiovasc. Thorac. Ann.*, vol. 28, no. 8, pp. 520–532, Oct. 2020, doi: 10.1177/0218492320957163.
- [19] L. Parker *et al.*, “A multi-modal computational fluid dynamics model of left atrial fibrillation haemodynamics validated with 4D flow MRI,” *Biomech. Model. Mechanobiol.*, vol. 24, no. 1, pp. 139–152, Feb. 2025, doi: 10.1007/s10237-024-01901-y.
- [20] A. Masci *et al.*, “Development of a Computational Fluid Dynamics Model of the Left Atrium in Atrial Fibrillation on a Patient Specific Basis,” presented at the 2017 Computing in Cardiology Conference, Sept. 2017. doi: 10.22489/CinC.2017.004-429.

- [21] M. Corti, A. Zingaro, L. Dede', and A. M. Quarteroni, "Impact of atrial fibrillation on left atrium haemodynamics: A computational fluid dynamics study," *Comput. Biol. Med.*, vol. 150, p. 106143, Nov. 2022, doi: 10.1016/j.combiomed.2022.106143.
- [22] L. Wang, Z. Wang, R. Fang, and Z.-Y. Li, "Evaluation of Stroke Risk in Patients With Atrial Fibrillation Using Morphological and Hemodynamic Characteristics," *Front. Cardiovasc. Med.*, vol. 9, p. 842364, Apr. 2022, doi: 10.3389/fcvm.2022.842364.
- [23] J. Dueñas-Pamplona *et al.*, "Reduced-order models of wall shear stress patterns in the left atrial appendage from a data-augmented atrial database," *Appl. Math. Model.*, vol. 130, pp. 713–727, June 2024, doi: 10.1016/j.apm.2024.03.027.
- [24] "Ahmed Qureshi *et al.*, "Modelling Virchow's Triad to Improve Stroke Risk Assessment in Atrial Fibrillation Patients," presented at the 2022 Computing in Cardiology Conference, Dec. 2022. doi: 10.22489/CinC.2022.378.
- [25] A. Qureshi *et al.*, "Mechanistic modelling of Virchows triad to assess thrombogenicity and stroke risk in atrial fibrillation patients," *Eur. Heart J. - Digit. Health*, vol. 3, no. 4, p. ztac076.2788, Dec. 2022, doi: 10.1093/ehjdh/ztac076.2788.
- [26] A. Darlington and M. D. McCauley, "Atrial Cardiomyopathy: An Unexplored Limb of Virchow's Triad for AF Stroke Prophylaxis," *Front. Cardiovasc. Med.*, vol. 7, p. 11, Feb. 2020, doi: 10.3389/fcvm.2020.00011.
- [27] Y. Sun *et al.*, "Finding low CHA2DS2-VASc scores unreliable? Why not give morphological and hemodynamic methods a try?," *Front. Cardiovasc. Med.*, vol. 9, p. 1032736, Jan. 2023, doi: 10.3389/fcvm.2022.1032736.

- [28] W. D. Bidgood, S. C. Horii, F. W. Prior, and D. E. Van Syckle, “Understanding and Using DICOM, the Data Interchange Standard for Biomedical Imaging,” *J. Am. Med. Inform. Assoc.*, vol. 4, no. 3, pp. 199–212, May 1997, doi: 10.1136/jamia.1997.0040199.
- [29] M. Gunawardhana, F. Xu, and J. Zhao, “How good nnU-Net for Segmenting Cardiac MRI: A Comprehensive Evaluation,” 2024, *arXiv*. doi: 10.48550/ARXIV.2408.06358.
- [30] A. Zingaro *et al.*, “A comprehensive stroke risk assessment by combining atrial computational fluid dynamics simulations and functional patient data,” Jan. 15, 2024, *Bioengineering*. doi: 10.1101/2024.01.11.575156.
- [31] Y. Zhu, L. Chen, W. Lu, Y. Gong, and X. Wang, “The application of the nnU-Net-based automatic segmentation model in assisting carotid artery stenosis and carotid atherosclerotic plaque evaluation,” *Front. Physiol.*, vol. 13, p. 1057800, Dec. 2022, doi: 10.3389/fphys.2022.1057800.
- [32] R. Kato *et al.*, “Pulmonary Vein Anatomy in Patients Undergoing Catheter Ablation of Atrial Fibrillation: Lessons Learned by Use of Magnetic Resonance Imaging,” *Circulation*, vol. 107, no. 15, pp. 2004–2010, Apr. 2003, doi: 10.1161/01.CIR.0000061951.81767.4E.
- [33] P. A. Yushkevich *et al.*, “User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability,” *NeuroImage*, vol. 31, no. 3, pp. 1116–1128, July 2006, doi: 10.1016/j.neuroimage.2006.01.015.

APPENDIX A




Figure A.1 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 2.




Figure A.2 DICOM Images for Atrial Fibrillation (AF) – Case 2





Figure A.3 DICOM Images for Sinus Rhythm (SR) – Case 2



(A) AF (B) SR
Figure A.4 ECG Comparison Between AF and SR – Case 2




Figure A.5 Highlighted regions indicate localization areas targeted during CT cardio-angiographic imaging for cardiovascular assessment of Case 3.




Figure A.6 DICOM Images for Atrial Fibrillation (AF) – Case 3




Figure A.7 DICOM Images for Sinus Rhythm (SR) – Case 3




Figure A.8 ECG Comparison Between AF and SR – Case 3




Figure A.9 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF – Case 2




Figure A.10 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 2




Figure A.11 Segmentations using nnU-Net AF and SR – Case 2




Figure A.12 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for AF – Case 3




Figure A.13 Colored Highlighting of Anatomical Regions of Interest via nnU-Net for SR – Case 3






Figure A.14 Segmentations using nnU-Net AF and SR – Case 3




(A) AF




(B) SR

Figure A.15 Mesh AF and SR – Case 2




(A) AF




(B) SR

Figure A.16 Mesh AF and SR – Case 3




(A) AF




(B) SR

Figure A.17 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 2



(A) AF



(B) SR

Figure A.18 Number of Triangles compared to the Triangles characteristic length (cm) AF and SR – Case 3




(A) AF – Case 2




(B) SR – Case 2

Figure A.19 Inlets / outlet which are for the CFD simulation for AF and SR – Case 2.




(A) AF – Case 3




(B) SR – Case 3

Figure A.20 Inlets / outlet which are for the CFD simulation for AF and SR – Case 3.




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.21 Velocity contours of the LA for Case 2 AF




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.22 Velocity contours of the LA for Case 2 SR




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.23 Velocity contours of the LA for Case 3 AF



(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30



Figure A.24 Velocity contours of the LA for Case 3 SR




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.25 Residence Time in the LA for Case 2 AF




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.26 Residence Time in the LA for Case 2 SR




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10




(C) Time = 0.5 s; Frame = 20




(D) Time = 0.75 s; Frame = 30




Figure A.27 Residence Time in the LA for Case 3 AF




(A) Time = 0 s; Frame = 0



(B) Time = 0.25 s; Frame = 10



(C) Time = 0.5 s; Frame = 20



(D) Time = 0.75 s; Frame = 30



Figure A.28 Residence Time in the LA for Case 3 SR