Medical imaging, shapes and statistics for stroke prediction in atrial fibrillation

Josquin Harrison

▶ To cite this version:

Josquin Harrison. Medical imaging, shapes and statistics for stroke prediction in atrial fibrillation. Medical Imaging. Université Côte d'Azur, 2024. English. NNT : 2024COAZ4027 . tel-04709084

HAL Id: tel-04709084
https://theses.hal.science/tel-04709084v1

Submitted on 25 Sep 2024

UNIVERSITÉ CÔTE D'AZUR

ÉCOLE DOCTORALE
SCIENCES ET TECHNOLOGIES
DE L'INFORMATION ET
DE LA COMMUNICATION

THÈSE DE DOCTORAT

Imagerie médicale, formes et statistiques pour la prédiction du risque d'accident vasculaire cérébral dans le cadre de la fibrillation atriale

Josquin Harrison
Centre Inria d'Université Côte d'Azur, Équipe Epione

Thése Dirigée par Dr Maxime Sermesant
Co-encadrée par Pr Hubert Cochet

Présentée en vue de l'obtention du grade de docteur en AUTOMATIQUE, TRAITEMENT DU SIGNAL ET DES IMAGES d'Université Côte d'Azur

Soutenue le 20 juin 2024

Devant le jury, composé de :

Pr Rasmus PAULSEN	DTU Danemark	Rapporteur
Pr Alistair YOUNG	King's College London	Rapporteur
Dr Xavier PENNEC	Inria	Examinateur
Pr Oscar CAMARA REY	UPF Barcelona	Examinateur
Dr Maxime SERMESANT	Inria	Directeur de thèse
Pr Hubert COCHET	Université de Bordeaux	Co-Encadrant

Inria

Epione
e-patient / e-medicine

liryc
L'institut des maladies du rythme cardiaque

Medical Imaging, Shapes and Statistics for Stroke Prediction in Atrial Fibrillation

Josquin Harrison
Centre Inria d'Université Côte d'Azur, Équipe Epione

Supervised by Dr Maxime Sermesant
Co-supervised by Pr Hubert Cochet

Presented to obtain the title of Doctor in AUTOMATIQUE, TRAITEMENT DU SIGNAL ET DES IMAGES of Université Côte d'Azur.

Defended on June 20th, 2024

Jury:

Pr Rasmus PAULSEN	DTU Danemark	Reviewer
Pr Alistair YOUNG	King's College London	Reviewer
Dr Xavier PENNEC	Inria	Examiner
Pr Oscar CAMARA REY	UPF Barcelona	Examiner
Dr Maxime SERMESANT	Inria	PhD Supervisor
Pr Hubert COCHET	Bordeaux University	Co-supervisor

Imagerie médicale, formes et statistiques pour la prédiction du risque d'accident vasculaire cérébral dans le cadre de la fibrillation atriale

Jury:

Président du Jury
Xavier Pennec, directeur de recherche, Inria

Rapporteurs
Rasmus PAULSEN, professeur des universités, DTU Danemark
Alistair YOUNG, professeur des universités, King's College London

Examinateur
Oscar CAMARA REY, professeur des universités, UPF Barcelona

Directeurs de thèse
Maxime Sermesant, directeur de recherche, Inria

Co-encadrant de thèse
Hubert Cochet, professeur des universités - praticien hospitalier, Université de Bordeaux

Abstract

Atrial Fibrillation (AF) is a complex heart disease of epidemic proportions. It is characterized by chaotic electrical activation which creates a haemodynamic environment prone to clot formation and an increase in risk of ischemic strokes. Although treatments and interventions exist to reduce stroke incidence, they often imply an increase in risk of other complications or consist in invasive procedures. As so, attempts of stratifying stroke risk in AF is of crucial importance for clinical decision-making. However, current widely used risk scores only rely on basic patient information and show poor performance. Importantly, no known markers reflect the mechanistic process of stroke, all the while more and more patient data is routinely available. In parallel, many clinical experts have hypothesized that the Left Atrium (LA) has an important role in stroke occurrence, yet have only relied on subjective measures to verify it. In this study, we aim at taking advantage of the evolving patient imaging stratification to substantiate this claim. Linking the anatomy of the LA to the risk of stroke can directly be translated into a geometric problem. Thankfully, the study and analysis of shapes knows a long-standing mathematical history, in theory as well as application, of which we can take full advantage. We first walk through the many facets of shape analysis, to realise that, while powerful, global methods lack clinically meaningful interpretations. We then set out to use these general tools to build a compact representation specific to the LA, enabling a more interpretable study. This first attempt allows us to identify key facts for a realistic solution to the study of the LA. Amongst them, any tool we build must be fast and robust enough for potentially large and prospective studies. Since the computational crux of our initial pipeline lies in the semantic segmentation of the anatomical parts of the LA, we focus on the use of neural networks specifically designed for surfaces to accelerate this problem. In particular, we show that representing input shapes using principal curvature is a better choice than what is currently used, regardless the architecture. As we iteratively update our pipeline, we further the use of the semantic segmentation and the compact representation by proposing a set of expressive geometric features describing the LA which are well in line with clinicians expectations yet offering the possibility for robust quantitative analysis. We make use of these local features and shed light on the complex relations between LA shape and stroke incidence, by conducting statistical analysis and classification using decision tree based methods. Results yield valuable insights for stroke prediction: a list of shape features directly linked to stroke patients; features that explain important indicators of haemodynamic disorder; and a better understanding of the impact of AF state related LA remodelling. Finally, we discuss other possible use of the set of tools developed in this work, from larger cohorts study, to the integration into multi-modal models, as well as opening possibilities for precise sensitivity analysis of haemodynamic simulation, a valuable next step to better understand the mechanistic process of stroke.

Keywords: Shape statistics, Cerebrovascular accident, Machine learning, Medical imaging, Atrial fibrillation.

Résumé

La Fibrillation Atriale (FA) est une maladie cardiaque complexe aux proportions épidémiques. Elle se caractérise par une activation électrique chaotique et crée un environnement hémodynamique propice à la formation de caillots et à l'augmentation du risque d'accidents vasculaires cérébraux (AVC) ischémiques. Bien qu'il existe des traitements et interventions préventives visant à réduire l'incidence d'AVC, ils impliquent une augmentation du risque d'autres complications médicales ou consistent en des procédures invasives. C'est pourquoi les tentatives de stratification du risque d'AVC dans la FA sont d'une importance cruciale pour la prise de décision clinique. Malgré cela, les scores de risque actuellement utilisés ne reposent que sur des informations élémentaires du patient et sont peu performants. Aucun marqueur connu ne reflète le processus mécanique de l'AVC, tandis que de plus en plus de données sont collectées de façon routinière sans être utilisées. Bien que de nombreuses études cliniques suggèrent que l'Atrium Gauche (AG) joue un rôle important dans l'occurence d'AVC, celles-ci ce basent sur des mesures subjectives pour le vérifier. En contraste, nous souhaitons exploiter l'avancée des méthodes de stratification d'imagerie médicale pour valider cette intuition. Pour cela, nous traduisons le lien existant entre l'anatomie de l'AG et le risque d'AVC en un problème géométrique, nous permettant ainsi de bénéficier d'un riche historique de ressources théoriques et appliquées sur l'analyse de formes. Nous passons en revue les nombreuses facettes de l'analyse de formes, et réalisons que, bien que puissantes, les méthodes existantes manquent d'interprétations cliniquement significatives. Nous nous basons néanmoins sur ces outils généraux pour construire une représentation compacte spécifique à l'AG qui permet une meilleure interprétation des résultats. Cette première tentative nous permet d'identifier les éléments clés d'une solution réaliste à l'étude de l'AG. Parmi eux, tout outil que nous construirons devra être suffisamment rapide et robuste pour être adapté à de potentiels larges études prospectives. Puisque l'étape cynétiquement déterminante du processus réside dans la segmentation sémantique des parties anatomiques de l'AG, nous nous concentrons sur l'utilisation de réseaux de neurones spécifiquement conçus pour les surfaces afin d'accélérer ce problème. En particulier, nous montrons que passer la courbure principale des formes en entrée des réseaux de neurones est un meilleur choix que ce qui est actuellement utilisé, quelle que soit l'architecture. Au fur et à mesure de l'amélioration itérative de notre pipeline, nous approfondissons l'utilisation de la segmentation sémantique et de la représentation compacte en proposant un ensemble de caractéristiques géométriques expressives décrivant l'AG – parfaitement alignées avec les attentes des cliniciens tout en offrant la possibilité d'une analyse quantitative robuste. Nous utilisons ces caractéristiques locales et mettons en lumière les relations complexes entre la forme de l'AG et l'incidence d'AVC, en effectuant une analyse statistique ainsi qu'une classification à l'aide de méthodes basées sur les arbres de décision. Les résultats fournissent des informations précieuses pour la prédiction d'AVC : une liste de caractéristiques de forme directement liées aux patients victimes; des caractéristiques qui expliquent d'importants indicateurs de troubles hémodynamiques; et une meilleure compréhension de l'impact du remodelage de l'AG lié à l'état de la FA. Enfin, nous discutons d'autres utilisations possibles des outils développés dans ce travail, de l'étude de cohortes plus importantes à l'intégration dans des analyses multimodales, en passant par une analyse potentielle de sensibilité précise de simulations hémodynamiques, une étape précieuse pour une comprehension exhaustive du processus mécanique de l'AVC.

Mots-clés: Statistiques de formes, Accident vasculaire cérébral, Apprentissage, Imagerie medicale, Fibrillation atriale.

Acknowledgments

First and foremost, I thank Maxime Sermesant for trusting me when he welcomed me in his team, and continuing to do so throughout my whole thesis, always encouraging me to pursue interesting research directions regardless of the time they may take. His support has been invaluable to me and I am grateful for it. I admire his ability to adapt to everyone he works with, and particularly the way he guided me throughout my entire PhD. I thank Hubert Cochet, for the significant amount of time he spent sharing his impressive medical knowledge, with myself and others. I am grateful to have benefited from his advice. Moreover, I thank him for teaching me that any good conference should feature a good brisket.

I thank the members of the jury: Rasmus Paulsen, Alistair Young, Xavier Pennec, and Oscar Camara Rey, for their help reviewing this manuscript and their participation in my thesis defence.

I would then like to thank Nicolas Ayache and the entire Epione team, for creating such a stimulating and friendly work environment. I thank Marco Lorenzi for being the one introducing me to them. I thank Xavier for inspiring the directions explored in this thesis, giving me invaluable insights into differential geometry, and welcoming me inside the g-stats team. I thank James Benn for taking so much time introducing me to nearly all the theoretical concepts present in this work. I would not have been able to carry out this project without him. Of course, I give a warm thank you to my present and past desk-mates, Zihao, Jairo, and Safaa, as well as Luís, Etrit, Harri, Olivier, and so many more, for making the time at Inria such a pleasant one, always full of interesting conversations. Lastly, I thank Isabelle and Nathalie for their reliability and constant help.

This project was dependant on many collaborations, that have been all fruitful thanks to many great people involved in them. I would not be able to cite everyone here, but I need to thank Buntheng and Anaelle from the IHU Lyric, Oscar and Jordi from UPF Barcelona, and Meyer Elbaz from Toulouse Hospital. They have not only been particularly important for this thesis, but have also welcomed me warmly every time we met. I hope the relationships developed along our collaborations will be longstanding.

I would like to finish by giving a shoutout to all the people that have stood by my side, outside of work. I give my warmest thanks to all of them. To my friends Omar El Nagar and Lyes Yahiaoui, and their families. I consider them brothers, and their families extensions of mine. To Dashiell, my brother, although his academic pedigree will inevitably be greater than mine, I couldn't let him beat me so easily and had to join a PhD program. To my dad who taught me to not be afraid of

constantly trying new things. To my mom for doing the hardest work of all and keeping me on the path of knowledge my whole life - I wouldn't even have the Baccalaureat if it wasn't for her. And to the newest member of my family, my wife Hava for her unwavering support but also her impactful council on this thesis.

Financial Support

This Work has been funded by the ERACoSysMed PARIS Project with contract number 15087, and has been supported by the French government through the National Research Agency (ANR) Investments in the 3IA Côte d'Azur (ANR-19-P3IA-000). The author is grateful to the OPAL infrastructure from Université Côte d'Azur for providing computing resources and support.

List of acronyms

AF	Atrial Fibrillation
CFD	Computational Fluid Dynamics
CNN	Convolutional Neural Network
CT	Computed Tomography
CV	Cross Validation
FEM	Finite Element Method
FR	Fibrillation Rhythm
HKS	Heat Kernel Signature
IPV	Inferior Pulmonary Veins
LA	Left Atrium
LAA	Left Atrium Appendage
LAAO	Left Atrium Appendage Ostium
LDDMM	Large Deformation Diffeomorphic Metric Mapping
LGBM	Light Gradient Boosting Machine
LIPV	Left Inferior Pulmonary Veins
LIPVO	Left Inferior Pulmonary Veins Ostium
LPV	Left Pulmonary Veins
LSPV	Left Superior Pulmonary Veins
LSPVO	Left Superior Pulmonary Veins Ostium
MCVC	Multi-Channel Variational Classifier
NN	Neural Network
ODE	Ordinary Differential Equation
PCA	Principal Component Analysis
PDE	Partial Differential Equation
PINN	Physics-Informed Neural Networks
PIDON	Physics-Informed DeepONet
PV	Pulmonary Veins
RFECV	Recursive Feature Elimination with Cross Validation
RIPV	Right Inferior Pulmonary Veins
RIPVO	Right Inferior Pulmonary Veins Ostium

ROC Receiver Operating Characteristic
RPV Right Pulmonary Veins
RSPV Right Superior Pulmonary Veins
RSPVO Right Superior Pulmonary Veins Ostium
SPV Superior Pulmonary Veins
SR Sinusal Rhythm
WKS Wave Kernel Signature

Contents

Chapter 1

Introduction

Contents

1.1 Thesis Context

The ERACoSysMed PARIS project[1] aims to exploit under-utilised, yet widely available, cardiac imaging data to quantitatively explore novel image-based biomarkers. The project is a collaborative effort between the Hamburg and Bordeaux hospitals, the Simula research laboratory, and Inria. Its ambition is to identify biomarkers that correlate with stroke in atrial fibrillation patients, and to prospectively outperform the current risk score by optimizing personalized treatment and clinical follow-up. The project aims to utilise existing medical records of patients with known clinical outcome, to tune and validate computer models and predictive machine learning methods in an iterative process. It focuses on two main problems: the development of a reliable haemodynamic simulation tool of the left atrium, conducted by Simula, and a robust methodology to provide

---

[1] ERACoSysMed PARIS project

quantitative tools to study the anatomy of the left atrium, conducted by Inria. The latter is the subject of this thesis.

1.2 Medical context

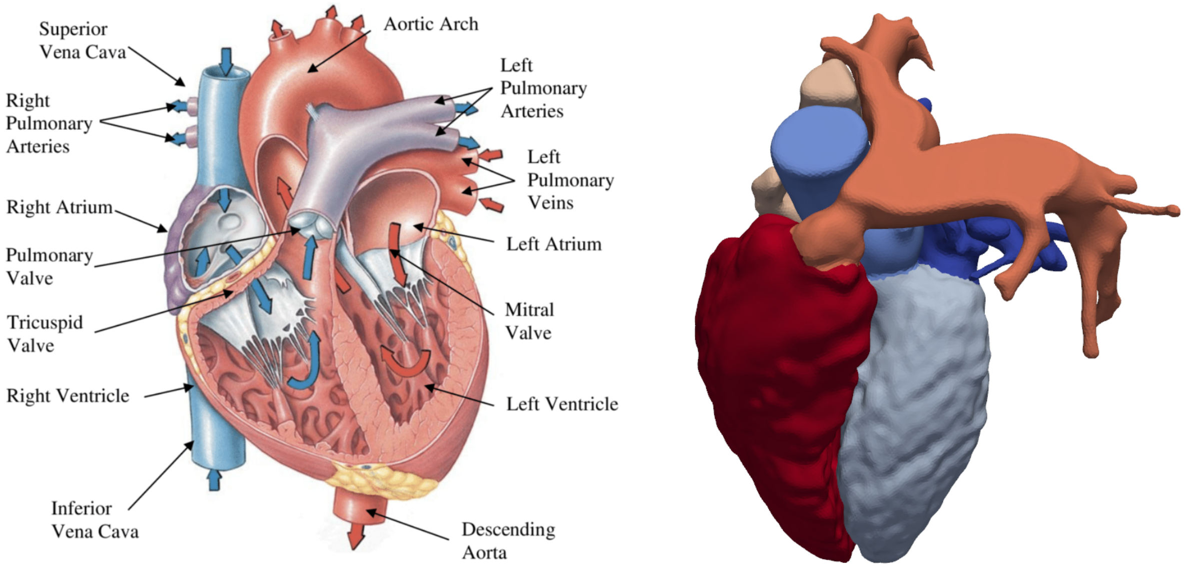

(a) Detailed sketch.

(b) 3D model.

Figure 1.1: Heart anatomy.

The Heart, the muscular organ that is tasked with pumping blood to the body of all mammals, has been studied by humans since ancient times. Ancient Egypt society considered the heart to contain the soul, a concept that is still symbolically present today. As the likes of the Greek Galen (129-216 AD), the Italian Da Vinci (1452-1519), or the Englishman Harvey (1578-1657), scientists have gradually uncovered the heart anatomy and its function throughout history. Today, the many aspects of the heart, from its anatomy, to its electrical conduction system and mechanical contraction cycle, are well understood independently. Indeed, we know now that the heart is composed of four chambers: the right atrium, right ventricle, left atrium and left ventricle, as shown in figure 1.1. The heart is in charge of the circulation of the blood in the body through blood vessels called arteries and veins. The right side of the heart receives the low-oxygenated blood from the body; pumps it to the lungs; receives it back on the left side through the pulmonary veins, into the left atrium, and through the left ventricle; and is finally pumped back out to the rest of the body through the aorta. The pumping is done via a mechanical contraction in the form of a twisting motion, itself driven by an electrical signal that synchronises the entire process. This complicated machinery is still at the heart of medical research today, and the quest to a unified understanding of the heart is still driving researchers across multiple disciplines, as showcased by the increasing push towards building a digital twin of the heart [36].

1.2.1 Cardioembolic strokes

As a consequence of its complex nature, the heart is also subject to many diseases, as each and any composing part of the heart can become pathological. In fact, Cardiovascular diseases (CVD) are the leading cause of death globally, accounting for an estimated 17.9 million lives each year, representing 32% of all global deaths. In particular, it is estimated that 85% of these deaths can be attributed to heart attack and stroke [158]. Stroke occurs when occlusions blocks blood supply to a part of the brain or blood vessels rupture within the brain. In either circumstance, neurological tissues sustain damage and can result in lasting brain damage, long-term disability, or death. Strokes are classified as either hemorrhagic or ischaemic in nature: on one hand, an hemorrhagic stroke occurs when blood from an artery suddenly begins bleeding into the brain, or weakened blood vessels rupture; on the other hand, an ischaemic stroke occurs when a blood clot, known as a thrombus, blocks or plugs an artery leading to the brain. The latter category accounts for an estimated two thirds of all strokes [89]. There are different underlying mechanisms that can cause an ischaemic stroke, among them, thrombi that occur in the heart are cause to a type of ischaemic stroke called cardioembolic stroke – which will be the subject of this thesis.

Epidemiology. Approximately 25% of all ischaemic cases are believed to be cardioembolic in origin, representing about 15% of all strokes [77]. Cardioembolic stroke is the most severe type of ischaemic stroke. They are associated with the highest risk of both early and late recurrence, a high probability of neurological dysfunction at hospital discharge, and the highest in-hospital mortality [77]. They also appear to occur more frequently with increasing age, as they account for 14.6% of cases below the age of 65 years, and this proportion goes up to 36% for patients aged 85 years and older [89]. Economically, the growing prevalence of strokes results in a staggering burden that is estimated to reach US$1 trillion in 2030 [57]. In particular, the costs associated to cardioembolic strokes amount to up to 40% above the average cost of strokes of other etiologies [23]. Moreover, secular trends show that this proportion has been increasing over the last decades [182].

Risk factors. Several variables can increase an individual's risk of suffering a stroke. On one hand, factors that are associated with all types of strokes have been well documented in medical research. Indeed, some modifiable risk factors include morbidities such as obesity, diabetes, high blood pressure, or high cholesterol, along with hygiene and behavioral choices such as smoking, physical inactivity, or poor diet [71]. Other non-modifiable risk factors based on genetic and environment have been identified. They include factors such as age, family history, anxiety and depression, or air pollution [85]. On the other hand, in the more specific case of cardioembolic stroke, various heart conditions will put a person at higher risk for developing one. In particular, the most frequent high-risk cardioembolic conditions include conditions such as recent myocardial infarction, systolic heart failure, mechanical valve prosthesis, or atrial fibrillation and are associated to an up to six-fold increased risk of stroke [89]. In this thesis, we focus on cardioembolic strokes in patients suffering from atrial fibrillation – representing the most common risk factor.

1.2.2 Atrial Fibrillation

The electrical system of the heart can fail in several ways, at different locations in the heart. This is generally referred to as arrhythmia. The most common form of arrhythmia concerns the atria, i.e. the top chambers of the heart, and is called Atrial Fibrillation (AF). AF is characterized by chaotic electrical activation followed by loss of atrial contraction, as illustrated in figure 1.2. In other terms, during atrial fibrillation, the heart's upper chambers beat chaotically and irregularly, resulting in an asynchronous and often very rapid heart rhythm. This has an impact on the blood flow dynamics, or haemodynamics, and creates an environment prone to a wide range of potential complications, most notoriously cardioembolic strokes. This complex disease is heterogeneous in nature: it can be paroxysmal or persistent; contain severe symptoms or be totally asymptomatic; and may self-terminate or become permanent and cause atrial remodelling due to the stress it induces on the heart [134].

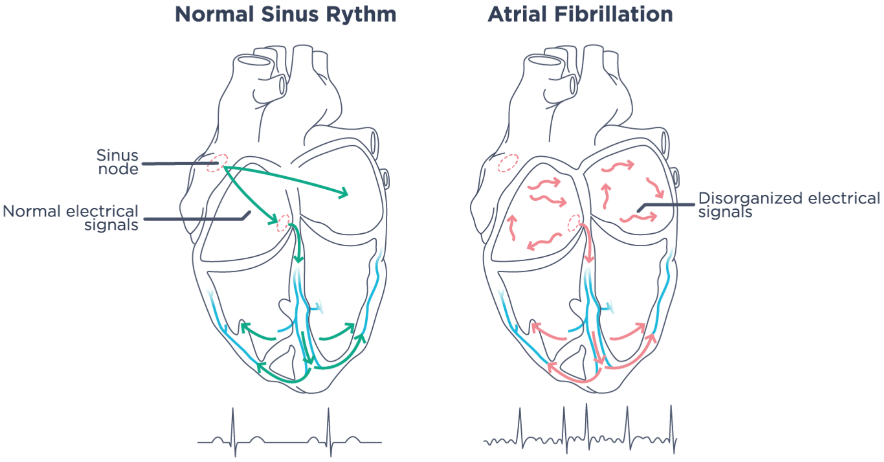

Figure 1.2: Atrial fibrillation[2].

Epidemiology. AF currently affects approximately 60 million people worldwide. Among them, 33 million are estimated to have a six-fold increase in risk of cardioembolic stroke specifically originating from the left atrium [53]. In fact, AF is the most common cardioembolic risk factor [136, 109, 3], and is also considered to be the most frequently encountered form of arrhythmia, occurring in approximately 5% of people aged 65 years and above [89]. More precisely, the prevalence of AF increases sharply from 0.1% among patients aged 55 years and below, to up tp 10% among those aged 80 years and above [62]. Projections from high-income countries predict the number of patients with AF to double, and the number of AF-related strokes to triple in the next few decades [62].

---

[2] Image taken from Withings

Risk factors. A broad range of modifiable, lifestyle-related risk factors and comorbidities are linked to the onset, persistence and progression of AF. These include hypertension, obesity, diabetes, alcohol and tobacco intake, or lack of exercise [53]. However, the risk factors most strongly associated with the development of AF are non-modifiable, with advancing age being the most important. Indeed, risk factors such as high blood pressure – which is responsible for approximately one-fifth of AF cases [53] – become more prevalent with age. Additionally, numerous studies have explored the genetic component of AF. Applying a polygenic risk score (PRS) to individuals at risk of AF has revealed a significant increased risk of AF in patients with high PRS [93].

Stroke prevention. In response to the increased cardioembolic stroke risk associated with AF patients – a condition strongly associated with non-modifiable risk factors – clinicians turn to preventive therapy. In particular, oral treatment for anti-coagulation is an effective preventive strategy that allows a 70% reduction of ischaemic stroke risk [97]. Although the overall benefit of oral anti-coagulants in this indication is clear, this treatment is accompanied by an increased risk of complications, as these molecules increase the risk of major bleeding. For drug-refractory patients, there exists surgery options, such as catheter ablation, to treat AF directly or the occlusion of Left Atrial Appendage (LAA) to prevent the formation of thrombus [91]. However, preventive therapy imply the capacity of identifying patients who are at risk, which can be a challenging task.

Stroke risk stratification and current challenges. In clinical routine, stratifying stroke risk is achieved by designing clinical scores tested on large cohorts of patients in order to measure their predictive power. Currently, the management of AF patients largely relies on the CHA_2DS_2-VASc score [110], a widely adopted stroke risk stratification tool used for helping decide when anti-coagulation therapy may be beneficial. Precisely, the CHA_2DS_2-VASc score is expressed as a weighted sum of points associated to the presence of seven risk factors collected from the patient's history and potential comorbidities, namely: (1) presence of congestive heart failure; (2) presence of hypertension; (3) presence of diabetes; (4) prior stroke or thromboembolism; (5) presence of vascular disease; (6) patient's age; (7) patient's sex. While its ease of use makes it attractive, it is a poor predictor, exhibiting a C-statistic of approximately 0.65 [54]. In addition the populations that benefit the most from anti-coagulation are not necessarily those with the most comorbidities [171]. More generally, the occurrence of stroke in AF patients is not well understood, and in turn, no known markers reflect the mechanistic process of strokes. As AF is gaining epidemic proportions and more and more patient-specific clinical data are routinely available, it is our goal to determine more precise and mechanism-aware stroke markers.

1.2.3 The Left Atrium

To pursue a more adequate stratification of stroke risk, we opt to focus on the way the heart works. As the blood is pumped in and out of each chambers, if at any point the heart cycle is broken due to AF, the contraction is lost and the blood starts to stagnate in the atria, eventually leading to a clot and a stroke. However, supposing the blood is sufficiently liquid, if the chamber that lost contractility was to be perfectly spherical, then the blood would have no place to stagnate and in turn no clot would form and no stroke would occur. Such observations have led clinical experts to take a closer look at the anatomy of the Left Atrium (LA) [133, 87, 51, 59] where most cardioembolic strokes originate under AF, with a specific focus on the shape of the Left Atrium Appendage (LAA) [184, 130, 15, 90, 135].

Anatomy of the LA. As illustrated in the 3D model in figure 1.3, the shape of the LA can be briefly summarised into the following components. First the 4 pulmonary veins (PV), through which the oxygenated blood from the lungs re-enters the heart, they are referred to in clockwise order starting from the bottom left as: Left Superior (LSPV), Left Inferior (LIPV), Right Inferior (RIPV) and Right Superior (PV). Second the body of the LA, the main part of the chamber where the blood is contained. Third, the mitral valve, the opening membrane that connects the LA to the left ventricle, although it will not appear throughout this work. Finally the appendage (LAA), a pouch connected to the LA body in similar fashion to the appendix. It is in this last region that clots are formed, earning the LAA the title of "the most lethal human attachment", as quoted from [86]. These observations are at the source of a now widely adopted surgical procedure to prevent clot formation: the previously mentioned occlusion of the LAA [91]. As we have a more accurate understanding of how cardioembolic strokes occur, we can use this information to draw an intuition as to how to better stratify stroke risk.

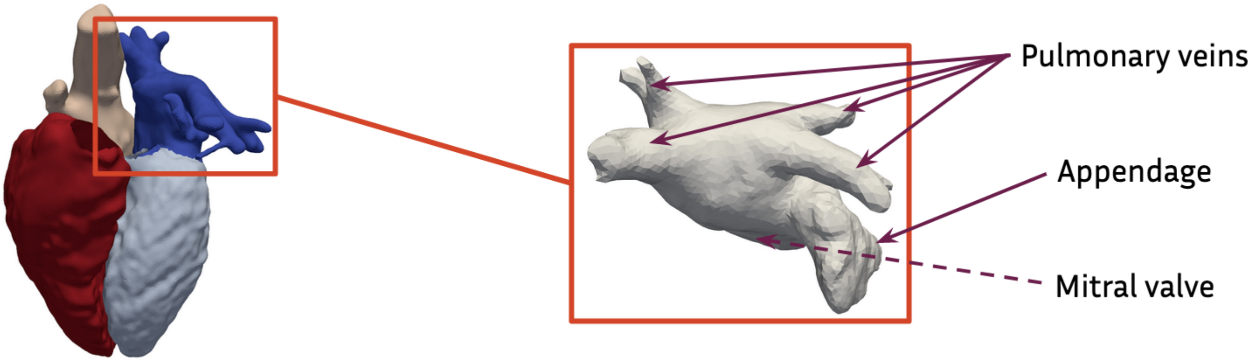

Figure 1.3: Model of the LA.

Anatomical heterogeneity As highlighted above, the shape of the LA has an impact on the occurrence of cardioembolic strokes. Aiming to further understand and quantify the extent of this impact emerges as a natural progression of the research in this direction. As much as people's physical traits, their organs can vary considerably between one another. This fact is known to be particularly prevalent in the anatomy of the LAA [184]. Although less documented, this variability also extends to the entire anatomy of the LA. For example PVs can be connected to the LA in multiple ways, with large differences in configuration, and may even converge into a common trunk that connects to the LA body. To highlight these phenomena in both the LA and LAA, we show some examples in figure 1.4.

This variability suggests that a successful classification of the morphology of the LA can help us gain a precise mechanistic understanding of stroke incidence, and thus derive relevant risk scores that can be more accurate than the currently employed methods. This line of research has been followed before, and has specifically focused on the LAA's morphology [184, 43], although the variability of the entire LA suggests that the PV and body may have an impact as well.

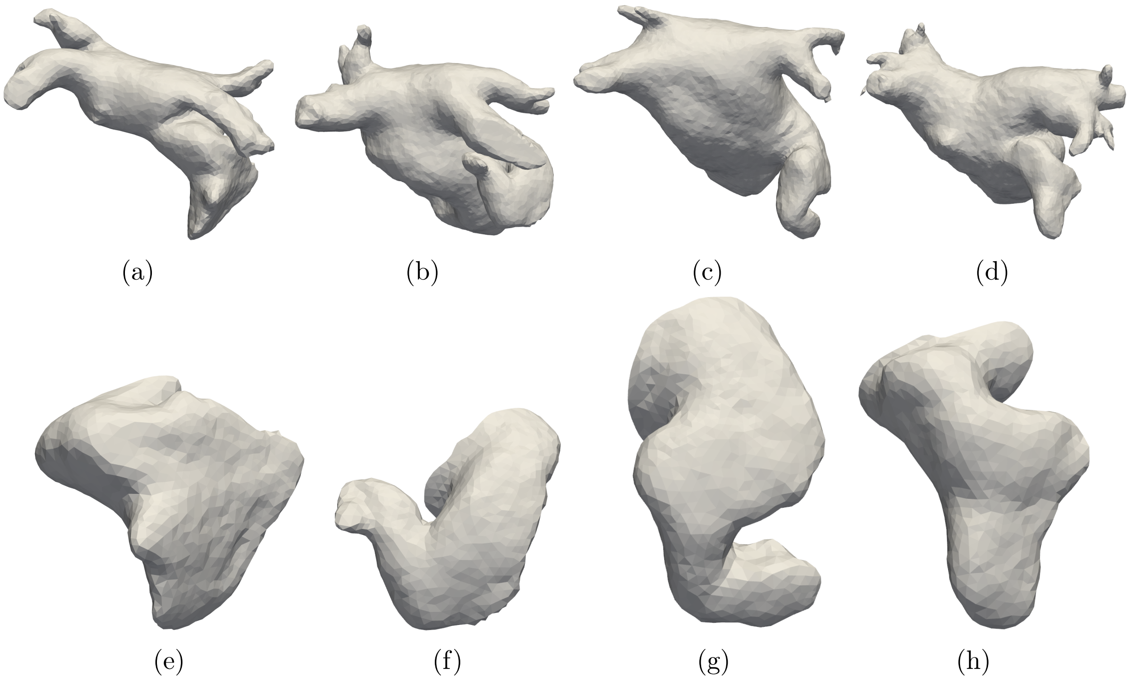

Figure 1.4: Examples of LAs (a to d) and their associated LAAs (e to h).

1.2.4 Medical imaging of the LA

The understanding of the heart's anatomy comes from our ability to observe it, as strongly implied in the previous section. Initially achieved solely through dissection, the advent of X-ray imaging, discovered by Roentgen in 1895, made it possible to look at a patient's body without opening it. With the growth of medical imaging came new techniques, and most famously Computed Tomography (CT), developed in 1971 by Hounsfield, which is now considered as the modern doctor's *truth machine* [19]. As of 2010, an estimated 5 billion imaging studies have been conducted worldwide, of which CT images represent the largest proportion among methods using ionising radiation techniques [152]. Widely used in cardiovascular imaging, the CT-scan, when augmented via a contrast agent, provides an accurate anatomical picture of the heart (see figure 1.5a) and is now routinely performed in every hospital centre [48]. While many other non-invasive measurements exist, such as electrocardiography, the use of wearable sensors or blood tests, they all give punctual information that varies greatly in time, even within a day. In opposition, the cardiac anatomy is not prone to such variation unless a disorder related remodeling occurs. As we introduced previously, medical experts have strongly hypothesised the shape of the LA to be a considerable factor of stroke occurrence, as the immense majority of clots occur in the LAA, and its anatomy is known to show high inter-individual variability. In an effort to better understand the link between AF-related strokes and the LA, important studies were conducted on LA haemodynamics [179] or LAA morphology [184, 43]. These studies have resulted in a widely used categorisation of the LAA, shown in figure 1.5b, demonstrating moderate association with stroke risk [43]. However, these works are mostly qualitative, and depend on a subjective classification that can differ greatly between practitioners.

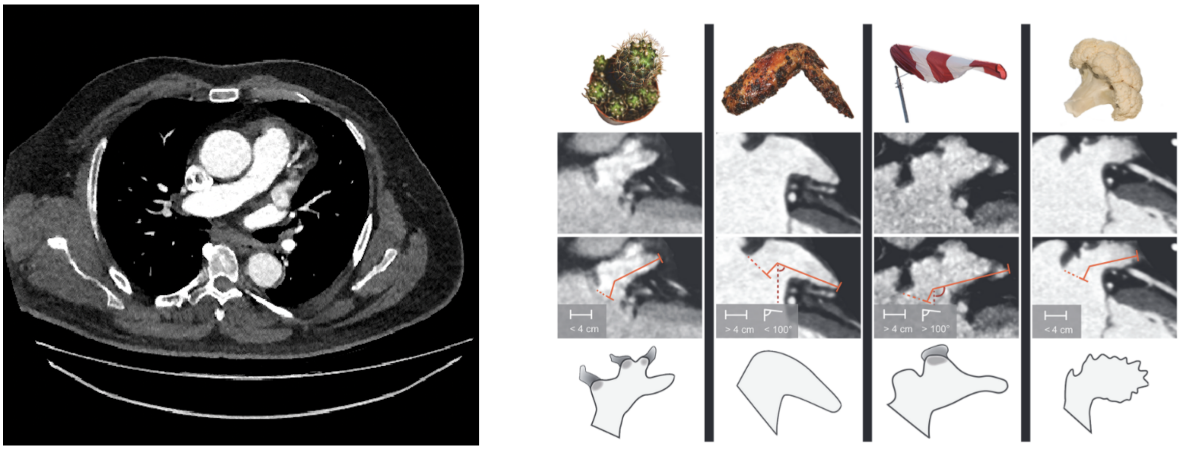

(a) CT-Scan slice.

(b) LAA classification proposed in [184]

Figure 1.5: What clinicians work with

The general medical context of this thesis can be resumed as follows: cardioembolic strokes are known to be the most deadly form of ischaemic strokes, they occur in the LA and are especially prevalent in AF patients. Current tools to precisely stratify stroke risk in AF patients are underdeveloped, nonetheless, the extreme variability of the LA's anatomy is an indicator of its potential into building more accurate risk markers.

1.3 Objectives and contributions

1.3.1 Thesis objectives

Following clinicians' intuition that the shape of the LA alone is a strong predictor of stroke occurrence, we forsake the multi-modal complexity of the heart, and focus solely on 3D models acquired from CT-scans. The general medical question at the centre of this PhD thesis is the following:

> **What is the relationship between the anatomy of the LA and the occurrence of cardioembolic strokes in AF patients?**

While this clinical question is straightforward, it requires an amount of methodological exploration.

As is typical in a medical research process, the path to achieve a complete answer and a deployable solution takes multiple stages, starting with exploratory and retrospective studies and followed by possibly multiple prospective studies of increasing extents, before its use in clinical routine. The work presented in this thesis consists of the initial exploratory stages. Precisely, it aims to provide a set of methods applicable for all future steps of the research process, along with first insights to the medical question defined above. In other terms, the main goal of this thesis is to define a relevant and thorough methodological approach to address the problem. Thus, we translate our problem into a methodological one, and largely focus on the following question:

How to study the shape of the LA and its impact on specific diseases?

Although this problem is more general than the medical question formulated above, we build all the methods and answers proposed in this work on a cohort of patients relevant to the problem of cardioembolic strokes in AF patients. To ensure for any answer to be clinically relevant, it is important to frame the exploration by some set of requirements.

Medical Requirements From earlier sections, we can already underline some important points to take into account for designing a relevant medical application:

1. The problem at the centre of this thesis being of medical nature, any proposed solution should bear high clinical relevance.
2. As medical research requires many steps, any relevant methodology designed on a small cohort of patients should be applicable to a potentially larger number of patients.
3. As the medical problem of this thesis is wide, any proposed method would need to focus on the anatomy of the LA rather than the specific medical phenomena.

Methodological Requirements. As we will explore different methods from the literature, and develop our own, we will need to keep these general bounds in mind. In fact, we translate the medical requirements into methodological terms to define what our end result should look like:

1. Any proposed method should yield expressive results, i.e. be interpretable.
2. Any proposed method need to be reusable at scale. Precisely, the tools developed need to be fully automated, fast and robust for processing large amounts of data in a realistic time.
3. The final pipeline should be of general purpose, i.e able to cover all potential shape-related follow up questions.

By the end of this work, we propose an automated pipeline that extracts an extensive number of LA shape features, with minimal conditions on the shape's discretisation, and in real time, enabling the study of the LA shape in a clinically coherent way. To get there, I propose to follow the exploratory route that lead us to it.

1.3.2 Thesis outline

Chapter 2. Dataset and preprocessing. This first chapter constitutes a short description of the starting point of this thesis. While it is always assumed that the basic object of study is the LA shape, it is not straightforward to construct a dataset of shapes. I show in this chapter how the LA anatomy can be extracted, i.e. segmented, from CT-scans, and what constitutes an efficient representation of shapes, in this case, meshes.

Chapter 3. Shape representations and their applications. Studying the shape of the LA, while being a clinical question, has a theoretical resonance to it. A shape as a mathematical object has known many definitions, and their study has generated many useful tools. To build an efficient pipeline, it is important to understand what these different tools can bring us. In this chapter, I focus on the many representations of shapes, taking a theoretical point of view from 3 widely used domains of shape analysis: operators, deformations and measures. Rather than giving a detailed and exhaustive exposition of each domain, I give a general geometric overview and select one central method per field that conveys the main ideas, as well as ground them in practice through use cases relevant to the study of the LA. Specifically, I introduce the generalised Laplacian for operators, the LDDMM framework for deformations and the currents for measures. This will help identify what is useful to help achieve the specific goal of this thesis, and what is not. In particular, global methods to study shapes are not well defined to directly use given our constraints, however, they offer tools valuable to reach a representation specifically tailored to the study of the LA shape.

Chapter 4. A first attempt at classification. The process of finding the right pipeline takes many iterations as we slowly define the fundamental steps needed in the process. In chapter 4, I focus again on the clinical problem, and give a first attempt at defining a way to answer it. I have recognised that a compact representation specific to the LA should give us the information to better understand it and propose a graph representation based on the anatomical labelling of the LA. I have also realised that any study of the heart should take into account its remodeling, and proposed a multi-modal learning type model to alleviate this, which I call the Multi-Channel Variational Classifier (MCVC). At this point, the available dataset was a fraction of its final version, making the MCVC particularly relevant. However as the dataset got bigger, I have observed that classical machine learning methods are sufficient, and thus better appropriate. In addition, thanks to the growing cohort and to clinicians' feedback, we have realised that the first approach contained two issues. First the automatic labelling is too slow for any larger studies, second the compact representation of the LA is not expressive enough. The 2 final chapters are dedicated to overcome these difficulties.

Chapter 5. Curvature input for deep learning on surfaces. As realised in the previous chapter, the automatic labelling part of the pipeline needs to be made faster. I have proposed to turn to neural networks to find the solution. In chapter 5, an overview is given of the recent development of deep learning methods specifically designed for surfaces. I have highlighted that the shape representation given as input to such models is overlooked, and have hypothesised that curvature is a better adapted choice. This idea is motivated from a geometric perspective, echoing chapter 3. To empirically highlight this fact, I have selected three specific neural network architectures, along with three commonly used benchmarking datasets for semantic segmentation and shape classification, in which principal curvature is compared against 3 widely used inputs. After successfully showing better performance from the principal curvature representation, the best method is applied to the anatomic segmentation of the LA shape.

Chapter 6. Geometric features In this final chapter, I focus on the expressivity issue of the representation developed in chapter 4. This is done by extracting specific features of the LA shape by fully leveraging the anatomic labels, from which we derived the graph representation, but also the ostia, the appendage and the body of the LA. Each individual object can be measured automatically in a variety of ways, constituting a large set of geometric features of the LA. The final step of this thesis is an extensive analysis of those features, relating them to cardioembolic stroke and showing that some features are largely correlated with blood stasis in the LAA. In addition I have provided more insight into the assumptions of chapter 4 concerning how the remodelling of the LA may affect its study. Precisely, it is shown that the AF state has the most impact on our study, and should be taken into account for a more thorough analysis. In summary, I have completed the pipeline, answering the methodological problem of the thesis, and provided the analysis that answers the clinical problem of this thesis.

Finally, I conclude by looking at possible future work that could elucidate the occurrence of stroke, and by showing how the tools we developed may help, and even serve other purposes.

1.3.3 Main contributions

Chapter 2 The automatic processing pipeline from CT-scan to mesh representation that I use throughout the thesis is shown in this short chapter. The pipeline is composed of two main parts. The first part corresponds to an automatic segmentation of the LA mask from the CT-scans. It is done using a neural network model developed through multiple iterations as a collaboration between Inria and the IHU Lyric, with objectives not limited to this thesis, and in which I only had a minor part. The second part is the post-processing pipeline of the segmentation masks, used to represent them as meshes. This pipeline also includes the automatic detection of segmentation errors and their correction. It is a personal work and contribution. Finally, I was charged with a large part of the retrieval of images and clinical information of the cohorts from both the Bordeaux and Toulouse hospital.

Chapter 3. The purpose of this chapter is to give an overview of the existing theoretical methods to represent and analyse shapes. However all use cases I have proposed are original work: the spectral shape analysis of the LA and the LAA in section 3.3; the diffeomorphic labelling of the LA in section 3.4.6; the motion generation pipeline based on parallel transport in section 3.4.7; and the attempts to solve the tri-Helmholtz PDE in section 3.5.3. Moreover, the motion generation pipeline in section 3.4.7 constitutes our part in a collaborative work to study the haemodynamic effect of AF on the LA. This work is still in preparation for submission:

> H. Kjeldsberg, **J. Harrison**, C. Reinbold, M. Sermesant, R. B. Schnabel, J. Sundnes, K. Valen-Sendstad: Comparative Analysis of Atrial Flow Patterns in Sinus Rhythm and Atrial Fibrillation Using High-Resolution CFD Simulations: Implications for Thrombogenesis *Paper in preparation for submission.* 2024.

Chapter 4. I have proposed a first pipeline to represent and analyse the LA shape. I contributed every element of the pipeline: the automatic labelling; the graph representation; and the neural network model, which is called MCVC. This work has been published at a peer-review methodological conference paper at MICCAI:

> **J. Harrison**, M. Lorenzi, B. Legghe, X. Iriart, H. Cochet, M. Sermesant: Phase independent latent representation for cardiac shape analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 537–546, Springer. 2021

And was disseminated through talks in clinical conferences:

> **J. Harrison**, M. Lorenzi, B. Legghe, X. Iriart, H. Cochet, M. Sermesant: Compact representation of atrial shape for thrombosis risk prediction in AF In: Heart Rhythm 19, no. 5, S458-S459. 2022.

> **J. Harrison**, M. Lorenzi, B. Legghe, X. Iriart, H. Cochet, M. Sermesant: Prediction of thrombosis in atrial fibrillation with compact atrial shape representation In: EHRA 2022.

Chapter 5. I have shown how principal curvature should be used as input to neural networks designed for surfaces, by showing empirical evidence of better performance through multiple experiments. This contribution is aimed to be published in the near future:

> **J. Harrison**, J. Benn, X. Pennec, M. Sermesant: Curvature representation for learning on surfaces. *Paper in preparation for submission.* 2024.

I also applied the newfound methods to the anatomic segmentation of the LA, which became a central part of the complete pipeline I built throughout the thesis.

Chapter 6. I have used the information and tools acquired from the previous chapters to propose a novel set of features describing the LA, computed automatically and in real-time speed. I contributed every methods used to derive these descriptors. I then followed a statistical analysis and a classification procedure of the geometric features using well established methods. I showed the relationship between the LA's anatomy and the occurrence of cardioembolic strokes. I also measured how the state of AF remodels the LA, and proposed a model aware of this to predict the occurrence of stroke, leading to a state-of-the-art method for stroke prediction. Moreover, I showed with striking precision how some features predict the presence of blood stasis in the LAA – an important marker of clot formation. Both the geometric features and the complete analysis will be submitted as a journal paper in a short time:

> **J. Harrison**, M. Elbaz, H. Cochet, M. Sermesant: Study of the atrial anatomy for the prediction of stroke through geometric features. *Paper in preparation for submission.* 2024.

Throughout the creation and definition of the complete set of geometric features, I have used them in collaboration with the lab from UPF Barcelona to conduct two haemodynamic studies of thrombus formation in the LA:

> J. Mill, **J. Harrison**, B. Legghe, A. L. Olivares, X. Morales, J. Noailly, X. Iriart, H. Cochet, M. Sermesant, and O. Camara: In-silico analysis of the influence of pulmonary vein configuration on left atrial haemodynamics and thrombus formation in a large cohort. In: International Conference on Functional Imaging and Modeling of the Heart, pp. 605–616, Springer International Publishing. 2021.

> J. Mill, **J. Harrison**, M. Saiz-Vivo, C. Albors, X. Morales, A. L. Olivares, X. Iriart, J. Noailly, H. Cochet, M. Sermesant, and O. Camara: The role of the pulmonary veins on left atrial flow patterns and thrombus formation. In: Scientific Reports 14, no. 1, 5860. 2024.

Additional contributions. In addition to the work directly relevant to this thesis, I have made minor contributions in other works related to the analysis of the LA. The first collaboration combines a deformation-based analysis of the LA shape in relation to arrhythmia and defines a set of features both clinical and shape related to derive a risk score of arrhythmia recurrence.

> S. Jia, H. Nivet, **J. Harrison**, X. Pennec, C. Camaioni, P. Jaïs, H. Cochet, and M. Sermesant: Left atrial shape is independent predictor of arrhythmia recurrence after catheter ablation for atrial fibrillation: A shape statistics study. In: Heart Rhythm O2, no. 6, vol. 2, pp. 622–632, Elsevier. 2021.

The second collaboration focuses on a stratification of patients to measure the risk of thrombus formation in the LAA. This is done via cluster analysis of features that are extracted from the shape of the LA, although with different methods than the ones I propose in chapter 6.

> S. Marta, J. Mill, **J. Harrison**, G. Jimenez-Pérez, B. Legghe, X. Iriart, H. Cochet, G. Piella, M. Sermesant, and O. Camara: Unsupervised machine learning exploration of morphological and haemodynamic indices to predict thrombus formation in the left atrial

appendage. In: International Workshop on Statistical Atlases and Computational Models of the Heart, pp. 200–210, Springer. 2022.

Chapter 2

Dataset and preprocessing

Contents

Abstract. The object of interest in this thesis is the anatomy of the Left Atrium (LA), yet it is not trivial to constitute a dataset of shapes. We establish the starting point of this thesis in this short chapter by presenting the cohort our work is designed on, and introducing the pre-processing pipeline we have developed to represent the LA as a mesh using segmentation masks predicted from CT-scans.

2.1 Dataset

The data used to build the methods and solutions in this thesis is a retrospective cohort, consisting of CT scans of both control patients and patients exhibiting the disease to study. Specifically, the available data consists of patients undergoing an LA Appendage (LAA) occlusion procedure. This implies that all patients of our dataset suffer from Atrial Fibrillation (AF), either in paroxysmal or persistent state, and all are considered at risk of stroke but could not follow an anticoagulant prescription due to high risks of bleeding.

Content. Throughout this work, the number of available data samples has varied, as patients from the Bordeaux Hospital and the Toulouse Hospital have been gradually included. By the end of this study, and after applying the preprocessing described in the next section, the dataset consisted of 404 volumetric CT scans with contrast, composed of 232 images from the Bordeaux Hospital and 172 images from the Toulouse Hospital. To define a meaningful *stroke positive* group of patient, we created a composite label combining past history embolism and/or the detection of LAA thrombus on CT scan. The proportion of *stroke positive* patients in the final dataset amounts to 39% of the 404 patients. The final dataset is composed of patients aged 76 ± 8, of which 68% are male. The patients have a CHA_2DS_2-VASc score of 4.4 ± 1.4 and a Body Mass Index (BMI) of 26 ± 5.

2.2 Automatic segmentation of the LA

Although CT-scans clearly outline the LA's geometry, they also contain a lot of additional information and *noise*. To be able to focus on the anatomical aspects of the LA, we first extract the shape from the image. This process is known as image segmentation and yields a mask on the images. This mask can focus on a single label, or it can consist of multiple labels. For instance, figure 2.1 illustrates a CT-scan slice from the dataset and its corresponding multi-label segmentation mask. In the case of a binary mask, it consists of a label 1 in the anatomical structure of interest and 0 elsewhere. The image segmentation of organs is traditionally performed by an expert's hand. It is a notoriously time consuming yet necessary procedure, even more so in our case. Thankfully, with the advent of deep learning and computer vision, this task can be automated efficiently.

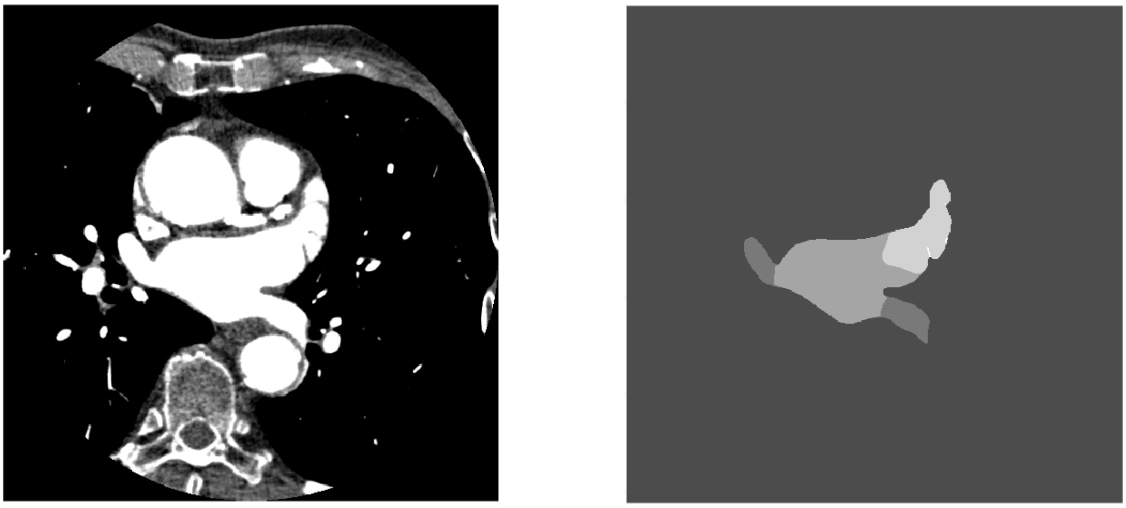
(a) CT-scan slice.

(b) Segmented mask slices.

Figure 2.1: Slice of a CT-scan and its segmentation.

2.2.1 Segmentation models

In an effort to enable the study of large cohorts we have set out to develop robust automatic segmentation models in collaboration with the IHU Liryc. We benefited from an access to more than

500 hand segmented images, and developed multiple generations of neural networks to extract the anatomy of the LA from the CT-scans.

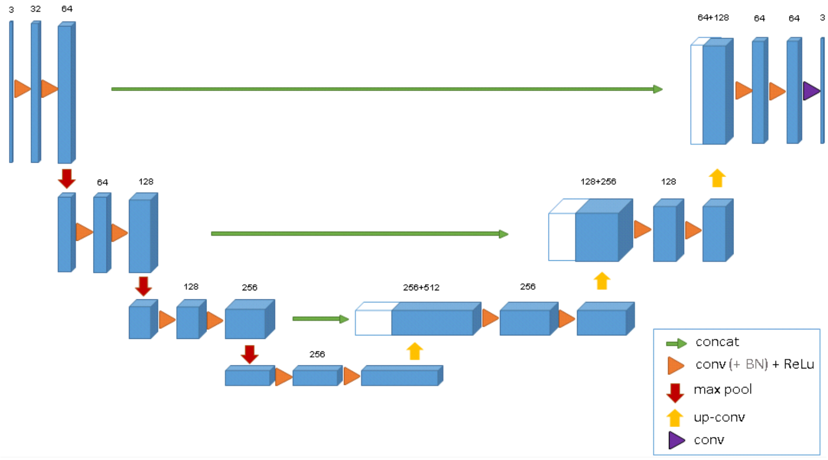

Figure 2.2: Overview of the 3D U-net, as introduced in [34].

3D U-net. An early version of the image segmentation model used in this thesis was based on the architecture of the 3D U-net [34], a deep neural network that learns to generate dense volumetric segmentations using few annotated 2D slices for training. The 3D U-net is based on the previous U-net architecture [151], but it takes 3D volumes as input and processes them with corresponding 3D operations: 3D convolutions, 3D max pooling, and 3D deconvolution layers. Figure 2.2 illustrates the network architecture. The input and output of the network are voxel tiles of the 3D image. The network consists of an *analysis* path and a successive *synthesis* path each with four resolution steps. The analysis path consists of a contracting encoder that analyses the whole image and is composed of four layers containing two $3 \times 3 \times 3$ convolutions each followed by a rectified linear unit (ReLU) and a $2 \times 2 \times 2$ max pooling with strides of two in each dimension. The synthesis path corresponds to an expanding decoder that produces a full-resolution segmentation, in which each layer consists of a $2 \times 2 \times 2$ deconvolution by strides of two in each dimension, followed by two $3 \times 3 \times 3$ convolutions each followed by a ReLU. Shortcut connections between layers of equal resolution in the encoder and the decoder allow the generation of essential high-resolution features. In the last layer, a $1 \times 1 \times 1$ convolution reduces the number of output channels to the number of labels. The model uses batch normalisation layers before each ReLU: each batch is normalised during training with its mean and standard deviation and global statistics are updated using these values. A crucial aspect of the architecture is the weighted softmax loss function. By assigning zero weights to unlabeled pixels, the model learns solely from labeled ones, enabling training on sparse annotations. The first model we use for the segmentation of the LA from cardiac images is directly built on this architecture. It uses successive 3D U-Nets and a proposed contour loss based on additional distance information to adjust the final segmentation [83]. Precisely, it consists of a dual 3D U-Net structure where a first U-Net roughly locates the LA and crops the image, and a second U-Net accurately segments the left atrium under higher resolution. This model is used for pre-processing the data used in the chapter 4.

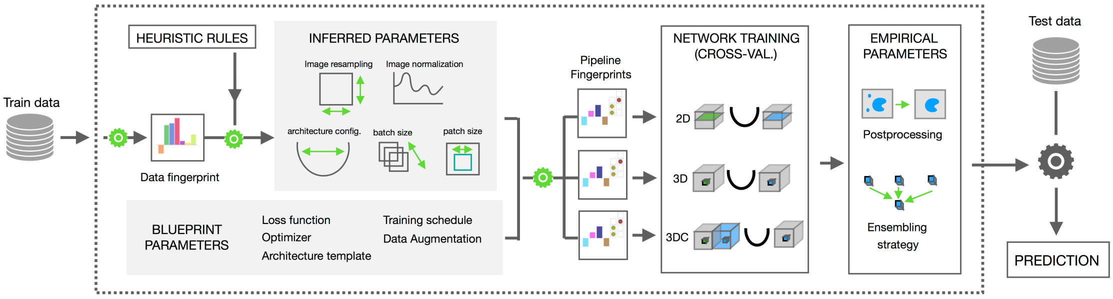

Figure 2.3: Configuration of the nnU-net, taken from [80].

nnU-net. The current version of the segmentation model we use is based on the nnU-net [79]. Similarly to the previous model, the nnU-net builds upon the U-net, but proposes a key innovation: it utilises a modular and configurable architecture, illustrated in figure 2.3, allowing automatic adaptation of the network to different medical imaging datasets and segmentation tasks. Traditionally, given a new problem, the solution needs to be manually designed and optimised. This process is prone to errors, and not very scalable. Instead, the nnU-net automatically adapts to a given dataset. Precisely, the nnU-Net systematically analyses the provided training data and creates a *dataset fingerprint*. It then creates a segmentation pipeline based on several U-net configurations (2D, 3D full resolution, 3D low resolution, 3D cascade), following a three-step recipe for each dataset: (1) the pipeline is initialized with fixed parameters that can be used in all cases. They correspond to the *blueprint parameters*, and include for example the nnU-net's loss function, learning rate and optimiser; (2) the *dataset fingerprint* is used to infer dataset specific parameters by following hard-coded heuristic rules. The inferred parameters are used to adapt the segmentation pipeline properties to the specific dataset, e.g. the depth of the U-net architecture and the pooling behavior; (3) empirical parameters are determined to select the best U-net configuration for the given dataset. The data used in chapters 5 and 6 is pre-processed using this model.

2.2.2 Limitations of automatic segmentation

The automatic segmentation of our dataset using these models have allowed us to identify several predictable segmentation errors linked to the LAA. The first errors the automatic segmentation generates consists of undesirable holes in the LAA, as seen in figure 2.4b. This issue is linked to what is known as an opacification defect: the image acquisition is performed after injection of a contrast agent, as the dye propagates in the heart, the shape of the LA becomes well contrasted with the rest. However, the contrast agent does not always successfully fill the LAA. This can be either due to the high concentration of trabeculations present there, or can be a sign of blood stagnation, also known as stasis. In fact, the presence of LAA thrombi is defined as a persistent stasis. This impacts the segmentation as it makes it harder for the model to distinguish a clear contour. While our segmentation model goes as far as to predict these opacification defects, the quality of the results remains unsatisfying. The second recurring error is due to the proximity of the Left Superior Pulmonary Vein (LSPV) – located right above the LAA – to the LAA (see figure 1.3), and results in a mask where the LAA and the LSPV are sometimes glued together, as shown in figure 2.4a. We deal with these problem in the next paragraph. Using the current version of the segmentation model on the final dataset, about 12% of the predicted segmentations had to be manually corrected. While the segmentation masks contain the geometrical information needed, they are not an optimal representation of shapes. To efficiently build, and exploit, surface processing tools, we decided to represent the LA as a polygon mesh [26].

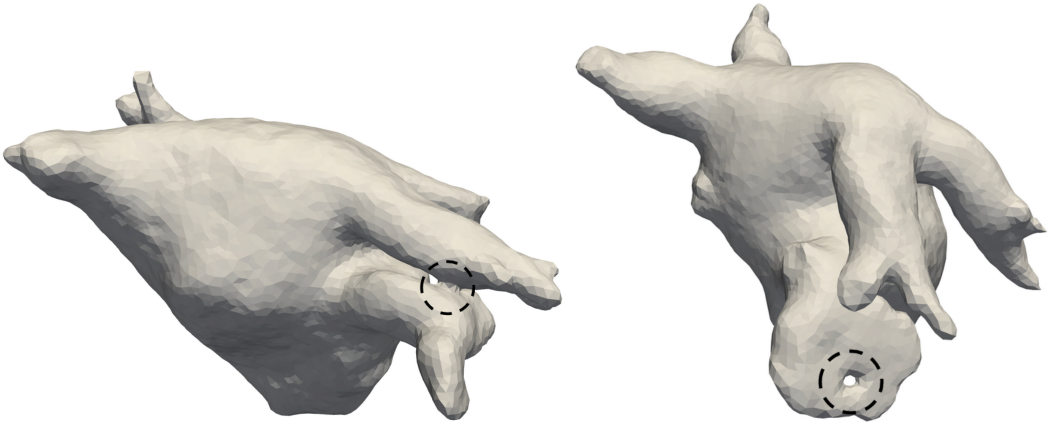

(a) Glued LSPV and LAA (b) Holes in the LAA

Figure 2.4: Common problems associated to the automatic image segmentation.

2.3 Meshing and cleaning

A mesh is a spatial discretisation of a continuous surface. It is defined as a collection of points, edges and faces connected together to build a series of polygons, representing a 3D surface. Overall, triangles are the most common choice of polygon in 3D mesh models because as they can be used to model 3D objects easily and accurately. In fact, we use a mesh representation of the LA for all the analyses explored in this thesis, and chose the polygons to be triangles. As examples, figures 1.4, 1.3, 1.1 are all meshes.

2.3.1 Mesh representation of the LA

All the methods we will build in this thesis are designed for meshes. As so, the quality of the meshing is paramount. A successful mesh representation should be smooth and closely approximate the underlying surface, locally the triangles forming the surface should all be similar in size and globally the mesh should be water-tight. Figure 2.5 provides both examples of a good and a bad quality mesh. Building a mesh fulfilling such conditions is not a straightforward task. In fact efficient meshing

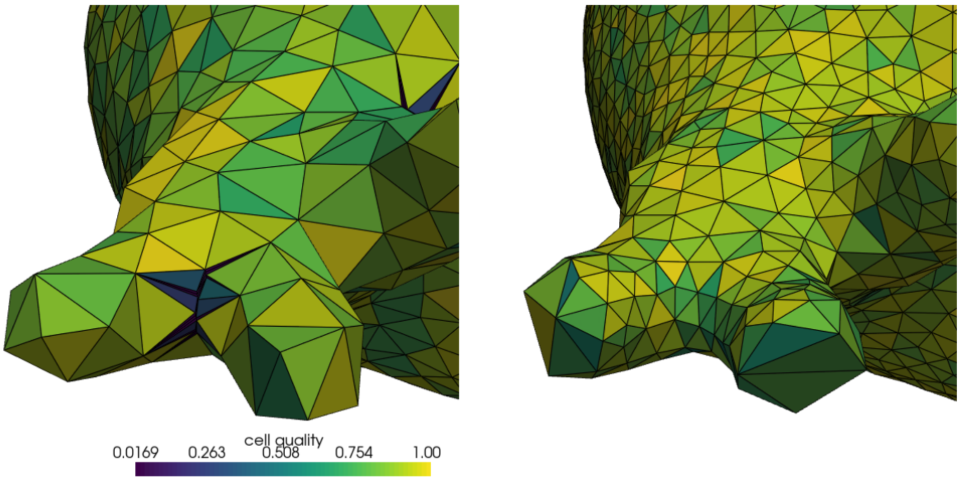

Figure 2.5: Illustrations of bad (left) and good (right) quality mesh.

procedure is an active research domain consisting of a large sub-field of computer graphics. We refer the readers to [26] for a complete exposition.

Proposed meshing pipeline. After an exploration phase, we have opted for a widely adopted method to represent the LA as a mesh, consisting in several steps. The first step is to extract a faithful approximation of the surface, that is, the boundary isosurface of the the binary mask. We do this using the marching cubes method [114], as seen in figure 2.6a. This yields a faithful but rough mesh. The second step is to re-mesh the previously obtained object to achieve triangle regularity and smoothness, as illustrated in figure 2.6b. This is done in our case via the `MMG` software [47], which we found gives the best results with minimum tuning. Finally, we need to ensure the *manifoldness*, or water-tightness, of the mesh. we do this via the repairing procedures proposed in [11], and implemented in `PyMeshFix`[1]. This entire pipeline is completely automatic and typically takes under 3 seconds to compute on a standard laptop computer.

Hyperparameters choice. As mentioned previously, all meshes used in this work are triangular meshes. Additionally, defining the average size of the meshes is another important decision to make: too coarse and valuable information is lost, too fine and the computation complexity increases significantly. We opt for an average of 5000 points per mesh, as illustrated in figure 2.7.

---

[1] https://pymeshfix.pyvista.org/index.html

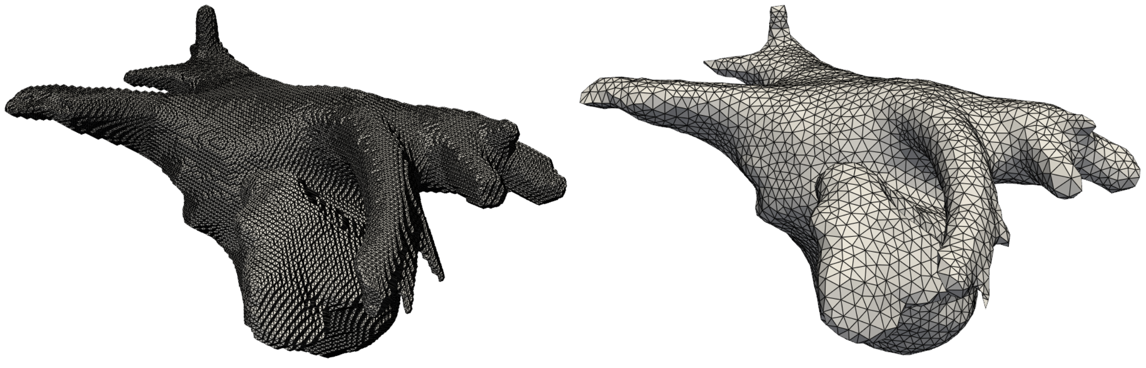

(a) Marching cubes result

(b) Smoothed mesh

Figure 2.6: The meshing pipeline. A rough contour is extracted via marching cubes, before being remeshed and smoothed.

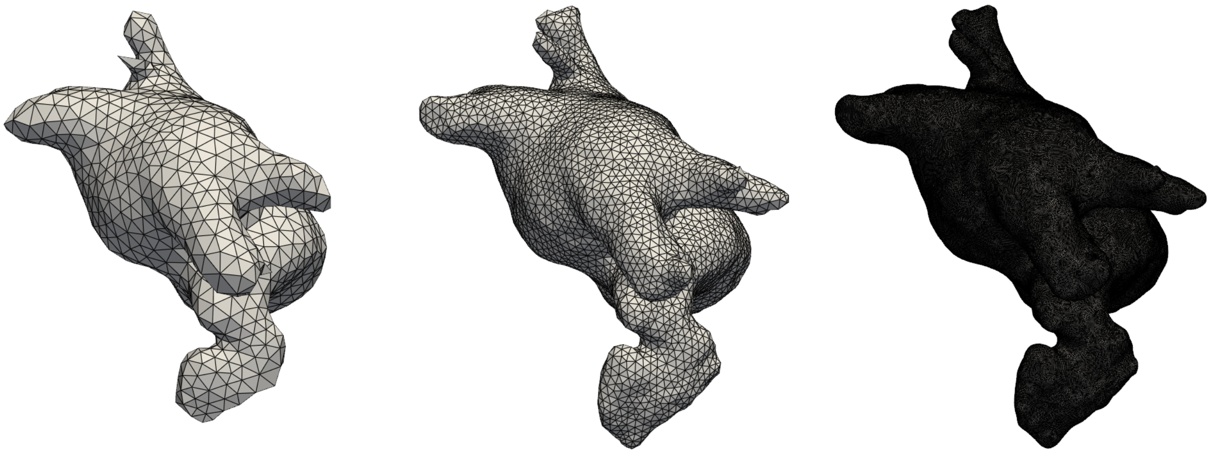

(a) 1k points, 2k triangles.

(b) 5k points, 10k triangles.

(c) 100k points, 400k triangles.

Figure 2.7: Illustration of the importance of mesh size.

2.3.2 Automatic segmentation error detection and correction

After reaching a guaranty for a good representation of the LA's surface, we are able to address the image segmentation problems identified in the previous section. As we described, the large majority of faulty segmentations are either due to a hole in the LAA or a joined LSPV and LAA. Both these problems can be detected automatically with a surface mesh representation.

Error detection. Both identified problems generates a hole in the surface, homeomorphic (i.e. of identical topology) to a torus (i.e. donut), while we know the LA contains none and should be homeomorphic to a sphere. Thus, the detection of an error in the segmentation is simply achieved by counting the number of holes in the resulting mesh. This number is more precisely called the genus of the surface, noted g, and can be computed using the Euler characteristic number, noted χ. Let V, E and F be the number of vertices, edges and faces of the mesh, then the genus g can be obtain from:

$$\chi = V - E + F$$
$$\chi = 2 - 2g$$

For a mesh of an LA, if $g = 1$, it denotes a problem in the segmentation.

Cleaning. While these problems are detected using meshes, they are best fixed using the initial binary mask representation. For this, we use morphological filters [49], that allow us to either erode or dilate the active part of the mask (where it is equal to 1) at its boundary. This is done by a structuring element of a prescribed size, in our case 2×2 pixels. A dilation followed by an erosion is called a morphological closing, and allows us to connect disjoint elements of the shape that are close to each other. The inverse operation, i.e morphological opening, will erase thin links of the shape. We apply closing to the mask to fill the hole in the LAA, and apply opening to disconnect the LAA and LSPV.

We apply the pre-processing pipeline described in this chapter to obtain the final processed cohort of 404 meshes and their relevant clinical information, that can be used for studying the link between the LA shape and strokes.

Chapter 3

Shape representations and their applications

Contents

Abstract. To study the impact of the Left Atrium's shape on stroke occurrence implies to study shapes. A problem that knows a rich body of work from both a theoretical and a practical standpoints. In this chapter, we propose to take a step back from the clinical problematic and introduce different notions of shapes, leading to useful representations that can help us in the goal of studying the Left Atrium.

3.1 Introduction

The clinical problem of this thesis is clear, it is up to us now to define what methods should be used to address it. The natural object that we will work with throughout the following chapters is the Left Atrium (LA) shape, in the form of surface meshes. Taking a step back from the medical context, we can see that understanding and analysing shapes has been subject to multiple fields of research, dating back to the likes of Euler and Gauss and is at the center of major mathematical domains such as, most famously, Riemannian geometry. Today, *digitised* shapes are ubiquitous, from computer-aided design and virtual worlds, to the present subject of medical image analysis. This presence has left us with a large amount of literature from which we can inspire ourselves for our own project. As we will see, shapes are not a straightforward object to work with, the shear size of research dedicated to it serves as testimony alone. As David Mumford explained in [131]: *shape is the ultimate non-linear sort of thing.* Their natural form is as embedded objects in space, defined by the space's own coordinate system, but to compare them we must ignore such coordinates as we do not want the *place* in space where they are embedded to have any impact. Their natural discretisation is as finite sets of points, but we want to ignore the number of points as to compare shapes rather than discretisations. The often large number of points needed to describe a given shape is yet another hurdle when working with shapes. The list of elements to consider when working with shapes can be easily continued, but we will stop here. To study shapes and therefore resolve one or more problems linked to it, researchers have set out to propose meaningful shape representations, from which distances and interpolations between shapes can be derived, and statistical tools can be used. From the many facets of shape representation, we will only go through a fraction of the work that has been done around this subject. Even then, we only focus on one or two versions of fundamental ideas of representations, when many other exist. Although some other facets of shape representation we ignore here will be introduced in chapter 5.3, we do not wish, nor pretend to be exhaustive. For instance, [81] reviews representations of which none will be touched upon in this thesis. Rather, we stay close to representations that have been historically linked to the medical field and highlight how geometry is the natural language to describe not only shapes, but also the space of shapes. As we take a tour of different theories of shapes, we will not forget the main problem by punctuating each chapter with practical use cases, of which all but one will be linked to the study of the LA.

Outline In this chapter we first propose, in section 3.3 to look at how operators defined on shapes contains valuable information to study them. In this context we show that classical and out-of-the-box methods such as spectral shape analysis is not sufficient to predict stroke from the LA's shape. In section 3.4 we show examples of using diffeomorphic deformation methods not to do a direct statistical analysis but rather to process and generate shapes for different tasks. As we seek to define correctly what it means for shapes to be *similar*, we propose in section 3.5 to follow methods from measure theory, and we attempt to implement an important link between two seemingly distinct views on shapes, and identify the problem that needs to be solved to link both in practice. As all these methods live in theory before practice, we first of all clear ourselves of some fundamental concepts and definitions in section 3.2

3.2 Mathematical context

It is difficult to conduct a proper analysis of shape, without going back to its initial development. Especially when the mathematical theory of shapes has been such a fruitful and longstanding one, from which most of the tools we will use stems from. It is also difficult to not go down the rabbit hole of definitions and theorems when introducing such fields. It would be besides the point of this thesis and the result would be a clumsy and cropped version of any of the many books dedicated to the subject. However, vocabulary belonging to differential geometry will eventually be thrown around within this thesis, and for the sake of clarity and simplicity, it is important to introduce what theoretical tools are at our disposal. This section is thus dedicated to introduce them in a loose manner with the goal of giving some perspective to the rest of this manuscript.

3.2.1 Riemannian geometry

At the centre of the theoretical study of shapes, lies Riemannian geometry. We will give a fast introduction to the fundamental concepts here, and will leave the thorough overview and precise construction to textbooks that have stood the test of time, such as [46] or [105], which we refer the readers too.

Manifolds A manifold in its barest form is a topological set, on which we would like to add structure to become a geometric object. The goal of this structure would be to enable the measurement of things, such as distances or angles, and more generally perform calculus. Because this is well understood in the case of euclidean vector spaces, it would be most desirable to import euclidean operations onto manifold. Which is where *smooth* manifolds comes in, they are a category of manifolds which are locally equivalent to a euclidean vector space. In other words, given a manifold M, we can attach at every point p of M a euclidean vector space, called the tangent space at p and noted T_pM. The combination of all these tangent spaces on M, is called the tangent bundle, noted TM. While this is a big step, we still need to perform computation on these objects, which leads us to Riemannian metrics.

Riemannian metrics To enable computations on vector spaces we need a metric, for the case of manifolds, we can equip it with a Riemannian metric g, defined as a smoothly varying family of inner products on the tangent bundle TM. Being an inner product on the tangent space T_pM for all $p \in M$, g enables us measurements of lengths and angles of tangent vectors. This also buys us the possibility to measure the length of a path, informally the generalisation of straight lines on manifold, more precisely a smooth map $\gamma : [0, 1] \to M$. This length is given by $L(\gamma) = \int_0^1 g(\dot{\gamma}(t), \dot{\gamma}(t))dt$, which in turn is used to define the distance between two points on M:

$$d(x, y) = \inf\{L(\gamma)|\gamma : [0, 1] \to M, \gamma(0) = x, \gamma(1) = y\} \tag{3.1}$$

for any $x, y \in M$

Connections A missing crucial component of this construction is to understand how to move a tangent vector v at p to the tangent space of another point q along a path γ. In a euclidean space this is easy, define a straight line by a direction vector w, and all we need to do is to translate v along the line by keeping its angle parallel with w. If we attempt the same operation on a surface embedded in a euclidean space, then the resulting translated vector v' wouldn't live in T_qM. To correctly perform this parallel transport we need a generalisation of the directional derivative, this is introduced by connections, that allows to define the covariant derivative which takes as input a tangent vector v at T_pM and a vector field w and outputs the vector $\nabla_w v(p)$ at T_pM. This tool simply defines a transported vector along a curve to be parallel if its covariant derivative along the curve is zero. Moreover, for any point $p, q \in M$, any tangent vector $v \in T_pM$ and any smooth curve γ between p and q, there exists a parallel vector field along the curve, called the parallel transport of v along γ. In fact, a Riemannian metric induces a canonical connection, called the Levi-Civita connection, as the unique connection satisfying the two following connections:

1. It is torsion free, i.e for any vector fields X and Y we have $\nabla_X Y - \nabla_Y X = [X, Y]$ where $[\cdot, \cdot]$ is the Lie bracket.
2. It preserves the Riemannian metric, that is

   $$\nabla_X g(Y, Z) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z)$$

   for all vector fields X, Y, and Z on M. Or equivalently, the parallel transport operator $\mathbf{\Pi}_p^q : T_pM \to T_qM$ along any smooth curve is an isometry.

Geodesics To have a complete pictures of what Riemannian geometry is, we still need to introduce a couple other quantities. Geodesics are curves whose tangent vectors are parallel in terms of a connection, that is, a smooth path $\gamma(t)$ is a geodesic of the metric g if it satisfies the ordinary differential equation

$$\nabla_{\gamma'(t)} \gamma'(t) = 0.$$

Choosing the Riemannian metric determined Levi-Civita connection, geodesics are then locally length minimising curve, or equivalently, that minimises the energy functional

$$\mathcal{E}(\gamma(t)) = \frac{1}{2}\int_0^1 g(\gamma'(t), \gamma'(t))\, dt.$$

Particularly, the distance defined in 3.1, is a geodesic distance, meaning the shortest path determined by it is a geodesic with respect to the Levi-Civita connection. In fact, the Hopf-Rinow theorem states that (M, d) is a complete metric space if and only if M is geodesically complete, i.e that solutions to the geodesic equation exists for all time.

Exponential maps For a point $p \in M$ and a direction $v \in T_pM$, the exponential map gives us the corresponding geodesic γ_p^v at time 1:

$$\begin{aligned} \exp_p : T_pM &\to M \\ v &\mapsto \gamma_p^v(1) \end{aligned}$$

The derivative at $v = 0$ is the identity map, which, by the inverse theorem, implies that exp is locally a diffeomorphism.

The derivative $D\exp_p$ of the exponential map at p is just the identity map of T_pM so that $\exp_p$ is a diffeomorphism from a small ball around zero in T_pM to a neighbourhood of p in M, by the inverse function Theorem. The importance of this statement is that the exponential map defines a system of coordinates around p called *normal coordinates*, in which the rays emanating from the origin in T_pM are mapped onto geodesics in M; that is, a point $v \in T_pM$ represents the coordinates of the point $\exp_p(v)$. In fact, the normal coordinates show that a Riemannian manifold is a smooth manifold modelled on its own tangent space.

Shapes In this context, the preferred way to define a shape is as an embedding. Given two smooth manifolds N and M, the set of smooth embeddings $\mathrm{Emb}(N, M)$ of N into M consists of C^∞ smooth maps

$$\varphi : N \to M \tag{3.2}$$

which are diffeomorphisms onto their images. The image $\operatorname{Im}\varphi(N)$ defines a submanifold of M, i.e a shape. φ defines the local coordinates of the shape embedded in M. When $M = \mathbb{R}^3$ and $N = S^1$ we have the set of closed embedded curves in $\mathbb{R}^3$; when $N = \mathbb{R}$ we have the set of open embedded curves in $\mathbb{R}^3$, when $N = \mathbb{S}^2$ we have the set of closed surfaces of genus zero in $\mathbb{R}^3$, when $N = \mathbb{T}^2$ we have the set of closed surfaces of genus 1 in $\mathbb{R}^3$, etc. For our case, we will restrict ourselves to surfaces with $M = \mathbb{R}^3$.

Infinite dimensional manifolds As we will see in a following section, Riemannian geometry can be applied not only to shapes but also to the space of shapes. However, the space of embeddings 3.2, is an infinite dimensional space. Thankfully, all of the geometric objects above lift to the setting of infinite dimensional smooth manifolds, so long as the Riemannian structure generates the same topology as the manifold [104]. When the Riemannian structure does not generate the manifold topology, which is often the case in applications, we can, and will, still talk about connections, geodesics, lengths, energy, distance, and exponential maps, but with the understanding that smoothness, existence, uniqueness, and non-degeneracy are not guaranteed and need to be established on a case-by-case basis. It is important to be aware of this since the absence of any one of these can greatly affect the results of an application.

3.2.2 Differential forms

Later on in this chapter, we will see how shapes can be represented through the action of integrating objects called differential forms. We give here a quick review of what they are, as to have a full mathematical context of the following sections, and we refer the readers to the friendly chapter on differential forms and integration in [9]. In fact, informally, differential forms are simply quantities that can be integrated. As a first example, consider a smooth function F over an interval in $\mathbb{R}$ and f its derivative: f such that $dF = fdx$, then we know the Newton-Liebniz formula:

$$\int_a^b dF = \int_a^b f(x)dx = F(b) - F(a)$$

Then $f(x)dx$ is called a 1-form. By extension, a 1-form defined on $\mathbb{R}^n$ can be written as $\omega = a_1(x)dx_1 + \cdots + a_n(x)dx_n$, where a_i are smooth functions. In fact, differential forms can be easily extended to manifolds:

Definition 3.2.1. *A differential form of degree 1, or 1-form, on a manifold M, is a smooth map*

$$\omega : TM \to \mathbb{R}$$

of the tangent bundle of M to the line, linear on each tangent space T_xM.

As 1-forms measure length, we can up the degree to measure areas with 2-forms, volumes with 3-forms and so on. This leads to the general definition of k-forms on manifolds.

Definition 3.2.2. *A differential form of degree k, or k-form, at a point x on a manifold M is a k-linear and antisymmetric function of k vectors of the tangent space T_xM of M at x. That is, for k vectors $\xi_1, \cdots, \xi_k$ in T_xM:*

$$\omega(\lambda'\xi_1' + \lambda''\xi_1'', \xi_2, \cdots, \xi_k) = \lambda'\omega(\xi_1', \xi_2, \cdots, \xi_k) + \lambda''\omega(\xi_1'', \xi_2, \cdots, \xi_k)$$
$$\omega(\xi_{i_1}, \xi_{i_2}, \cdots, \xi_{i_k}) = (-1)^v\omega(\xi_1, \xi_2, \cdots, \xi_k)$$

where

$$v = \begin{cases} 0 & \textit{if the permutation } (i_1, ..., i_k) \textit{ is even} \\ 1 & \textit{if the permutation } (i_1, ..., i_k) \textit{ is odd} \end{cases}$$

Differential forms come along with an entire calculus, called exterior calculus, to manipulate them. The first operation attached to differential form is the wedge product $\wedge$, the extension of exterior product to forms. For a p-form ω^p and a k-form ω^k, the wedge product $\omega^p \wedge \omega^k$ is a $(p+k)$-form, that is:

1. Skew-commutative: $\omega^p \wedge \omega^k = (-1)^{pk} \omega^k \wedge \omega^p$
2. Distributive: $(\lambda_1 \omega_1^p + \lambda_2 \omega_2^p) \wedge \omega^k = \lambda_1 \omega_1^p \wedge \omega^k + \lambda_2 \omega_2^p \wedge \omega^k$.
3. Associative: $(\omega^p \wedge \omega^k) \wedge \omega^l = \omega^p \wedge (\omega^k \wedge \omega^l)$

To further illustrate the relation between the wedge product and forms, consider the flat manifold $\mathbb{R}^n$ with the coordinate system $(x_1, \cdots, x_n)$, for a point x, the n 1-forms $dx_1, \cdots, dx_n$ form a basis of the space of 1-forms on $T_x\mathbb{R}^n$. From this it follows:

Theorem 3.2.1. *Every differential k-form on $\mathbb{R}^n$ can be uniquely decomposed into:*

$$\omega^k = \sum_{i_1 < \cdots < i_k} a_{i_1, \cdots, i_k}(x) dx_{i_1} \wedge \cdots \wedge dx_{i_k}.$$

Where $a_{i_1, \cdots, i_k}(x)$ are smooth functions on $\mathbb{R}^n$.

For an intrinsic point of view, we may introduce the k^{th} exterior power of a vector space V, noted $\Lambda^k(V)$, defined by the space spanned by elements of the form $x_1 \wedge x_2 \wedge \cdots \wedge x_k$, where $x_i \in V$ for $1 \leq i \leq k$. Recall that a real Hilbert space H is a vector space with an inner product, noted (x, y) for $x, y \in H$, satisfying:

$$\begin{aligned}
(ax_1 + bx_2, y) &= a(x_1, y) + b(x_2, y), \text{ for all } a, b \in \mathbb{R} \\
(x, y) &= (y, x) \\
(x, x) &\geq 0 \\
(x, x) &= 0 \text{ iff } x = 0
\end{aligned}$$

By fixing an element $y \in H$, then F_y defined by $F_y(x) = (y, x)$ assigns to each $x \in H$ a real number and is a linear bounded functional on H. The space of all linear bounded functional of a Hilbert space H is called its dual space and is noted H^*. The following Riesz representer theorem clarifies the type of linear bounded functional one can find on Hilbert spaces:

Theorem 3.2.2. *For every bounded linear functional $F \in H^*$ on a real Hilbert space H, there exists a unique $y \in H$ such that:*

$$F(x) = (y, x), \text{ for all } x \in H$$

The Riesz operator $\mathcal{R} : H^* \to H$ is an isometric isomorphism.

For a Riemannian manifold (M, g) and the Lebesgue measure on the M, the space of differential p-forms $\Lambda^p(M)$ is the dual space to the tangent bundle TM. The isometric isomorphism defined by Riesz is given in this case by the musical isomorphisms $\flat : TM \to T^*M$ and $\sharp : T^*M \to TM$. Specifically, for a 1-form ω, $\omega^\sharp$ is the vector field defined by $\omega(X) = g(\omega^\sharp, X)$ for every vector field X, and $\flat$ its inverse. This makes it possible to obtain an L^2 metric on smooth differential 1-forms $\Lambda^1(M)$ as:

$$(\alpha, \beta)_{L^2} = \int_M g(\alpha^\sharp, \beta^\sharp) d\mu$$

We can extend this definition to $\Lambda^p(M)$ as:

$$(\alpha, \beta)_{L^2} = \int_M g^p(\alpha, \beta) d\mu. \tag{3.3}$$

for any p-forms $\alpha, \beta \in \Lambda^p(M)$. Where g^p is a pointwise metric on $\Lambda^p T_x^* M$:

$$g^p(\alpha_1 \wedge \cdots \wedge \alpha_p, \beta_1 \wedge \cdots \wedge \beta_p) = \sum_\pi \text{sign}(\pi) g(\alpha_1^\sharp, \beta_{\pi(1)}^\sharp)(x) \cdots g(\alpha_p^\sharp, \beta_{\pi(p)}^\sharp)(x).$$

Where π range over the set of permutations of $\{1, \cdots, p\}$.

Aside the fact that duality is a powerful mathematical tool, this is useful for the visualisation (and discretisation) of differential forms. In our case, 3D spaces, this means we can conceive 1-forms and 2-forms as vector fields and 0-forms and 3-forms as scalar field.

3.3 Spectral shape analysis

As we have seen previously, the general goal of Riemannian geometry is to extend calculus to manifolds. This is done by working with a metric g, that in turn characterises the manifold, or surface in the case of shapes. In this section we will see how the extension of operators can help representing shapes in a concise manner, as well as manipulating them efficiently. Within the general field of spectral shape analysis, we will only overview how the extension of the Laplace operator can help us represent shapes and compute distances on them. We refer the readers to the review paper [177] for a complete overview of this field.

3.3.1 The Laplace-Beltrami Operator

The Laplace operator notoriously shows up in many places of the euclidean world, and therefore it is natural to transfer it to the Riemannian setting. To define it properly, recall that the Laplacian of a scalar function f is:

$$\Delta f = \text{div}(\text{grad} f) \tag{3.4}$$

The generalised version, also called the Laplace-Beltrami operator, uses the Riemannian divergence and the Riemannian gradient. Consider the local parametrisation φ of a manifold M, we can write the metric g component-wise as $g_{ij} = \langle \partial_i \varphi, \partial_j \varphi \rangle$. If we denote g^{ij} as the components of the inverse of g, then the gradient of a scalar function f on M is defined in local coordinates and with the Einstein notation:

$$(\text{grad} f)^i = g^{ij}\partial_j f$$

And the divergence of a vector field X is defined thanks to the Levi-Civita connection, in local coordinates again:

$$\nabla \cdot X = \frac{1}{\sqrt{|g|}}\partial_i(\sqrt{|g|}X^i)$$

When combining both together we get the Laplace-Beltrami operator:

$$\Delta f = \frac{1}{\sqrt{|g|}}\partial_i(\sqrt{|g|}g^{ij}\partial_j f) \tag{3.5}$$

This fundamental operator has many uses when working on surfaces. We can see that the Laplace-Beltrami operator contains the metric in its formulation, which in turn contains a lot of valuable information about the *shape* of the manifold. Amongst the possible ways to extract what the laplacian tells us about the surface, one is via its eigen decomposition, ensured by the spectral theorem and computed by solving:

$$\Delta f = \lambda f \tag{3.6}$$

In an even more concise way, one can ask how well the eigenvalues of the Laplace operator (or spectrum) can describe a shape. This area of research was most notably started by Kac in [88], and knows a long history of results, theoretical and applied. As it turns out, it is impossible to fully reconstruct a shape from its spectrum, implying that different shapes can have the same spectrum. However, such examples are hard to find in the wild. In fact, it is possible to reconstruct, at the minimum, the dimension, total curvature and volume of a shape from its spectrum. As such similar shapes are hard to find in a set of surfaces representing, e.g anatomies, it makes the Laplace spectral decomposition a suitable tool to discriminate them. This idea was initially proposed in [150], and is alluring for shape analysis because of some particular useful properties. First of all, the Laplace-Beltrami operator 3.5 does not depend on any choice of coordinates, meaning that any representation stemming from it is intrinsic and two isometric shapes will have the same spectrum. A desirable property when studying shapes, as we do not want rigid transformations between shapes to impact their analysis. Second of all, we wish to compare shapes that are discretised in different ways. In our case, the mesh representation for each LA is computed to fit their shape optimally, and therefore do not have the same number of points, edges or faces. Because eigenvalues are a set of real scalars, as long as we decompose each shapes with a fix number of n eigenvalues, we can represent and compare them in $\mathbb{R}^n$. This out of the box representation method makes it a perfect candidate for a first attempt at predicting stroke from the shape of the LA. We do so in the next section, but first let's note another important tool that stems from the Laplace-Beltrami operator: the computation of distances on a surface.

An important tools when processing meshes is the possibility to compute distance between two points. In fact, this will be an important element for computing geometric quantities of the LA, in chapter 6. There exists a long list of methods to do so, for instance if a mesh can be viewed as a graph, and one can apply the famous Dijkstra algorithm [44] to find the shortest path between two points. A more efficient way to do so is by considering the surface as a manifold M, translate the problem as finding the geodesic distance ϕ between any two points, and find a PDE easy to solve that yields the distance. In this vein, the preveiling method is through the Eikonal equation $|\nabla\phi| = 1$, which directly yields the geodesic distance, as initially proposed in [95]. Although this method knows a lot of efficient implementations, it requires solving the equation from scratch each time we want to compute a distance on the surface and it is difficult to parallelise. To alleviate this, [38] proposes the *heat method* leveraging on robust and sparse representation of the Laplace operator on M, and computing the geodesic via the three following steps:

1. Integrate the heat flow $\dot{u} = \Delta_M u$ for a fixed time t
2. Evaluate the vector field $X = -\nabla u/|\nabla u|$.
3. Solve the Poisson equation $\Delta_M \phi = \nabla \cdot X$

Therefore keeping in store a lot of information to have faster successive computation of geodesic on a given manifold M. We will use this specific method to define several geometric features of the LA in chapter 6. It is good to note that this method can also help parallel transport vectors across the surface or compute the log map at a given source point.

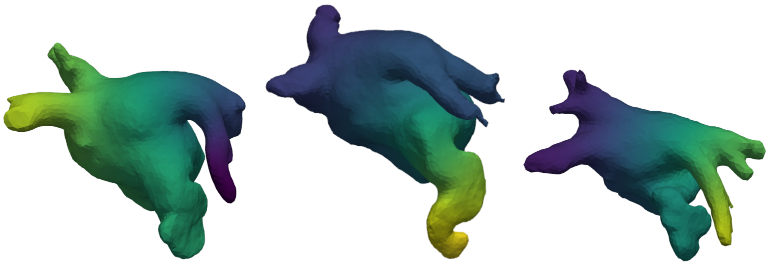

Figure 3.1: First eigenfunction of the Laplacian of different LAs.

3.3.2 Use case: a LAA study via the Laplace spectrum

As stated above, as a widely used technique to study cohorts of surfaces, spectral shape analysis can be tempting for our problem. Indeed, this out of the box method can be used with minimum pre-processing being isometry and discretisation invariant. What is left to know is whether the distance between shapes offered by this method has any relevant anatomical meaning. In practice, faithful eigenvalues rely on a faithful Laplace approximation. We use the discrete laplacian proposed in [163], and the python implementation attached to it[1], for its compatibility with a wide range of surface representations, and keep this method throughout the manuscript. Such a method yields a sparse matrix of the representation of the laplacian L as well as the sparse mass matrix M, which

---

[1] https://github.com/nmwsharp/robust-laplacians-py

are used to find the eigenvalues through the generalised eigenvalue problem:

$$Lx = \lambda Mx$$

which we solve using the `scipy` sparse linear algebra library, for the first 128 (ordered) eigenvalues. We apply this method on the available cohort of LAs, however we already see an important setback. The orthonormal basis are hardly related between different LAs, in turn, the m-th eigenfunctions of (nearly) any two LAs do not describe the same anatomical variations, as shown in 3.1. Which finally impacts their eigenvalues as they will describe the magnitude of different aspects in the anatomy, making a direct comparison impossible. However, all is not lost. An important part in the stratification of stroke risk from the LA shape is actually focused on the LAA, rather than the entire chamber, due to the fact that most clot formations are formed in the LAA. We can separate the LAA with labelling techniques proposed later in the manuscript, and compute the Laplacian's spectrum on them directly. While there are still discrepancies between different LAAs eigenfunctions, the spectral basis seem to coincide better in between shapes. However, it is still hard to determine which variations of the shape are explained by each eigenfunctions, as shown in 3.2. This is a notoriously hard problem in any cases. We may still see if this representation has any predictive power regarding stroke risk.

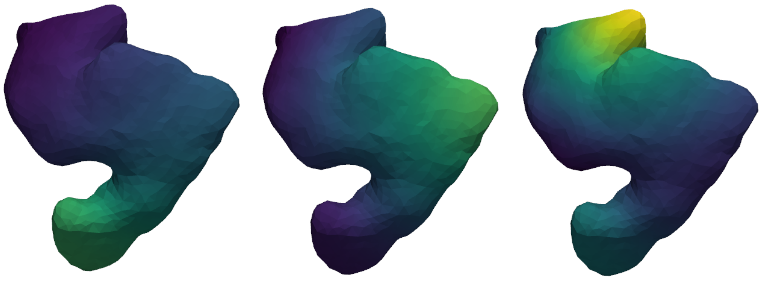

Figure 3.2: First 3 eigenfunctions of the Laplacian of an LAA

We propose two small studies to highlight how these representations may be used in our case. First we conduct a K-Means cluster analysis, to understand whether the LAA's spectra have natural groups regardless of stroke incidence. We find via the silhouette score that the dataset is optimally separated in two groups. However, it is hard to interpret how these two groups differ. In addition, it seems these two groups have little to do with stroke incidence. We visualise this using the TriMap dimension reduction method proposed in [4], and show both the cluster groups and the stroke groups in 3.3.

In a second time we propose a supervised study, in which we apply a random forest to learn a classification between patients with and without stroke. To pick the number of eigenvalues, we perform it on a growing number of them, up to 128. We find 91 first eigenvalues to yield best results and perform 5 fold cross-validation. We report the test results in table 3.1 and compare it to the same procedure on the LA. We also add the classification proposed in [100], which is a classification

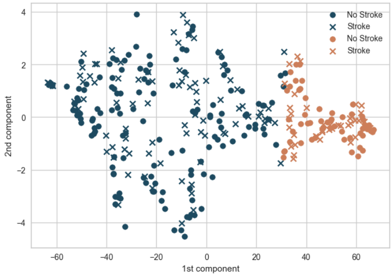

Figure 3.3: TriMap of LAA Laplace spectrum

of the LAA done by clinicians into 4 types called chicken-wing, cauliflower, windsock, and cactus. Typically 3 of the 4 types are considered at risk, that is, all but chicken-wing. Although better than the other proposed methods, either the spectrum representation is not well adapted for our problem, or it is impossible to distinguish an LA at risk from a normal one. In either case we do not have a lot of leg room for follow up analysis. The spectrum representation does not allow for interpretation, a n-th eigenvalue of x amount has hardly any meaningful clinical importance. Although we leave it at that, this will serve as our baseline for shape-related stroke stratification in the later chapters of this thesis.

	Accuracy	F1	Precision	Recall	AUC
LAA *type* [100]	.568	.588	.647	.507	.558
LA spectrum	.606	.336	.542	.246	.548
LAA spectrum	.655	.426	.677	.315	.621

Table 3.1: Spectral results

3.3.3 Final remarks

The idea that a surface can be fully represented through the spectrum of an operator has generated a lot of interest. Many papers have suggested other, and perhaps better, operators instead of the Laplacian, such as a Dirac operator [112], an overview of these is presented in [177]. In parallel, other representations of surfaces based on operators exist, a notable proposition based on the Laplacian is the Heat Kernel Signature (HKS) in [169], which leverages the eigendecomposition of the Laplacian to approximate the heat kernel, and derives multi-scale representations of the surface from it. This representation will re-appear in chapter 5.

3.4 Diffeomorphometry

To study the variations of shapes within a cohort, one can compute these variations directly. This idea dates back to the biologist D'Arcy Thompson, who set out to measure the evolution of species by modelling it as a geometric transformation in [40]. This theory of patterns was later formalised by Grenander in [64], and the definition and study of such transformations was introduced by Michor and Mumford in [132], and is now generally referred to as diffeomorphometry. As we expose the

geometric ides of this field, we take time introducing and using the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework. However it is important to note the existence of other methods surrounding diffeomorphometry, for instance the intrinsic view introduced by the SRNF method [82], and its Sobolev counterpart [73].
To make things mathematically meaningful, recall that *shape*, in our case, are submanifolds of $\mathbb{R}^3$. The transformations we seek out need to be anatomically relevant, in mathematical terms they imply smooth orientation preserving diffeomorphism.

3.4.1 Diffeomorphisms

Considering our shapes to live in the three dimensional Euclidean space $\mathbb{R}^3$ with the standard Euclidean inner product $\langle \cdot, \cdot \rangle$, a diffeomorphism is a smooth bijective map

$$\eta : \mathbb{R}^3 \to \mathbb{R}^3$$

with smooth inverse η^{-1} and positive Jacobian determinant $\det D\eta > 0$. The set, and smooth manifold, of all orientation preserving diffeomorphisms of $\mathbb{R}^3$ is denoted

$$\mathcal{D}\left(\mathbb{R}^3\right) = \{\eta : \mathbb{R}^3 \to \mathbb{R}^3 : \eta \,\text{and}\, \eta^{-1} \,\text{are smooth, with}\, \det D\eta > 0\}.$$

The tangent space to $\mathcal{D}\left(\mathbb{R}^3\right)$ at the identity diffeomorphism is the vector space of smooth vector fields on $\mathbb{R}^3$ equipped with the Lie bracket

$$T_e\mathcal{D} = \{X : \mathbb{R}^3 \to T\mathbb{R}^3 : X \,\text{is}\,\text{smooth}\}.$$

The tangent space to $\mathcal{D}\left(\mathbb{R}^3\right)$ at a point $\eta \in \mathcal{D}\left(\mathbb{R}^3\right)$ is given by the set of smooth sections of the pullback bundle $\eta^* T\mathbb{R}^3$; in particular, the set of smooth vector fields covering η:

$$T_\eta\mathcal{D}\left(\mathbb{R}^3\right) = \{X : \mathbb{R}^3 \to T\mathbb{R}^3 : \pi_{\mathbb{R}^3} \circ X = \eta\}.$$

where $\pi_{\mathbb{R}^3} : T\mathbb{R}^3 \to \mathbb{R}^3$ is the tangent bundle projection map.

The identity map e is a diffeomorphism; the composition of two diffeomorphisms is smooth, and again, a diffeomorphism; and the operation of inversion $\eta \mapsto \eta^{-1}$ is smooth. In particular, $\mathcal{D}\left(\mathbb{R}^3\right)$ has a group structure, and therefore, in addition to being a smooth manifold, is a Lie Group. If $R_\eta : \mathcal{D}\left(\mathbb{R}^3\right) \to \mathcal{D}\left(\mathbb{R}^3\right)$ denotes composition on the right by a fixed diffeomorphism $\eta \in \mathcal{D}\left(\mathbb{R}^3\right)$ then its differential

$$dR_\eta : T_e\mathcal{D} \to T_\eta\mathcal{D}$$

is given by composition on the right with η

$$dR_\eta \cdot X = X \circ \eta.$$

If $L_\eta : \mathcal{D}\left(\mathbb{R}^3\right) \to \mathcal{D}\left(\mathbb{R}^3\right)$ denotes composition on the left by a fixed diffeomorphism $\eta \in \mathcal{D}\left(\mathbb{R}^3\right)$ then its differential

$$dL_\eta : T_e\mathcal{D} \to T_\eta\mathcal{D}$$

is given by

$$dL_\eta \cdot X = D\eta \cdot X.$$

Remark 3.4.1. *Although we are endowing the space $\mathcal{D}\left(\mathbb{R}^3\right)$ with a geometric structure, this space is in fact only a smooth infinite dimensional* Frechet *manifold and calculus in Frechet manifolds is quite involved due to the failure of the inverse function Theorem. Furthermore, $\mathbb{R}^3$ is an open manifold and we need to consider appropriate decay conditions at infinity for both the diffeomorphisms and the vector space of vector fields. In this thesis we will not give a detailed treatment of calculus in Frechet spaces, and we will not go into any real or functional analytic details to make our definitions and objects mathematically rigorous; instead, having acknowledged the inherent analytic difficulties we will proceed formally with the given definitions and refer the reader to [76] for a detailed mathematical construction of the diffeomorphism groups on $\mathbb{R}^n$, on open subsets Ω, or on compact manifolds M.*

This covers what we mean by transforming a shape in $\mathbb{R}^3$. That is, we can smoothly deform surfaces S in $\mathbb{R}^3$, with diffeomorphisms η of $\mathbb{R}^3$; the new surface is the image of S, respectively, under the diffeomorphism:

$$\eta(S) = \{\eta(x) \in \mathbb{R}^3 : x \in S\}.$$

3.4.2 Riemannian metrics on the Diffeomorphisms group

Just like on finite dimensional manifolds we can equip the infinite dimensional Frechet manifold $\mathcal{D}\left(\mathbb{R}^3\right)$ with a Riemannian metric. This is most easily done by specifying an inner product $(\cdot,\cdot)(e)$ on the Tangent space at the identity $T_e\mathcal{D}$, also called the Lie algebra. We can then extend it to a metric $(\cdot,\cdot)(\eta)$ on the remaining tangent spaces $T_\eta\mathcal{D}$ by right-invariance. This is a short-hand way of saying that the value of $(\cdot,\cdot)(\eta)$ on two tangent vectors X_η and Y_η in $T_\eta\mathcal{D}$ is given by the value of $(\cdot,\cdot)(e)$ on their right translates $\tilde{X}_e = dR_{\eta^{-1}} \cdot X_\eta = X_\eta \circ \eta^{-1}$ and $\tilde{Y}_e = dR_{\eta^{-1}} \cdot Y_\eta = Y_\eta \circ \eta^{-1}$:

$$(X_\eta, Y_\eta)(\eta) := \left(dR_{\eta^{-1}} \cdot X_\eta, dR_{\eta^{-1}} \cdot Y_\eta\right)(e) = \left(X_\eta \circ \eta^{-1}, Y_\eta \circ \eta^{-1}\right)(e).$$

L^2 metric We can define an L^2 inner product at the identity using the Euclidean inner product $\langle\cdot,\cdot\rangle$

$$(X, Y)(e) = \int_{\mathbb{R}^3} \langle X, Y\rangle \, d\mu$$

which can be extended to a right-invariant metric on all of $\mathcal{D}\left(\mathbb{R}^3\right)$ as explained above. Recall that for everything to work, we need to consider appropriate conditions at infinity. Nothing in this context ensures that the integral is finite and by extension that the metric is well defined. One often talks about completions of smooth compactly supported vector field in L^p where metrics are well defined. But we will ignore these technical difficulties and proceed with the definition of Sobolev metrics, as if there are no convergence problems.

Sobolev metric If $\Delta = -\left(\partial_x^2 + \partial_y^2 + \partial_z^3\right)$ denotes the usual – positive definite – Laplace operator on $\mathbb{R}^3$ then we can lift it to an operator on vector fields by applying it component-wise: if $X = f \cdot \partial_x + g \cdot \partial_y + h \cdot \partial_z$ for functions f, g, and h then

$$\Delta X = \Delta f \cdot \partial_x + \Delta g \cdot \partial_y + \Delta h \cdot \partial_z.$$

Denoting by Δ^k the application of Δ k times, where $k = 0$ denotes the identity operator I, we define a Sobolev H^s metric at the identity by

$$(X, Y)_s (e) = \sum_{k=0}^{s} \int_{\mathbb{R}^3} \langle X, \Delta^k Y \rangle \, d\mu.$$

Integration by parts shows that the vector Laplacian is self-adjoint so that the H^s metric is symmetric and it does not matter on which vector field it is applied under the integral (of course, outside of the integral $\langle X, \Delta^k Y\rangle \neq \langle \Delta^k X, Y\rangle$). While the L^2 metric measures the average of the inner product between X and Y over $\mathbb{R}^3$, the Sobolev H^s metric measures the derivatives of the field as-well - this can have geometric consequences. The Sobolev metric can also be extended to a right-invariant metric $\mathcal{D}\left(\mathbb{R}^3\right)$ following the given recipe. It will be convenient to rewrite the Sobolev metric as

$$(X, Y)_s (e) = \int_{\mathbb{R}^3} \langle X, A^s Y \rangle \, d\mu$$

where $A^s = \sum_{k=0}^{s} \Delta^k$.

Instead of the vector Laplace operator one can more generally consider any self-adjoint operator $L : T_e\mathcal{D} \to T_e\mathcal{D}$ and define an inner product at the identity by

$$(X, Y)_L (e) = \int_{\mathbb{R}^3} \langle X, L \cdot Y \rangle \, d\mu \tag{3.7}$$

and extend this to a right-invariant metric on $\mathcal{D}\left(\mathbb{R}^3\right)$ following the given recipe.

3.4.3 Geodesics on the Diffeomorphisms group

In finite dimensional Riemannian geometry, a smooth Riemannian metric automatically generates a smooth Levi-Civita connection with which one can define geodesics. In infinite dimensions, this is not the case and certainly not straightfoward in the Frechet setting. One typically needs to prove smoothness of the metric, and connection on a case by case basis, and then conclude local existence and uniqueness of geodesics. For the diffeomorphism groups with Sobolev type metrics, which includes the L^2 metric, this was done in [124]. Since smooth diffeomorphisms are dense in the Sobolev class their results restrict to our Frechet setting of diffeomorphism groups and we can use their results directly. It is quite complicated to write down an explicit expression for the connection ∇ of a right-invariant metric on $\mathcal{D}\left(\mathbb{R}^3\right)$ but recall from section 1 that geodesics are equivalently characterised as local extrema of the energy functional

$$\mathcal{E}\left(\eta(t)\right)_0^t = \frac{1}{2}\int_0^t \left\|\partial_s\eta(s)\right\|^2\left(\eta(s)\right)ds.$$

The velocity vector $\partial_t\eta(t) \in T_{\eta(t)}\mathcal{D}$ belongs to the tangent space of $\mathcal{D}\left(\mathbb{R}^3\right)$ at each point $\eta(t)$, but we can right-translate this velocity vector field back using $dR_{\eta^{-1}(t)}$ to obtain a curve $v(t)$ in $T_e\mathcal{D}$:

$$v(t) = \partial_t\eta(t) \circ \eta^{-1}(t).$$

This allows us to rewrite the energy functional as

$$\mathcal{E}\left(\eta(t)\right)_0^t = \frac{1}{2}\int_0^t \left\|dR_{\eta^{-1}}(s)\cdot\partial_s\eta(s)\right\|_s^2(e)\,ds. \tag{3.8}$$

and express the geodesic equations as an equation in $T_e\mathcal{D}$. [124] carries out the derivations for explicit geodesic equations. Once we solve the velocity vector $v(t)$, it is possible to obtain the geodesic curve $\eta(t)$ by solving the flow, or reconstruction, equation

$$\partial_t\eta(t) = v(t) \circ \eta(t). \tag{3.9}$$

The curve $\eta(t) \in \mathcal{D}\left(\mathbb{R}^3\right)$ is a family of diffeomorphisms and, at each t, gives us a diffeomorphism with which we can transform/deform $\mathbb{R}^3$ and any shape occupying space in it.

3.4.4 The shape matching problem

With some geometry of the diffeomorphism group in hand we can now state the central problem of diffeomorphometry. Given a source shape S and a target shape $\tilde{S}$ we want to find a diffeomorphism η of $\mathbb{R}^3$ taking S close to $\tilde{S}$ within a prescribed matching accuracy ϵ, but we also want η to be as close to the identity as possible. The entire construction above is to show that we have metrics and their geodesics, as well as formulas allowing us to solve them, that allow us to construct a geodesic-like curve in $\mathcal{D}\left(\mathbb{R}^3\right)$ whose end point fulfils the requirements of diffeomorphometry as thought up by D'Arcy Thompson. In this language, the shape matching is expressed through the constrained optimsiation problem

$$\operatorname{argmin}_{\eta\in\mathcal{D}(\mathbb{R}^3)}\left(d_*(e,\eta) + d\left(\eta(S),\tilde{S}\right)\right) \tag{3.10}$$

where d_* is the Riemannian distance function induced by the right-invariant metric $(\cdot,\cdot)_*$ on $\mathcal{D}\left(\mathbb{R}^3\right)$ and d is a measure of similarity between two given shapes, which we will talk about in the next section. This is the general method of the Large Deformation Diffeomorphic Metric Mapping (LDDMM). We introduced several classes of metrics on the diffeomorphism group in section 2.3, but this was only for concreteness. When phrased in general terms one sees that the optimisation solution is quite flexible, allowing for many different choices of metrics, including those learned from explicit examples. Given an explicit metric on the diffeomorphism group, and given an explicit similarity measure d, one can find the solution to the optimisation problem by solving a system of partial differential equations on $\mathbb{R}^3$ involving the geodesic equation of the given right-invariant metric plus an inhomogeneity constraint depending on the similarity measure. Describing everything from the geometry of the diffeomorphism group also brings to the table a whole set of tools to extend the framework for the analysis of shapes. Typically, a cohort of n shapes can be described by a template shape or atlas T, computed as the mean of the set of shapes, called the Fréchet mean, and a set of diffeomorphism $\{\varphi_i\}_{i\leq n}$ computed as geodesics. This representation allows us to analyse things such as the main modes of variation from this template, or the development of learning-based algorithms such as clustering, classification or regression. In the next section, we explain how the diffeomorphism between two shapes is computed in practice.

3.4.5 Practical construction of diffeomorphisms

Now that we have introduced the geometric view on the deformation-based shape analysis, we can now see how to construct the discretised version of the geodesic-like diffeomorphism between two shapes, in the context of the LDDMM framework. Although the Riemannian metric on diffeomorphism groups call for self-adjoint differential operators, they are notoriously hard to work with. Instead, the large majority of studies use kernel operators, which offer an algebraic view and yielding positive definite matrices, easier to work with. Precisely, kernels are positive semi-definite scalar functions, that can be attached to a Hilbert space via the following theorem by Aronszajn [10]:

Theorem 3.4.1. *For any kernel k on $\mathbb{R}^3$, there exists a Hilbert space $(H, \langle\cdot,\cdot\rangle)$ and a mapping $\phi : \mathbb{R}^n \to H$ such that for all $x, y \in \mathbb{R}^3$*

$$k(x, y) = \langle\phi(x), \phi(y)\rangle$$

In fact, when k is positive definite, then it is a *reproducing kernel.* That is, it uniquely defines a Hilbert space H. This space is called a RKHS, and the kernel acts as the inner product on H. We will not go into any details regarding this facet of LDDMM, and we refer the readers to [52] for a comprehensive exposition.

Given the RKHS structure and a kernel of choice k, shape S is defined via a set of N control points $(c_k^{(t)})_{k=1,\ldots,N}$. Then the velocity vector $u(t)$ of equation 3.9 takes the form

$$v_t(x) = \sum_{k=1}^{N} k(x, c_k^{(t)})\mu_k^{(t)} \tag{3.11}$$

where the $(\mu_k^{(t)})_{k=1,\ldots,N}$ are a set of momentum vectors attached to each control point c_k at a given time t. In the rest of this work we choose k to be a Gaussian kernel, defined as $k(x, y) = \exp(-\frac{\|x-y\|^2}{\sigma^2})$, although other kernels exist, their impact have been studied in [122]. We can further define the associated norm to all such vector field in the RKHS as

$$\|v\|_K^2 = \sum_{i,j} K(c_i, c_j)\mu_i^T \mu_j. \tag{3.12}$$

The energy functional 3.8 of a given diffeomorphism takes the simpler form

$$\int_0^1 \|v_t\|_K^2 dt$$

And the equations driving the diffeomorphism can be redefined into a Hamiltonian system of ODEs, written here in matrix notation:

$$\begin{cases} \dot{c}^{(t)} = K(c^{(t)}, c^{(t)})\mu^{(t)} \\ \dot{\mu}^{(t)} = -\frac{1}{2}\nabla K(c^{(t)}, c^{(t)})\mu^{(t)T}\mu^{(t)} \end{cases} \tag{3.13}$$

This rewritting is particularly suitable for the use of numerical schemes in direct applications. In the case of shape matching, i.e registering a source shape S into a target shape S'. The optimal transformation which we stated as equation 3.10, is rewritten as optimising a cost function in terms of control points c and momenta μ of the form:

$$C(c, \mu) = d(\phi_1.S, S')^2 + ||v_0||_K^2 \tag{3.14}$$

Optimisation of C is done via gradient descent with the help of automatic differentiation. While the distance d is often simply defined as the L_2 norm between landmarks, the lack of point-to-point correspondance between shapes renders it void. To alleviate this problem [172] proposed a *currents*-based distance built with kernels, yielding the inner-product between two surfaces S and S' of the form:

$$\langle [S], [S'] \rangle = \sum_{i,j} K(c_i, c_j')n_i^T n_j'. \tag{3.15}$$

Where c_i (resp. c_j') are the control points of S (resp. S') and n_i (resp. n_j') their normals and K is a kernel of choice. Although effectively solving the discretisation discrepancies, the kernel-based current distance can induce deformations that collapse shapes together when normals within a small neighbourhood have opposite orientation. A distance based on Varifolds was proposed in [32] that ignores normal orientation:

$$\langle [S], [S'] \rangle = \sum_{i,j} K(c_i, c_j')\frac{(n_i^T n_j')^2}{\|n_i\|\|n_j\|}. \tag{3.16}$$

Without going into any further details, we will dedicate the next section to expose the importance of a correct definition of shape similarity. Let us now see how such constructions can help us.

3.4.6 Use Case 1: Automatic labeling

The LDDMM setting for constructing diffeomorphism has direct applications related to our own problematic. Throughout the manuscript, our goal is to extract relevant geometric descriptors of the LA. Towards this idea, we wish to leverage on the anatomical parts of the LA, however it is not straightforward to identify them automatically on a mesh. This identification translates into labelling, or semantic segmentation, of the LA mesh into 6 different classes, defined by 6 anatomical parts of interest: the 4 PVs, the body and the LAA. While learning-based methods may sound attractive, we would need labels to train, which are not known to us. Another solution is to leverage on the diffeomorphic shape matching method explained previously. If we can represent our entire cohort as one template shape T and a set of deformations from this template $\{\phi_i\}_{i\leq n}$, where n is the size of our cohort, then we would only need to hand-label the template. That is, we could approximate any shape S_i by $\phi_i.T$ and interpolate the labels carried out by warping T with ϕ_i onto S_i, as illustrated in 3.4. It is important to note that the LDDMM framework yields a deformation of the ambient space, that is, it warps the space in which the shape lives, rather than intrinsically transforming the shape. This offers an important advantage, the deformation can be applied to any discretisation of the shape, we may compute it on a coarse mesh and apply it on a very fine mesh, optimising the computation considerably, however this extrinsic view also requires careful attention to the placement of the shapes. If two shapes live far away from each other in the ambient space, the resulting transformation of matching one to the other would be a combination of a rigid and an elastic transformation, which would result in a big computational overhead. In addition, the numerical scheme is not well adapted to such large deformation, and would yield erroneous results. As our LA dataset does not share a common coordinate reference, we have to address this problem. To do so, we follow the classical methodology by performing rigid registration before applying the deformation procedure. Therefore we apply the ICP algorithm, reviewed in [156], which ensures convergence under the assumption of good initial transformation. This initial transformation can be, e.g aligning the principal axes of inertia, or Procrustes analysis (see [63]). We found that the choice of initial transformation has a big impact on the final alignment, to ensure robustness we apply both examples and pick the method that yields the smallest Hausdorff distance post ICP. Now that all meshes are in a common reference system, we may compute the template representation of our LA cohort, also called an atlas model. Recall we aim at finding a *mean* shape T and a collection of diffeomorphisms $\{\phi_i\}_{i\leq n}$, this is done as described in the previous section, although the cost function 3.14 is of the more general form:

$$C(T, c, \mu_{i=1,\dots,n}) = \sum_i d(\phi_i.T, S_i)^2 + \sum_i ||v_i||_K^2 \tag{3.17}$$

In this setting, the template is a shape initially picked from the dataset and smoothly deformed to minimise 3.17. As it turns out, the resulting mean shape is strongly conditioned by the starting point, this may be due to the distances between all shapes being very close w.r.t the entire space of diffeomorphisms, yielding sub-threshold distances very fast. For this reason we need to be careful when choosing this initial shape. To do so, we find the 5 shapes which the volume is the closest to the mean volume of the dataset, and manually pick the shape with the straightest LAA.

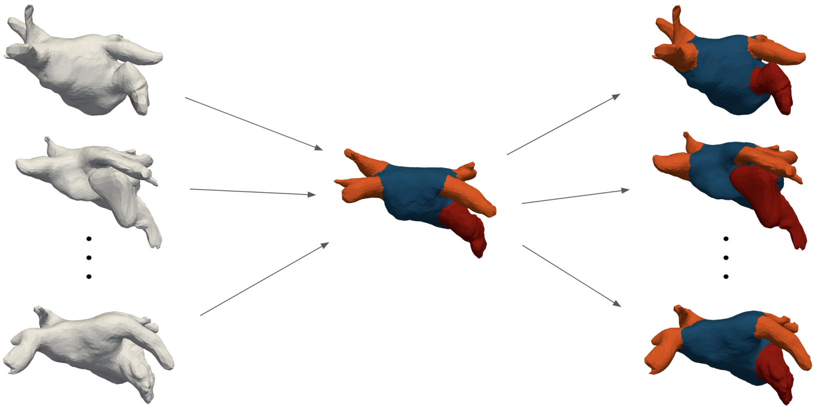

Figure 3.4: Labelling pipeline.

We may now give a summary of the labelling pipeline in the correct order:

1. Centre the LA shapes to have zero mean
2. Pick the template's starting point as the mesh which volume is closest to the mean and LAA's is the straightest
3. Apply the rigid registration procedure to the entire dataset with the template as target
4. Compute the atlas model of the LA cohort.
5. Hand label the resulting template
6. Warp the template by each deformation ϕ_i and interpolate the labels from $\phi_i.T$ to S_i.

Implementation The LDDMM framework has been efficiently implemented in the open-source software `deformetrica`, see [25]. We use the L-BFGS algorithm for optimisation of the cost function 3.17 and the 4th order Runge-Kutta scheme for the ODE system 3.13 with the Gaussian kernel width set to 10. We also employ the use of the currents kernel, which we will discuss in the following section 3.5, with kernel width of 10. All other parameters are kept from the default setting, we run the registration procedure for 300 iterations. Regarding the rigid registration pipeline we use the implementations proposed in `trimesh`. All other tasks from volume computation to hand labelling and manually correcting the labelling were implemented with the help of `pyvista`.

Results The rigid registration procedure takes under 30sec to compute for each mesh, and the LDDMM registration procedure takes about 30min per mesh. After labelling, we corrected each mesh, yielding a ground truth labelling of the LA enabling us to measure the success of the pipeline. As this is a *human-in-the-loop* process, we show both the mean metrics as well as the worst case results in table 3.2, although the average results are high, yielding a generally fast hand correction, the worst cases are sub-par. As our main goal is to automatically process LA shapes to extract robust shape features, we will propose several other methods along the manuscript, namely, a *multi-atlas* version of the pipeline in 4, as well as a different route, based on neural networks in 5. Nevertheless, this first step gives a good feeling of the task at hand, and will inform our later and better versions.

	Accuracy	f1 score	Precision	Recall
Mean	0.926 ±0.058	0.925 ±0.058	0.956 ±0.057	0.924 ±0.058
Min	0.557	0.558	0.559	0.557

Table 3.2: Labelling pipeline results.

3.4.7 Use Case 2: Motion generation

Although we do not use the Riemannian structure offered by the LDDMM framework for statistical shape analysis, we can fully take advantage of it for generating motion out of static LA shapes from a single example. This use case will highlight how to use the concept of parallel transport to do so, and how it can be relevant in studying cardiac properties. Specifically, we set out to study the haemodynamic impact of fibrillation on the LA, in collaboration with the computational physiology group of Simula[2] who developed a robust Computational Fluid Dynamics (CFD) tool to simulate cardiovascular flow in a moving domain, proposed in [98], and both the Sant Pau Hospital of Barcelona, and the Hamburg Hospital that provided the dataset.

Context Studying the impact of fibrillation on the Left Atrium (LA) would require a complete spatio-temporal model. We start with 44 meshes of LAs hand segmented from CT-Scans. To overcome the lack of motion available for each patient, several methods exist. One solution is to use Electro-mechanical models to simulate motion. Although recent work has established good results on coupling sinusal rhythm (SR) electrical signal with mechanical models (e.g [14]), the high complexity of the shape combined with the thinness of the LA wall makes it a hard problem to solve. Instead of having a full electro-mechanical model, it is also possible to build a model solely for the computation of boundary conditions aimed at CFD simulations, [37] proposed to add a lumped-parameter 0D circulation model to tune the state of the rhythm from sinusal to fibrillated before building a full electro-mechanical model, although powerful, this type of approach requires calibration of each component relying on assumptions that would require better medical knowledge (especially from more imaging data) to validate. In order to make the least amount of assumptions, we will rely on two surrogate models of SR and fibrillation rhythm (FR) movement acquired and segmented from Cine MRI. To personalise the temporal models to each shape in our cohort, we rely on parallel transport in the context of the LDDMM framework.

---

[2] https://www.simula.no/research/research-departments/computational-physiology

Parallel transport Keeping the notation of the previous section, we set out to leverage the Riemannian structure of the LDDMM framework to parallel transport a diffeomorphism along another diffeomorphism. In other words, consider three shapes S, $\bar{S}(=\phi.S)$ and $S'(=\phi'.S)$, and define c and μ as the parametrisation of ϕ. Then the parallel transport of μ, c through ϕ' would generate the translated map ϕ_T, yielding the translation of S' via ϕ as $\bar{S}' = \phi_T.S'$. Although theoretically sound, a robust numerical scheme tailored to the LDDMM framework, as an adaptation of the pole ladder scheme, was only recently proposed in [67] and applied to cardiac motion in [68], which we refer the reader for a complete exposition. To apply these set of tools to the problem of generating motion of a shape from a template model, consider the set of shapes $(S_i)_{i=0,\dots,n}$ representing the evolution in time of the template model, where n is the number of time-steps. We can represent this set as a template shape S_0 and a set of deformations $(\phi_i)_{i=1,\dots,n}$ where $\phi_i.S_0 = S_i$. We refer to this as an atlas representation of the movement. In order to translate such motion to a shape S', we only need to compute the transformation ϕ' matching S_0 to S' and translate the set of maps $(\phi_i)_{i=1,\dots,n}$ through it, i.e parallel transport the set of parameters $(\mu_i, c_i)_{i=1,\dots,n}$ through ϕ'. A graphical abstract of the pipeline is shown in 3.5.

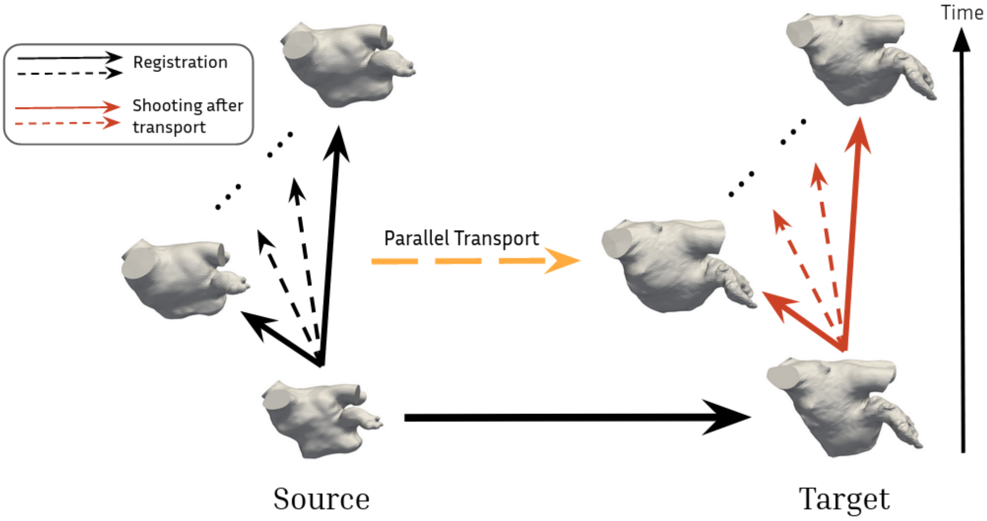

Figure 3.5: Graphical abstract of the parallel transport pipeline used to generate motion.

Implementation and results We follow the same implementation as in the previous use case to compute the atlas representation with an added step.To lighten the computational burden of computing all deformations, we subsample each shape at computation time by remeshing a coarser version with the `mmg` software [3], and apply the deformations to the fine meshes. We follow the parallel transport implementation proposed in [68] which was built on `deformetrica`. In addition we scale the parallel transport w.r.t volume change, following suggestions made in [66]. This technique showed

---

[3] http://www.mmgtools.org

a clear improvement when transporting the movement to LAs with a big volume difference compared to the surrogate shape. The resulting motions, although close to realistic, had a trembling factor, problematic for the SR models. To alleviate this problem we kept half of the meshes describing the movements (from n=88 to n=44) and resampled them by registering shapes between each time steps and keeping the mid-point of the deformation. This created a much smoother overall movement, see 3.6 for a volume curve comparison. Since LAs from the cohort were taken at different times in the heart cycle, we created an atlas for each RR time with base template at the closest timestep possible, once the computations were made we reordered every movements to start at the same time point (RR=0). Each model were validated by ensuring the volume curves are close to the surrogate model volume curves for both SR and FR movements, as well as ensuring the movement was qualitatively good. We found 5 out of the original 44 shapes to have unrealistic movements, defined by an unnatural collapse of the shape when going from diastole to systole. Finally, for the 39 kept model, we compute the deformation matrix to represent the endocardium boundary conditions in the CFD simulations. The analysis of the CFD simulations are still underway.

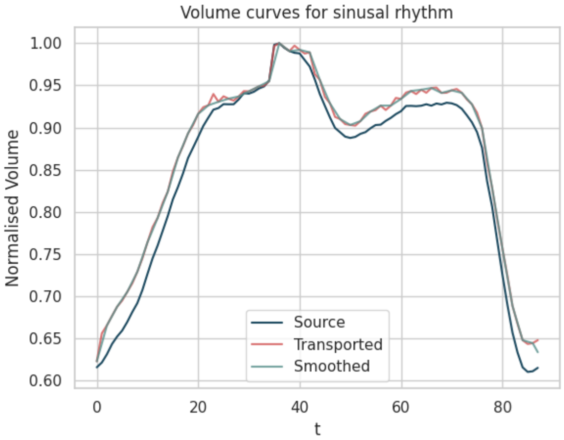

Figure 3.6: Normalised volume curve showing the surrogate motion compared with a the motion transported to another shape with and without smoothing.

3.4.8 Final remarks

In this section we showed how measuring shape variations as diffeomorphic deformations, can be utilised in out quest to understand the LA's properties. This concept, dating back to D'Arcy Thompson, and developed by a long list of researchers, is well established in the field of shape analysis. As previous work have used the LDDMM framework for shape analysis, even directly related to the LA shape [84], we may be tempted to pursue stroke prediction directly, however important setbacks does not allow to do so. The first element of those setbacks appears when constructing the atlas representation of the LA cohort: the PVs of the mean shape may not be mapped directly to the veins of other LA shapes, this is highlighted by the worst case results of table 3.2, implying that the deformations, while being diffeomorphic, are not anatomically relevant. A second indication that the

LDDMM framework is not well adapted to directly study the LA is in the estimation of the mean shape, as discussed previously, it is strongly dependant on the initial shape to compute it. This makes it impossible to have robust representation of mean shapes, as well as main mode of variation within a dataset. This problem may also be due to the fact that measured distances are small, as a cohort of similar shapes are very close to each other w.r.t the entire group of diffeomorphisms. In addition, work has shown that the LDDMM framework may not be adequate to classification tasks, e.g [159]. Although supported by biologists of the like of D'Arcy Thompson, the deformation of a shape into another, while a pleasant idea, may not encode efficiently their difference. Despite this, its efficient construction of shape interpolation can be particularly helpful for in processing shapes, offering the possibility to solve important preprocessing steps like automatic labelling and motion generation. The first application gives us an entry point into the automatic processing and study of the LA, which will be extensively used and refined throughout the manuscript. The second application fills up a hole to enable unbiased haemodynamic studies of the heart, an important perspective of this manuscript.

3.5 Measuring shape similarity

In the previous chapter we set out to compute deformations that serve as efficient proxies for plausible anatomical changes of shapes. In its theory we left an important aspect in the background: How to measure the similarity between shapes. Indeed, this information is present in equations 3.10 and 3.14, and is a crucial part of Diffeomorphometry, but also as its one entire subject which can be used alone for the analysis of shape. Recall that a shape is defined by a smooth C^∞ embedding $\varphi : N \to M$. It is only natural to consider the space of all such embeddings $\mathrm{Emb}(N, M)$ as the space of shapes. However this space is not *parametrisation invariant.* Given the point-set image $\mathrm{Im}\,\varphi(N)$, we can find many other embeddings with the same point-set image: we could first permute all of the points of N and then apply f again - the point-set image in $\mathbb{R}^3$ is still the same but with the internal points permuted. Since we want to compare the shapes, and not the maps generating them, we identify maps with the same point-set images. The permutations of points of N are precisely diffeomorphisms of N to itself - $\mathcal{D}(N)$ (again, smooth bijective maps with smooth inverses). The quotient

$$U(N, M) = \mathrm{Emb}(N, M)/\mathcal{D}(N) \tag{3.18}$$

is a smooth manifold and principal $\mathcal{D}(N)$-bundle which can be thought of as the set of a smooth submanifolds in M of type N - the point sets of shapes in M having the topology of N [132]. Given two subsets A and B of $\mathbb{R}^3$ we can define the distance between them as

$$d_H(A, B) = \max \left\{ \sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A) \right\}$$

where $\sup_{a \in A} d(a, B)$ is the largest distance from a point a to the set B (which is, itself, the smallest distance between a and points of B), and similarly for $\sup_{b \in B} d(b, A)$. This is the Hausdorff distance:

two sets are ϵ-close in space if every point in A is within an ϵ distance of B and if every point in B is within an ϵ-distance of A. If S and $\tilde{S}$ are two embedded surfaces in $\mathbb{R}^3$ then a distance between them is $d_H(S, \tilde{S})$. Unfortunately, this is not so useful as a measure of a shape since two sets can be everywhere close in space while exhibiting different behaviour on a smaller scale. For example, if S is a sphere and $\tilde{S}$ surface oscillating back and forth across S with 0 mean displacement from S and amplitude ϵ, then these two surfaces will be close in the Hausdorff distance - they have on average the same shape but a different form on a smaller scale. Capturing these finer scale differences in shape is important in many areas of science, including biology and medical imaging where important differences of interest live at finer scales. These finer differences can be captured by the difference in the tangent planes of both surfaces. Although there are many different ways of doing this, we can give the following example, while not necessarily representing the best or even most complete method, it illustrates the concept and the heart of the shape similarity problem. We start an embeddings $f : N \to \mathbb{R}^3$. The differential of an embedding maps the tangent space of N to its image inside the tangent space of $\mathbb{R}^3$

$$Df : TN \to T\mathbb{R}^3.$$

We can then parameterise any two shapes via $f(x) = (f_1(x), f_2(x), f_3(x)) \in \mathbb{R}^3$ and $g(x) = (g_1(x), g_2(x), g_3(x)) \in \mathbb{R}^3$, and say they are close if both the images f and g and their differentials Df and Dg are close:

$$\sum_i \sup_{x \in N} |f_i(x) - g_i(x)|^2 + \sum_{j,k} \sup_{x \in N} \left| \frac{\partial f_j}{\partial x_k} - \frac{\partial g_j}{\partial x_k} \right| < \epsilon$$

where the second term is the Frobenius norm of the differentials as Jacobian matrices, however we can take any variant of these norms. The problem with this is obtaining an explicit expression for a parameterisation of the entire surface, which is usually not possible, or computationally expensive. The metric d_H is a zero-th order measure of similarity between two shapes, in the sense that it can detect global similarity and proximity in space, but not similarity at a finer scale as the example above described. A distance measure describing the difference between normal vectors to surfaces is a first order measure and will detect the finer scale behaviour described above. A distance measure describing differences in the rates at which normal vectors change between two shapes is a second order measure - typical measures of this kind of behaviour in surfaces involve the principal curvatures κ_1 and κ_2, Gaussian curvatures K, or mean curvatures H of a surface. Each higher order measure will detect differences at finer and finer scales. This can all be phrased mathematically in terms of distances measuring the differences between k-th order derivatives of the embeddings defining the shapes (or measures on the Jet spaces of embeddings), and we will call a shape distance measuring this k-th order behaviour an SJ^k measure (short for shape-jet measure of order k). As it turns out, the best place to look for such shape encoding, is in *measure theory.* In the following sections we will look at a common way to encode shapes as measures before looking at a theoretical approach to find an SJ^k measure embedding of shapes.

3.5.1 Shapes encoded as measures

Through the construction of a shape space and the examples above, note that we wish to have two important properties to hold, regardless the representation: parametrisation invariance and discretisation invariance. This alone points us to measures as a convenient represesentation as it contains both properties by construction. The notion of encoding surfaces as measures dates back to [56], with a goal of building a bridge between geometry and measure theory, a domain appropriately called geometric measure theory, we refer the readers to [128] for a complete view. Recall that a measure, in its barest definition, attributes a mass to a subset $S \subset M$:

$$\mu : S \subset M \mapsto \mu(S) \in \mathbb{R} \cup \{\infty\}$$

The most direct use of a measure to represent a shape is via the dirac measure

$$\delta_x = \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{otherwise} \end{cases}$$

and a shape S in M can be directly represented in its discrete, and thus computer-friendly, setting as a weighted point cloud and associate it to the measure:

$$\sum_{i=1}^{n} \mu_i \delta_{x_i} : S \subset M \mapsto \sum_{x_i \in S} \mu_i \in \mathbb{R}$$

This representation automatically buys us the parametrisation invariance, as shapes have been turned into an unordered set and comparison of shapes in this context relates to the Monge-Kantorovich problem: How to find the optimal transportation of one distribution of mass to the other. One problem remains, both distributions must have the same total mass, consider two shapes with their own discretisation, then both will be represented by measures as

$$\alpha = \sum_{i=1}^{n} \alpha_i \delta_{x_i}, \qquad \text{and} \qquad \beta = \sum_{i=1}^{m} \beta_i \delta_{y_i}.$$

With a careless choice of weights, for instance $\alpha_i = \beta_i = 1$, the total mass would be different for both cases. Thankfully, we can alleviate this problem easily by *mollifying* the shape, that is, setting $\alpha_i = \frac{1}{n}$ and $\beta_i = \frac{1}{m}$. Effectively, solving the discretisation invariance problem and turning the spaces of shapes into the space of density measures.

Remark 3.5.1. *The currents and varifolds alluded to in the previous section, are measures as well, which can be written in this context as* $\sum_{i=1}^{n} \alpha_i \delta_{(x_i, \vec{n_i})}$, *where* $\vec{n}$ *denotes the normal. In which case the currents as viewed by [172] is a measure on* $\mathbb{R}^3 \times \mathbb{S}^2$ *and the varifolds as viewed by [32], that consider unoriented normals are measures on* $\mathbb{R}^3 \times \mathbb{S}^2_{\pm}$. *In both cases they carry the idea expressed earlier of capturing finer difference between shapes by comparing the tangent planes of surfaces rather than solely their position in space.*

The space of density measures This space has been extensively studied from the perspective of optimal transport (OT) [173], and knows a broad range of applications, from machine learning to economics, and even as a fast substitute for diffeomorphometry, we refer the reader to [58] for a full view on its current development. To keep a geometric perspective, we may formalise things a little bit more. Normalised smooth positive densities are more generally given by volume forms μ on (our space of embedding) M, that has fixed total volume of $\mu(M) = 1$, we denote the space of such measures as

$$\mathrm{Dens}(M) = \left\{ \nu \in \Lambda^n(M) \ : \ \nu > 0, \ \int_M d\nu = 1 \right\} \tag{3.19}$$

It is shown in [129] that this space is *equivalent* to the space of diffeomorphisms of M quotiented by the space of volume-preserving diffeomorphism, or volumorphisms of M, which we note $\mathcal{D}(M)/\mathcal{D}_\mu(M)$. In similar fashion as in the Diffeomorphometry section, the main lines of study for such spaces are related to finding a correct metric on it. On $\mathcal{D}(M)$ we have seen that two major choice are L^2 and Sobolev metrics. When constrained (or descended to) $\mathcal{D}(M)/\mathcal{D}_\mu(M)$, they generate the OT distances. For the L^2, we get the famous Wasserstein distance which computes the minimum energy needed to transport the mass of one density to the other, written in this context:

$$d_W(\alpha, \beta) = \inf_\eta \int_M d_M(x, \eta(x)) d\mu.$$

Where η denotes a diffeomorphism such that $\eta\alpha = \beta$ and d_M denotes the L^2 metric on M, see [16] for more on this version of OT. More recently, the (right-invariant) Sobolev metric version of OT, the $\dot{H}^1$ metric, was derived in [94], and yields a closed-form Riemannian distance on $\mathrm{Dens}(M)$:

$$d_{\dot{H}^1}(\alpha, \beta) = \sqrt{\mu(M)} \arccos \left(\frac{1}{\mu(M)} \int_M \sqrt{\frac{d\alpha}{d\mu} \cdot \frac{d\beta}{d\mu}} d\mu \right) \tag{3.20}$$

As one could have guessed from the arccos function present, it is further showed in [94] that the space of densities endowed with the $\dot{H}^1$ metric is isometric to an infinite sphere, a particular pleasing structure for both theoretical and practical developments.

Remark 3.5.2. *L^2 optimal transport was also considered as a full fledge framework for Diffeomorphometry, due to its efficient computation. However the flow generated by geodesics from the Wasserstein distance does not preserve the topology of the surface. This property cannot be ignored in deformation-based shape analysis, and no clear solution has been proposed to alleviate this problem. On the other hand $\dot{H}^1$ optimal transport was dismissed even faster for shape analysis, due to the geodesics between shapes being a simple fading of one shape to the other. However, the authors of [94] point out its potential in information geometry, as they show the $\dot{H}^1$ metric is the infinite dimensional counter part to the Fisher-Rao metric, a centre piece of information theory.*

After this quick overview of the well studied and well structured space of measure, and its efficiency to represent shapes, we may focus on the problem of finding an SJ^k measure for representing shapes, that is, a representation that accurately captures and describes the higher order differential properties of a shape.

3.5.2 Sobolev currents

To probe the world of measure theory in its full generality requires us to turn to differential forms. Although we have introduced them briefly in section 3.2.2, we did not explicit what they are used for. If we look at them in the simplest way, differential k-forms in $\mathbb{R}^3$ are a *type* of function which assigns a weighted length ($k = 1$) to an infinitesimal vector based at a point x, a weighted area ($k = 2$) to an infinitesimal parallelogram based at a point x, and a weighted volume ($k = 3$) to an infinitesimal parallelepiped based at a point x. We use differential forms to calculate weighted lengths of curves, weighted areas of surfaces, and weighted volumes of solids, and this is done through integration

$$\begin{aligned}
\int_L \alpha &:= \lim_i \sum_i \alpha_{x_i}(L'(x_i)) \\
\int_A \omega &:= \lim_i \sum_i \omega_{x_i}(T_1(x_i)), T_2(x_i)) \\
\int_V \mu &:= \lim_i \sum_i \mu_{x_i}(E_1(x_i)), E_2(x_i), E_3(x_i));
\end{aligned}$$

here L' is the unit tangent vector to the curve L and the right-hand side is a Riemann sum in which weighted lengths are assigned to the tangent vector at sample points along the curve and then added; saving a full blown theory of measure, this is precisely what integration is. Similarly for a surface A: a weighted area is assigned to the square/parallelogram spanned by an orthonormal basis of the tangent space to A at sample points x_i, and then summed. The pattern should be set, but to be clear the vectors E_i form an orthonormal basis for the tangent space to a three dimensional submanifold V in $\mathbb{R}^3$. In a sense, as integration is the fundamental operation of measure theory, differential forms are the fundamental objects of measure theory. We can now easily introduce formally what a current is. Since differential forms are objects that we integrate, we can treat the integral and it's domain (either a curve, a surface, or a solid) as something that *acts* on differential forms whose degree k matches the dimension of the domain of integration; that is, for a fixed k-dimensional submanifold N of $\mathbb{R}^3$ a k-current is a functional on the set of smooth differential k-forms $\Lambda^k\left(\mathbb{R}^3\right)$ which assigns to each k-form a real number

$$\begin{aligned}
\int_N : \Lambda^k\left(\mathbb{R}^3\right) &\to \mathbb{R} \\
\alpha &\mapsto \int_N \alpha
\end{aligned}$$

Following the efficiency of measures to encode shapes, currents have been used to represent shapes in various ways. We have already seen the version derived in [172], that was used to offer a discretisation invariant similarity measure between shapes in the context of LDDMM Diffeomorphometry. Although it may be interesting to consider them for the direct analysis of shapes. For instance [18] proposed to leverage on Finite Elements Method (FEM), for the representation of curves, which showed good comparative results in [159]. As it turns out, currents are a natural solution for describing an SJ^k embedding of a shape, for an arbitrary k and therefore yield a distance between shapes that measures their k-th order behaviour. The central idea, proposed in [17], is to consider Sobolev metrics on the space of differential forms, which naturally carries a k-th order measure, and identify what currents (i.e shapes) look like in this space. For the rest of this section, we allow ourselves to get more involved into the computations in order to follow along [17] and fully highlight the ideas behind it. We start at the beginning, and define a shape as an element from a Sobolev H^s embedding of a closed p-dimensional manifold N to the torus $\mathbb{T}^n$:

$$\mathcal{E}_N^s = \{\varphi : N \to \mathbb{T}^n \, : \, \varphi \text{ is an embedding}\}$$

Recall that we can define an L^2 metric on the space of smooth differential p-forms $\Lambda^p(\mathbb{T}^n)$. To move towards Sobolev metrics we first introduce new operations on the space of differential forms:

- The differential of forms, d, maps a p-form to a $p+1$-form, and the co-differential δ maps $p+1$-forms to p-forms, they can be related via the following equality, for $\alpha \in \Lambda^p(\mathbb{T}^n)$ and $\beta \in \Lambda^{p+1}(\mathbb{T}^n)$:
$$\int_{\mathbb{T}^n} g^{p+1}(d\alpha, \beta) d\mu = \int_{\mathbb{T}^n} g^p(\alpha, \delta\beta) d\mu$$
- The hodge star $\star$ maps p-forms to $(n-p)$-forms, and is defined by $g^p(\alpha, \beta)\mu = \alpha \wedge \star\beta$ and relates d and δ via:
$$\delta = (-1)^{n(p+1)+1} \star d \star$$
with the property that $\star\star = (-1)^{p(n-p)}$.
- With these operators we can define the Hodge-Laplacian operator as $\Delta := d\delta + \delta d$, which commutes with d, δ and $\star$.

We may now move away from the L^2 metric of equation 3.3 by defining the Sobolev metric on $\Lambda^p(\mathbb{T}^n)$:

$$(\alpha, \beta)_{H^k} = \int_{\mathbb{T}^n} g^p(A^k \alpha, \beta) d\mu$$

with:

$$A^k = \sum_{r=0}^{k} a_r \Delta^r.$$

Where we assume $k > \frac{n}{2} + 1$ to ensure all components are at least C^1 differentiable, this is what gives us the kth-order sought after. The current map in this context is defined by $\mathcal{C} : \mathcal{E}_N^s \to H^k(\Lambda^p(\mathbb{T}^n))^*$ as:

$$\mathcal{C}(\varphi)(\alpha) = \int_{\varphi(N)} \alpha.$$

It can be shown that the current map $\mathcal{C}(\varphi)$ is a bounded linear functional on $H^k(\Lambda^p(\mathbb{T}^n)$ which ensures a unique Riesz representer of $\mathcal{C}(\varphi)$ in $H^k(\Lambda^p(\mathbb{T}^n)$, denoted $\beta_{\mathcal{C}(\varphi)}$ that satisfies:

$$\mathcal{C}(\varphi)(\alpha) = (\beta_{\mathcal{C}(\varphi)}, \alpha)_{H^k}$$

To better identify what the riesz representer of a shape actually is, we can decompose it and focus on each part individually. The hodge decomposition separates $\beta_{\mathcal{C}(\varphi)}$ into 3 parts as:

$$\beta_{\mathcal{C}(\varphi)} = d\sigma + \delta\gamma + h$$

where h is a harmonic form, α is called an exact $(p-1)$-form and γ a co-exact $(p+1)$-form. Let us now focus on each component by projecting the Riesz representer to the respective subspaces. The projection onto the co-exact Sobolev space of forms is done via the projection operator $P = \delta\Delta^{-1}d$ and the Hodge start $\star$ as $P \star \beta_{\mathcal{C}(\varphi)}$. It can be shown that the projection is precisely given by:

$$P \star \beta_{\mathcal{C}(\varphi)} = \star d\sigma \tag{3.21}$$

Moreover, for any $(n-p)$-form α we have:

$$\begin{aligned}
(P \star \beta_{\mathcal{C}(\varphi)}, \alpha) = (\beta_{\mathcal{C}(\varphi)}, \star P\alpha)_{H^k} &= \int_{\varphi(N)} \star P\alpha \\
&= (-1)^{n^2-p^2+1} \int_{\varphi(N)} d \star \Delta^{-1} d\alpha \\
&= (-1)^{n^2-p^2+1} \int_{\partial\varphi(N)} \star \Delta^{-1} d\alpha \\
&= 0
\end{aligned}$$

Where we make use of the self-adjointness and commutativity of each operator, as well as the Stokes theorem, and the final result is due to the fact that $\varphi(N)$ has empty boundary. As this holds for any form α and combining it with equation 3.21, we get $d\sigma = 0$. Similarly, we can project onto the space of harmonic forms via the orthogonal projection Q and get the following result:

$$(Q\beta_{\mathcal{C}(\varphi)}, \alpha)_{H^k} = (\beta_{\mathcal{C}(\varphi)}, Q\alpha)_{H^k} = \int_{\varphi(N)} Q\alpha = \int_{M_\varphi} dQ\alpha = 0$$

Where again we make use of the Stokes theorem, and the fact that $Q\alpha$ is a harmonic form, therefore by definition $dQ\alpha = 0$. M_φ denotes the interior of the embedding φ. In other words, the Riesz representer of a shape is given by an exact form γ, which we decompose as $\gamma = \Omega \cdot \mu$, where Ω is an H^{k+1} function and μ is the Lebesgue measure. We can therefore write, for any $\varphi \in \mathcal{E}_N^s$:

$$\beta_{\mathcal{C}(\varphi)} = \delta(\Omega_\varphi \cdot \mu).$$

The sole non-zero component, γ, is in fact positive form, which can be normalised to live in the space

$\mathrm{Dens}(\mathbb{T}^n)$. In other words, an SJ^k embedding of the shape can be entirely identified by a density measure, and analysed through the prism of optimal transport, either through the Wasserstein distance or $\dot{H}^1$ distance. This identification is written in detail in the following theorem.

Theorem 3.5.1. *Denote* $\zeta_k = \sum_{\xi \in \mathbb{Z}^n/\{0\}} \frac{1}{(2\pi|\xi|)^{2k}}$. *Suppose that* $k > \frac{n}{2} + 1$*, so that* $\zeta_k < \infty$ *and let scalar coefficients* $(a_0, \cdots, a_k)$ *such that* $2a_0 < \frac{a_k}{\zeta_k}$. *Then each shape* $\mathcal{C}(\varphi)$ *is represented by a unique* H^k *probability measure:*

$$p_\varphi = \frac{1}{\int_{\mathbb{T}^n} \Omega_\varphi \cdot \mu} (\Omega_\varphi \cdot \mu) > 0, \tag{3.22}$$

obtained by the normalised solution of the PDE:

$$A^k \Omega_\varphi = \chi_{M_\varphi} \tag{3.23}$$

Where χ_{M_φ} *denotes the characteristic function of the interior of the embedding.*

This result gives us, at once: an SJ^k embedding with full control on the number of orders at which we want to compare shapes; a formal encoding of a shape into a density measures, perhaps more robust than the *mollification* method, on which we can simply use the $\dot{H}^1$ distance and its analytic formulation for fast computations; a straightforward way to build this representation, thanks to its PDE formulation of 3.23.

Remark 3.5.3. *This holds only for co-dimension one shapes. Although the authors of [17] derived a specific representation for shapes of dimension 1, other cases would measures representing shapes equal to zero.*

Remark 3.5.4. *Without following the derivations above, it might sound surprising that a shape can be mapped onto an infinite dimensional sphere. However the geometric intuition behind it can help make sense of it alone. First, the current map is a linear functional, which can be seen geometrically as an infinite dimensional vector. Second, the Riesz representer ensures us the unique representation of shapes, in other words, none of these infinite dimensional vectors are colinear. Given the correct basis, this means that they all point to a different infinite dimensional direction. We can then normalise them, and represent them as points on an infinite dimensional sphere.*

The ingredients to apply this are straightforward. χ_{M_φ} is exactly the binary mask, which in medical imaging is the natural representation of a shape when segmented from an image. All that is left is to find an efficient method to solve the linear PDE 3.23. However, this part is deceptively hard. Although the PDE is linear, the degree k is always high, which, as we highlight in the next section, renders most methods unusable.

3.5.3 The setbacks of Poly-Helmholtz PDEs

The crux of the rather elegant representation of shapes proposed above lies in the efficient resolution of a k-th order linear PDE, given by 3.23. The theorem 3.5.1 allows us to use $k = 3$ at the minimum, and set $a_0 = a_k = 1$ and all other $a_i = 0$. This gives us the simplest possible version of the PDE for shape representation as:

$$(I - \Delta^3)u = f \tag{3.24}$$

The construction being on $\mathbb{T}^n$ implies periodic boundary conditions on the domain Ω. This problem is often referred to as a tri-Helmholtz equation, or Poly-Helmholtz for an arbitrary k. As we set out to find an efficient way to solve this equation, we can start by taking advantage of the periodic domain by trying spectral methods, which are techniques that naturally support periodic boundary conditions.

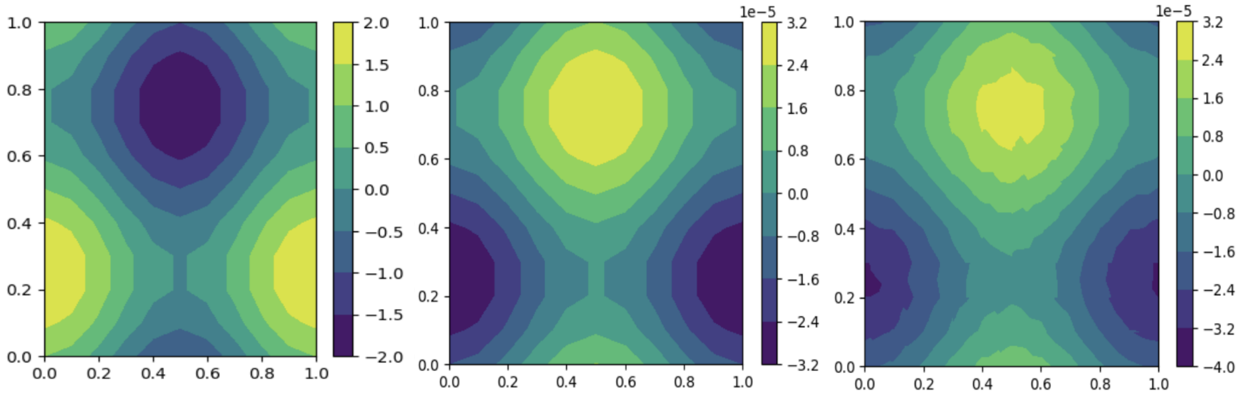

Figure 3.7: Simple solution via spectral methods. f (left), u (middle), $\hat{u}$ (right).

Spectral methods The central tool to this group of methods is to represent functions by their spectral coefficients, computed in a basis of choice, in which the solutions can be solved efficiently. The efficiency of spectral methods to solve a wide range of PDE problems has led to a large body of work studying them. As we will not go into any details regarding the theory surrounding it, we refer the readers to [27]. In this method, a function f is decomposed as

$$f(x) = \sum_n f_n^\phi \phi_n(x)$$

and a linear PDE of the form

$$Lu = f$$

is written as findind the spectral coefficients u_n^ϕ of u in

$$\sum_n \langle \psi_i | L\phi_n \rangle u_n^\phi = \sum_n \langle \psi_i | \phi_n \rangle f_n^\phi. \tag{3.25}$$

To leverage on the periodic boundary conditions, we choose ϕ as the Fourier basis. The entire experiment was done using the `Dedalus` package [30], which implements the tau method, a generalisation of 3.25. We follow their recommandations regarding all parameters to obtain a robust solver specific to problem 3.24. To make sure that the solver yields correct results, we conduct two separate experiments in 2d. First we test it on functions of the type

$$f(x,y) = \cos(ax+by) + \sin(cx+dy) \tag{3.26}$$

where a, b, c, d scalar coefficients, for which we have a closed form answer of the type

$$u(x,y) = \frac{\cos(ax+by) + \sin(cx+dy)}{1-(a^2+b^2)^3} \tag{3.27}$$

On this experiment the spectral solver yields good results, faithful to the closed-form. See for example the case where $a = b = 1$ and $c = d = 0$ in figure 3.7. We can now turn to cases where f is a 2d mask of a shape, a problem closer to our final application. We use the leaf dataset for this experiments, proposed by CITE, and preprocessed by CITE, and follow the exact same procedure as before. In this case, we realise that the solver fails to yield exact results. In fact, the solver always yields the same Gaussian distribution on the domain of the mask, regardless of the input shape, see figure 3.8. We believe this may be due to one of two problems. The most likely is the *sharp* boundary of shape when represented as a mask, this makes the f, or the right hand side (RHS) of the PDE not smooth enough to use Fourier decomposition. A second reason may simply be the bad approximation of the operator Δ^3 via Fourier differentiation. Nonetheless, we are not aware of any spectral-based methods specifically taylored to tackle one or both of those problems, and we rather propose to explore other methods altogether.

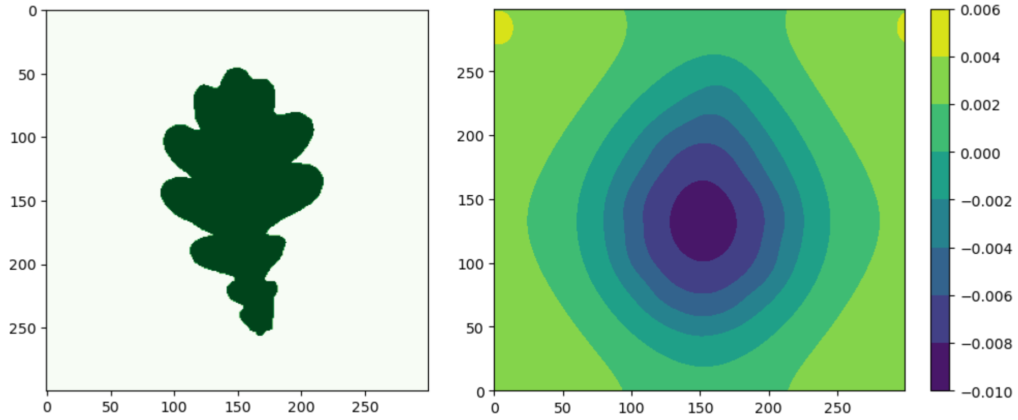

(a) A 2d leaf.

(b) The failed solution.

Figure 3.8: Spectral methods fail on shapes.

Finite elements With the failure of standard spectral methods, we turn ourselves to the most widely used techniques for solving PDEs, finite element methods. Again, we will not get into any details of how FEM works, nor give any presentation of the current state of this domain, we refer the readers to [33] for this. In concept, FEMs solve PDE problems by discretising, or meshing, the domain Ω and looking for a solution function in a finite dimensional counterpart V to the original infinite dimension function space. This finite dimensional counterpart is defined by basis functions of choice, often polynomials, and the discretised domained is built with fundamental elements of choice, often triangles. Finally, FEM concerns itself by solving a weak (or variational) formulation of the PDE. That is, instead of looking for $u \in V$ such that

$$Lu = f,$$

the problem is written in weak form, in which the goal is to find $u \in V$ such that

$$a(u, v) = \langle f, u \rangle, \quad \forall v \in V$$

Where a is a bilinear form. In our case the weak formulation can be derived using classical techniques, and 3.24 becomes, in the 2d case:

$$a(u, v) = \sum_{j=0}^{k} \binom{k}{j} \langle \frac{\partial^k v}{\partial^{k-j} x \partial^j y}, \frac{\partial^k u}{\partial^{k-j} x \partial^j y} \rangle.$$

For the method to be theoretically sound, conditions are required on both the discretised domain and the finite dimensional function space. To this end, the number of degrees of freedom of the polynomial basis function are directly linked to the number of orders of the PDE. For instance, a fourth-order PDE would require, at minimum, a quintic polynomial with 18 parameters for a triangular element [33]. This renders the use of classical methods difficult for our problem, containing at the minimum 6 orders of differentations. To alleviate this problem researchers suggested to relax the condition on the domain, by using non-conforming elements. In turn, this can accommodate for polynomials of much lower degrees, for instance, [175] show that polynomials with degrees as low as 2 can accommodate for orders as high a 6 in the 2d case. In this context, [183] proposed a non-conforming element that seem to fit the problem at hand 3.24, with guarantees for $2m$-th order PDE on $\mathbb{R}^n$ if $m = n + 1$. This element was further studied for the specific case of tri-Helmholtz equations, i.e 3.24, in [21], and used in a shape-related context in [22]. We took benefit from this study and it's attached implementation within the `Firedrake` package [70], to apply it to our problem. Similar to the spectral methods, we first tested the method on known pairs (f, u), before trying on 2d shapes. Again, the results on pairs yield good results, but attempts to find the proper representation for shapes suffered from some numerical collapse. Giving the same Gaussian like distribution over the domain, regardless of the RHS shape (see 3.9), as for spectral methods. Although probably due to the lack of smoothness of the RHS, it is unclear why this happens. Non-conforming elements approach to FEM is still a young field, and basic numerical stability might be to fault in this case. Moreover, there seems to be a growing interest in developing efficient for higher order PDE, and more and more methods have been proposed within FEMs, as well as in growing variants of FEM. For instance, non-local operator approachs [149], and virtual element methods [8, 41, 75] show good promises for poly-Helmholtz equations.

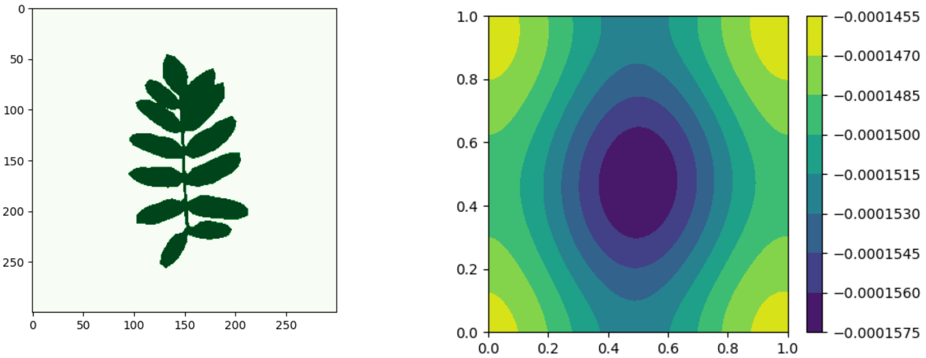

(a) A 2d leaf.

(b) The failed solution.

Figure 3.9: Non-conforming FEM on 2d shapes.

Neural Operators To make this attempt of solving poly-Helmholtz PDEs more exhaustive, we focus on a last set of tools, recently developed and aimed at taking advantage of the effectiveness of Neural Networks (NN) to approximate functions. To this end, specific NN architectures and learning procedures were designed specifically for the goal of solving PDEs. In concept, given a linear PDE of the form

$$Lu = f,$$

a NN will be trained to learn the inverse operator

$$\begin{aligned} G_\theta : \mathcal{F} &\to \mathcal{U} \\ f &\mapsto u \end{aligned}$$

where θ denotes the NN's parameter, which are commonly optimised by minimising a relative L^2 loss. Networks designed for this purpose are generally referred to as *Neural Operators* [101]. The two main neural operator architectures are the Fourier Neural Operators (FNO) [106] and the Deep Operator Networks (DeepONet) [115]. The most widely used architecture, the DeepONet, attempts to learn in an instance-wise manner. That is, the DeepONet accepts as input the values of f on the whole domain, as well as a query point y at which to evaluate u. Both inputs are given to separate networks (respectively, the *branch net* and the *trunk net*), and the outputs of each networks, both of a fixed dimension p, are multiplied component-wise and summed, to yield the scalar $G_\theta(f)(y)$. This requires to have the same discretisation throughout the dataset. In addition, although it has shown good results, is notoriously long to optimise and requires a lot of data. In the FNO case, each layer performs a Fourier decomposition of the input, applies a linear transformation on the Fourier modes and maps back to the input space via inverse Fourier transform. This general methods has the desirable property of being discretisation invariant, and naturally accommodates for periodic boundary conditions. For these reasons, we attempt a first experiment by generating a set of 1000 pairs (f, u) from equations 3.26 and 3.27, discretise into 40×40 grids, and train a model by following the authors recommendations. Results, however, are sub-par. The mean relative error does not go beneath 0.8 and qualitative results show that the FNO mostly yields a blurred version of the input (see 3.10). We are unsure as to why the FNO architecture is unable to yield better results, and did not attempt to run the experiment on 2D leaves as we are not able to get the full supervised dataset (f, u).

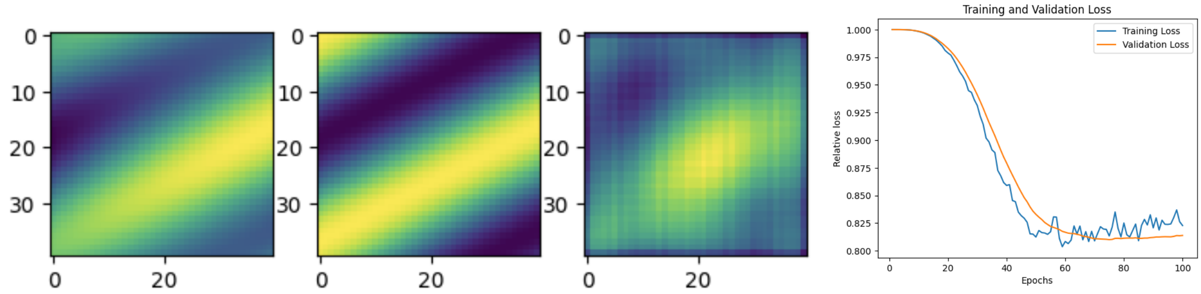

(a) An example: f (left), u (middle), and $\hat{u}$ (right).

(b) Relative absolute error loss.

Figure 3.10: FNO results.

One way to both add numerical robustness and alleviate the problem of accessible data, is to incorporate the operator from the original PDE into the network. This concept, first introduced as Physics-Informed Neural Networks (PINN) in [148], proposes to trace back the results to the RHS of the PDE and compare it with the initial f. Staying in the theme of neural operators, authors have proposed a physics-informed FNO, named PINO [107], and a physics-informed DeepONet (PIDON) [176]. For PINO, the authors proposed to perform Fourier differentiation on the output of the model, and incorporate it in the loss. That is, if we denote by RAE the relative absolute error, then the full loss to minimise if of the form

$$\mathrm{RAE}(u, G_\theta(f)) + \mathrm{RAE}(f, L_{\mathcal{F}} G_\theta(f))$$

For the Tri-Helmholtz scenario, this is not applicable, as performing Fourier differentiation 6 times accumulates to much error to have a faithful approximated operator $L_{\mathcal{F}}$. On the other-hand PIDON proposes to use automatic differentiation of the NN-approximated operator, which is known to be more faithful than pure numerical differentiation. Although the PIDON learning procedure adds numerical stability, it adds a considerable computational complexity, as we need a total of 24 back-propagations to compute the full Δ^3, in addition to the initial problems inherent to DeepONet for the Tri-Helmholtz problem, making unpractical in any *real world* setting. we still performed our basic test over the dataset created for the FNO, as to get a sense of a performance ceiling of neural operators for the case of high-order linear PDEs. As it turns out, results are rather good (see 3.11), which makes us think that, as the field grows, possibly faster and more robust architecture will arise as a potential candidate for our specific problem.

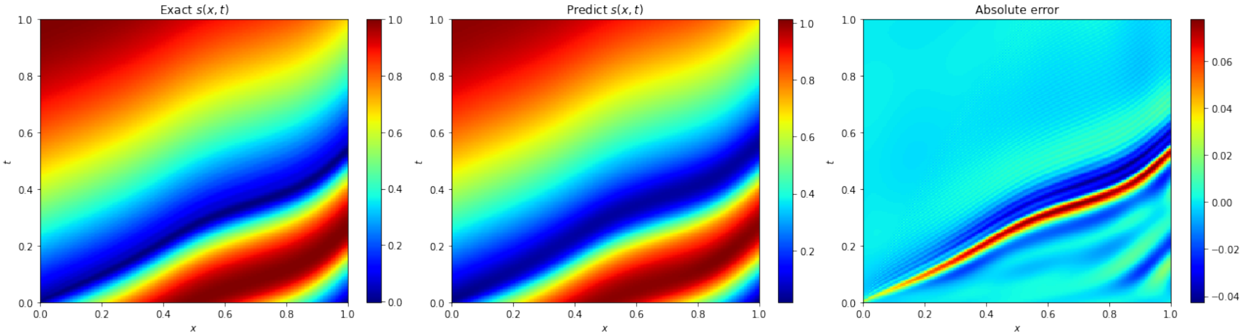

Figure 3.11: PIDON. Example of result: solution (left), prediction (middle), absolute error (right).

3.6 Conclusion

In this first chapter, we explored several ways to tackle shape analysis. Far from exhaustive, the three parts of the chapter highlight major domains focused on the measurement of shapes. As we have seen, the necessarily omnipresent geometrical tools can be applied on different levels. In the first part, we followed the main construction underlying spectral analysis methods, which focuses on efficient representation of shapes through the spectrum of a given operator defined on the shape. We then saw how geometry can be used not only on shapes, but on the space of diffeomorphisms acting on shapes, the center ideas from Diffeomorphometry. Although more theoretically involved, we stayed as close as possible to a geometric interpretation of these infinite dimensional spaces, and yet how they can be directly applied to studying shapes, especially in cases of anatomical shapes. In the final part, we focused on the space of shapes directly, and shifted the interest on finding the best distance directly on it. We showed that this problem relates to finding an embedding of shapes that captures differences at multiple scales, and that the language of measure theory helps us defining it. Following this idea we showed that Sobolev currents is the natural way to capture such differences, and we followed the construction of [17], that give a natural distance on such spaces. The crux of this theory, however, renders it hard to apply in practice, as it relates to high-order PDEs that are yet to be solved efficiently, highlighting the fact that the step from theory to application is often not straightforward. Throughout the chapter we showcased how such methods may be applied to the analysis of the LA. Ranging from direct analysis through out of the box methods such as shape spectrum classification, to multiple tools that help into providing an automatic process for many steps to come, either through fast distance computation on the LA, anatomic labeling or even motion generation. While this chapter can be distinguishable from the rest of the manuscript by its more involved theoretical aspects, a clear take-away message can be drawn out. Although advanced mathematical tools have been created for analysing and measuring shapes, they may not be appropriate as direct method for the statistical shape analysis of the LA. This is due to the lack of expressiveness with regards to the parts of the LA that are correlated with stroke, as well as the difficulty to have clear interpretation of such methods. Indeed, at this stage we do not know whether the LA has any impact whatsoever on stroke incidence, yet the previous attempt only shows an absence of positive answer, which is not the same as a negative answer. On the other hand, advanced and general methods stemming from this type of theory can help us automatically preprocess shapes, which in turn enables us to build tools down the line specific to the LA that are both expressive and interpretable. Which is precisely what we set ourself to do throughout the rest of this manuscript.

Chapter 4

A first attempt at classification

Contents

Abstract. We set out to design automatic and interpretable tools to study the LA. In this chapter, rather than directly using general shape representations as presented in the previous chapter, we propose to represent the left atrium as a geometric graph: a compact representation robust to noise and easy to study with standard tools from statistics and machine learning. Moreover, we identify a problem when studying the left atrium, and by extension the heart, linked to the variability of image acquisition times in the heart cycle among patients, leading to a high heterogeneity within the cohort. To accommodate this, we propose MCVC, a novel multi-channel variational neural network architecture.

4.1 Introduction

At the risk of being redundant with the introductory chapter, let us recall some general facts about this study. Atrial fibrillation (AF) is a complex cardiac disease that dramatically increases an individual's risk of developing a cardioembolic stroke. While preventive therapy exists, stroke prediction tools are limited, and image based markers are not well understood. This lack of knowledge is due on one hand, to the complexity of analysing and comparing the shape of the Left Atrium (LA); and on the other hand, to the unavailability of relevant data, which limits the complexity of the models we can use. The data availability problem is particularly present in this chapter, as at this stage only a small fraction of the full cohort of patients later processed is available. In addition, conducting data analysis in cardiology exacerbates the small dataset problem as the heart cycle renders impossible to compare images taken at systole with those taken at diastole time. To palliate this issue, we set out to design tools that enables such studies. This chapter constitutes a first attempt at addressing this problem and sets the direction for the rest of this thesis. In this context, we first propose a graph representation of the LA, to focus on the impact of pulmonary veins (PV) and LA Appendage (LAA) positions, resulting in a simple, easy to analyse object. Secondly, we propose the Multi-Channel Variational Classifier (MCVC) for handling and classifying heterogeneous datasets based on the consistent representation of each dataset in a common latent space. We show that such a model is analogous to a meta-classifier, where each dataset is characterised by specific projection in a common latent space, while sharing the same separating boundary. Our final goal is to predict stroke risk from the anatomy of the LA. Thus, we apply this model to the graph representation of the LA and interpret the outputs of the model to give novel time-dependant bio-markers related to PV and LAA configurations for the prediction of stroke.

4.2 Graph representation of the LA

As we advance from general shape representation, we aim to find a representation tailored to the LA. What we are looking for is a compact representation on which we can apply traditional machine learning algorithms. As the final goal of this analysis is to explore underlying links between stroke and PV and LAA positions and orientations in the LA, we introduce the concept of a lightweight graph representation, containing, in theory, all the information needed. To do so automatically, we first build upon the ideas from section 3.4.6 for the automatic labelling of anatomical parts, before computing the graph with a simple set of rules. Let us first recall the information about our dataset and the pre-processing pipeline

4.2.1 Dataset and pre-processing pipeline

Our study was performed on 3D Computational Tomography scans (CT-scans) and the corresponding clinical data from 107 patients suffering atrial fibrillation, of which 64 are labelled stroke positive – a composite criterion composed of the information corresponding to the detection of LAA thrombus on the CT-scan and/or past history embolism. In particular, our database is composed of 50 patients with CT-scans are acquired during systole (of which 27 are stroke positive) and 57 patients with scans acquired in diastole (of which 37 are stroke positive). Cardiac segmentation was conducted automatically with a 3D U-Net Neural Networks (NN) as proposed by [83], and then hand-corrected by experts. All the data was acquired by the Bordeaux University Hospital. From the available segmentation masks we use the open-source package `MMG`[1] for meshing the shapes. First we apply a marching cube algorithm, giving a rough meshing of the surface, the mmg3d algorithm is then applied with specific parameters to keep the number of triangles under 2500.

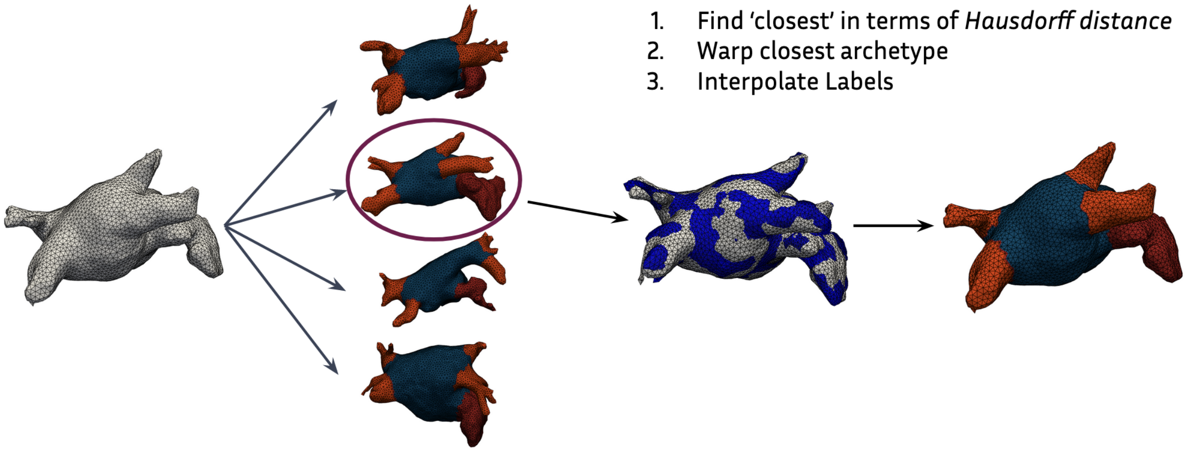

Figure 4.1: Overview of the automatic segmentation using multiple archetypes.

4.2.2 Diffeomorphic labelling

To study the impact of PVs and LAA geometry in clot formation, we propose to represent the LA as a graph similar to its centre-line. To do so, we first need to identify the different anatomical parts of the LA: we use the LDDMM framework method for LA labelling, as introduced in section 3.4.6, where we automatically label the PVs and LAA of every mesh using an *atlas* representation. This representation allows to fully represent a population of size n by an average shape T and a set of deformations $\{\varphi_i\}_{i\leqslant n}$. After hand-labelling the mean shape T, we warp the labels through the deformations in order to label every patient. This general method, although straightforward, can lead to significant errors. For instance, consider two similar blob-like shapes, each with a large bump in a different place. In order to match the two shapes, the shortest path computed via the LDDMM framework would eliminate the bump of the first shape and grow a new one, rather than stretch the surrounding of the bump to match the second. This phenomenon can happen with PVs in the LA context, leading to incorrect segmentations. Although this whole procedure is designed to minimise the number of man hours required to segment the anatomical parts of the LA, it is still necessary to make sure the obtained labels are accurate, while spending a minimal amount of time on their correction.

As a solution, we simply propose to use multiple hand-labelled atlases rather than a single one. We start with 4 *archetypal* shapes selected among the entire dataset, that represent generally different classes of LAs, for example: common trunk Left PV (LPV), small LAA, big LAA, etc.

---

[1] https://www.mmgtools.org/

Similarly to the initial method, each target shape to label is then rigidly transformed to fit each archetype. We then find the closest archetype to the target shape by computing the Hausdorff distance between them, and use it for registration and label interpolation. Figure 4.1 depicts a visualisation representation of this pipeline. The pipeline outside of this change stays the same. We use `trimesh` implementation of the ICP algorithm for rigid registration and the `deformetrica` software [25] for the diffeomorphic registration, with both the deformation and current kernels widths set to 10mm. Finally the interpolation between the warped templates and the input shape is done with `pyvista`.

Automatic post-processing We realise some errors remain at the final interpolation step. Precisely, when the LPVs are in the common trunk configuration, the labels of both the Left Superior PV (LSPV) and Left Inferior PV (LIPV) may be glued together. Thankfully, we can detect and clean them automatically. Indeed, we detect this issue by checking that all parts containing the same label have as sole neighbour the LA body. When this is not the case, we reassign the boundary of both glued labels to be part of the body. The second problem we encountered was due to the closeness of the LSPV and the LAA. The interpolation may then label some parts of the LAA as LSPV. Again, we can automatically check that every anatomical part – every sub-shape containing one label– is a single connected component. When not the case, we considered the smallest(s) connected component(s) to be an error, and reassigned it's value to its neighbouring label. A visualisation of both errors can be seen in figure 4.2.

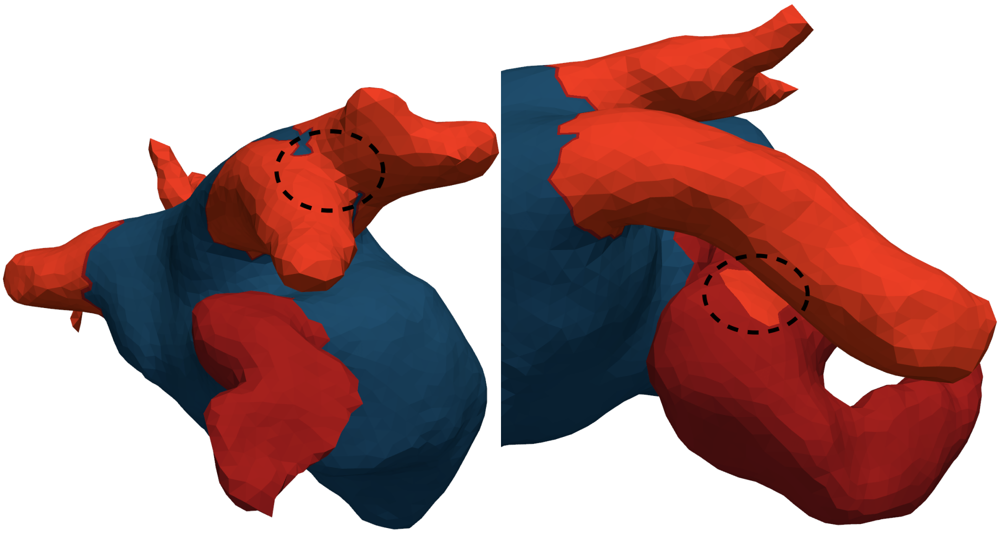

(a) Common LPVs are glued together. (b) Labels get interpolated from LSPV to LAA.

Figure 4.2: Illustration of errors occuring in the interpolation step of the labeling pipeline.

Results Following this multi-atlas pipeline, we get better overall results on the entire dataset than when using the initial method, as observed in table 4.1. In fact, the added steps in the pipeline also make it faster. Although we perform more rigid registration steps (one for each archetype), the diffeomorphic transformation is the computational crux of the method, and thanks to the archetypes, it takes nearly half the amount of iterations to reach an acceptable surface matching error in comparison to the initial method (from ~ 150 steps to ~ 80 steps). In turn, the full process of getting clean hand-corrected labels by experts is also considerably sped up due to the automatic labels being closer to ground truth for each shape. In the future, we add to the archetypes 16 randomly selected labelled shapes, to increase the performance further without a big increase in computational cost.

Method		Accuracy	f1 score	Precision	Recall
Initial	Mean	0.926 ±0.058	0.925 ±0.058	0.956 ±0.057	0.924 ±0.058
	Min	0.557	0.558	0.559	0.557
Multi-Atlas	Mean	0.969 ±0.035	0.969 ±0.038	0.972 ±0.022	0.969 ±0.035
	Min	0.848	0.833	0.901	0.848

Table 4.1: Comparison of the labelling pipelines performances.

4.2.3 Graph extraction

Graphs as shape representers have been studied in different settings, as a tool to apply deep learning methods to meshes [29], or as an augmented skeleton representation of shapes for general shape analysis [143]. We pursue the simplest form of geometric graphs: a fixed number of nodes labelled with coordinates and connected to each other. To achieve this graph representation, we extract the centre of mass of each of the 6 labels (i.e. 4 PVs, LAA and body of the LA) as well as the centre of each junction between labels, representing the 5 ostia. These points define the graph nodes. Each label centre node is connected to the corresponding ostium centre node, and the resulting branches are connected to the LA body centre. Because the LAA is a more complex shape than veins, we extend it with 2 nodes, resulting in 13 nodes in total. This is done by cutting the LAA with the plane defined by the LAA centre node and the inverse direction of the branch it is attached to; and computing the barycenter of the cut LAA. We iterate this process twice to get the 2 new nodes. This simple construction illustrated in figure 4.3, is in theory the most concise option that incorporates vein and LAA orientation and configuration. As we will handle the set of coordinates attached at each node, we need to ensure they are all in a unified coordinate reference. To do so, we set the body centre as the origin, and fix the x-axis as the direction from the centre to the LAA ostium. The LIPV Ostium (LIPVO) is then projected on the plane defined by the new x-axis to define the y-axis. Although we can now use any method from machine learning or statistics to compare sets of 13 3D coordinates, we propose a way to classify datasets with invariance to a prescribed heterogeneity.

4.3 Multi-channel variational classifier

Now that we have drastically simplified the shapes representing the LA, we could pursue standard statistical analysis methods. However, we are faced with an important hurdle due to the heterogeneity

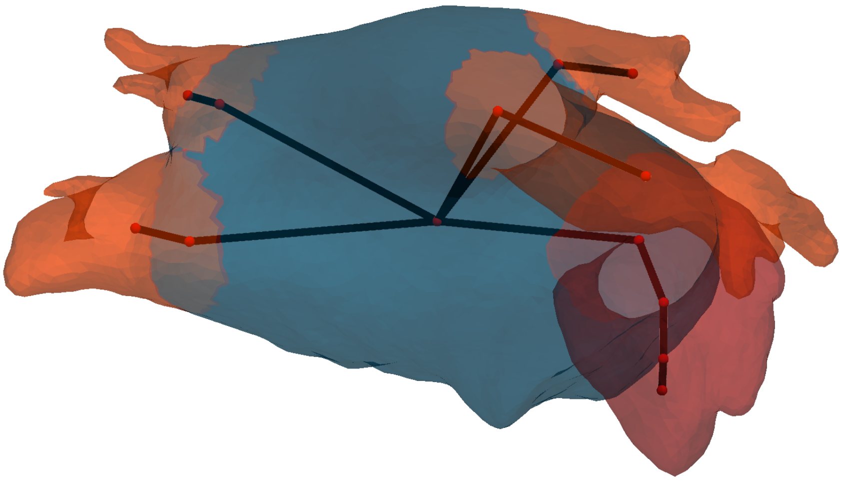

Figure 4.3: Example of an LA graph.

of the dataset, making it difficult to use such methods. In cardiology, images are taken at different times in the heart cycle, using different imaging systems, and for different tasks. Developing analysis methods allowing to integrate information across heterogeneous datasets is therefore a necessity for enabling statistical studies. In this section we follow this idea.

Our approach is shown, in concept, in figure 4.4. Inspired by [7], we propose a multi-channel variational formulation to merge multiple heterogeneous datasets into a common latent representation. Instead of following the classical learning procedure consisting in training one model per dataset, we train a single model that learns to represent multiple datasets in a shared latent space, and then uses the latent features of each subset for classification. The framework we propose realises a combination of multi-task learning [153] and meta Learning [74]: the common representation across datasets imposed by the shared latent space induces homogeneity in the latent projections of each subset. In addition to having a unified representation, this also enriches the amount of data provided to train a classifier for automated image-based diagnosis, and minimise the numbers of parameters to train, a valuable bonus when data is scarce. In addition, the learned latent distribution of the joint representation is a meta parameter that constitutes an excellent prior for future datasets as it is robust to multiple datasets sharing a common task. Finally, although we propose this framework as a solution to learning with heterogeneous datasets, it entirely accommodates for learning with multiple modalities.

We propose to use such a model to explore underlying links between Pulmonary Veins (PV), Appendage (LAA) positions and orientations in the Left Atrium and stroke. In the following sections we formulate the classification methodology within a supervised framework where the common representation is improved in terms of Kullback-Leibler divergence. We then apply the model to the joint analysis of LA Graphs where the data is split in systole and diastole subsets.

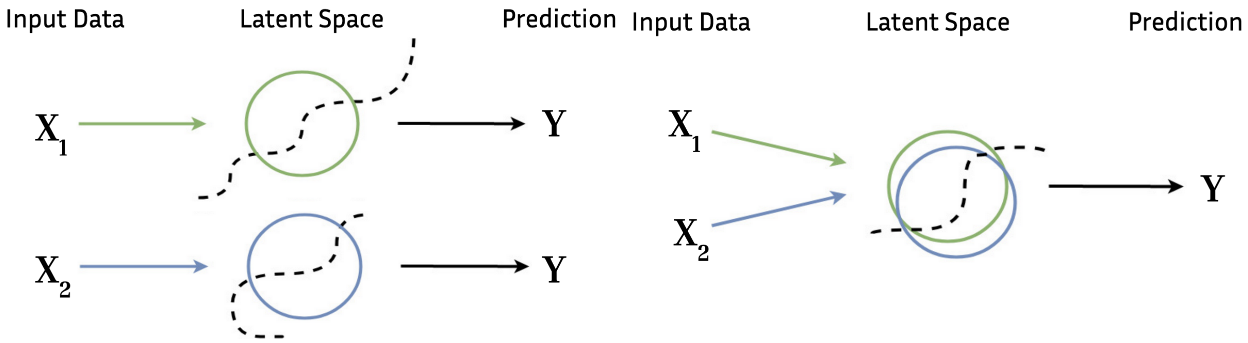

(a) Classical learning procedure. (b) MCVC learning procedure.

Figure 4.4: Visualisation of the training procedure in the classical regime and the MCVC regime.

4.3.1 Kullback-Leibler formulation

Remark: Until the end of this section, we will denote as label the ground truth of a prediction task, rather than the anatomic labels previously mentioned.

In order to learn a unified latent representation of multiple heterogeneous datasets, we build on Variational Autoencoders (VAE) [96] and assume the existence of a single latent variable z generating the multiple datasets. Let's denote $D = \{D_k\}_{k=1}^N$, the collection of N datasets with respective dimensions d_k, and $(x, y) = \{x_k, y_k\}_{k=1}^N$ a set of pairs of observations x_k and labels y_k from every dataset. Let $z \in \mathbb{R}^d$ be the latent variable shared by all elements of (x, y), where d is the dimension of the latent space such that $d \ll \inf \{d_k | 1 \leqslant k \leqslant N\}$. Our aim is to estimate the label distribution across all datasets, that is $p(y|x)$, using a representation z shared by all datasets D_k, such that $z \sim p_k(z|x_k, y_k)$ for any k.

To alleviate the intractability of $p_k(z|x_k, y_k)$, we use variational inference and introduce a set variational approximations $q_{\phi_k}(z|x_k, y_k) \in \mathcal{Q}$, where $\phi = \{\phi_k\}_{k=1}^N$ are the inference parameters for each dataset. In turn, the common representation can be achieved by forcing the approximations $q_{\phi_k}(z|x_k, y_k)$, for each k to be as close as possible to the common posterior $p_\theta(z|D)$ on average. By using the Kullback-Leibler divergence, this problem translates to:

$$\underset{q \in \mathcal{Q}}{\operatorname{argmin}} \, \mathbb{E}_N [\mathcal{D}_{KL}(q_{\phi_k}(z|x_k, y_k) || p_\theta(z|D))] \tag{4.1}$$

Where θ denotes the common classification parameters. Because $p_\theta(z|D)$ is intractable, we cannot directly solve this optimisation problem. Instead, we can derive a lower bound to (4.1) by expanding the Kullback-Leibler divergence:

$$\mathcal{D}_{KL}[q_{\phi_k}(z|x_k, y_k) || p_\theta(z|D))] = \int_{\mathbb{R}^d} q_{\phi_k}(z|x_k, y_k)[\ln q_{\phi_k}(z|x_k, y_k) - \ln p_\theta(z|D)] dz \tag{4.2}$$

Using Bayes' theorem, we can now rewrite the divergence to:

$$\begin{aligned}\mathcal{D}_{KL}[q_{\phi_k}(z|x_k,y_k)||p_\theta(z|D))] =&\mathcal{D}_{KL}[q_{\phi_k}(z|x_k,y_k)||p_\theta(z|x))] \\ &- \mathbb{E}_{z\sim q_{\phi_k}}[\ln p_\theta(y|z,x)] + \ln p_\theta(y|x)\end{aligned} \tag{4.3}$$

Rearranging we get

$$\begin{aligned}\ln p_\theta(y|x) - \mathcal{D}_{KL}[q_{\phi_k}(z|x_k,y_k)||p_\theta(z|D))] =&\mathbb{E}_{z\sim q_{\phi_k}(z|x_k,y_k)}[\ln p_\theta(y|z,x)] \\ &- \mathcal{D}_{KL}[q_{\phi_k}(z|x_k,y_k)||p_\theta(z|x))]\end{aligned} \tag{4.4}$$

Finally, we want to impose this constraint over all datasets, by supposing that every dataset is conditionally independent, we have the following evidence lower bound:

$$\begin{aligned}\mathcal{L}(x,y,\theta,\phi) =&\frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_{z\sim q_{\phi_k}(z|x_k,y_k)}\left[\textstyle\sum_{k=1}^{N}\ln p_\theta(y_k|z,x_k)\right] \\ &- \mathcal{D}_{KL}[q_{\phi_k}(z|x_k,y_k)||p_\theta(z|x))]\end{aligned} \tag{4.5}$$

Maximising the lower bound $\mathcal{L}$ is therefore equivalent to optimising the initial problem (4.1). The distribution $p_\theta(y_k|z,x_k)$ of shared parameters θ is learned from the latent space and acts as a classifier. In addition the learned distribution p_θ is a meta-parameter that contains information from every datasets in D.

In practice, $p_\theta(y_k|z,x_k)$ and each $q_{\phi_k}(z|x_k,y_k)$ are approximated through NNs. Precisely, each $q_{\phi_k}(z|x_k,y_k)$ is implemented by a separate encoder, and a single network acting as a decoder implements $p_\theta(y_k|z,x_k)$. We assume that $\mathcal{Q}$ is the Gaussian family, parameters θ and ϕ are initialised randomly, the optimisation is done by stochastic gradient descent, using an adaptive learning rate with an Adam optimiser and back-propagation. Unlike variational auto-encoders, the reconstruction objective in (4.5) is over the labels y, which is what transforms the traditional decoder into a classifier. Moreover having more than one encoder impacts the reconstruction loss which becomes a cross reconstruction of the labels from every dataset. This constraint forces the encoders to identify a common latent representation across all datasets.

4.3.2 Synthetic experiments

The aim of our synthetic experiments[2] is to highlight the possibilities offered by our method. For the sake of interpretability, we chose to use a relatively simple parameterisation for our model, consisting of a NN with three fully connected layers as our encoder and 2 fully connected layers for the classifier. All activation functions are ReLUs, and a Softmax function is used for classification. The synthetic

---

[2] The code is available at `https://github.com/Inria-Asclepios/mcvc`

data was generated with the `make-classification` function from the `scikit-learn`[3] package, by generating clusters of points for multi-label classification. We then generate similar datasets by applying transformation and adding noise to the initial problem. Thanks to this method, we can generate a high variety of independent datasets where the features have different distributions, but sharing the same target space.

We generate a collection of 5 independent datasets of dimensions varying between 10 and 20, sharing three labels. The latent dimension is set to 4 and we train our model jointly across datasets, as well as on each dataset separately. On 100 experiments, results show a generally increased overall accuracy score on test data when the datasets are trained jointly. With a 90% confidence level, the accuracy for training with MCVC lies in $[0.83, 0.99]$ with a median of 0.92, while the accuracy for separate training lies in $[0.73, 0.94]$ with a median of 0.86. About the same amount of epochs are required for convergence. Moreover, we can see in figure 4.5 that training multiple datasets jointly leads to a coherent latent space shared between labels (figure 4.5b), while training datasets independently and trying to represent the latent spaces of each dataset in a common space leads to poor consistency of the latent representation across datasets (figure 4.5a).

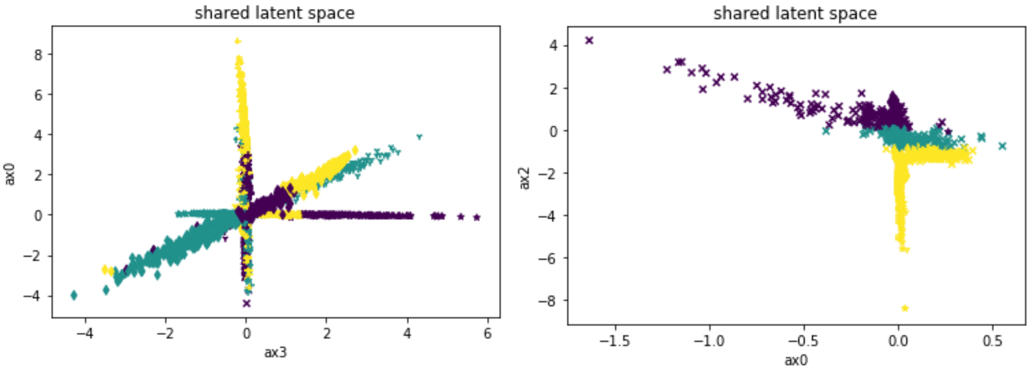

(a) 5 independent latent representations projected in a common space.

(b) Shared latent representation of the 5 datasets obtained with our method.

Figure 4.5: Latent space after convergence on 5 datasets with 3 labels. Our model provides a unified representation (fig. 4.5b), not achievable when datasets are trained independently (fig. 4.5a).

4.3.3 Noisy labels

In the shared latent space between the collection of datasets D, some subspace is allocated for labels that may not belong to a given dataset D_k. This brings robustness to the classification when there is an uncertainty on the ground truth labels. In fact, if we assign a wrong label to a given dataset D_k the model is able of re-assigning the observation to one of the existing one, which is enforced by the constraint of obtaining a coherent latent space. To illustrate this point we created four datasets

---

[3] https://scikit-learn.org/

sharing three common possible labels; then, for one dataset we modified the third label into a fourth unseen one. In figure 4.6 we show the evolution of the joint latent space throughout training. As the test accuracy drops from 0.88 to 0.75 while the fourth label completely disappears in the predictions, these observations illustrate the preference of the model for obtaining a coherent latent space over aiming for a high-accuracy.

(a) After 10 epochs. (b) After epoch 50. (c) After convergence.

Figure 4.6: Evolution of joint latent space with noisy labels. The wrong label (red) disappears completely from predictions to become a coherent label (green).

4.4 Application to the LA

Our clinical database is composed of 2 subsets : 50 patients in systole (of which 27 are thrombus, i.e. stroke, positive) and 57 patients in diastole (of which 37 are thrombus positive). Thus we use two encoders to which we feed the nodes of the graphs previously obtained in order to enable a joint analysis of the dataset. The labels are 1 if the patient is stroke positive and 0 otherwise. As such, the label classes are balanced for both datasets. We train the MCVC model with the same architecture and hyper-parameters as for synthetic data. On a 10-fold cross-validation (CV), the model yields an average test accuracy of 0.72, for the diastolic set 0.74, and for the systolic set 0.7. In contrast, attempts to classify the subsets independently suffer from mode collapse; even a careful hyper-parameters fine-tuning results in very poor accuracy scores, reaching at best 0.63 in both cases. This highlights the robustness of the model and the clear advantage of joint analysis. Figure 4.7 shows the shared latent space with the systole subset as ($\times$) and diastole subset as ($\bullet$). We see a clear separation of classes as well as a common representation of the subsets: systole and diastole are well grouped together for each class. When we attempt to classify the complete dataset using a single network, i.e. without any splitting of the data and disregarding the considerable changes in shape during the cardiac cycle, we obtain slightly worse accuracy results of 0.74. In addition, such a model is less interpretable as important features for a given class can be contradictory in between subsets.

As an additional baseline we performed PCA followed by logistic regression with cross validation and grid search on the number of principal components. We observed a much lower accuracy both on the complete dataset or individual subsets (0.65 at diastole, 0.71 at systole and 0.66 on the complete dataset).

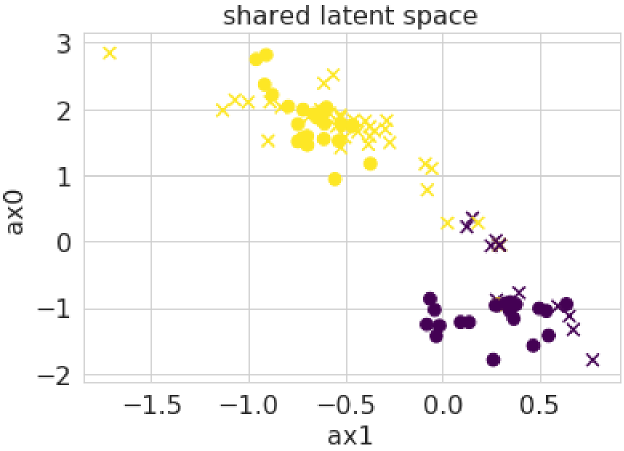

Figure 4.7: Illustration of the shared latent space learned by the MCVC on our dataset. The systole subset (×) and diastole subset (•) are homogeneously represented, while the classes (colors) are well separated.

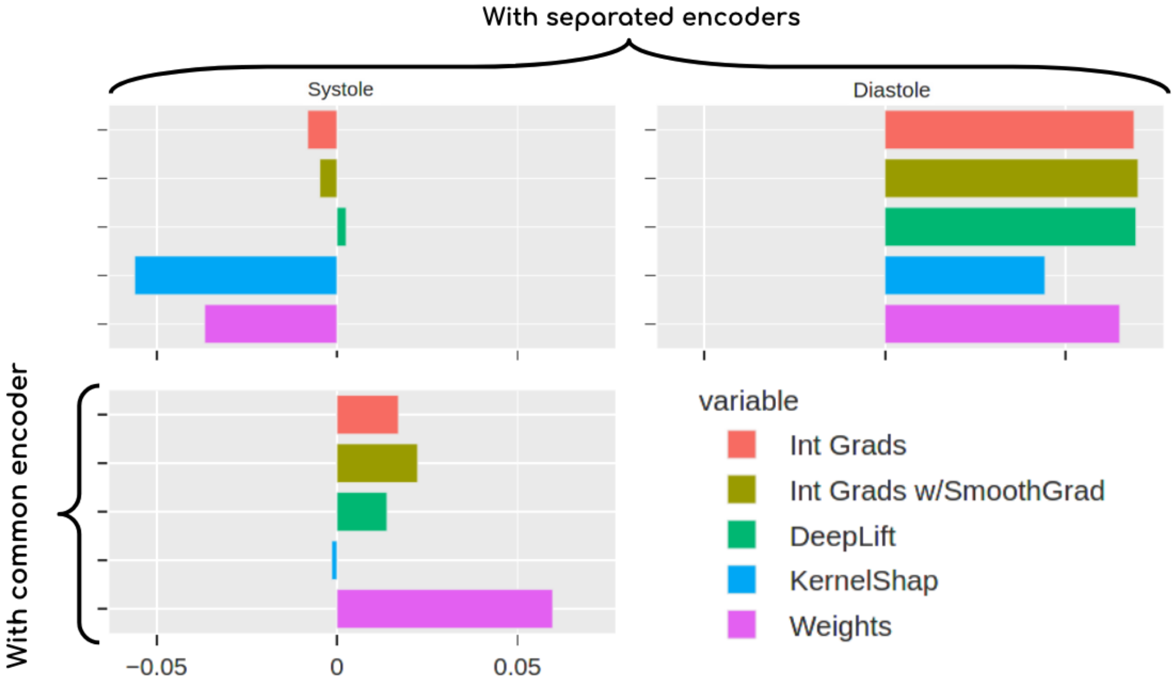

Figure 4.8: Feature importance scores of the RIPV center's x-axis when training jointly against using the MCVC.

Results interpretation. In order to evaluate the interpretability of our model, and investigate the expressiveness of the LA representation we have created, we compare results from 3 interpretation algorithms implemented in the `captum`[4] python package: Integrated Gradient [170], DeepLIFT [165],

---

[4] https://captum.ai/

and KernelSHAP [116]. For each fold of the CV, we run these algorithms on samples that are classified with more than a 95% certainty. We then look at the mean value of feature attribution scores. We illustrate an example of analysis in figure 4.8, which shows the importance scores on the x value of the RIPV centre point when the dataset is split (systole and diastole) and when it is trained commonly (i.e not split). In particular, it highlights the benefits of splitting the datasets, and therefore of the MCVC, in ensuring clinical coherence. At the contrary, we see that the model without splitting the data disregards this feature, when it is in fact of significant importance for the diastole subset. We perform this analysis on all features in an attempt to highlight possible stroke bio-markers.

In this spirit, we also perform a visual analysis of the feature importance, as shown in figure 4.9. We represent in black the atlases of the population at systole and diastole; in blue, the important feature values for control patients; and in red, the ones for stroke positive patients. At diastole, the model seems to focus on the left PVs, their ostia being closer to and rotated towards the interior – compared to the rest of the frame in stroke positive cases. In addition, for those cases, all veins intersections with the LA body are more horizontal. This *fold* could impact the haemodynamic environment of the LA, which plays an important part in clot formations. At systole, the left PV and the angulation of the appendage are the focus. For thrombus positive cases, PVs are rotated towards the interior – with the LIPV being on top and aligned with the LAA, and the LAA ostium tends to be closer to the centre and the LAA lower. Finally we see again the PVs tendency to *folded*. On top of bringing interpretability to the MCVC results, this analysis also highlights the importance of analysing systole and diastole images separately.

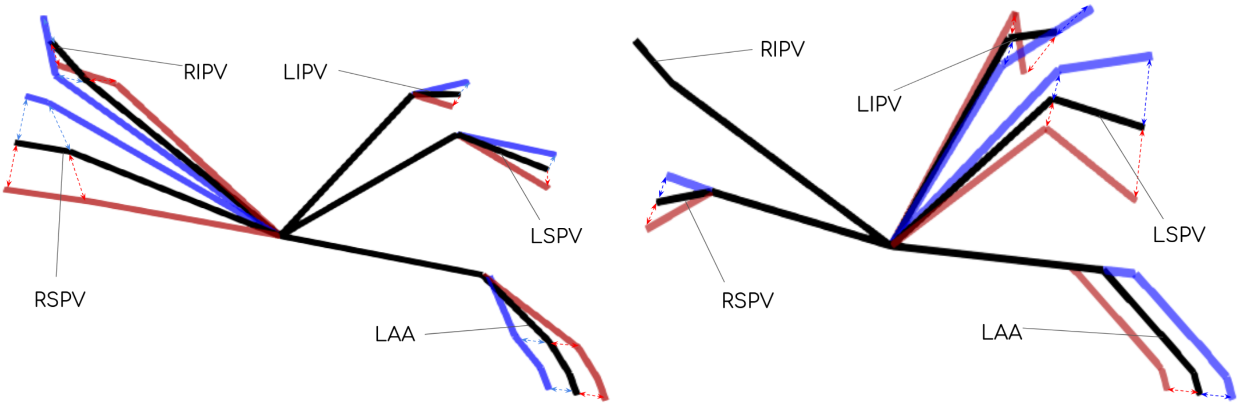

(a) Diastole (b) Systole

Figure 4.9: Visualisation of important features for predicting presence (red) and absence (blue) of thrombosis, both at diastole and systole.

4.5 Conclusion

In this work, we have provided a graph representation of the LA to analyse possible image-based bio-markers. In order to enable joint analysis of systole and diastole graphs, we presented a new method at the crossroad between multi-task learning and meta learning to tackle the joint analysis of multiple heterogeneous datasets. By leveraging on the idea that the whole is better than its parts, we proposed a classification scheme with good interpretation properties of the latent space highlighted in the study of LA graphs. We believe that the coherent latent space inherited from our model makes it possible to have deep NN as encoders while conserving the interpretability of simpler models. We aim at further exploiting this property by applying the method to joint analysis of datasets containing much more heterogeneity. Finally, we believe the lightweight graph representation can be added in a more complete and multi-disciplinary study of the LA.

Drawbacks. Although we have put our best effort in this chapter to sell both the graph representation and the MCVC model, we can now see in the future of this thesis, where more data is available, and better techniques are developed. As so, it is important to highlight some drawbacks of the current pipeline. First, the labelling procedure takes no less than 30 minutes per shape – this is problematic if or when the methods developed in this thesis are to be applied to cohorts of thousands of patients. Second, the graph representation of the LA, although compact, may be too simplistic. In addition, this representation is still not expressive enough, a set of 13 3d coordinates is difficult to clearly interpret. Finally, we realised that with more data, the MCVC model looses it's edge to simpler classification method on the separate sub-groups, dropping the need to unify the entire dataset. Moreover, the systole-diastole split turns out not be the most relevant one. In fact, the imaging procedure is now systematically done at systolic times: the diastole timed images of our dataset are rather old ones, and in the future, almost no such examples will be seen.

Future directions. While the drawbacks stated above seem severe, they enable us to better identify the future directions to take towards an end-to-end pipeline for the systematic analysis of the LA shape. Moreover, we consider that the methods proposed in this chapter are valid preliminary tools, and we believe it is worth building on them and fixing the problems we identified rather than starting over. Let us list, in concept, what can be done. First of all, the computational cost of labelling shapes needs to be addressed. Although we started this study with raw CT images, thanks to the application of the current pipeline, we are now in possession of clean surface meshes with ground truth anatomic labels. This allows us to turn to learning-based supervised methods that could offer a substantial speed-up of the anatomical labelling. In parallel, a number of improvement have been recently proposed to apply deep learning on surfaces, making it even more relevant for us to pursue this idea. In fact, we dedicate the chapter 5 to this idea.

The second major problem is linked to the expressively of the representation. In the same spirit as the graph method, we realise we could extract much more information from the anatomic labels in the form of geometric features. Quantities that contain information about the shapes in a more straightforward fashion, that can be better interpreted, and potentially directly computed by radiologists on CT-scans. This line of research will be followed in the chapter 6 of this manuscript, where we will see how it improves our understanding of the relationship between stroke and LA anatomy. In addition, these new features will help us better capture the heterogeneity of the dataset.

Indeed, the proposal to split the data into systole and diastole subgroups is difficult to substantiate. This is due to the lack of tools available to measure such heterogeneity. Thanks to the geometric features, we will see that although the idea to take in account the LA's heterogeneity is a good one, the split should be regarding the state of the AF, rather than the phase of the heart cycle.

Finally, although we realised the MCVC model is not the most relevant when sufficient amount of data is available for every subgroup, we suggest that it can still be relevant in a fully multi-modal setting, which is the next important step for preventive stroke risk – although outside of the scope of this thesis. In fact, variational models for multi-modal learning have recently known a peak in interest, with the development of mixture-of-experts [164], or product-of-experts [102] methods.

Chapter 5

Study of shapes with Neural Networks

Contents

Abstract. In this chapter we take a closer look at how neural networks adapted to surfaces can help us for the study of the LA, and more specifically for the semantic segmentation of its anatomical parts. While we have already presented the theoretical setting for the study of shapes, we expose this chapter in a self-contained manner, and reintroduce the field, although with a more practical eye. We show that while the design of efficient neural network architectures adapted to surfaces has gone a long way, the representation of the input shape itself is often neglected, or over-complicated. We propose to scale back and use curvature as the input. We explain this proposition through the Shape Operator and highlight its improving effect in conjunction with different architectures on segmentation as well as classification tasks. Finally, we successfully apply the best methods to the labeling problem.

5.1 Introduction

Surfaces are a natural representation for many real world objects ranging from organs to archaeological artifacts, as well as a central tool in virtual environments such as games, or Computer-Aided Design. This ubiquity has resulted in a large of work dedicated to mathematical methods developed for the efficient use of surfaces, as well as their analysis. The mathematical study of surfaces was consolidated by Gauss, who showed that the bending of a surface in space could also be measured intrinsically by the lengths of curves on the surface and the angles at which curves intersect. This intrinsic point of view was taken up by Riemann who generalised the ideas to higher dimensions, which led to the field of Riemannian geometry. The *Shape Operator* is the main tool linking the intrinsic geometry of a surface with its bending in ambient space and contains the main geometric invariants of a surface, namely principal curvature, principal directions, and mean and Gauss curvatures. Extensively studied by O'Neill in [138], the shape operator has become an efficient practical tool in surface processing. Among other things, it is used to define local tangent frames, making it the backbone of many computational frameworks, such as the computation of surface *features*, like creases [141].

The goal of *traditional* computational surface analysis is to find a representation of a surface that is expressive enough to capture details relevant for the problem, or task, at hand, while being computationally light-weight. To this end, many surface representations, with varying properties, have been proposed, often much more complex than curvature, and able to capture the surface structure at different scales.

The effectiveness of convolutional Neural Network (CNN) in image processing opened new doors to surface processing. The design of efficient convolution-like operations to adapt neural networks (NN) to surfaces, alleviated the need for complex representations, to the point where most state of the art architectures use extrinsic vertex coordinates as input, letting the NN models learn the surface structure at multiple scales. While some attempts were made to use well known representations as inputs, yielding some increase in performance, the general consensus is that the model should be able to learn it by itself. While neural networks are efficiently capturing the features at multiple scales, the use of a local representation more expressive than extrinsic coordinates, and more natural to interpret, should naturally improve its performance. The optimal choice should be somewhere between coarse extrinsic vertex coordinates, and more complex indirect representations.

We hypothesize that the optimal choice for such a local representation is curvature, as it characterizes the surface up to isometry (location and orientation in space), and it is purely local which allows the neural network to decide on more general surface features. In addition it is fast to compute, leaving very little computational overhead in any scenario. This chapter tests our hypothesis against three widely used representations of surfaces in three state of the art learning architectures. The three representations we test against are extrinsic coordinates, the Heat Kernel Signature (HKS) [169] and the SHOT descriptor [160] and the three learning architectures we test the representations in are DiffusionNet [161], DeltaNet [181] and PointNet++ [147]. We first draw a state of the art of learning on surfaces in section 5.2, we then give a quick review of curvature defined from the shape operator in section 5.3, and finally we conduct extensive experiments in section 5.4 that shows how principal curvature enhances any state of the art model in different tasks.

5.2 Related work

The first step for surface processing is usually to discretise it as a mesh or point cloud, particularly useful for visualisation or rendering. From these starting point, novel representations have been derived in an effort to provide tools for different surface related tasks. These tasks range from surface matching, semantic segmentation, classification, or even shape retrieval. The general trend was to find compact descriptors of a shape which could be then compared within a dataset. A long list of such descriptors exist, among them signature-based descriptors, such as Heat Kernel Signature (hks) [169] or Wave kernel signature [13], proved to be particularly efficient. Closely related are histogram methods, which are often combined with signatures to provide expressive representation such as the SHOT [160] or the Echo [126] descriptors. Geometric measure theory has also been a source of inspiration for developing efficient representation, for example geometric currents [18] leveraging on finite elements, or kernel-based currents [172] and varifolds [32] tailored for shape deformation. Such representation can be used in conjunction with classical statistical analysis tools, e.g [126, 160, **?**], although they are often building blocks for specialised methods on surfaces such as LDDMM [185], functional maps [139] or spectral-based analysis [177].

With the advent of deep learning, many methods previously stated were re-written with the help of neural networks, e.g Deep functional maps [111], ResNet-LDDMM [5]. Representation of surfaces themselves were proposed as neural networks, e.g DeepSDF [142] or DeepCurrents [140]. As Convolutions proved particular effective when learning on images, i.e structured grid, some work proposed voxel-based solutions to the study of surfaces [119], or representing surfaces as geometric images [166], on which convolutions can be applied. A second generation of *Geometric Deep Learning* has focused on building network architecture specifically tailored to work directly on surfaces, i.e meshes or point clouds. From a point cloud perspective, Point Net [146], and its extension Point Net ++ [147] consider the surface as a set of points by applying set operations on them, DGCNN [178] apply a convolution-like operation on *dynamic graphs* constructed layer-wise. MeshCNN [72] on the other hand fully leverage the mesh structure to develop operation unique to triangulations, in a similar vein GaTr [28] proposes to represent geometric data in an Algebra of choice and design an architecture with operations belonging to this algebra. An effort to have efficient generalisation of convolutions on surfaces was proposed by [118], although the lack of global coordinates creates ambiguity in local operations. To alleviate this problem, a large body of work has proposed *rotation equivariant* operations, namely GemCNN [42] augmenting Graph Neural Networks, or Field Convolutions [125]. Finally, recent models propose to bypass the problem of generalising convolutions by focusing on well defined operations on Surfaces, such as discrete exterior calculus in Hodge Net [167], heat diffusion in Diffusion Net [161] or a suit of known operators in DeltaNet [181].

As model architectures included more and more knowledge of shapes, the need for a better representation to input those models has decreased. Outside of models that contain operations specific to the structure of choice (e.g [72, 28]), most models naturally accepts as input the coordinates at every point. Some papers propose to augment the model by inputting higher dimensional descriptors such as the ones previously mentioned (e.g HKS, WKS, SHOT), initially design for a more *direct* analysis. Such proposition can be seen in Diffusion Net, where HKS interacts well with the diffusion part of the architecture, or PointNet++ were they combine HKS, WKS, Gaussian curvature through concatenation and PCA. However, later work has dismissed this idea (e.g [125]), citing the results of DiffusionNet [161] that show no great improvement from coordinates to HKS.

As methods for learning on surfaces has evolved, we suggest that a better input representation is a simpler one, yet should be more expressive than coordinates. We propose that we can scale back to the simplest differentiable invariant of a surface: curvature.

5.3 The Shape Operator, Curvature, and Congruence

A loose description of curvature of a surface S at a point p may be the measurement of the *bending* in the vicinity of p. Viewed extrinsically, this is the variation of the surface normal N about this point, a precise method to quantify this rate of change is done by the Shape operator.

Definition 5.3.1. *For a point p of a surface $S \subset \mathbb{R}^3$, and any tangent vector v to S at p, the shape operator is a linear operator defined as*

$$\begin{aligned} \mathbf{S}_p : T_p(S) &\longrightarrow T_p(S) \\ v &\longmapsto -\nabla_v N \end{aligned}$$

where $T_p(S)$ is the tangent plane of S at p, and N is a unit normal vector field on a neighborhood of p in S.

In other words, the shape operator $\mathbf{S}_p$ tells us how the normal vector changes as we move in S, in the direction of v from p. One possible way to visualise the Shape operator is through the Gauss map, which identifies each point $p \in S \subset \mathbb{R}^3$ with its unit normal vector N_p, now thought of as a point in $\mathbb{S}^2$. The Shape operator is then the differential of the Gauss map at p and is a tangent vector to $\mathbb{S}^2$ at the image N_p of p, as in 5.1.

The operator $\mathbf{S}_p$ is linear for each $p \in S$, and self-adjoint in the Euclidean inner product $\langle \cdot, \cdot \rangle$:

$$\langle \mathbf{S}_p(v), w \rangle = \langle v, \mathbf{S}_p(w) \rangle,$$

and can therefore be represented by a symmetric 2×2 matrix $[\mathbf{S}_p] : T_pS \to T_pS$ at each point $p \in S$. It's well-known that symmetric matrices admit a complete system of orthonormal eigenvectors (e_1, e_2) spanning the space on which they act. The matrix representation $[\mathbf{S}_p]$ with respect to the basis

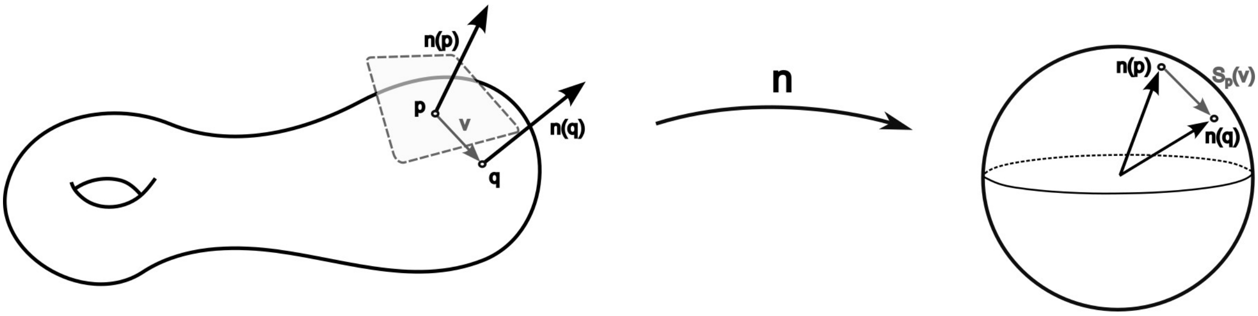

Figure 5.1: The shape operator visualised via the Gauss map.

(e_1, e_2) has the simple form:

$$[\mathbf{S}_p] = \begin{pmatrix} \kappa_1 & 0 \\ 0 & \kappa_2 \end{pmatrix}. \tag{5.1}$$

Definition 5.3.2. *Let S be a surface in $\mathbb{R}^3$, p a point in S, $\mathbf{S}_p$ the shape operator at p and $[\mathbf{S}_p]$ its matrix representation.*

1. *The eigenvalues $\kappa_1(p)$ and $\kappa_2(p)$ of $[\mathbf{S}_p]$ at p are the* principal curvatures *of S at p, and their corresponding eigenvectors e_1 and e_2 are the* principal directions*;*
2. *The* Gauss curvature *K_p of S at p is the product $\kappa_1(p) \cdot \kappa_2(p)$ of the principal curvatures;*
3. *The* mean curvature *H_p is the average $\frac{\kappa_1(p)+\kappa_2(p)}{2}$ of the principal curvatures*

The Gauss and Mean curvatures can be equivalently interpreted as the determinant and half the trace of $[\mathbf{S}_p]$, respectively.

The importance of these quantities is that they completely describe the shape of a surface - which we'll explain momentarily - and it turns out that the Gauss and Mean curvature generate all possible differential invariants of a surface [65, 137]. In particular, Gauss and Mean curvature are fundamental characteristics of the shape of a surface, and the inclusion of the higher order invariants they generate into a representation could even improve the results shown here.

5.3.1 Congruence

To explain how Gauss and Mean curvature completely describe the shape of a surface we need a few definitions. An *isometry* of $\mathbb{R}^3$ is a map $F : \mathbb{R}^3 \to \mathbb{R}^3$ whose differential preserves the angles between tangent vectors at every point of $\mathbb{R}^n$:

$$\langle v, w \rangle_p = \langle DF_p \cdot v, DF_p \cdot w \rangle_p, \quad \forall v,\, w \in T_p\mathbb{R}^3.$$

If g_S is the Riemannian inner product induced on TS by the Euclidean inner product $\langle \cdot, \cdot \rangle$ then an isometry between two surfaces is a map $\eta :\, S \to \overline{S}$ whose differential preserves the angles between tangent vectors to S:

$$g_S(v,w) = g_{\overline{S}}(D\eta \cdot v, D\eta \cdot w).$$

Every isometry F of $\mathbb{R}^3$ restricts to an isometry of surfaces $F|_S = \eta : S \to F(S)$, but the converse need not be true, unless an additional hypothesis on the shape operators is satisfied.

Two surfaces S and $\overline{S}$ are *congruent* if there exists an isometry $F : \mathbb{R}^3 \to \mathbb{R}^3$ such that $F(S) = \overline{S}$; that is, congruent surfaces are surfaces which differ only in their location and orientation in space. It's more-or-less clear that the shape operators $\mathbf{S}$ and $\overline{\mathbf{S}}$ of two congruent surfaces are related by

$$DF_p \cdot \mathbf{S}_p(v) = \overline{\mathbf{S}}_{F(p)} \left(DF_p \cdot v \right), \quad \forall v \in T_pS;$$

in particular, the matrices $[\mathbf{S}_p]$ and $[\overline{\mathbf{S}}_p]$ are conjugate to one another via $[DF_p]$. It turns out (Theorem 9.2 [138]) that if there exists an isometry $\eta : S \to \overline{S}$ such that

$$D\eta_p \cdot \mathbf{S}_p(v) = \overline{\mathbf{S}}_{\eta(p)} \left(D\eta_p \cdot v \right), \quad \forall v \in T_pS$$

- i.e. the matrices $[\mathbf{S}_p]$ and $[\overline{\mathbf{S}}_p]$ are conjugate to one another via the matrix representation $[D\eta_p]$ - then there exists an isometry $F : \mathbb{R}^3 \to \mathbb{R}^3$ such that $F|_S(S) = \eta(S) = \overline{S}$, and the two surfaces are congruent. The conclusion of this brief mathematical digression is that two congruent surfaces have the same intrinsic geometry and shape in space, and two surfaces with the same intrinsic geometry and shape in space are congruent.

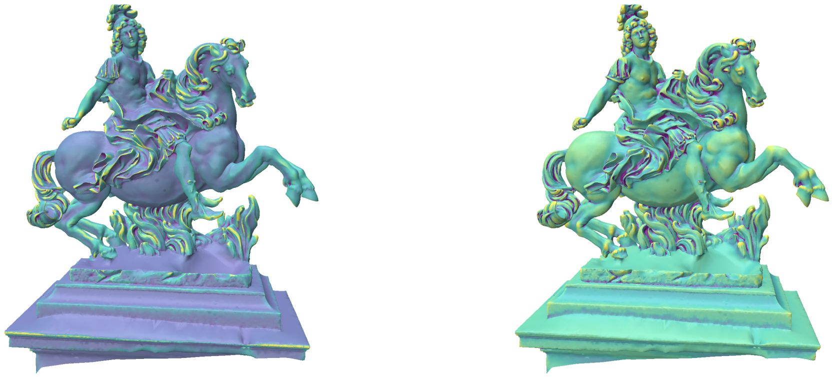

(a) κ_1 (b) κ_2

Figure 5.2: Principal curvature visualisation of a Louis XIV statue.

5.3.2 Discrete curvature

As well as being an important theoretical tool, curvature is a central notion in mesh processing. A large body of work has been dedicated to estimating its discrete counterpart. Among them, many methods propose to infer Gaussian curvature directly, such as in [121], or involve the use of Geometric Measure Theory [55], as in [35, 103]. Interestingly, many efficient methods propose to first discretise the shape operator in order to compute the Gaussian curvature from it. This is done either directly on the mesh triangles, such as in [155], or by first locally fitting a function to the surface, and then computing explicitly the shape operator. To get a better feel for why this is a natural construction of the shape operator, consider a surface S, given a point $p \in S$, then the surface around p can be parameterized as $X(u, v)$ with $(u, v) \in \mathbb{R}^2$. The inner product at T_pS, also called the first fundamental form, is then given for any two tangent vectors v, w by:

$$(v, w)_p = v^T \begin{pmatrix} E & F \\ F & G \end{pmatrix} w,$$

where $E = \langle \partial_u X, \partial_u X \rangle$, $F = \langle \partial_u X, \partial_v X \rangle$ and $G = \langle \partial_v X, \partial_v X \rangle$. And the surface normal at p can be defined as

$$n = \frac{\partial_u X(u, v) \times \partial_v X(u, v)}{|\partial_u X(u, v) \times \partial_v X(u, v)|}.$$

We can now define the second partial derivatives of X on the normal direction n, a quantity called the second fundamental form, noted $\mathbb{I}$:

$$\mathbb{I} = \begin{pmatrix} L & M \\ M & N \end{pmatrix}$$

where $L = \langle \partial_{uu} X, n \rangle$, $M = \langle \partial_{uv} X, n \rangle$, and $N = \langle \partial_{vv} X, n \rangle$. The partial derivatives of the surface normal can then be expressed via the Weingarten equations, in terms of the components of the first and second fundamental form:

$$\partial_u n = \frac{FM - GL}{EG - F^2} \partial_u X + \frac{FL - EM}{EG - F^2} \partial_v X$$
$$\partial_v n = \frac{FN - GM}{EG - F^2} \partial_u X + \frac{FM - EN}{EG - F^2} \partial_v X$$

This enables us to write the matrix form of the shape operator at p as:

$$[\mathbf{S}_p] = (EG - F^2)^{-1} \begin{pmatrix} LG - MF & ME - LF \\ ME - LF & NE - MF \end{pmatrix}$$

From these derivations, it becomes interesting to find good local parametrisation of surfaces, that is, a bi-variate scalar function f such that:

$$X(u, v) = (u, v, f(u, v))$$

The shape operator can then be easily derived from the first and second derivatives of f. An efficient way to find such functions is via osculating jets, proposed in [31]. For the following experiments, we use a multi-scale version of this, proposed in [141], in which the shape operator is computed by using varying size of neighborhoods around a point, yielding a robust method for estimating curvature on a mesh.


(a) Gaussian curvature. (b) HKS, 1st component. (c) SHOT, 1st component.

Figure 5.3: A visualisation of the three representation on the *Armadillo man*.

5.4 Experiments

We will test the principal curvature representation against the three most commonly used representation: the HKS [169], the SHOT descriptor [160], and the extrinsic coordinates. All representations will be tested and ranked for three different architectures: DiffusionNet [161], PointNet++ [147], and DeltaNet [181]. Finally, we pick three tasks of varying complexity to measure the impact of each method: human segmentation [117], molecular segmentation [20], and shape classification [108]. An exemple of each dataset is shwon in figure 5.4

Outside of point coordinates, the two most commonly used representations, in conjunction with NN architectures, are the HKS and SHOT representations. As explained in 5.2, the HKS is a purely instrinsic representation derived from the Laplace operator defined in chapter 3.3, and constitutes the most widely used signature-based method to represent shapes. The SHOT representation is a descriptor mixing signature and histogram-based methods to describe shapes, and is therefore an extrinsic representation. As they belong to two different classes of surface representations we believe they are the most adequate for benchmarking our proposed curvature representation. We visualise the first dimension of each representation in figure 5.3, where we can clearly see that the HKS already captures features of the shape at a bigger scale, while the curvature is only describing local variations. The first component of the SHOT descriptor contains more noise, however we see how it picks up a combination of local and global features.

We regard DiffusionNet as the state of the art in NN architecture, as it shows the most promising results on general benchmark tasks. In addition, they have shown very little difference in their results when changing the input from coordinates to HKS, making it the *harder* test for our representation. PointNet++ have been designed as a general method to process shapes arising from many situation, in a controlled environment as is our case, but also from segmented images, as encountered in the autonomous driving field [147]. As so they have the least geometric structure to describe surface: all one needs is a point set. We believe that in this case, adding better surface information to the input may greatly enhance the performance of the model. The authors of PointNet++ have already touched on this subject, recommending a linear combination of HKS, WKS and Gaussian curvature, followed by a PCA projection, leading to a 64 dimensional feature per point. We aim to show that a 2d (or even 1d) input of curvature information is more relevant for a smaller computational cost. DeltaNet proposes an architecture intrinsic to surfaces by design, by combining four operators defined on the surface: Laplacian, divergence, curl and norm. Most papers that propose other surface descriptors rather than coordinates as input, do it solely to have an intrinsic representation of the surface. We believe that curvature not only gives isometry invariance (section 5.3), but is also more robust numerically and better performance from our curvature representation in this architecture would support our belief.

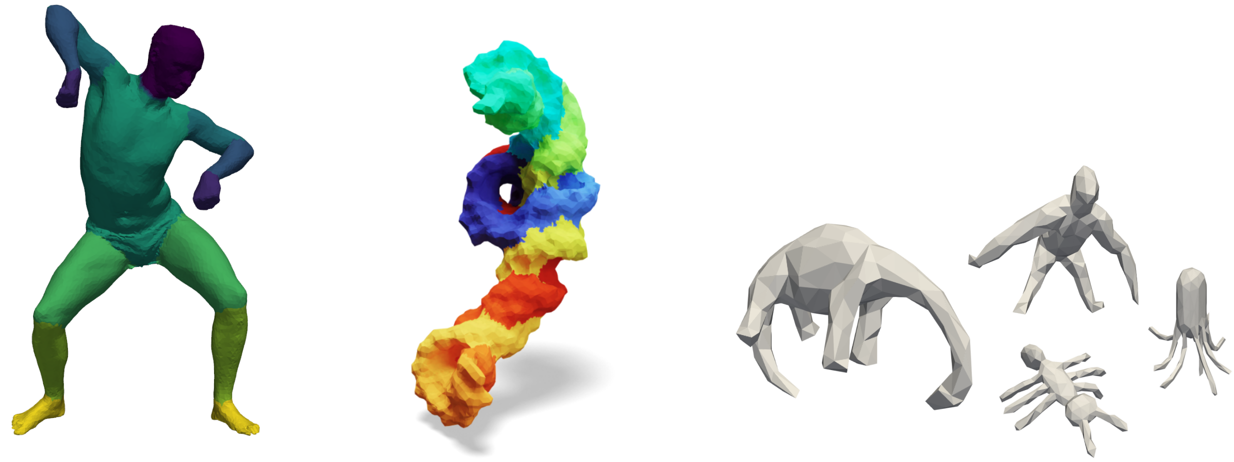
(a) Human poses [117]. (b) Molecules [20]. (c) Shrec'11 [108].

Figure 5.4: Samples of the segmentation and classification datasets used for experiments.

5.4.1 Implementation

The performance of each representation is strongly dependent on the chosen implementation. We have tried to be as fair as possible by not developing our own implementations and instead use methods which have already been tried, tested, and validated in the literature.

For calculating the discrete principal curvatures via quadratic surface fitting, we have used `igl`'s implementation with a fixed neighbourhood radius of 5; the Gaussian curvature κ is then computed directly as the product of κ_1, κ_2.

HKS depends on the Laplacian, and we have used the method implemented in `robust-laplacian` based on [162] - this is also consistent with what is used in Diffusion Net. The eigen decomposition is then performed with `scipy`.

For the SHOT representation, we use the implementation in the `pcl` library, which computes 352 features per vertex, in this case we normalise all shapes and use a ball of radius .1; all other parameters are left untouched.

Regarding the neural networks, we use the implementations made publicly available by the authors, modifying only when needed to accommodate more than just coordinates as input. We also use the same parameters proposed in each paper when they are known, which we detail for every task below.

5.4.2 Time Complexity

As a first experiment we compute[1] the time complexity per number of points for each representation, see figure 5.5. HKS and curvature are both efficient for meshes with up to ∼100k points, although curvature is consistently faster, even for larger meshes (tested up to ∼500k points), displaying the very small overhead incurred by the use of curvature in surface models.

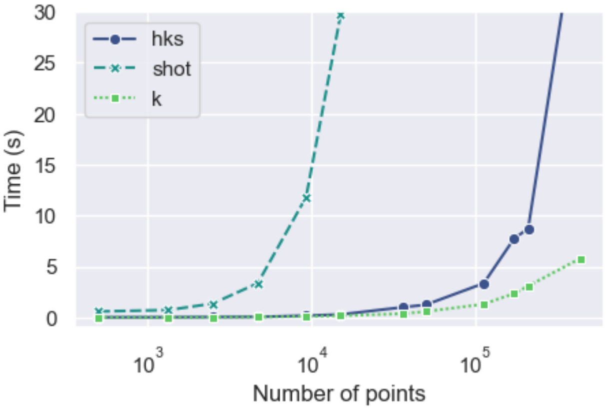

Figure 5.5: Time of computation for each representation with respect to the number of points in a mesh.

5.4.3 Human anatomy segmentation

We first attempt to segment the human parts from the composite dataset proposed in [117], containing samples from other human dataset, namely FAUST [24], SCAPE [6], Adobe [1], MIT [174], and SHREC07 [60]. As in the initial paper, We use Shrec'07 dataset as test set. As proposed in [180], we differ by evaluating on vertices rather than faces. For PointNet++ and DeltaNet we resample each shape to 1024 points, and we leave the meshes untouched for DiffusionNet, as per the experiments conducted in each paper. We optimise the negative log-likelihood for 100 epochs, with the ADAM optimiser and a scheduler step every 20 epochs. We run the experiment 5 times and pick the parameters from the last epoch for evaluation. We report the mean test accuracy in table 5.1.

The results highlights the assumption that better representation lead to better performance. PointNet++ shows the greatest improvement when moving away from coordinates, this is due its architecture having the least geometric information. The better results come from the principal curvature,

---

[1] Computed on an Apple M2 chip

	xyz	shot$_{16}$	shot$_{64}$	hks	κ_1, κ_2	κ
Point Net ++	69.6	71.4	72.4	78.1	**80.6**	74.5
Delta Net	72.4	58.1	66.2	69.2	**86.8**	60.0
Diffusion Net	94.7	95.0	95.0	95.1	**97.4**	95.4

Table 5.1: Human part segmentation results.

which shows how expressive this representation is. This assumption may be even more expressed by looking at the Delta Net experiment, where κ_1, κ_2 greatly outperforms all other methods, however coordinates are better than any other more complex representations. DiffusionNet may show that it is agnostic to the type of input, as long as it loosely describe the shape, however the difference is still significant. To further demonstrate the impact of good representation, even in the case of DiffusionNet, we show in 5.6 the worst cases for xyz , κ_1, κ_2 and HKS inputs. The clear improvement may be even more important than general accuracy in some cases, e.g with human-in-the-loop type corrections.

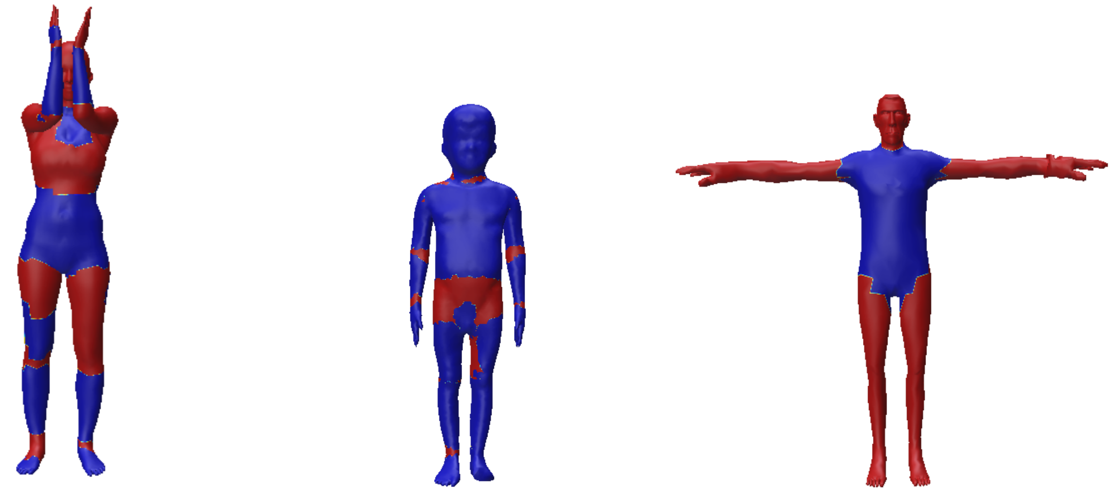

(a) xyz. 29.2% accuracy. (b) κ_1, κ_2. 80.9% accuracy. (c) *hks*. 12.3% accuracy.

Figure 5.6: Human part segmentation with DiffusionNet. Worst cases for different representations, blue shows the correct prediction, red the error.

5.4.4 Molecular segmentation

The molecular dataset made available by [20] and first proposed in [145], can be considered a harder segmentation task then the Human part dataset. With a wider range of shapes in the form of RNA molecules, and its 260-way part segmentation task. We perform random 5 fold cross validation with a train-test ratio of 80-20, the same splits were used for all cases. We resample all meshes to 2048 points, except in the case of DiffusionNet where we kept the original discretisation. We run the model for 200 epochs and compute the test on the last epoch. Table 5.2 shows the mean test accuracy results.

	xyz	shot$_{16}$	shot$_{64}$	hks	κ_1, κ_2	κ
Point Net ++	35.4	70.9	71.9	70.2	**72.4**	69.0
Delta Net	29.2	45.6	**56.5**	49.6	55.5	29.2
Diffusion Net	82.6	88.4	89.1	85.6	**89.4**	84.0

Table 5.2: Molecular segmentation results.

Again, we see a great improvement when adding a better representation of the surface in the case of PointNet++, going from failing in the case of coordinates to outperforming DeltaNet, with principal curvature having the best results. DiffusionNet shows a non-negligable jump in performance as well. Although SHOT descriptor outperform other representation in the case of DeltaNet, the general performance of this architecture is underwhelming. We believe this is due to the accumulation of errors in the discretisation of surface operators used. Indeed, one layer computes a chain of 6 operators on the surface, since the RNA shapes are very irregular, the error for each operator could be significant.

5.4.5 Classification

Finally, we propose to compare representations in the context of classification. This final experiment should show whether or not geometrically informative inputs interacts well with pooling-type operations. We choose the widely adopted baseline Shrec11, proposed in [108]. It is a 30-way classification dataset with 20 shapes per category. Their are many methods for training and evaluation used across the literature. We choose the simplified mesh dataset and the harder version of training, using only 10 samples per class and evaluating on the rest, and do this with 5-fold cross validation procedure. We train for a 100 epochs with a scheduler step at 50 epoch and optimise the cross-entropy loss with a label smoothing factor of 0.2. Mean test accuracy are showed in table 5.7a.

	xyz	shot$_{16}$	shot$_{64}$	hks	κ_1, κ_2	κ
Point Net ++	71.5	69.8	60.7	60.8	85.7	**96.2**
Delta Net	75.7	54.9	60.4	98.6	90.1	**98.8**
Diffusion Net	80.3	52.6	67.4	98.9	94.2	**99.1**

(a) Results.

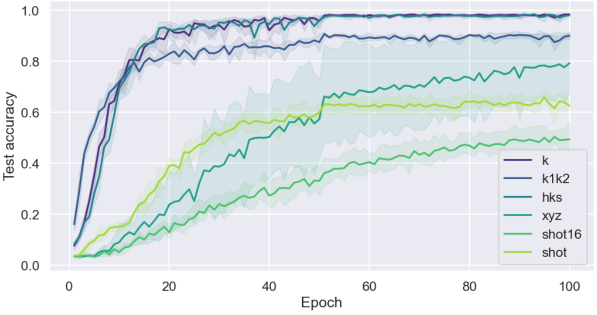

(b) DiffusionNet: Training evolution of the test accuracy with 95% confidence interval per representations.

Figure 5.7: Shrec'11 classification.

Yet again we see a great improvement when turning to better representations, even more so when using gaussian curvature κ. This indicates that gaussian curvature interacts better with pooling operations present in classification architectures. Additionally, figure 5.7b shows that all geometric representations yield less variability across each folds, while HKS, gaussian curvature κ, and principal curvature κ_1, κ_2 converge much faster than all others.

5.5 Application to the Left Atrium

Keeping the main thing, the main thing, we can now go back to the LA. Everything developed and tested above was to speed up the anatomic labeling pipeline of LA meshes, we can now focus on that. Our ever expanding dataset is now composed of 232 segmented LA originating from the Bordeaux hospital center, these shapes were transformed into meshes and labelled with the multi-atlas method proposed in section 4. However a new batch of data from the Toulouse hospital is now available, therefore we set out to train a semantic segmentation model on the Bordeaux dataset and apply it on the Toulouse dataset. To do so we pick the best performing model that we have found, the DiffusionNet coupled with principal curvature, and train it using the same procedure and hyper-parameters as in the previous experiments. As a result, we significantly improve general performances. In table 5.3, we show comparative results of both LDDMM methods, computed on the Bordeaux dataset, and the new DiffusionNet method, computed on the Toulouse dataset. The net gain of this method is three fold. First, the intrinsic representation provided by the principal curvature eliminates the need for any rigid registration pre-processing, which gets rid of human intervention at this step. Second, the worst case is drastically improved (see 5.3), which means we have, perhaps nearly, reached the point were the pipeline can be run entirely without a human in the loop. Finally, the inference time for a given shape is around 5 seconds on a standard laptop CPU, coming down from around 40 minutes for the LDDMM case, which makes the pipeline from image to labelled mesh usable in real time scenarios.

Method		Accuracy	f1 score	Precision	Recall
LDDMM	Mean	.926 ±.058	.925 ±.058	.956 ±.057	.924 ±.058
	Min	.557	.558	.559	.557
Multi-Atlas LDDMM	Mean	.969 ±.035	.969 ±.038	.972 ±.022	.969 ±.035
	Min	.848	.833	.901	.848
DiffusionNet w/ κ_1, κ_2	Mean	.992 ±.007	.991 ±.007	.992 ±.007	.992 ±.007
	Min	.964	.964	.967	.964

Table 5.3: Labelling pipeline results.

With this successful attempt at efficiently labelling the LA meshes, it is natural to ask how these methods would perform to directly classify the cohort for automatic stroke prediction, alleviating the need for any further analysis. Unfortunately, the model fails to learn in this scenario and systematically places all shapes into the same class, even when carefully fine-tuning the hyper-parameters. This strengthens our belief that to respond to this question, it is the important to extract meaningful and precise geometric features first and perform a thorough analysis on them, as it is the only way to decidedly say and quantify whether or not the shape has any impact on stroke occurrence.

5.6 Conclusion

In this work we showed that curvature should be the representation of choice when it comes to processing surfaces with neural networks. Indeed, as neural networks are able to learn surface features at multiple scales, the need for complex representations as inputs is not relevant anymore, however curvature remains a better local representation than extrinsic coordinates. In all experiments, curvature showed better performance than any other choice of input, both qualitatively and quantitatively, and can be compute with minimal overhead. It's combination with PointNet++, which has the least prior information about the surface, showed that it can help the network better understand the surface structure. When combined with DeltaNet, which contains only intrinsic operation, the improvement indicates that curvature gives more than just a rigid transformation invariance. Even in the case of DiffusionNet, where the diffusion operation seems to interact nicely with any representation, curvature as input showed significant improvements. Following these empirical results, we believe curvature should be the standard practice when using models to learn on surfaces. We successfully applied these newly found concepts to the semantic segmentation of the LA's anatomical parts, although showed that attempting direct classification with such model does not give any meaningful results. Yet again, the black-box nature of neural networks prevents us from probing it and giving any final answer to the main problematic of this thesis. In the following and final chapter, we propose the last missing step to LA representation pipeline in the form of expressive geometric feature. Which will finally enable us to have full insights into the relationship between LA shape and stroke.

Chapter 6

Geometric features of the LA

Contents

Abstract We now show a full class of features and analyse them. Leading to an expressive patient stratification of shape related strokes. We go further and show that there are in fact a distinct difference between patients with paroxysmal and persistent forms of AF in terms of shapes. Finally, we highlight that the opacification defect on arterial phase, a strong haemodynamic feature, can be predicted well with our features.

6.1 Introduction

In the previous chapters, we have followed well established method for global shape analysis, showing that we need to build representations more specific to the Left Atrium (LA). To this end, we proposed a better control over the shape by automatically labelling the anatomical parts of the LA, and leverage on it to derive the sought after representation. In chapter 4, we propose a two parts pipeline for the analysis of the LA shape: First a diffeomorphism-based labeling method of the anatomical parts of the LA; Second a geometric graph, also represented in figure 6.3, as to solely extract the information about vein and LA Appendage (LAA) configuration. Although yielding good preliminary results, the lack of data made it difficult to conclude in any meaningful way. The increase availability in data, from ~ 100 patients to ~ 400 patients, helped us realise that both part of the pipeline needed improving. Chapter 5 was motivated by the need for a more efficient labeling pipeline, we successful results we focus in this chapter on more expressive set of features representing the LA. For instance, the LA graph was initially represented in its raw form, as a set of 13 3D coordinates composing the nodes of the graph, while it can be used in a more straightforward approach to measure angles and orientations directly. These values, by being more expressive, should help the general performance, but also help understand important features. In fact, the main criticism we received from clinicians from our previous work was regarding the interpretability of results. Naturally, the sole tool of medical experts to observe the LA is imaging, for this reason, any important information, should, to some extent, be drawn and measured directly by radiologists, on a 3D imaging medium of choice, in our case CT-scans.

Unsurprisingly, this is precisely what is being done in most related work. For instance, the previously mentioned classification of LAA morphology into 4 types, proposed in [184], can be done directly by the clinician's eye, although at the risk of variability in the LAA classification choice between one expert to the other. Other papers have proposed to study the LAA with more objective measures, such as the area and volume [130, 15], orientation [90], or even more precisely the shape of LAA Ostium (LAAO) [135]. An overview of such methods can be found in [39]. Although the LAA has seemingly been the primary focus surrounding the LA anatomy, some index relating the body of the LA to cardioembolic strokes have been studied, such as the body volume [87], or sphericity [133]. The pulmonary veins (PV), on the other hand, have only been studied recently, and mostly from a haemodynamic perspective [51, 59]. Overall, the large body of work surrounding this subject contains a clear direction in research: the goal is to find a set of local descriptors that best explicit the variability of the LA's morphology, in order to study their link with a specific pathology, in our case, cardioembolic strokes.

This fits perfectly in our pipeline. Now that we can control the LA shape efficiently via anatomic labeling, we can extract as many geometric features as needed, with the important added value that everything can be done automatically, opening doors to large scale studies with minimal human-hour. Moreover, the robustness of the features we will extract is only dependent on the robustness of the semantic segmentation and not on their own computations. In contrast with other studies, we can analyse a large set of local features, describing many parts of the LA, rather than focus on a small subgroup, enabling the possibility to pick up more complex relationship within the anatomy.

In the next section we show how such features can be computed by leveraging not only on the

LA graph, but on the complete anatomical labels, reaching a total of 66 uncorrelated features. This concludes the pipeline described as the main goal at the beginning of this thesis. We follow by performing several statistical analysis of those geometric features, which leads us to a better understanding of the impact of the LA shape on stroke occurrence, but also of the remodelling that the LA undergoes when the Atrial Fibrillation (AF) becomes persistent, and finally the link between the presence of blood stasis in the LAA and its shape. This study thus constitutes the provided answer to the main medical problematic of this thesis.

6.2 Feature extraction

The anatomical labels of the LA makes it possible to automatically measure a significant amount of geometric quantities describing it. To do so we use the intermediary step of computing objects describing specific anatomical elements of the LA, before extracting the geometric features from them. Those objects are: the graph representation proposed in section 4, or more precisely the set of 13 3D coordinates composing it; the ostia of the LAA and each PV, defined as the intersection between every label and a curve in 3D; and the LA body and LAA themselves, representing surface meshes with holes that we can identify. Let us now describe the features that we derive from each object.

6.2.1 Some post-processing

To be as robust as possible, we do some post-processing to correct errors potential errors from the automatic. We use the same *label cleaning* pipeline proposed in chapter 4. This will ensure that all labels defined a single connected component and are all neighbours only to the body. We can however make a distinction with the errors generated by the diffeomorphic labeling procedure initially employed. First the problem of connected Left Superior Pulmonary Veins (LSPV) and Left Inferior Pulmonary Veins (LIPV), is almost non existent using the latest method, as it happened only twice over the total 404 mesh. Second, the presence of small mislabelled patch happened more often, and wasn't a problem restricted to the LSPV and LAA. Nonetheless, both problems are solved automatically. In addition, we wish to have smooth label boundaries, or ostia. We want smooth boundaries in order to have features that are not dependent on the irregularity of the mesh edges. To do so we replace each point of the extracted curve with the pairwise mean of each neighbour. We have found that by repeating this procedure twice yields the desired smoothness. A graphical representation of this procedure can be seen in figure. 6.2, and results can be seen in figure 6.1. We can now focus on the feature extraction.

6.2.2 From the geometric graph

The graph itself was initially used as a raw representation to input in a prediction model in chapter 4. However, while containing a lot meaningful information, this set of 3d coordinates lacks expressive power. Instead of using it directly, however, we can use it to extract relevant features, giving clearer

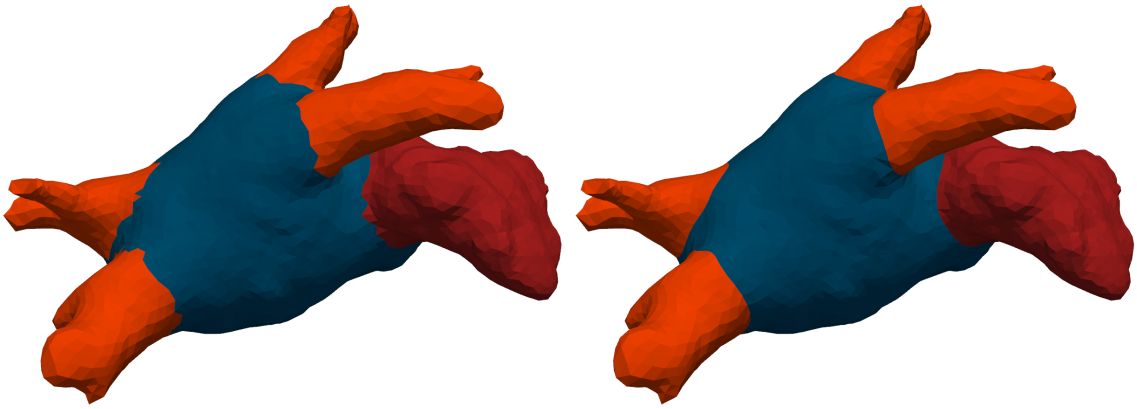

(a) Labels with rough boundaries.

(b) Smoothed labels.

Figure 6.1: Smoothed boundaries

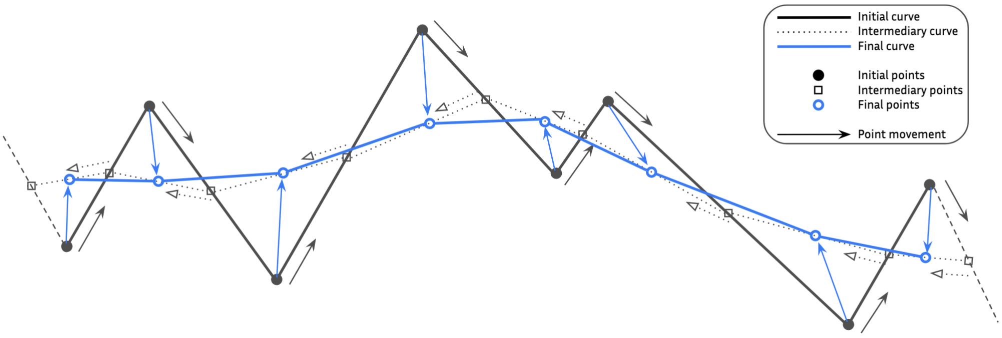

Figure 6.2: Boundary smoothing procedure.

information about the shape. In fact, we can measure angles, orientations, distances and alignments, solely from the node coordinates of the graph and the configuration of its branches. Although the goal is too extract as much features as possible, we list here only a set of selected features that will become relevant in later analysis. Some of them are depicted in 6.3.

Folds To measure the angles of the veins with respect to the rest of the LA, we compute the angle between the branch of a given vein with the branch linking the centre of the LA and the vein's ostium. We refer to these as the PV's *fold*, they are represented in 6.3a as ϕ_1, ϕ_2, ϕ_3 and ϕ_4.

Distances In addition to angles, we can use the graph to measure distances. Among the relevant ones, we can compute the distance between ostia, specifically the distance between right PV (RPV) ostia and the one between left PV ostia. They are shown in 6.3b as L_α and L_β.

Alignments Although a rather vague word, we refer as *alignment* of PVs and LAA as a way to measure how asymmetrical the LA is. This is again information contained in the graph, which we can extract through multiple features. For example, define the vector described by the RSPVO and RIPVO, and the vector described by the LSPVO and LIPVO, then we can measure the angle

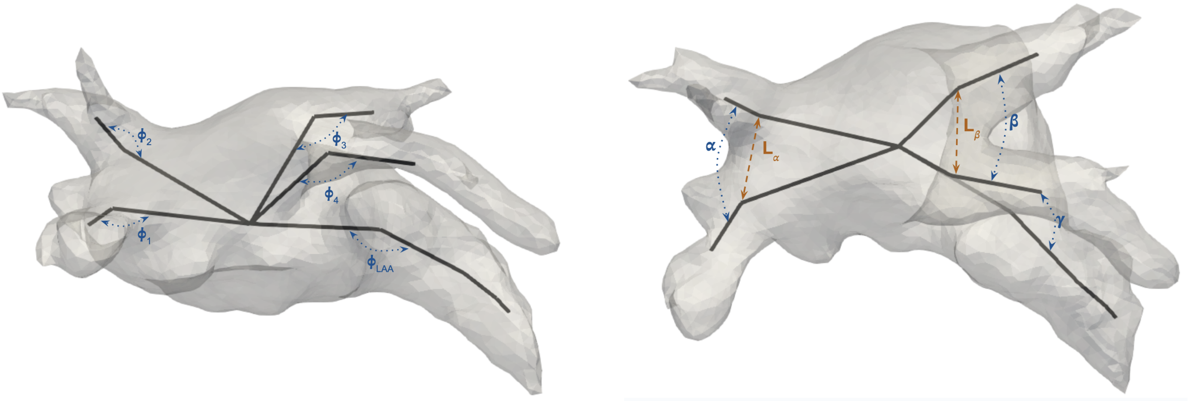

(a) Angles. (b) Distances and alignments.

Figure 6.3: Some relevant features extracted from the geometric graph.

between them, effectively describing how different the vein blocks are connected to the LA, we refer to this value as *PV alignment*.

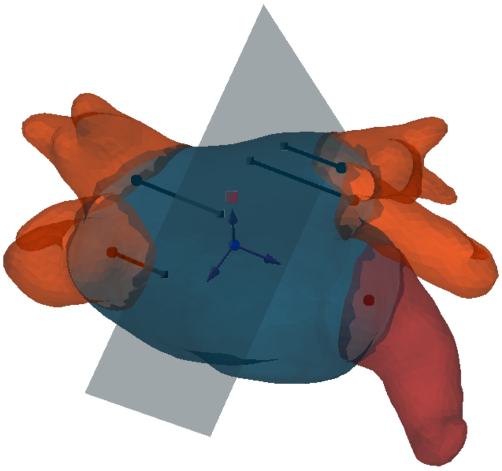

Figure 6.4: Local coordinates.

Local coordinates While the *folds* described above can measure one type of insertion angle of veins, the measure of vein orientation can be done more precisely. To do so we first propose a local coordinate system entirely defined with ostia points and the LA centre point. We do this in the following steps:

1. The *x-axis* is defined by the vector with the LA centre as origin and pointing to the LAA.
2. We then project the vein ostia onto the plane defined by the LA centre and the normal in the LAA direction; the *y-axis* is defined by the direction from the la centre to the mean of the projected ostia.
3. Finally, the z direction is the cross product between the two previously defined vectors.

Figure 6.4 illustrates this construction. Thanks to this, we can now rethink how to define the PV orientations and configurations. We first define a vein direction as the normalised vector with the vein's ostium as origin and the vein's centre point as direction. We then perform change of basis to have each vein direction described in the local coordinate system. These new directions now lie on the unit sphere and can be re-written in spherical coordinates (r, θ, ϕ), with the magnitude r can be ignored, yielding the two angles θ, ϕ as descriptors of the vein's orientations. We can also represent each vein as a point on the unit sphere $\mathbb{S}^2$, helping us visualise their distribution, as seen in fig. 6.5. To go further, we propose to represent each LA as one point on $\mathbb{S}^2$ via the Fréchet mean of all 4 veins, as showed in fig. 6.5d. We then call *vein configuration* of an LA as the spherical coordinates of these points.

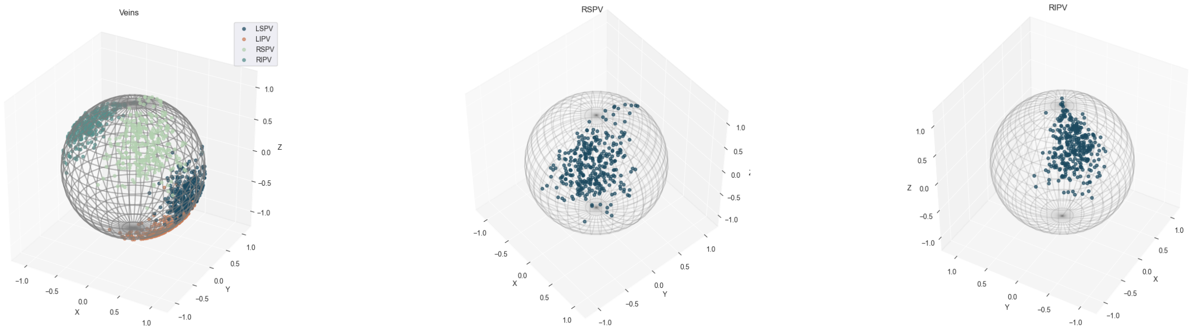

(a) All veins.
(b) RSPV.
(c) RIPV.

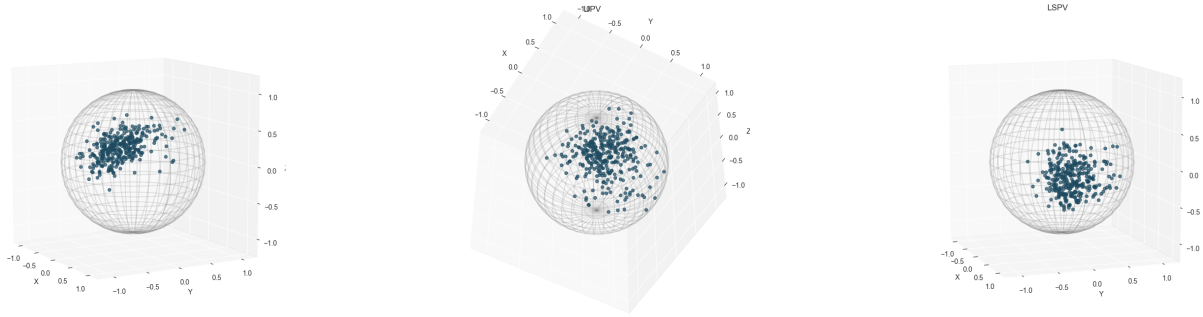

(d) Frechet means.
(e) LIPV.
(f) LSPV.

Figure 6.5: Vein representations.

An intermediate vein study We propose to take a closer look at this representation of PV. First we can visualise each vein on the sphere for the entire dataset. This highlights the important *up-down* shift that right PVs have, while the left PVs distributions are much more clustered. Going further, we can perform a study of such representation through geometric statistics, and leverage on the `geomstats` package [123]. We perform geometric k-means for clustering directly on $\mathbb{S}^2$, as shown in 6.6a. The lack of global coordinates on a manifold leads us to represent the data in the tangent space at a specified point. The typical base point for this is the Fréchet mean of the dataset. We also need to choose the reference frame in this tangent space, again, we follow standard practice by using the two principal components of a PCA analysis, called Tangent PCA, shown in 6.6b. While we may successfully find 2 cluster, we do not find any link between cluster and stroke, as seen in 6.6c. We still add the (x, y)-coordinates in the tangent space to our set of features, as they may turn out to be meaningful when combined with others.

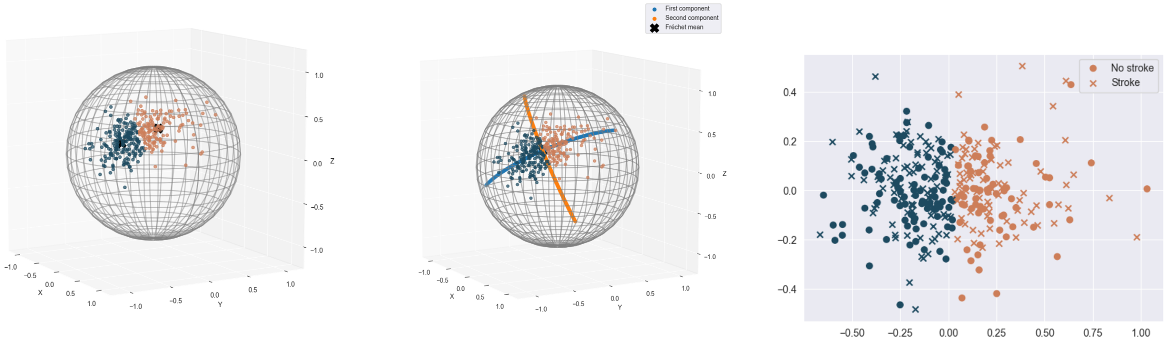

(a) Clustering.

(b) TPCA.

(c) Tangent space.

Figure 6.6: Vein study.

Other examples The features previously described are not the complete list of computed features. Although we do not wish to describe them all, we can add here the ones that appear in figure 6.3: α and β are the angles between RPVs and LPVs respectively; ϕ_{LAA} is the angle described by the LAA w.r.t the body centre, measured at every point of the LAA nodes; and γ is the angle describing the alignment between the LSPV and the LAA. Although the list of computed feature does not stop there, all individual properties of the LA shape we are trying to measure is contained in the features we overviewed.

6.2.3 From curves

Although the LA graph is a valuable source of information, we can go beyond it and further leverage on the anatomic labels. To this end, we focus in this section on the ostia curves, that is, in our setting, the polyline defined by the set of connected edges of each label's boundary, see 6.7a. For these curves, we propose to measure the length, and the total curvature, that is, the sum of the curvature computed by finite differences, visualised in 6.7b. Since the curve is computed from the smoothed boundary in 3d, it does not lie on a 2d plane, however, to capture additional features, we approximate it as a 2d curve by projecting it to the closest plane. This plane is defined via a base point and direction: the mean of all vertices of the polyline; the mean of each point's binormal vector (from the standard Frenet frames construction). We can now define the maximum diameter or the minimum and maximum radii, see 6.7c. For good measure, we also fit the ellipse that best approximate the projected closed curve and extract the small and long axis distances as well as the ellipse eccentricity.

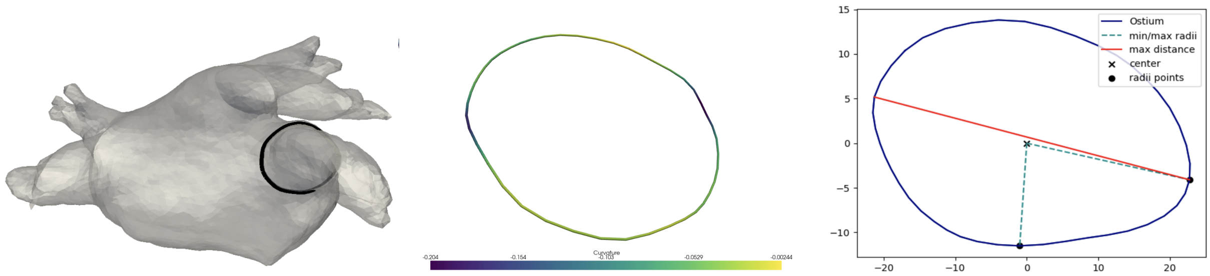

(a) Ostium. (b) Ostium curvature. (c) Ostium 2D projection.

Figure 6.7: Ostium features.

6.2.4 From the surface

The last subset of features we will describe are extracted from surfaces of anatomical parts. Precisely the body surface and the LAA surface. Although many "trivial" quantities have been computed, we show here how to extend to more elaborate measurements.

Common quantities We can first compute common quantities, such as volumes and surface area, as well as their Body to LAA ratios. We can also compute the Surface Area to Volume ratio (SA:V) of the LAA, a quantity often measured in biology [61] and physics [144], which we will turn out to be an excellent predictor of blood stagnation in the LAA (see section 6.3.5).

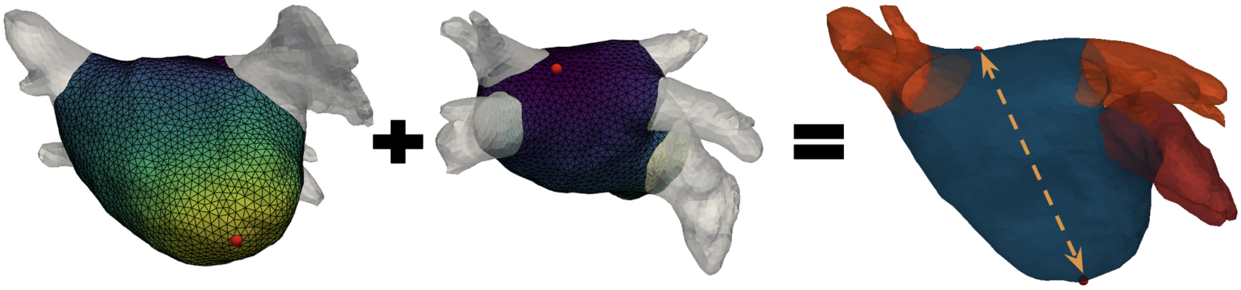

Figure 6.8: Body length computation.

Distance based Although the labels already gives us good control over the LA mesh, we are still missing some key geometric indicators. Specifically, we identified three of them that could help us extract meaningful information from the shape. They are the tip of the LAA – the furthest point of the LAA surface; the body's apex – the lowest point of the body surface; and the body's ceiling – the highest point of the body surface. Although not directly indicated in the labels, nor computable with the local coordinates, we can rely on distances on the surface to extract all three information. That is, we can redefine each element as a specific distance based on the ostia points as follow:

1. The body apex is the furthest point from all the PV ostia.
2. The body ceiling is the closest point from all the PV ostia.

3. The LAA tip is the furthest point from the LAA ostium.

To clarify, the PV ostia distance is achieved by computing the mean every distance map from every points of every ostium to the entire LA body. The LAA ostium distance is achieved by computing the mean of every distance from every points of the LAA ostium to the entire LAA. An ostium is typically ~ 50 points, and the body contains ~ 4000 points, making it important to have an efficient distance estimation. Thankfully, the vector heat method introduced in [38], and described in section 3.3, provides an efficient way to compute repeat distances on a given shape, making it possible to compute both distances in real time. Once the maps computed and the 3 points acquired, we can extract additional features. The first one is the body length, which we define as the distance between apex and ceiling, as shown in figure 6.8. The LAA tip can be used for multiple purposes. We first use it to define the LAA *surface length*, by computing the geodesic length between the point with minimum LAA ostium distance and the LAA tip, as shown in figure 6.9a. We also append the LAA tip to the LA graph, in order to make the LAA branch of the graph a more faithful description of the LAA, similar to a center-line. This is then used to define the LAA *length*, as the sum of the edges of the LAA branch. We can further exploit distances from ostia by considering distance lines on the surface. We compute 12 equidistant lines between the LAA Ostium and its tip – with respect to the LAA Ostium distance map–, as well as 2 lines in the other direction on the LA body, as shown in figure 6.9b. This is then used to compute the area of the slice prescribed by each line (see figure 6.9c), and define the *squeeze* (resp. *release*) factor, as the ratio between the largest line before (resp. after) the ostium and the ostium. We refer to the mean of the squeeze factor and the release factor as the bottleneck factor.

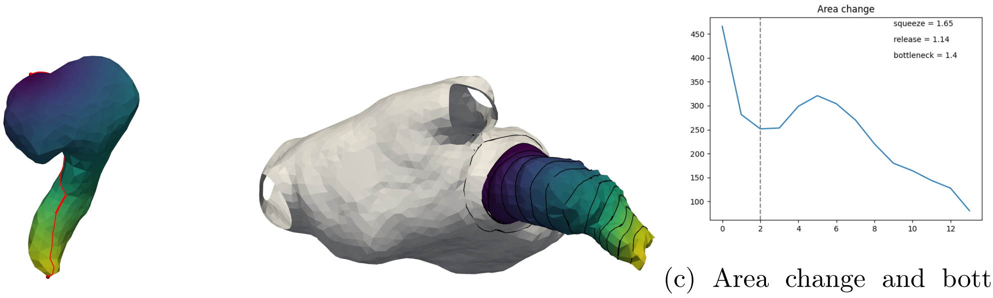

(a) The tip of the LAA.

(b) Distance lines.

(c) Area change and bottleneck.

Figure 6.9: Exploiting surface distances.

Other features We also extract other features, such as the length between the LAAO and the first intersection with the LAA surface in the LAAO normal direction, which we call the characteristic intersection, as well as the ratio between this characteristic intersection and the distance to the tip of the LAA. Finally, we add the first 4 eigenvalues of the laplace decomposition of the LAA, as described in 3.3, also called the Shape's DNA [150].

6.2.5 End note

We showed how to extract a large variety of features from the LA, highlighting how having control over the LA shape in the form of anatomic labels enables us to construct a specific and expressive representation of the LA. The definition of all the features above were motivated by the goal of predicting stroke, however they can help in many other contexts, from defining meaningful contour lines, to capturing the movement of the LA in a dynamic study. We hope this set of tools will help pursue multiple lines of research surrounding the LA in the future.

6.3 Statistical Analysis

With the previous section constituting the final component of the set of tools we built to analyse the LA shape. It is now possible to focus on the clinical question of this thesis: **What is the relationship between the anatomy of the LA and the occurrence of cardioembolic stroke?**

To do so, we will analyse the geometric features through a more classical approach than in section 4, by using tools from statistical analysis as well as the decision-tree based method Light Gradient Boosting Machine (LGBM) proposed in [92] and its Python API[1]. Their are two main arguments for this decision. First, the goal of extracting features is to increase the expressive power of the shape representation, but also to allow for better interpretability. While *explainable* AI is an active field of research, the simpler models are still the most trustworthy, and should still be the first methods to try [154], and although gradient boosting methods may not be the most easy to interpret, they are still a better choice than neural networks. Second, the considerable increase in our dataset, from ~ 100 patients to ~ 400 patients, and from a unique Hospital centre to both Bordeaux and Toulouse hospitals, renders the use of architectures like the MCVC inadequate, as highlighted in the drawbacks section of chapter 4. We give the complete characteristics of the cohort from both centres in table 6.1.

In the following sections we pursue several directions. First we attempt to discriminate between the stroke and control groups of our geometric feature dataset using both univariate statistical analysis and an LGBM classifier. We then realise that the remodelling of the LA when AF becomes persistent [134], has a significant impact on the shape, which we measure using the same discriminating methods as above. This leads us to separate the cohort in two groups, the patients suffering from paroxysmal AF and the ones suffering from persistent AF, and we conduct independent analysis on both of them. Finally, we postulate that the presence of blood stasis in the LAA – as detected by an opacification defect observed on CT scans at the arterial phase – are strong biomarkers of haemodynamic perturbations, and by extension stroke risk. We show that this marker can be

---

[1] https://github.com/microsoft/LightGBM

efficiently predicted by shape alone.

	Bordeaux ($n = 232$)	**Toulouse** ($n = 172$)	**Total** ($n = 404$)
Age	75 ± 9	77 ± 8	76 ± 9
Gender	67%	79%	68%
BMI	26 ± 5	26 ± 5	26 ± 5
Paroxysmal AF state	64%	47%	56%
LVEF	56 ± 9	57 ± 9	57 ± 9
CHA_2DS-VASc	4.6 ± 1.4	4.5 ± 1.3	4.4 ± 1.4
HAS-BLED	$3.2 \pm .9$	$3.3 \pm .9$	$3.3 \pm .9$
Cardioembolic stroke	41%	37%	39%
LAA Thrombus	3%	6%	5%
Other stroke	3%	4%	4%
Stroke positive	41%	37%	39%

Table 6.1: Baseline characteristics of patient cohort.

6.3.1 A note on baselines

The reason for our work is to get a more precise understanding of stroke risk by leveraging on imaging, a growing commodity for patient stratification. To measure how successful, or unsuccessful, the shape may be to predict stroke, we must define a correct baseline. We give here a discussion on what should be a good comparison methods.

CHA_2DS-VASc score Clinicians typically rely on the CHA_2DS-VASc score to assess the risk of cardioembolic events, with a score bigger or equal to 2 implying a high risk leading to predictive treatment. While it is clear that the CHA_2DS-VASc score should be our main baseline, it is less straightforward to choose *which* version of it. Indeed, as we are studying a retrospective cohort, we should consider the value of the CHA_2DS-VASc score before cardioembolic stroke. Precisely, the **S** in CHA_2D**S**, is the target to predict. We leave the choice to the reader by showing the baseline performance of the CHA_2DS-VASc both pre- and post-stroke in table 6.2.

Another concern regarding this score is the decision threshold used to diagnose a patient at risk. The cohort of patients contained in our dataset are all undergoing LAA occlusion, that is, they were all considered at high risk of stroke. As clinicians are using the CHA_2DS-VASc to assess this risk, every patient has a score already higher than one. If we use such threshold, then we obtain a Recall of 1, but very poor accuracy, see table 6.2. We may want to define a higher threshold, for instance a threshold at 5 yields the highest accuracy, which we give as well in table 6.2. To ensure a better understanding of how good the CHA_2DS-VASc is for predicting stroke, we additionally show the precision-recall curve and the ROC curve in figure 6.10.

Shape-based baselines To add more perspective on our method's performance, we also want to compare it with method that also leverage on the anatomy of the LA. In fact, some clinicians

	Threshold	Accuracy	f1	Precision	Recall	AUC
CHA_2D-VASc	1	.395	.552	.383	.992	.513
	5	.599	.431	.463	.403	.560
CHA_2D**S**-VASc	1	.380	.548	.378	1.00	.502
	5	.708	.659	.628	.694	.705

Table 6.2: CHA_2DS-VASc results.

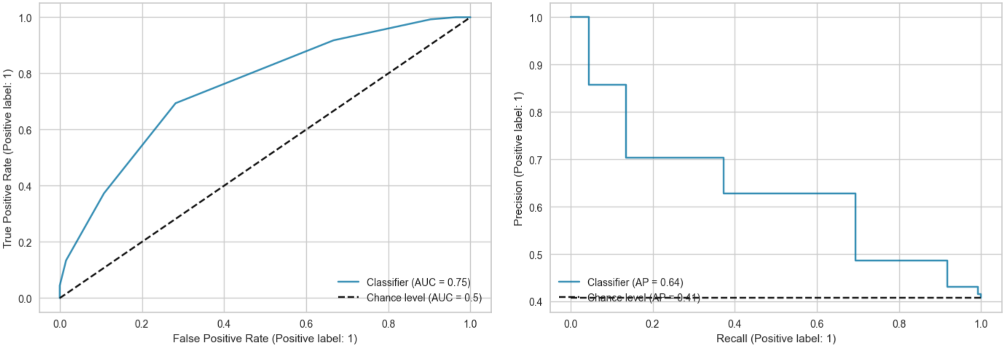

(a) CHA_2D**S**-VASc ROC. (b) CHA_2D**S**-VASc Precision Recall Curve.

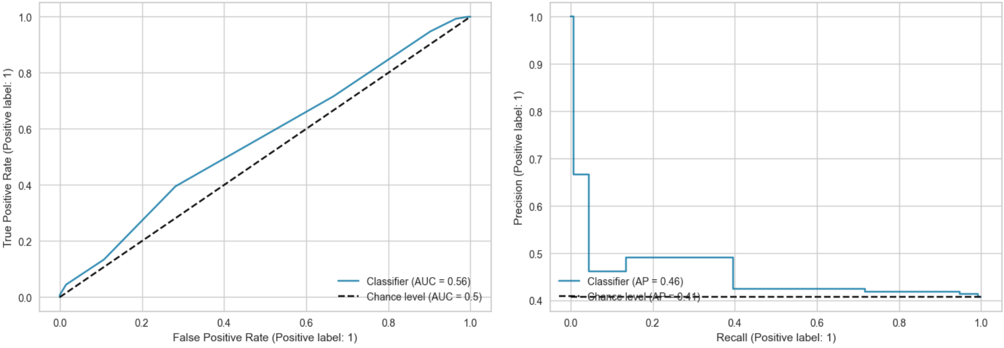

(c) CHA_2D-VASc ROC. (d) CHA_2D-VASc Precision Recall Curve.

Figure 6.10: CHA_2DS-VASc performance.

already use shape-based information in their decision making. The most widely use marker is the body volume [87], which we propose as a first shape-based baseline. This is typically measured with the LA Volume index (LAVi), that is, the LA volume corrected by surface area, and expressed in mL/m^2 [127]. We found that a threshold at 62 mL/m^2 best discriminates control and stroke groups, as seen in table 6.3. It is interesting to note that a LAVi is considered normal under 32 mL/m^2, while we found no patients with values under 32, further highlighting the high risk characteristic of the cohort. More recently, clinicians have used a classification of the LAA into 4 groups, as proposed in [100]. The 4 groups are generally referred to as chicken wing, windsock, cauliflower, and cactus, with chicken wing being the classification considered as low risk, and all others high risk. This method knows some limitations in terms of performance, but also in the choice of group assignment as different radiologist will give different classification. We were still able to have experts classify the LAA of 200 patients from our cohort, which add to the baselines 6.3. Finally, we add the very first method proposed in this thesis from chapter 3.3, that is, the eigenvalues of the Laplace decomposition, also called the Shape-DNA method proposed in [150], which we refer to as *spectrum*. This last addition gives more perspective as to how the specific LA representation we construct significantly helps the performance compared to well established, out-of-the-box methods. Now that we have setup a context of what can be a better stroke predictor, let us see if our approach brings more or better information for the clinician's decision.

	Accuracy	F1	Precision	Recall	AUC
LAVi	.609	.247	.590	.156	.540
LAA *type* [100]	.568	.588	.647	.507	.558
LA *spectrum*	.606	.336	.542	.246	.548
LAA *spectrum*	.655	.426	.677	.315	.621

Table 6.3: Shape-based method baselines.

6.3.2 Initial analysis

To study the link between our LA features and cardioembolic stroke we will conduct both univariate statistical analysis and classification using LGBM. In any cases, we start with a preprocessing step of feature selection.

Feature selection Motivated by finding the best possible shape descriptor, the set of all features computed in the previous section contains a lot of redundancy, as to make sure we were measuring the correct shape *behavior*. To prevent bias and increase performance, we perform a first round of feature selection. To do so we compute the correlation matrix of all features, and group all features that are highly correlated together, that is, with a correlation coefficient bigger than 0.8, as shown in 6.11. This gives 22 groups of features, with a total of 59 features. For each group, we select the feature with the highest Pearson correlation coefficient with the target stroke label. This first sweep leaves us with 66 uncorrelated LA descriptors. Regarding the LGBM classification task, we perform a second round of feature selection, as is the standard practice in machine learning. After testing several methods such as MRMR [45] or factor analysis [99], we select the Recursive Feature Elimination method [69] with Cross-Validation (RFECV), as it generally outperforms the other methods in the following analysis.

Univariate statistical analysis Let us now take a first look into the features. For each of the 66 features we perform a standard student t-test, and consider a feature to be significant when its p-value was under 0.05. Out of all features, 6 are significant predictors of stroke risk, among them 5 are related to the LAA and its ostium, and 1 is related to vein configuration. Precisely, the ratio between the tip and characteristic intersection of the LAA (Tip:CI), is the most significant. We may interpret this feature as measuring how *long* the LAA is after turning. The volume of the LAA (LAA Vol) and the area of the ostium (LAAO Area) also have an impact, the bigger they are the higher the risk, which supports previous knowledge that size impacts stroke occurrence. The following two

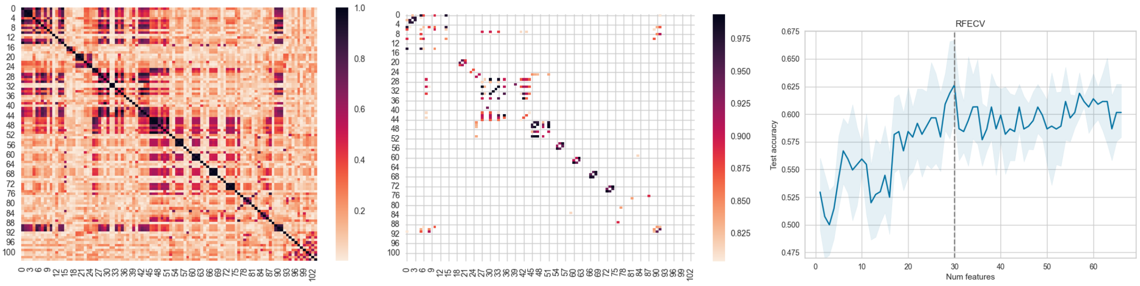

(a) Feature correlation matrix. (b) Highly correlated ($> .8$). (c) RFECV over 5 folds.

Figure 6.11: Correlation between features.

significant features are measures of how irregular the LAA is, first its total absolute curvature (LAA $\int \kappa$) shows that the more local irregularity the higher the risk, second the LAA release factor (LAA RF), shows that, rather counter intuitively, the bigger the LAA expands after the LAAO the smaller the risk. Finally, the last significant feature (IPV:SPV) is the ratio between the length between the Inferior PVs (IPV) and the length between the Superior PVs (SPV), which may be seen as a measure of asymmetry in the LA. Precisely if the IPV length is smaller than the SPV length the risk is higher. The summary is given in table 6.4. While these results might sound attractive, they must be taken with a grain of salt. The gold standard of a p-value < 0.05 does not tell the full story, we can see that for each feature there is a big crossover between both groups. In addition, we can use dimension reduction methods, namely UMap [120], to visualise the distribution of the dataset (see fig. 6.13b), and realise there is not much of a linear separation between them.

features	p-value	statistics	Control mean	Stroke mean
LAA Tip:CI	.007	-2.702	.455 ± .107	.48 ± .131
LAAO Area	.011	-2.546	514 ± 162	564 ± 235
LAA Vol	.013	-2.494	11448 ± 4927	13024 ± 7757
LAA $\int \kappa$	.027	-2.219	197 ± 33	205 ± 37
LAA RF	.028	-2.203	.854 ± .240	.913 ± .286
IPV:SPV	.029	2.181	1.01 ± .093	.995 ± .087

Table 6.4: Significant features.

Non-linear classification To identify a possibly more complex relationship between the LA anatomy and stroke risk, we turn to machine learning methods and perform classification between the stroke and control groups via an LGBM classifier. For this process, all features are normalised and we perform RFECV with 5 folds (see 6.11c) for feature selection, as explained above, giving best performance for 30 features. Regarding hyper-parameters, we set the number of leaves to 100 and the maximum number of bins at 60, all other parameters are kept as default. We perform 5 fold cross validation and report everything as mean test results.

We give the main result in table 6.5. This method yields considerably better balance between all 5 performance measures compared to baselines (see tables 6.2, 6.3) indicating that the model learns

	Accuracy	F1	Precision	Recall	AUC	p-value
LGBM	.686	.591	.637	.554	.666	.014

Table 6.5: Geometric feature results.

meaningful information. The precision-recall and ROC curve (fig. 6.12) enforces this fact as well. The accuracy is slightly better than other baselines although still underwhelming, and the recall – or true positive rate – is low, considering we would rather give early treatment to more people than not enough.

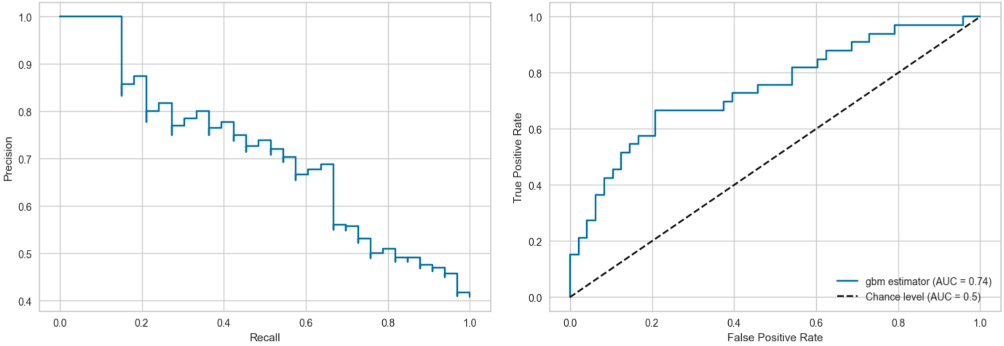

(a) Precision-Recall.

(b) ROC.

Figure 6.12: Geometric features performance curves.

Although initial results are generally underwhelming, we can still focus on model interpretation using the gain importance method and its implementation in the `LGBM` package. This feature importance is defined as the difference in the objective cost function – in our case binary cross-entropy – when training the model with and without a given feature, in other words, the gain in performance each specific feature yield. The gain measure is the feature importance we will use throughout the rest of this chapter. The 4 most important features have considerably more weight than the others ranging from 120 to 160, while the 5th one is under 110 and all other ones have a score of under 100. In contradiction with the univariate study, the most important features are all related to PV configuration rather than the LAA. Precisely, the angle of the LSPV insertion, the total absolute curvature of the RSPV ostium, the area of the RIPV ostium and the angle between left and right PVs, which we refer to as PV alignment, as seen in table 6.13a. Although we can interpret the model to an extent, it is rather dangerous to give too much precision at this point, as the model fails to give a decisive classification between control and stroke groups.

6.3.3 Rhythm counts

Our first study, detailed in chapter 4, highlighted an important factor to take into account when studying cardiac shape, that is, the anatomic non-linear transformation of the heart. Our first assumption was to look at the natural variations from the heart cycle, as the LA knows a non negligible transformation between systolic phase and diastolic phase. As our cohort grew, we realised

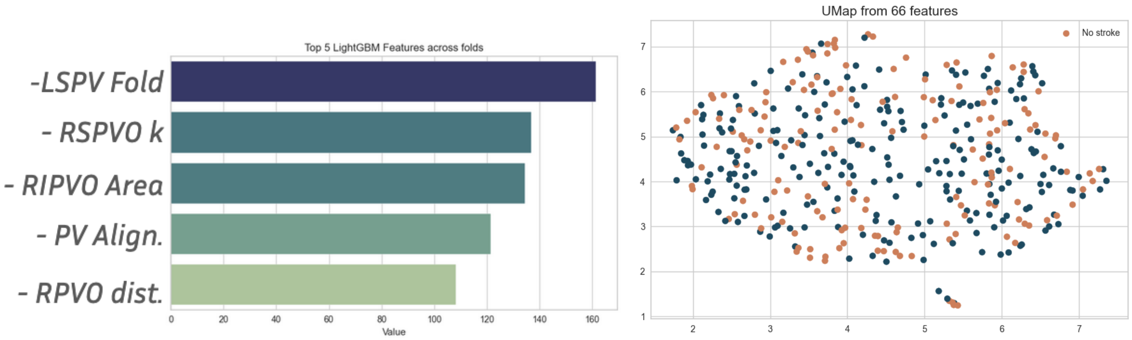

(a) Top 5 important features.

(b) UMap.

Figure 6.13: Feature importance and projection.

that, rather than the systole/diastole split, the state of the AF was a remodelling of bigger importance. Typically, when a patient suffers AF recurrently, the fibrillation will become persistent. This change in AF state forces the LA to undergo a morphological change, a remodelling that, naturally, is picked up by our geometric features. In fact, we can show a distinct difference between the two groups. By conducting student t-tests, we found 25 out of the 66 features to be significant, with p-values ranging from 10^{-17} to 0.0007, we show the first 6 and their related distribution, estimated via the KDE method in fig. 6.14.

features	p-value	statistics
Body Vol	$< 10^{-17}$	-8.65
LAAO small axis	$< 10^{-11}$	-7.28
LAA SA:V	$< 10^{-11}$	7.16
LAA size	$< 10^{-10}$	-6.7
RPVO dist.	$< 10^{-9}$	-6.64
RSPVO area	$< 10^{-9}$	-6.4

(a) Top 6 significant features.

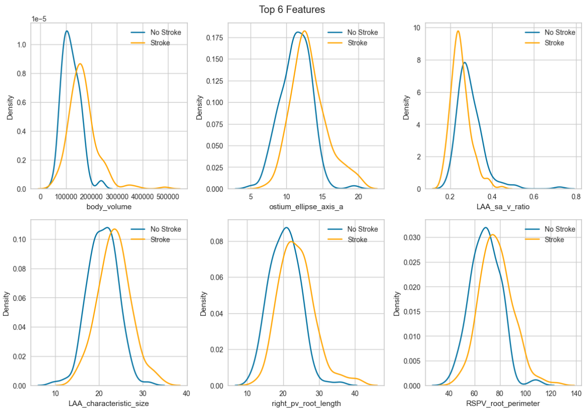

(b) KDE of top 6 features.

Figure 6.14: Geometric features are impacted by AF state.

	Accuracy	F1	Precision	Recall	AUC	p-value
LGBM	.833	.841	.853	.829	.834	0.0

Table 6.6: AF state is well predicted with geometric features and LGBM.

We can go further and predict it via the LGBM model with convincing results, shown in 6.6. As in our initial analysis, we can look at feature importance, see fig. 6.15a. It is clear that the main morphological change is in the size of the Body, however angles of PVs are also impacted, highlighting that the transition between paroxysmal and persistent AF is a non-linear remodelling of the LA. It is also notable that a simple dimension reduction plot can already show a good separation between the two groups, as shown in fig. 6.15.

Having noticed that the heart's remodelling affects our analysis, rather than deciding which group should be separated based on patient information, we may decide to take an unsupervised route and look for an optimal separation. However, after conducting cluster analysis using the k-means method, we can show – via the silhouette score – that the ideal number of cluster is 2 (searched between 2 and 15). Although these 2 clusters are different from the persistent/paroxysmal split, this separation showed less performance in the next steps of our analysis. More so, they are less interpretable than the persistent/paroxysmal split, leading us to assume that it is the most important bias, if not the only one, to take in account in our study. To the best of our knowledge, this information is never incorporated in other studies, and we suggest it should be when attempting to build any predictor from the shape of the LA in the presence of AF.

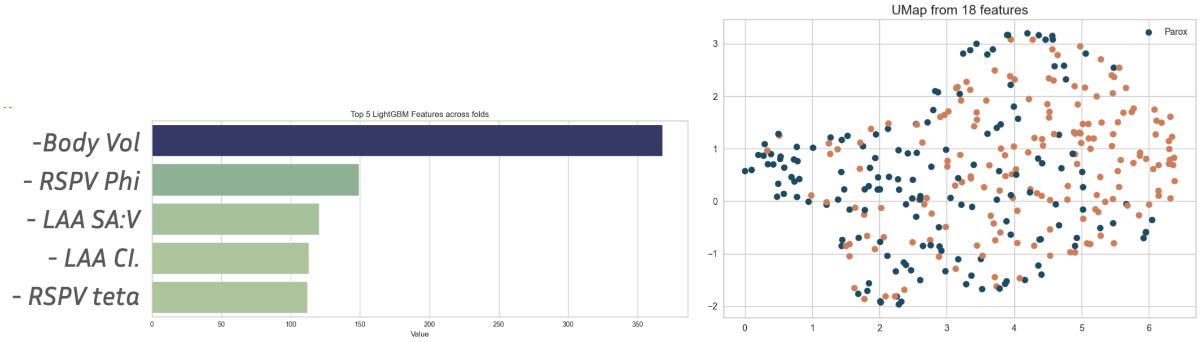

(a) Top 5 important features.

(b) UMap of dataset.

Figure 6.15: Feature importance and visualisation of AF state.

6.3.4 Rhythm aware analysis

With the knowledge that the cohort should be split into the persistent and paroxysmal groups, we can refine our initial analysis. In this section, we conduct the same analysis as previously (section 6.3.2), although done independently for the persistent group and the paroxysmal group. Table 6.7 summarises the data balance between each group.

	Control	Stroke
Persistant	105	70
Paroxysmal	90	63
All	238	166

Table 6.7: Dataset recap.

Univariate statistical analysis Performing the univariate analysis (table 6.8), we can see that considerably more features are significant in the persistent group. Although some initially relevant features remain, specifically LAA Tip:CI, LAAO area and LAA volume, the rest has now disappeared.

However, LAA SA:V can be considered a closely related measure to total absolute curvature ($\int \kappa$) and release factor (RF). Amongst the newly appearing features we can see relations to vein configuration, most notably linked to the RIPV angle, in both the persistent and the paroxysmal cases. The distance between IPVs is also significant, although it is unclear if it is a sign of LA size or LA asymmetry. However, the significance of the ratio between LA body and LAA volumes indicates more clearly a degree of asymmetry – or unbalance – between both part. Once again, although the statistical test may indicate some significance, the strong overlap between control and stroke groups forbids us to make any conclusive remark as to which feature should be selected to predict stroke occurrence. Although we can not yet measure the impact of the LA shape on stroke risk, the lack of features for the paroxysmal group may indicate a problem regarding labels, rather than an absence of connection between shape and stroke. Indeed, it is not always possible to identify the origin of a stroke, as such, the stroke target label may consist of strokes other than cardio-embolic in nature. We suspect that the proportion of patients suffering from a non-cardiac stroke is bigger in the paroxysmal group, in turn explaining the number of significant feature discrepancy between both groups.

Group	Features	p-value	statistics	Control mean	Stroke mean
Persistant	RIPV ϕ	.0004	3.34	-.623 ± .337	-.782 ± .287
	LAA Vol	.006	-2.74	12820 ± 5021	15593 ± 8333
	LAAO Area	.01	-2.59	554 ± 157	634 ± 250
	IPV Dist.	.011	-2.54	52.7 ± 7.32	56.0 ± 9.8
	LAA Tip:CI	.013	-2.48	.442 ± .085	.478 ± .104
	RSPV Fold	.029	-2.19	2.84 ± .164	2.89 ± .151
	LAAO Eccentricity	.036	-2.11	.494 ± 1.79	1.54 ± 4.58
	Body:LAA Vol	.042	-2.04	171965 ± 46852	191127 ± 77210
Paroxysmal	RIPV Fold	.01	-2.59	2.70 ± .251	2.80 ± .193

Table 6.8: Significant features of both groups.

Group	Accuracy	F1	Precision	Recall	AUC	p-value
Persistent	.771	.724	.714	.734	.762	.001
Paroxysmal	.71	.651	.667	.636	.697	.025
AF aware	.741	.677	.691	.685	.73	.013
Initial	.686	.591	.637	.554	.666	.014

Table 6.9: Results on subsets, against initial method.

LGBM classification Having grown our cohort to $\sim$ 400 patients, we can disregard the need for complex learning methods to unify heterogeneous groups, as presented in 4, and rather focus on classifying both groups independently. Once again, we do so by training an LGBM classifier. Although we keep the models hyper-parameters from the previous study, we perform per-group RFECV feature selection. We report the results in table 6.9 and precision-recall and ROC curves in 6.16. As expected, results show a clear improvement, most notably in the persistent group. In any case, the better balance between performance measures highlights a more stable method to predict stroke. We can also build a larger prediction model by combining both LGBM classifier, which we refer to as *AF aware*. As before we perform a feature importance analysis, and show the top 5 most important features in figure 6.17. In contrast with the initial model, the model uses a wider range of geometric features for both group, as it not only focuses on PV features but also takes in account features from the LAA and LAAO. We can note that the right PVs are more important than the left branch, echoing the vein study we proposed in section 6.

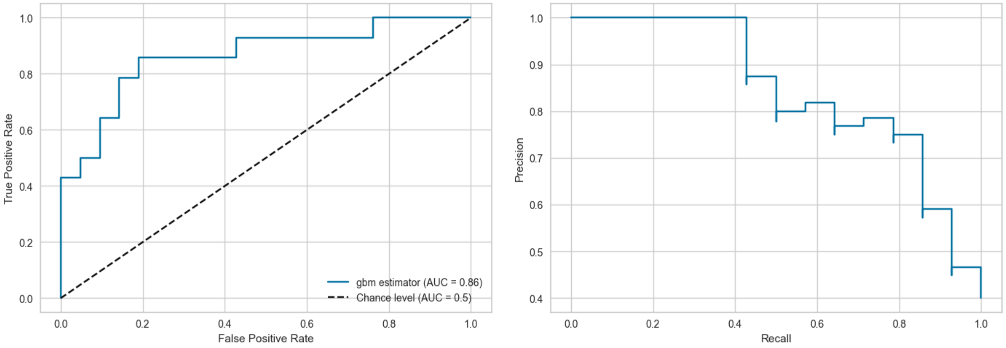

(a) Persistent ROC. (b) Persistent Precision Recall.

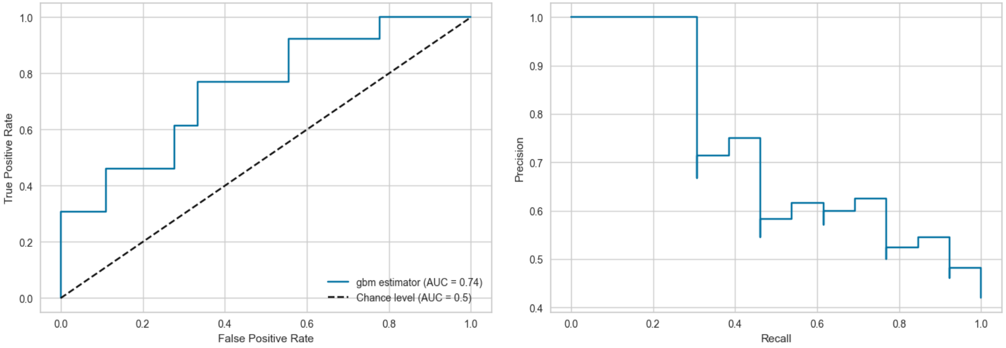

(c) Paroxysmal ROC. (d) Paroxysmal Precision Recall.

Figure 6.16: LGBM performance per subgroup.

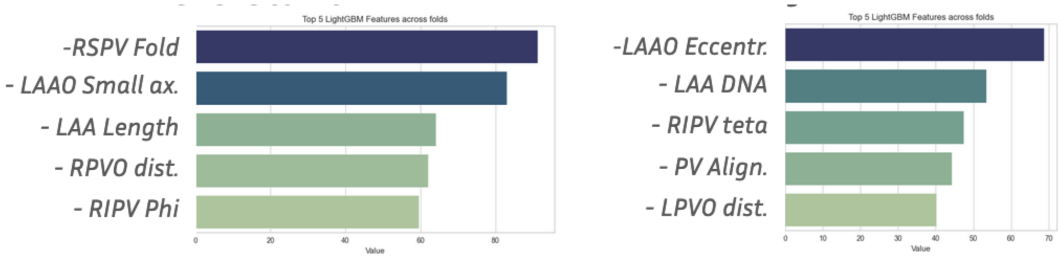

(a) Persistent group. (b) Paroxysmal group.

Figure 6.17: Feature importance per subgroup.

Final benchmarks The final model is a combination of both LGBM models trained independently on the persistent and paroxysmal subsets, which we call *AF aware*, as previously mentioned. In this concluding section, we compare the LGBM AF-aware model with the following baselines, either currently widely used by clinicians or proposed in earlier stages of the thesis:

1. $\text{CHA}_2\text{DS-VASc}$ score, we show both the most realistic case, considering the score before the occurrence of stroke with a decision threshold at 1 (noted $\text{CHA}_2\text{D-VASc}$), and the best case, with stroke accounted for and a threshold at 5.
2. The LAA-classification proposed in [100], with the chicken types being a stroke marker vs the rest.
3. The LAA spectrum method as introduced in the very first analysis of LA shape of this thesis in chapter 3.3. We also add *awareness* of the AF type by combining models trained on the persistant and paroxysmal groups independently.
4. Both our LGBM model and the Shape-DNA method, whithout taking in account the persistant and paroxysmal groups.

The results, given in table 6.10, show that our model, the LGBM aware method, stands out from all other methods in terms of performance, while having a similar recall to the best version of the $\text{CHA}_2\text{D}\mathbf{S}\text{-VASc}$ score. We also believe results will improve when conducting a bigger prospective study with a wider range of patients.

Method	Accuracy	F1	Precision	Recall	AUC
$\text{CHA}_2\text{D-VASc}$	.395	.552	.383	**.992**	.513
$\text{CHA}_2\text{D}\mathbf{S}\text{-VASc}$	.708	.659	.628	**.684**	.705
LAA *type* [100]	.568	.588	.647	.507	.558
LAA *spectrum*	.655	.426	.677	.315	.621
LAA *spectrum* (aware)	.673	.528	.643	.461	.648
LGBM	.686	.591	.637	.554	.666
LGBM (aware)	**.741**	**.677**	**.691**	**.685**	**.73**

Table 6.10: Final benchmarking.

6.3.5 A study of Opacification defect

The previous section focuses on the most straightforward way to study the relationship between the LA anatomy – in the form of features – and cardioembolic stroke. That is, by measuring the predictive power of the geometric features with respect too the stroke label. However, as explained earlier, the available cohort used to conduct this analysis is considered at high risk of stroke, as all patients are screened for a potential LAA occlusion procedure. This can also be highlighted by high mean $\text{CHA}_2\text{DS-VASc}$ score, where all patients have a minimum score of 1, as well as a high LAVi. While a larger prospective study will clarify the relationship between LA anatomy and cardioembolic events, we propose to follow a secondary study by focusing on a softer objective. Radiologists can detect the presence of blood stagnation in the LAA by defining it as a large enough observed opacification defect on the image. In fact, when the opacification defect is persisting through the arterial and venous phase, the blood stagnation is considered to be an LAA thrombus. This defect on the arterial phase is an important indicator on the haemodynamic environment in the LA, which can be a valuable information when attempting to prevent strokes. In fact, is has been shown that this filling defect at the early phase of a CT-scan is an important marker of cardioembolic stroke [113, 78].

We suggest that the filling defects can be considered a relevant surrogate of stroke occurrence to predict. Because we are studying a retrospective cohort, we only have access to past stroke history, sometimes occurring more than a decade before the CT-scan was taken. In which case the LA may have been considerably remodelled since the embolic event. Moreover, patients classified as *stroke negative* may experience stroke in the week following the CT-scan. In contrast, the filling defect is a marker readily accessible at the time of the CT scan, making it a faithful representation when focusing on the LA anatomy. We aim to see if the geometric features have any predictive power with respect to this opacification defect, giving additional insight into the impact of the shape on the haemodynamic environment, and by extension stroke occurrence.

Detection Although the presence of an opacification defect on arterial phase for each patient was collected in the Bordeaux cohort, the Toulouse Hospital did not systematically report it. However, the final version of CT scan segmentation model, introduced in chapter 2, makes it possible not only to segments the LA shape, but also to label the defect region. To distinguish between small defects – possibly due to trabeculations – and the presence of a blood stasis, we set a threshold on the total volume of the labelled opacification defect region. To define the threshold, we perform 5 fold cross validation on the Bordeaux dataset, and found 1481mm^3 to be the best value with results reported in table 6.11. This automatic detection method labelled 42% of the population as suffering from blood stasis, against the 38% defined by experts. We then applied this threshold to the Toulouse segmentations to define an blood stasis label on the entire cohort.

	Accuracy	f1	Precision	Recall	AUC
nnU-Net	.860	.838	.814	.865	.860

Table 6.11: nnU-Net opacification defect detection with a threshold at 1481 mm^3.

Prediction After successfully detecting the presence of blood stagnation in the LAA on the entire cohort, we employ the same analysis pipeline as previously, to understand if the geometric features of the LA can predict its presence. Results from univariate analysis show that 32 out of the 66 features are statistically significant. Of them, the most predictive are all directly related to the LAA and its relationship with the LA body as shown in table 6.18a and figure 6.18b.

	Accuracy	f1	Precision	Recall	AUC	p-value
LGBM	.838	.794	.781	.806	.832	2e-5

Table 6.12: Blood stasis prediction from geometric features.

The LGBM classification model also show positive results (table 6.12), with a mean test accuracy of .838. Interestingly the feature importance analysis shows that the LAA SA:V is by far the most

features	p-value	t-statistic
LAA Size	$< 10^{-28}$	-12.2
LAA SA:V	$< 10^{-22}$	10.4
Body:LAA Vol.	$< 10^{-20}$	-9.72
LAAO rad.	$< 10^{-18}$	-9.2
Body Vol.	$< 10^{-18}$	-9.03
LAAO area	$< 10^{-15}$	-8.32

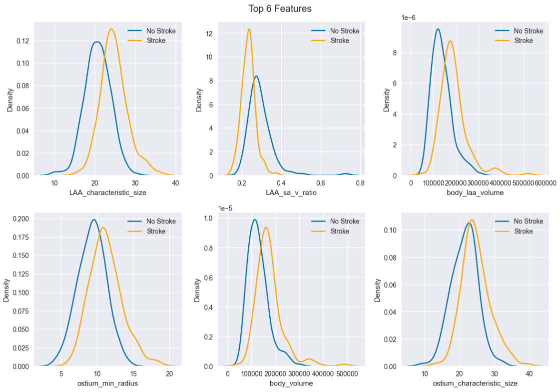

(a) Top 6 significant features.

(b) KDE of top 6 features.

Figure 6.18: Blood stasis prediction statistical results.

important (fig. 6.19), indicating that this quantity could be a better method for patient diagnosis than any known LAA classification. In addition, this quantity can be precisely measured while typical LAA classification strongly rely on the clinicians appreciation. It is also interesting to note that the SA:V is measure that is often use in biology [61] and physics [144]. We look forward to see further results in broader studies regarding this feature. Finally, a 2D dimension reduction visualisation via the tSNE method also shows good separation between both groups.

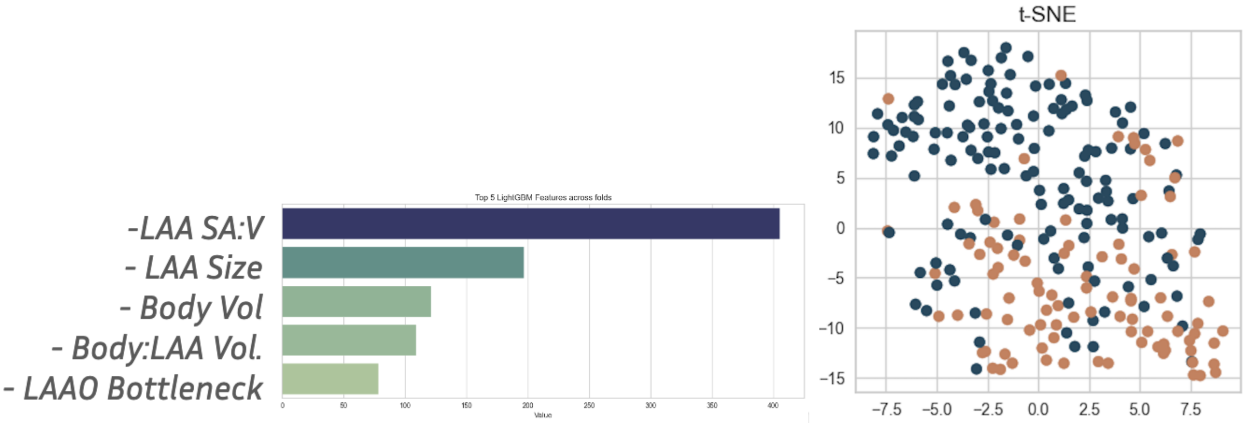

(a) Feature importance.

(b) tSNE.

Figure 6.19: Blood stasis feature importance and visualisation.

6.4 Conclusion

In this final chapter, we first completed the pipeline we set out to build from the beginning, through the direct computation of geometric features of the LA shape. This pipeline enabled us to stay close to clinicians expectations, and constitutes a tool that paves the way to a better understanding of the relation between LA shape and cardioembolic stroke. In the second half, we have realised how this may be the case, by conducting both statistical analysis and supervised classification of a cohort constituted from patients of both the Bordeaux and Toulouse Hospital. This analysis lead us to discover several important fact. First, the heterogeneity of the dataset, precisely the one induced by the remodelling of the LA due to AF state, is an important factor to take in account when studying such cohorts. Second, and perhaps most central, the shape of the LA appears to have an important impact on the occurrence of stroke. In fact, we were able to predict stroke considerably better than the $\text{CHA}_2\text{DS-VASc}$, the most widely used score by clinicians today. In addition, we showed that a strong indicator of Thrombosis, the presence of blood stasis in the LAA, can be efficiently predicted by our geometric features, and we identified the LAA SA:V as the predominant predictor of such phenomenon. However, while more interpretable than any other previously suggested method, we were not able to precisely define a fixed set of features that predict the occurrence of stroke. Even more, any interpretation we gave in this chapter should be taken with a grain of salt, as the study has been conducted on a particularly at-risk cohort of patients, and may not be relevant to the total population potentially screened for such risks.

As we are aware that this cohort may not be clinically meaningful for stroke prevention, it successfully served as material to develop a fully automated method to extract stroke risk features from CT scans. These can now be applied in large prospective studies on AF patients in which the risk/benefit ratio of preventive medicine is still unclear. With the overall positive results, and the growing availability of CT-scans, we hope such future prospective studies may help better identify both the precise features that are most predictive of stroke, and their direct impact on blood flow, in order to fully understand the mechanistic impact of cardioembolic strokes.

Chapter 7

Conclusion and Perspectives

Throughout this thesis, we worked towards answering the medical question:

> **What is the relationship between the anatomy of the Left Atrium (LA) and the occurrence of cardioembolic strokes in Atrial Fibrillation (AF) patients?**

This question, although grounded in real world problems, hides the methodological challenge of measuring shapes and their differences. Such questions, as we have seen, have been addressed within numerous domains, ranging from purely theoretical to purely practical.

7.1 Main contributions

The open-endedness of the question allowed us to explore a significantly large spectre of existing methods instead of exclusively focusing on one. We have dedicated the chapter 3 to this. This allowed us to better understand the underlying problems of studying the LA shape, and fix a clear goal as to how we should answer such problems. That is, any tools we develop need to: (1) yield an expressive representation of the shape; (2) be easily interpretable by clinicians; and (3) should scale to potentially thousands of images. In other words, tools should be relevant to the clinical world. As so, known *out-of-the-box* methods for global shape analysis are not well adapted to the problem. It turns out the variation of shapes they measure are not expressive enough and the results are difficult to interpret. Thus, we decided to build methods specific to the LA aiming to answer any problem related to it, rather than applying general methods to specific clinical problems. In chapter 4 we propose a first version of such tools, in which we realise labeling the anatomical parts of the LA is the first step towards building expressive representation. We also saw that the heterogeneity of the dataset stemming from the nature of the heart should be taken in account. This stage allowed us to confirm that mesh labeling followed by geometric feature extraction was the way to go in order to respond to all our problems. However, to make the pipeline fast enough for larger scale applications, we needed to address the computational cost of mesh labeling, as this was the crux of the process. We dedicated the chapter 5 to this particular problem, and proposed an improvement to neural network methods designed for surfaces, in the form of a better representation of the input shapes via principal curvature. With these new techniques, we were able to bring the anatomic labeling procedure down from around $40min$ to around $5sec$ per mesh, along with better performance, and minimal preprocessing. Armed with a way to label LA meshes efficiently, we were able to focus on the final step of our pipeline, extract geometric features, proposed in chapter 6. With now total control over the shapes, we took full advantage of the anatomical labels and extracted 104 geometric features (with some redundancies) that measure quantities of anatomic parts such as volume, orientations, skewness or relative configurations between the different parts of the LA, etc. This final set of features representing the LA shape marks the endpoint of the automatic tools we proposed to tackle the initial clinical problem. In summary, we have constructed a pipeline for translating segmentation masks of the LA, to shape features, as shown in 7.1, that takes under 20 seconds to compute on a laptop computer. Our pipeline can be adjoined to an automatic CT-scan segmentation model to make it completely end-to-end, with the added value that an expert can be in the loop at each step of the process. In contrast to other existing approaches that entirely rely on modern deep learning methods, albeit yielding promising results, we use these methods punctually and for the purpose of automating tedious tasks, which are always verifiable by clinicians.

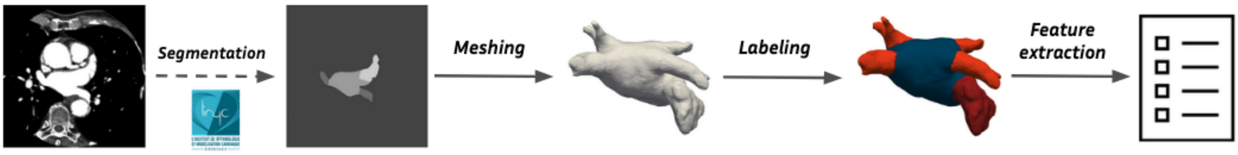

Figure 7.1: Overview of the final end-to-end pipeline, from CT-scan to geometric features.

Thanks to the data gathered from the Hospitals of Bordeaux and Toulouse, we were able to apply our pipeline to a cohort of 404 AF patients, and give elements of responses to the clinical problem of interest. Initial attempts at direct classification of geometric features between stroke and control patients, although better than the current widely used methods, have not yield overwhelming results. However, we recognised that the state of the AF (i.e. persistent or paroxysmal) has a big impact on the shape features. After measuring this impact, we derived an AF state-aware classification by simply splitting the full dataset in subgroups, and realised significantly better results. When interpreting the results, we saw that Pulmonary Veins (PV) are as important as the LAA in the occurrence of cardioembolic stroke, which contrasts with the classical lines of research focusing solely on the LAA. Finally, we realised that the geometric features we have extracted are excellent predictors of blood stasis in the LAA, itself being an important marker of cardioembolic events. This prediction is mostly explained via the Surface Area to Volume ratio (SA:V) of the LAA, which may be used alone, and is easily computable by cardiac radiologists directly. Overall, we believe that we have successfully demonstrated that a link exists between the LA shape and the occurrence of stroke, and have given elements of answers as to what part of the shape has a mechanistic impact on strokes. In addition, our pipeline is specific to the LA but can be used to study any problem linked to it. With this in mind, we can look at the limitations and perspectives of this thesis.

7.2 Limitations

Although several limitations have been underlined in detailed parts of this work, we can highlight three main limitations.

Scope of the cohort. The entirety of this thesis is considered a preliminary study, aiming to give a first answer to an intuition from clinical experts. This entails the use of a retrospective dataset collected for a separate purpose. In fact, the cohorts from both the Bordeaux and Toulouse hospitals, were constituted of patients undergoing an LAA occlusion procedure, making it a *harsh* dataset of patients at high risk of stroke which were drug-refractory with additional risk of major bleeding. This translates into a limited range of subjects, especially control patients, and gives a skewed view on who gets screened for stroke risk. Addressing this limitations implies a full prospective study with a wider range of patients.

Interpretation. Although we have dedicated a considerable amount of time to keep the methods close to clinicians' expectations, we were still unable to precisely pinpoint specific bio-markers of stroke and rarely saw one feature stand out as most important. This is possibly due to a complex relationship between geometric features, and how they interact with each other. In addition, the dataset being only composed of high risk patients made it difficult to classify, and in turn difficult to identify features that are most important. As it is the case for the previous limitation, access to larger cohort will help solving this issue.

Pipeline automation. Although the pipeline is well on its way to be used in real-world and real time scenarios, an expert in the loop is still required to correct errors from the automatic image segmentation and the anatomical labeling. Regarding the image segmentation, we have briefly talked about ways to detect the most occurring errors that manifest in the form of a hole in the LAA, or an Left Superior PV (LSPV) glued to the LAA. While we can detect the presence of either one, we do not discriminate between them, and automatic methods to repair them differ. We could however fix these problems automatically, as holes in the LA are considered *handles* and glued LSPV-LAA are considered *tunnels* of the mesh, and methods exist to identify them [12]. Regarding the semantic segmentation of the mesh, we realise that all problems are regarding the identification of LAA ostium, this is a notoriously hard part of the anatomy to define, and to this day, there are no ways to automatically repair it. However, we believe the straightforward solution would be to have access to more data to improve the quality of the semantic segmentation.

With these limitations in mind, let us define some important future directions to this work.

7.3 Perspectives

As in most exploratory studies, the number of follow up questions are greater than the ones answered. We conclude this thesis by outlining a handful of possible future work that could follow this thesis.

Haemodynamic simulations The interpretation concerns we expressed in the limitations, were linked to better understanding the interaction between the shape, in its full complexity, and the mechanistic occurrence of stroke. This interaction can be described by analysis the blood flow through the prism of the LA shape. In other words, perform a haemodynamic simulation study of the impact of the geometric features on flow quality. Thankfully, the pipeline we proposed is based on automatic control of the shape of the LA. We can leverage on the labels and features we obtained to perform precise handle-based deformations with techniques such as proposed in [157] and [168], yielding a wide range of possible deformations, of which we have already implemented a considerable amount, as seen in figure 7.2. In addition, we can generate cardiac motion for any LA shape thanks to the parallel transport methods showcased in chapter 3.4.7. The combination of both tools would enable us to generate an entire cohort of dynamic meshes to pursue a sensitivity analysis of the impact of features on haemodynamic flow, similar to [50]. We are looking forward to making this study possible thanks to collaboration with teams from UPF Barcelona and Simula, that have developped robust Haemodynamic models in the context of the PARIS project.

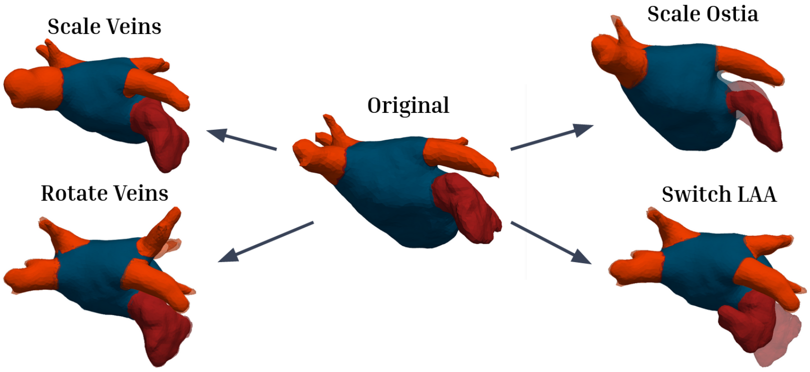

Figure 7.2: Modifying the shape via mesh processing.

Finer details The geometric features we captured describes the major parts of the LA shape. Although most of the main variation in the shapes are captured, it may be interesting to add information at smaller scales. For instance, some LA anatomies contain small veins in addition to the main four pulmonary veins. We could easily adapt our feature extraction process to accommodate for them by introducing them in the anatomical labelling framework. Looking at an even finer scale, representing the LA as a surface constitutes a simplification of its anatomy. It does not account for the LA wall or the trabeculations in the LAA (see figure 7.3), yet they can have an important haemodynamic impact. Recent advances in medical imaging, such as novel photon counting CT systems, would allow us to capture elements at such finer scales.

(a) Trabeculae scheme, taken from [2].

(b) Segmented 20μm image of the LAA.

Figure 7.3: Advances in medical imaging will allow us to capture finer details of the shapes.

Multi-modality As in any complex disease, a complete solution will never be found in a single place. The heart's function rely on many moving parts, and understanding stroke occurrence involves understanding how those moving parts interact with each other. While we focused on the shape of the LA, our pipeline should be one facet of a larger set of information, extracted from other modalities, allowing clinicians to better diagnose patients. Such modalities can include electric signal from 12-lead ECGs, 3D motion from Cine-MRIs, blood flow from ultra-sounds, and many more. Thankfully, the results from a combination of studies, including this one, has led to a new project: The **RHU Talent**. It will incorporate multiple modalities in a prospective study, encompassing a larger number and wider range of patients, with better follow-up. We believe it is the natural direction to take following the work of this thesis.

Bibliography

[1] Adobe: Adobe mixamo 3d characters. Retrieved from Mixano.com (2016) 84

[2] Alkhouli, M., Di Biase, L., Natale, A., Rihal, C.S., Holmes, D.R., Asirvatham, S., Bartus, K., Lakkireddy, D., Friedman, P.A.: Nonthrombogenic roles of the left atrial appendage: Jacc review topic of the week. Journal of the American College of Cardiology **81**(11), 1063–1075 (2023) 116

[3] Alshehri, A.M.: Stroke in atrial fibrillation: Review of risk stratification and preventive therapy. Journal of family & community medicine **26**(2), 92 (2019) 4

[4] Amid, E., Warmuth, M.K.: Trimap: Large-scale dimensionality reduction using triplets. arXiv preprint arXiv:1910.00204 (2019) 33

[5] Amor, B.B., Arguillère, S., Shao, L.: Resnet-lddmm: advancing the lddmm framework using deep residual networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(3), 3707–3720 (2022) 77

[6] Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM SIGGRAPH 2005 Papers. pp. 408–416. Association for Computing Machinery (2005) 84

[7] Antelmi, L., Ayache, N., Robert, P., Lorenzi, M.: Sparse multi-channel variational autoencoder for the joint analysis of heterogeneous data. In: International Conference on Machine Learning. pp. 302–311. PMLR (2019) 66

[8] Antonietti, P.F., Manzini, G., Mazzieri, I., Scacchi, S., Verani, M.: The conforming virtual element method for polyharmonic and elastodynamics problems: a review. The virtual element method and its applications pp. 411–451 (2022) 56

[9] Arnol'd, V.I.: Mathematical methods of classical mechanics, vol. 60. Springer Science & Business Media (2013) 28

[10] Aronszajn, N.: Theory of reproducing kernels. Transactions of the American mathematical society **68**(3), 337–404 (1950) 39

[11] Attene, M.: A lightweight approach to repairing digitized polygon meshes. The visual computer

26, 1393–1406 (2010) 20

[12] Attene, M., Campen, M., Kobbelt, L.: Polygon mesh repairing: An application perspective. ACM Computing Surveys (CSUR) **45**(2), 1–33 (2013) 114

[13] Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: 2011 IEEE international conference on computer vision workshops (ICCV workshops). pp. 1626–1633. IEEE (2011) 77

[14] Augustin, C.M., Fastl, T.E., Neic, A., Bellini, C., Whitaker, J., Rajani, R., O'Neill, M.D., Bishop, M.J., Plank, G., Niederer, S.A.: The impact of wall thickness and curvature on wall stress in patient-specific electromechanical models of the left atrium. Biomechanics and modeling in mechanobiology **19**(3), 1015–1034 (2020) 43

[15] Bai, W., Chen, Z., Tang, H., Wang, H., Cheng, W., Rao, L.: Assessment of the left atrial appendage structure and morphology: comparison of real-time three-dimensional transesophageal echocardiography and computed tomography. The international journal of cardiovascular imaging **33**, 623–633 (2017) 6, 90

[16] Benamou, J.D., Brenier, Y.: A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. Numerische Mathematik **84**(3), 375–393 (2000) 49

[17] Benn, J., Marsland, S.: The measurement and analysis of shapes: An application of hydrodynamics and probability theory. Annals of Global Analysis and Geometry **62**(1), 47–70 (2022) 51, 53, 59

[18] Benn, J., Marsland, S., McLachlan, R.I., Modin, K., Verdier, O.: Currents and finite elements as tools for shape space. Journal of Mathematical Imaging and Vision **61**, 1197–1220 (2019) 50, 77

[19] Bercovich, E., Javitt, M.C.: Medical imaging: from roentgen to the digital revolution, and beyond. Rambam Maimonides medical journal **9**(4) (2018) 7

[20] Berman, H.M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., Bourne, P.E.: The protein data bank. Nucleic acids research **28**(1), 235–242 (2000) 82, 83, 85

[21] Bock, A., Cotter, C.: A note on error analysis for a nonconforming discretisation of the tri-helmholtz equation with singular data. arXiv preprint arXiv:2101.04181 (2021) 56

[22] Bock, A., Cotter, C.J., Kirby, R.C.: Planar curve registration using bayesian inversion. Computers & Mathematics with Applications **159**, 155–172 (2024) 56

[23] Bogiatzi, C., Hackam, D.G., McLeod, A.I., Spence, J.D.: Secular trends in ischemic stroke subtypes and stroke risk factors. Stroke **45**(11), 3208–3213 (2014) 3

[24] Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh

registration. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3794–3801 (2014) 84

[25] Bône, A., Louis, M., Martin, B., Durrleman, S.: Deformetrica 4: an open-source software for statistical shape analysis. In: Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings. pp. 3–13. Springer (2018) 42, 64

[26] Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Lévy, B.: Polygon mesh processing. CRC press (2010) 19, 20

[27] Boyd, J.P.: Chebyshev and Fourier spectral methods. Courier Corporation (2001) 54

[28] Brehmer, J., De Haan, P., Behrends, S., Cohen, T.S.: Geometric algebra transformer. Advances in Neural Information Processing Systems **36** (2024) 77, 78

[29] Bronstein, M.M., Bruna, J., Cohen, T., Veličković, P.: Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. arXiv preprint arXiv:2104.13478 (2021) 65

[30] Burns, K.J., Vasil, G.M., Oishi, J.S., Lecoanet, D., Brown, B.P.: Dedalus: A flexible framework for numerical simulations with spectral methods. Physical Review Research **2**(2), 023068 (2020) 54

[31] Cazals, F., Pouget, M.: Estimating differential quantities using polynomial fitting of osculating jets. Computer aided geometric design **22**(2), 121–146 (2005) 81

[32] Charon, N., Trouvé, A.: The varifold representation of nonoriented shapes for diffeomorphic registration. SIAM journal on Imaging Sciences **6**(4), 2547–2580 (2013) 40, 48, 77

[33] Ciarlet, P.G.: The finite element method for elliptic problems. SIAM (2002) 55, 56

[34] Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19. pp. 424–432. Springer (2016) 17

[35] Cohen-Steiner, D., Morvan, J.M.: Restricted delaunay triangulations and normal cycle. In: Proceedings of the nineteenth annual symposium on Computational geometry. pp. 312–321 (2003) 80

[36] Coorey, G., Figtree, G.A., Fletcher, D.F., Snelson, V.J., Vernon, S.T., Winlaw, D., Grieve, S.M., McEwan, A., Yang, J.Y.H., Qian, P., et al.: The health digital twin to tackle cardiovascular disease—a review of an emerging interdisciplinary field. NPJ digital medicine **5**(1), 126 (2022) 2

[37] Corti, M., Zingaro, A., Quarteroni, A.M., et al.: Impact of atrial fibrillation on left atrium haemodynamics: A computational fluid dynamics study. Computers in Biology and Medicine

150, 106143 (2022) 43

[38] Crane, K., Weischedel, C., Wardetzky, M.: Geodesics in heat: A new approach to computing distance based on heat flow. ACM Transactions on Graphics (TOG) **32**(5), 1–11 (2013) 32, 97

[39] Cresti, A., Camara, O.: Left atrial thrombus—are all atria and appendages equal? Cardiac Electrophysiology Clinics **15**(2), 119–132 (2023) 90

[40] D'Arcy, W.T.: On growth and form. Cambridge Univ. Press **1**(6), 7 (1942) 34

[41] Dassi, F., Mora, D., Reales, C., Velásquez, I.: Mixed variational formulations of virtual elements for the polyharmonic operator (- δ) n. Computers & Mathematics with Applications **158**, 150–166 (2024) 56

[42] De Haan, P., Weiler, M., Cohen, T., Welling, M.: Gauge equivariant mesh cnns: Anisotropic convolutions on geometric graphs. arXiv preprint arXiv:2003.05425 (2020) 77

[43] Di Biase, L., Santangeli, P., Anselmino, M., Mohanty, P., Salvetti, I., Gili, S., Horton, R., Sanchez, J.E., Bai, R., Mohanty, S., et al.: Does the left atrial appendage morphology correlate with the risk of stroke in patients with atrial fibrillation? results from a multicenter study. Journal of the American College of Cardiology **60**(6), 531–538 (2012) 7, 8

[44] Dijkstra, E.W.: A note on two problems in connexion with graphs. Numerische Mathematik **1**(1), 269–271 (1959) 32

[45] Ding, C., Peng, H.: Minimum redundancy feature selection from microarray gene expression data. Journal of bioinformatics and computational biology **3**(02), 185–205 (2005) 101

[46] Do Carmo, M.P., Flaherty Francis, J.: Riemannian geometry, vol. 2. Springer (1992) 25

[47] Dobrzynski, C.: MMG3D: User guide. Ph.D. thesis, Inria (2012) 20

[48] Donal, E., Lip, G.Y., Galderisi, M., Goette, A., Shah, D., Marwan, M., Lederlin, M., Mondillo, S., Edvardsen, T., Sitges, M., et al.: Eacvi/ehra expert consensus document on the role of multi-modality imaging for the evaluation of patients with atrial fibrillation. European Heart Journal–Cardiovascular Imaging **17**(4), 355–383 (2016) 7

[49] Dougherty, E.R., Lotufo, R.A.: Hands-on morphological image processing, vol. 59. SPIE press (2003) 22

[50] Dueñas-Pamplona, J., Garcia, J.G., Castro, F., Munoz-Paniagua, J., Goicolea, J., Sierra-Pallares, J.: Morphing the left atrium geometry: A deeper insight into blood stasis within the left atrial appendage. Applied Mathematical Modelling **108**, 27–45 (2022) 115

[51] Durán, E., García-Villalba, M., Martínez-Legazpi, P., Gonzalo, A., McVeigh, E., Kahn, A.M., Bermejo, J., Flores, O., Del Álamo, J.C.: Pulmonary vein flow split effects in patient-specific simulations of left atrial flow. Computers in biology and medicine **163**, 107128 (2023) 6, 90

[52] Durrleman, S., Prastawa, M., Charon, N., Korenberg, J.R., Joshi, S., Gerig, G., Trouvé, A.: Morphometry of anatomical shape complexes with dense deformations and sparse parameters. NeuroImage **101**, 35–49 (2014) 39

[53] Elliott, A.D., Middeldorp, M.E., Van Gelder, I.C., Albert, C.M., Sanders, P.: Epidemiology and modifiable risk factors for atrial fibrillation. Nature Reviews Cardiology **20**(6), 404–417 (2023) 4, 5

[54] van der Endt, V.H., Milders, J., Penning de Vries, B.B., Trines, S.A., Groenwold, R.H., Dekkers, O.M., Trevisan, M., Carrero, J.J., van Diepen, M., Dekker, F.W., et al.: Comprehensive comparison of stroke risk score performance: a systematic review and meta-analysis among 6 267 728 patients with atrial fibrillation. Europace **24**(11), 1739–1753 (2022) 5

[55] Federer, H.: Curvature measures. Transactions of the American Mathematical Society **93**(3), 418–491 (1959) 80

[56] Federer, H.: Geometric measure theory. Springer (2014) 48

[57] Feigin, V.L., Brainin, M., Norrving, B., Martins, S., Sacco, R.L., Hacke, W., Fisher, M., Pandian, J., Lindsay, P.: World stroke organization (wso): global stroke fact sheet 2022. International Journal of Stroke **17**(1), 18–29 (2022) 3

[58] Feydy, J.: Analyse de données géométriques, au delà des convolutions. Ph.D. thesis, Université Paris-Saclay (2020) 49

[59] García-Isla, G., Olivares, A.L., Silva, E., Nuñez-Garcia, M., Butakoff, C., Sanchez-Quintana, D., G. Morales, H., Freixa, X., Noailly, J., De Potter, T., et al.: Sensitivity analysis of geometrical parameters to study haemodynamics and thrombus formation in the left atrial appendage. International journal for numerical methods in biomedical engineering **34**(8), e3100 (2018) 6, 90

[60] Giorgi, D., Biasotti, S., Paraboschi, L.: Shape retrieval contest 2007: Watertight models track. SHREC competition **8**(7), 7 (2007) 84

[61] Glazier, D.S.: A unifying explanation for diverse metabolic scaling in animals and plants. Biological Reviews **85**(1), 111–138 (2010) 96, 110

[62] Go, A.S., Hylek, E.M., Phillips, K.A., Chang, Y., Henault, L.E., Selby, J.V., Singer, D.E.: Prevalence of diagnosed atrial fibrillation in adults: national implications for rhythm management and stroke prevention: the anticoagulation and risk factors in atrial fibrillation (atria) study. Jama **285**(18), 2370–2375 (2001) 4

[63] Gower, J.C.: Generalized procrustes analysis. Psychometrika **40**, 33–51 (1975) 41

[64] Grenander, U.: General pattern theory: a mathematical theory of regular structures. Journal of the Royal Statistical Society-Series A Statistics in Society **158**(3), 629 (1995) 34

[65] Guggenheimer, H.W.: Differential geometry. Courier Corporation (2012) 79

[66] Guigui, N., Moceri, P., Sermesant, M., Pennec, X.: Cardiac motion modeling with parallel transport and shape splines. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). pp. 1394–1397. IEEE (2021) 44

[67] Guigui, N., Pennec, X.: Numerical accuracy of ladder schemes for parallel transport on manifolds. Foundations of Computational Mathematics **22**(3), 757–790 (2022) 44

[68] Guigui, N., Pennec, X.: Parallel transport, a central tool in geometric statistics for computational anatomy: Application to cardiac motion modeling. In: Handbook of Statistics, vol. 46, pp. 285–326. Elsevier (2022) 44

[69] Guyon, I., Weston, J., Barnhill, S., Vapnik, V.: Gene selection for cancer classification using support vector machines. Machine learning **46**, 389–422 (2002) 101

[70] Ham, D.A., Kelly, P.H.J., Mitchell, L., Cotter, C.J., Kirby, R.C., Sagiyama, K., Bouziani, N., Vorderwuelbecke, S., Gregory, T.J., Betteridge, J., Shapero, D.R., Nixon-Hill, R.W., Ward, C.J., Farrell, P.E., Brubeck, P.D., Marsden, I., Gibson, T.H., Homolya, M., Sun, T., McRae, A.T.T., Luporini, F., Gregory, A., Lange, M., Funke, S.W., Rathgeber, F., Bercea, G.T., Markall, G.R.: Firedrake User Manual. Imperial College London and University of Oxford and Baylor University and University of Washington, first edition edn. (5 2023). https://doi.org/10.25561/104839 56

[71] Hankey, G.J.: Population impact of potentially modifiable risk factors for stroke. Stroke **51**(3), 719–728 (2020) 3

[72] Hanocka, R., Hertz, A., Fish, N., Giryes, R., Fleishman, S., Cohen-Or, D.: Meshcnn: a network with an edge. ACM Transactions on Graphics (ToG) **38**(4), 1–12 (2019) 77, 78

[73] Hartman, E., Sukurdeep, Y., Klassen, E., Charon, N., Bauer, M.: Elastic shape analysis of surfaces with second-order sobolev metrics: a comprehensive numerical framework. International Journal of Computer Vision **131**(5), 1183–1209 (2023) 35

[74] Hospedales, T., Antoniou, A., Micaelli, P., Storkey, A.: Meta-learning in neural networks: A survey. arXiv e-prints pp. arXiv–2004 (2020) 66

[75] Huang, X.: Nonconforming virtual element method for 2 m th order partial differential equations in r n with m> n. Calcolo **57**(4), 42 (2020) 56

[76] Inci, H., Kappeler, T., Topalov, P.: On the regularity of the composition of diffeomorphisms. arXiv preprint arXiv:1107.0488 (2011) 36

[77] Inogés, M., Arboix, A., García-Eroles, L., Sánchez-López, M.J.: Gender predicts differences in acute ischemic cardioembolic stroke profile: Emphasis on woman-specific clinical data and early outcome—the experience of sagrat cor hospital of barcelona stroke registry. Medicina

60(1), 101 (2024) 3

[78] Inoue, T., Suematsu, Y.: Poor enhancement pattern of left atrial appendage in cardiac computed tomography is associated with stroke in persistent atrial fibrillation patients. Journal of Thoracic Disease **11**(8), 3315 (2019) 109

[79] Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021) 18

[80] Isensee, F., Jäger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: Automated design of deep learning methods for biomedical image segmentation. arXiv preprint arXiv:1904.08128 (2019) 18

[81] Iyer, N., Jayanti, S., Lou, K., Kalyanaraman, Y., Ramani, K.: Three-dimensional shape searching: state-of-the-art review and future trends. Computer-Aided Design **37**(5), 509–530 (2005) 24

[82] Jermyn, I.H., Kurtek, S., Klassen, E., Srivastava, A.: Elastic shape matching of parameterized surfaces using square root normal fields. In: European conference on computer vision. pp. 804–817. Springer (2012) 35

[83] Jia, S., Despinasse, A., Wang, Z., Delingette, H., Pennec, X., Jaïs, P., Cochet, H., Sermesant, M.: Automatically segmenting the left atrium from cardiac images using successive 3d u-nets and a contour loss. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 221–229. Springer (2018) 17, 62

[84] Jia, S., Nivet, H., Harrison, J., Pennec, X., Camaioni, C., Jaïs, P., Cochet, H., Sermesant, M.: Left atrial shape is independent predictor of arrhythmia recurrence after catheter ablation for atrial fibrillation: A shape statistics study. Heart Rhythm O2 **2**(6), 622–632 (2021) 45

[85] Johansson, A., Drake, I., Engström, G., Acosta, S.: Modifiable and non-modifiable risk factors for atherothrombotic ischemic stroke among subjects in the malmö diet and cancer study. Nutrients **13**(6), 1952 (2021) 3

[86] Johnson, W.D., Ganjoo, A., Stone, C.D., Srivyas, R.C., Howard, M.: The left atrial appendage: our most lethal human attachment! surgical implications. European journal of cardio-thoracic surgery **17**(6), 718–722 (2000) 6

[87] Jordan, K., Yaghi, S., Poppas, A., Chang, A.D., Mac Grory, B., Cutting, S., Burton, T., Jayaraman, M., Tsivgoulis, G., Sabeh, M.K., et al.: Left atrial volume index is associated with cardioembolic stroke and atrial fibrillation detection after embolic stroke of undetermined source. Stroke **50**(8), 1997–2001 (2019) 6, 90, 100

[88] Kac, M.: Can one hear the shape of a drum? The american mathematical monthly **73**(4P2), 1–23 (1966) 31

[89] Kamel, H., Healey, J.S.: Cardioembolic stroke. Circulation research **120**(3), 514–526 (2017) 3, 4

[90] Kanmanthareddy, A., Reddy, Y.M., Vallakati, A., Earnest, M.B., Nath, J., Ferrell, R., Dawn, B., Lakkireddy, D.: Embryology and anatomy of the left atrial appendage: why does thrombus form? Interventional Cardiology Clinics **3**(2), 191–202 (2014) 6, 90

[91] Katsanos, A.H., Kamel, H., Healey, J.S., Hart, R.G.: Stroke prevention in atrial fibrillation: looking forward. Circulation **142**(24), 2371–2388 (2020) 5, 6

[92] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: Lightgbm: A highly efficient gradient boosting decision tree. Advances in neural information processing systems **30** (2017) 98

[93] Khera, A.V., Chaffin, M., Aragam, K.G., Haas, M.E., Roselli, C., Choi, S.H., Natarajan, P., Lander, E.S., Lubitz, S.A., Ellinor, P.T., et al.: Genome-wide polygenic scores for common diseases identify individuals with risk equivalent to monogenic mutations. Nature genetics **50**(9), 1219–1224 (2018) 5

[94] Khesin, B., Lenells, J., Misiolek, G., Preston, S.C.: Geometry of diffeomorphism groups, complete integrability and optimal transport. arXiv preprint arXiv:1105.0643 (2011) 49

[95] Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. Proceedings of the national academy of Sciences **95**(15), 8431–8435 (1998) 32

[96] Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) 67

[97] Kirchhof, P.: Integrated care of patients with atrial fibrillation: the 2016 esc atrial fibrillation guidelines (2017) 5

[98] Kjeldsberg, H.A., Sundnes, J., Valen-Sendstad, K.: A verified and validated moving domain computational fluid dynamics solver with applications to cardiovascular flows. International Journal for Numerical Methods in Biomedical Engineering **39**(6), e3703 (2023) 43

[99] Kline, P.: An easy guide to factor analysis. Routledge (2014) 101

[100] Korhonen, M., Muuronen, A., Arponen, O., Mustonen, P., Hedman, M., Jäkälä, P., Vanninen, R., Taina, M.: Left atrial appendage morphology in patients with suspected cardiogenic stroke without known atrial fibrillation. PLoS One **10**(3), e0118822 (2015) 33, 34, 100, 101, 108

[101] Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., Anandkumar, A.: Neural operator: Learning maps between function spaces with applications to pdes. Journal of Machine Learning Research **24**(89), 1–97 (2023) 57

[102] Kutuzova, S., Krause, O., McCloskey, D., Nielsen, M., Igel, C.: Multimodal variational autoencoders for semi-supervised learning: In defense of product-of-experts. arXiv preprint

arXiv:2101.07240 (2021) 74

[103] Lachaud, J.O., Romon, P., Thibert, B., Coeurjolly, D.: Interpolated corrected curvature measures for polygonal surfaces. In: Computer Graphics Forum. vol. 39, pp. 41–54. Wiley Online Library (2020) 80

[104] Lang, S.: Differential and Riemannian manifolds. Springer Science & Business Media (1995) 28

[105] Lee, J.M.: Smooth manifolds. Springer (2012) 25

[106] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895 (2020) 57

[107] Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. ACM/JMS Journal of Data Science (2021) 58

[108] Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoua, G., Suetens, P.D., et al.: Shape retrieval on non-rigid 3d watertight meshes. In: Eurographics workshop on 3d object retrieval (3DOR). Citeseer (2011) 82, 83, 86

[109] Lin, H.J., Wolf, P.A., Kelly-Hayes, M., Beiser, A.S., Kase, C.S., Benjamin, E.J., D'Agostino, R.B.: Stroke severity in atrial fibrillation: the framingham study. Stroke **27**(10), 1760–1764 (1996) 4

[110] Lip, G.Y., Nieuwlaat, R., Pisters, R., Lane, D.A., Crijns, H.J.: Refining clinical risk stratification for predicting stroke and thromboembolism in atrial fibrillation using a novel risk factor-based approach: the euro heart survey on atrial fibrillation. Chest **137**(2), 263–272 (2010) 5

[111] Litany, O., Remez, T., Rodola, E., Bronstein, A., Bronstein, M.: Deep functional maps: Structured prediction for dense shape correspondence. In: Proceedings of the IEEE international conference on computer vision. pp. 5659–5667 (2017) 77

[112] Liu, H.T.D., Jacobson, A., Crane, K.: A dirac operator for extrinsic shape analysis. In: Computer Graphics Forum. vol. 36, pp. 139–149. Wiley Online Library (2017) 34

[113] Liu, X.W., Yan, L.q., Wang, G.s., Sun, Y.h., Bao, W.j., Zhang, H.w., Yang, L., Chen, C., Li, M.Y., Yang, L., et al.: Left atrial appendage filling defects restricted to the early phase of cardiac computed tomography is significantly associated with ischemic stroke. Clinical Imaging **98**, 16–21 (2023) 109

[114] Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: Seminal graphics: pioneering efforts that shaped the field, pp. 347–353. Association for Computing Machinery (1998) 20

[115] Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via deep-

onet based on the universal approximation theorem of operators. Nature machine intelligence **3**(3), 218–229 (2021) 57

[116] Lundberg, S., Lee, S.I.: A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems **30**, 4765–4774 (2017) 72

[117] Maron, H., Galun, M., Aigerman, N., Trope, M., Dym, N., Yumer, E., Kim, V.G., Lipman, Y.: Convolutional neural networks on surfaces via seamless toric covers. ACM Trans. Graph. **36**(4), 71–1 (2017) 82, 83, 84

[118] Masci, J., Boscaini, D., Bronstein, M., Vandergheynst, P.: Geodesic convolutional neural networks on riemannian manifolds. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 37–45 (2015) 77

[119] Maturana, D., Scherer, S.: Voxnet: A 3d convolutional neural network for real-time object recognition. In: 2015 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 922–928. IEEE (2015) 77

[120] McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426 (2018) 102

[121] Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. In: Visualization and mathematics III. pp. 35–57. Springer (2003) 80

[122] Micheli, M., Glaunes, J.A.: Matrix-valued kernels for shape deformation analysis. arXiv preprint arXiv:1308.5739 (2013) 39

[123] Miolane, N., Guigui, N., Le Brigant, A., Mathe, J., Hou, B., Thanwerdas, Y., Heyder, S., Peltre, O., Koep, N., Zaatiti, H., et al.: Geomstats: a python package for riemannian geometry in machine learning. Journal of Machine Learning Research **21**(223), 1–9 (2020) 94

[124] Misiołek, G., Preston, S.C.: Fredholm properties of riemannian exponential maps on diffeomorphism groups. Inventiones mathematicae **179**, 191–227 (2010) 37, 38

[125] Mitchel, T.W., Kim, V.G., Kazhdan, M.: Field convolutions for surface cnns. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10001–10011 (2021) 77, 78

[126] Mitchel, T.W., Rusinkiewicz, S., Chirikjian, G.S., Kazhdan, M.: Echo: Extended convolution histogram of orientations for local surface description. In: Computer Graphics Forum. vol. 40, pp. 180–194. Wiley Online Library (2021) 77

[127] Møller, J.E., Hillis, G.S., Oh, J.K., Seward, J.B., Reeder, G.S., Wright, R.S., Park, S.W., Bailey, K.R., Pellikka, P.A.: Left atrial volume: a powerful predictor of survival after acute myocardial infarction. Circulation **107**(17), 2207–2212 (2003) 100

[128] Morgan, F.: Geometric measure theory: a beginner's guide. Academic press (2016) 48

[129] Moser, J.: On the volume elements on a manifold. Transactions of the American Mathematical Society **120**(2), 286–294 (1965) 49

[130] Mügge, A., Kühn, H., Nikutta, P., Grote, J., Lopez, J.A.G., Daniel, W.G.: Assessment of left atrial appendage function by biplane transesophageal echocardiography in patients with nonrheumatic atrial fibrillation: identification of a subgroup of patients at increased embolic risk. Journal of the American College of Cardiology **23**(3), 599–607 (1994) 6, 90

[131] Mumford, D.: Pattern theory: the mathematics of perception. arXiv preprint math/0212400 (2002) 24

[132] Mumford, D.B., Michor, P.W.: Riemannian geometries on spaces of plane curves. Journal of the European Mathematical Society **8**(1), 1–48 (2006) 34, 46

[133] Nakamori, S., Ngo, L.H., Tugal, D., Manning, W.J., Nezafat, R.: Incremental value of left atrial geometric remodeling in predicting late atrial fibrillation recurrence after pulmonary vein isolation: a cardiovascular magnetic resonance study. Journal of the American Heart Association **7**(19), e009793 (2018) 6, 90

[134] Nattel, S., Burstein, B., Dobrev, D.: Atrial remodeling and atrial fibrillation: mechanisms and implications. Circulation: Arrhythmia and Electrophysiology **1**(1), 62–73 (2008) 4, 98

[135] Nucifora, G., Faletra, F.F., Regoli, F., Pasotti, E., Pedrazzini, G., Moccetti, T., Auricchio, A.: Evaluation of the left atrial appendage with real-time 3-dimensional transesophageal echocardiography: implications for catheter-based left atrial appendage closure. Circulation: Cardiovascular Imaging **4**(5), 514–523 (2011) 6, 90

[136] Oladiran, O., Nwosu, I.: Stroke risk stratification in atrial fibrillation: a review of common risk factors. Journal of community hospital internal medicine perspectives **9**(2), 113–120 (2019) 4

[137] Olver, P.J.: Classical invariant theory. Cambridge University Press (1999) 79

[138] O'neill, B.: Elementary differential geometry. Elsevier (2006) 76, 80

[139] Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. ACM Transactions on Graphics (ToG) **31**(4), 1–11 (2012) 77

[140] Palmer, D., Smirnov, D., Wang, S., Chern, A., Solomon, J.: Deepcurrents: Learning implicit representations of shapes with boundaries. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18665–18675 (2022) 77

[141] Panozzo, D., Puppo, E., Rocca, L.: Efficient multi-scale curvature and crease estimation. Proceedings of Computer Graphics, Computer Vision and Mathematics (Brno, Czech Rapubic **1**(6) (2010) 76, 81

[142] Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous

signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019) 77

[143] Pizer, S.M., Fletcher, P.T., Joshi, S., Thall, A., Chen, J.Z., Fridman, Y., Fritsch, D.S., Gash, A.G., Glotzer, J.M., Jiroutek, M.R., et al.: Deformable m-reps for 3d medical image segmentation. International journal of computer vision **55**, 85–106 (2003) 65

[144] Planinšič, G., Vollmer, M.: The surface-to-volume ratio in thermal physics: From cheese cube physics to animal metabolism. European Journal of Physics **29**(2), 369 (2008) 96, 110

[145] Poulenard, A., Rakotosaona, M.J., Ponty, Y., Ovsjanikov, M.: Effective rotation-invariant point cnn with spherical harmonics kernels. In: 2019 International Conference on 3D Vision (3DV). pp. 47–56. IEEE (2019) 85

[146] Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017) 77

[147] Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. Advances in neural information processing systems **30** (2017) 76, 77, 82

[148] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational physics **378**, 686–707 (2019) 58

[149] Ren, H., Zhuang, X., Rabczuk, T.: A higher order nonlocal operator method for solving partial differential equations. Computer Methods in Applied Mechanics and Engineering **367**, 113132 (2020) 56

[150] Reuter, M., Wolter, F.E., Peinecke, N.: Laplace–beltrami spectra as 'shape-dna'of surfaces and solids. Computer-Aided Design **38**(4), 342–366 (2006) 31, 97, 101

[151] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18. pp. 234–241. Springer (2015) 17

[152] Roobottom, C., Mitchell, G., Morgan-Hughes, G.: Radiation-reduction strategies in cardiac computed tomographic angiography. Clinical radiology **65**(11), 859–867 (2010) 7

[153] Ruder, S.: An overview of multi-task learning in deep neural networks. arXiv e-prints pp. arXiv–1706 (2017) 66

[154] Rudin, C.: Please stop explaining black box models for high stakes decisions. Stat **1050**(26), 457 (2018) 98

[155] Rusinkiewicz, S.: Estimating curvatures and their derivatives on triangle meshes. In: Proceed-

ings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. pp. 486–493. IEEE (2004) 81

[156] Rusinkiewicz, S., Levoy, M.: Efficient variants of the icp algorithm. In: Proceedings third international conference on 3-D digital imaging and modeling. pp. 145–152. IEEE (2001) 41

[157] Rustamov, R.M.: Multiscale biharmonic kernels. In: Computer Graphics Forum. vol. 30, pp. 1521–1531. Wiley Online Library (2011) 115

[158] Şahin, B., İlgün, G.: Risk factors of deaths related to cardiovascular diseases in world health organization (who) member countries. Health & Social Care in the Community **30**(1), 73–80 (2022) 3

[159] Salili-James, A., Mackay, A., Rodriguez-Alvarez, E., Rodriguez-Perez, D., Mannack, T., Rawlings, T.A., Palmer, A.R., Todd, J., Riutta, T.E., Macinnis-Ng, C., et al.: Classifying organisms and artefacts by their outline shapes. Journal of the Royal Society Interface **19**(195), 20220493 (2022) 46, 51

[160] Salti, S., Tombari, F., Di Stefano, L.: Shot: Unique signatures of histograms for surface and texture description. Computer Vision and Image Understanding **125**, 251–264 (2014) 76, 77, 82

[161] Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: Discretization agnostic learning on surfaces. ACM Transactions on Graphics (TOG) **41**(3), 1–16 (2022) 76, 77, 78, 82

[162] Sharp, N., Crane, K.: A Laplacian for Nonmanifold Triangle Meshes. Computer Graphics Forum (SGP) **39**(5) (2020) 83

[163] Sharp, N., Crane, K.: A laplacian for nonmanifold triangle meshes. In: Computer Graphics Forum. vol. 39, pp. 69–80. Wiley Online Library (2020) 32

[164] Shi, Y., Paige, B., Torr, P., et al.: Variational mixture-of-experts autoencoders for multi-modal deep generative models. Advances in neural information processing systems **32** (2019) 74

[165] Shrikumar, A., Greenside, P., Kundaje, A.: Learning important features through propagating activation differences. In: International Conference on Machine Learning. pp. 3145–3153. PMLR (2017) 71

[166] Sinha, A., Bai, J., Ramani, K.: Deep learning 3d shape surfaces using geometry images. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VI 14. pp. 223–240. Springer (2016) 77

[167] Smirnov, D., Solomon, J.: Hodgenet: Learning spectral geometry on triangle meshes. ACM Transactions on Graphics (TOG) **40**(4), 1–11 (2021) 77

[168] Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Symposium on Geometry processing. vol. 4, pp. 109–116. Citeseer (2007) 115

[169] Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: Computer graphics forum. vol. 28, pp. 1383–1392. Wiley Online Library (2009) 34, 76, 77, 82

[170] Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: International Conference on Machine Learning. pp. 3319–3328. PMLR (2017) 71

[171] Svendsen, J.H., Diederichsen, S.Z., Højberg, S., Krieger, D.W., Graff, C., Kronborg, C., Olesen, M.S., Nielsen, J.B., Holst, A.G., Brandes, A., et al.: Implantable loop recorder detection of atrial fibrillation to prevent stroke (the loop study): a randomised controlled trial. The Lancet **398**(10310), 1507–1516 (2021) 5

[172] Vaillant, M., Glaunes, J.: Surface matching via currents. In: Biennial international conference on information processing in medical imaging. pp. 381–392. Springer (2005) 40, 48, 50, 77

[173] Villani, C., et al.: Optimal transport: old and new, vol. 338. Springer (2009) 49

[174] Vlasic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. In: Acm Siggraph 2008 papers, pp. 1–9. Association for Computing Machinery (2008) 84

[175] Wang, M., Xu, J.: Minimal finite element spaces for 2m-th-order partial differential equations in rn. Mathematics of Computation **82**(281), 25–43 (2013) 56

[176] Wang, S., Wang, H., Perdikaris, P.: Learning the solution operator of parametric partial differential equations with physics-informed deeponets. Science advances **7**(40), eabi8605 (2021) 58

[177] Wang, Y., Solomon, J.: Intrinsic and extrinsic operators for shape analysis. In: Handbook of numerical analysis, vol. 20, pp. 41–115. Elsevier (2019) 30, 34, 77

[178] Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. ACM Transactions on Graphics (tog) **38**(5), 1–12 (2019) 77

[179] Watson, T., Shantsila, E., Lip, G.Y.: Mechanisms of thrombogenesis in atrial fibrillation: Virchow's triad revisited. The Lancet **373**(9658), 155–166 (2009) 8

[180] Wiersma, R., Eisemann, E., Hildebrandt, K.: Cnns on surfaces using rotation-equivariant features. ACM Transactions on Graphics (ToG) **39**(4), 92–1 (2020) 84

[181] Wiersma, R., Nasikun, A., Eisemann, E., Hildebrandt, K.: Deltaconv: anisotropic operators for geometric deep learning on point clouds. ACM Transactions on Graphics (TOG) **41**(4), 1–10 (2022) 76, 77, 82

[182] Winter, Y., Wolfram, C., Schaeg, M., Reese, J.P., Oertel, W.H., Dodel, R., Back, T.: Evaluation of costs and outcome in cardioembolic stroke or tia. Journal of neurology **256**, 954–963 (2009) 3

[183] Wu, S., Xu, J.: Nonconforming finite element spaces for 2m-th order partial differential equations on rn simplicial grids when m=n+1. Mathematics of Computation **88**(316), 531–551 (2019) 56

[184] Yaghi, S., Chang, A.D., Akiki, R., Collins, S., Novack, T., Hemendinger, M., Schomer, A., Mac Grory, B., Cutting, S., Burton, T., et al.: The left atrial appendage morphology is associated with embolic stroke subtypes using a simple classification system: a proof of concept study. Journal of cardiovascular computed tomography **14**(1), 27–33 (2020) 6, 7, 8, 90

[185] Younes, L.: Shapes and diffeomorphisms, vol. 171. Springer (2010) 77