UNIVERSITÉ DE MONTRÉAL

ATLAS CONSTRUCTION FOR MEASURING THE VARIABILITY OF COMPLEX ANATOMICAL STRUCTURES

HERVÉ LOMBAERT
DÉPARTEMENT DE GÉNIE INFORMATIQUE ET GÉNIE LOGICIEL
ÉCOLE POLYTECHNIQUE DE MONTRÉAL

THÈSE PRÉSENTÉE EN VUE DE L'OBTENTION
DU DIPLÔME DE PHILOSOPHIÆ DOCTOR (Ph.D.)
(GÉNIE INFORMATIQUE)
JUIN 2012

UNIVERSITÉ DE MONTRÉAL

ÉCOLE POLYTECHNIQUE DE MONTRÉAL

Cette thése intitulée:

ATLAS CONSTRUCTION FOR MEASURING THE VARIABILITY OF COMPLEX ANATOMICAL STRUCTURES

présentée par: LOMBAERT Hervé
en vue de l'obtention du diplôme de: Philosophiæ Doctor
a été dûment acceptée par le jury d'examen constitué de:

M. PAL Christopher J., Ph.D., président
Mme CHERIET Farida, Ph.D., membre et directrice de recherche
M. GRADY Leo, Ph.D., membre et codirecteur de recherche
M. DAHDAH Nagib, M.D., membre
M. SIDDIQI Kaleem, Ph.D., membre

"Learning never exhausts the mind"

— Leonardo da Vinci (1452–1519)

DEDICATION

Aos meus queridos pais
(To My Loving Parents)

ACKNOWLEDGMENTS

This thesis represents a personal achievement that has been made possible through the help and support from many encounters. Notably, I am deeply thankful to all members of my *Ph.D. committee* who took patience to read and evaluate this manuscript. I am, of course, extremely grateful to my director, *Farida Cheriet*, who gave me all the Freedom to pursue my dreams, fulfill my curiosities, experiment and develop all thoughts that eventually led to the young researcher I am today. My scientific development is also greatly due to my codirectors, *Leo Grady*, and previously *Yiyong Sun*, who knew how to precisely stimulate my curiosity and intellect. They both introduced me, since my beginnings at *Siemens Corporate Research in Princeton*, to this captivating world of research where imagination is the limit. I also owe a lot to *Chenyang Xu* who has always been a great mentor and a personal role model all along these years, *Frank Sauer* and *Gareth Funka-Lea* for opening doors to new possibilities, and *Jonathan Polimeni*, from the *Martinos Center at the Harvard Medical School*, whose amazing knowledge greatly inspired me.

During my thesis, I also had the privilege to collaborate with *INRIA Sophia Antipolis*, in a stunning team led by *Nicholas Ayache* who always trusted me, gave me an insightful eye on research, and introduced me to great researchers. Among them are *Jean-Marc Peyrat*, also a great friend, who was extremely helpful all along this collaboration, *Hervé Delingette*, with whom I had great brainstorming sessions, *Xavier Pennec*, who is a great source of inspiration and certainly deeply peaked my curiosity on mathematics, *Maxime Sermesant*, with whom I had great discussions on cardiac anatomy and functions, and incredible students such as *Marco Lorenzi* and *Christoph Seiler*, who inspired me with their rigorous thoughts on Demons. My stay in France also brought me to collaborate with the *CREATIS* team led by *Isabelle Magnin*, and in particular, with *Pierre Croisille* who introduced me to a whole new fascinating world of cardiac microstructures and diffusion tensor imaging, and who always had time to meet me, be it in Lyon, Monaco, or Montreal, as well as with *Patrick Clarysse* and *Stanislas Rapacchi*, who were always very helpful in our collaboration.

The world of research is an intriguing environment where work and friendship are very closely intermingled. Throughout my thesis, I had the privilege to meet great researchers and friends, in Montreal, *David, Cedric, Chantal, Fantin, Fouzi, Jeremie, Jonathan*$^{\times 2}$*, Lama, Luc, Madeleine, Mathias, Najat, Pascal, Pascal, Philippe*, *Rafik* for dragging me into the

IEEE student branch, *Rola, Samuel, Sébastien, Séverine, Thomas, Yasmin*, in Princeton, *Alban, Ali, Brad, Christoph, Christophe*, *Daniel* for being a great roommate and for introducing me to the Levelsets, *David, Guillaume*$^{\times 2}$, *Julien*$^{\times 2}$, *Kalpit, Kanwal, Klaus, Marie-Pierre, Nicolas*$^{\times 2}$, *Oliver, Rui, Tony, Wolfgang*, and on the French Riviera, *Andrew, Barbara, Benoit, Erik, Erin*, *François and Vincent* for their pride, *Grégoire, Hugo*, *Isabelle* for her incredible help, *Jan, Jatin, John, Jerome, Ken*, *Kristin* for liquorice, *Liliane, Marine, Michael, Mikael, Nicolas*, *Olivier* for our 200km bike tour in Corsica with *Jean-Marc*, *Rocio, Stefan, Stephan, Stéphanie*, *Tommaso* for his crash course on Demons, *Vikash, Wassim*.

On a personal note, I must thank my close and deep friends who gave me the courage to pursue the Ph.D. when times were darker, *Ali, Antoine*, *Arnaud* for his amazing kindness, *Bruno*, *Gérald and Marthe* for their unstinting support, Laurent for his insightful thoughts and help, *Mazen*, *Sophie* for her patience, as well as *Catherine pour sa touche de folie qui m'est chère*, and above all, my *parents and family*, *meus pais e familia que sempre me apoiaram em meus esforços*.

The research undergone during this thesis has been funded through the sponsorship of *Siemens Corporate Research*, as well as the *NSERC Alexander Graham Bell Canada Graduate Scholarship* (CGS) and several travel grants (*INRIA, EGIDE, NSERC, IAPR, IPMI*).

RÉSUMÉ

La recherche sur l'anatomie humaine, en particulier sur le cœur et le cerveau, est d'un intérêt particulier car leurs anomalies entraînent des pathologies qui sont parmi les principales causes de décès dans le monde et engendrent des coûts substantiels. Heureusement, les progrès en imagerie médicale permettent des diagnostics et des traitements autrefois impossibles. En contrepartie, la quantité phénoménale de données produites par ces technologies nécessite le développement d'outils efficaces pour leur traitement. L'objectif de cette thèse est de proposer un ensemble d'outils permettant de normaliser des mesures prélevées sur différents individus, essentiels à l'étude des caractéristiques de structures anatomiques complexes.

La normalisation de mesures consiste à rassembler une collection d'images dans une référence commune, aussi appelée construction d'atlas numériques, afin de combiner des mesures provenant de différents patients. Le processus de construction inclut deux étapes principales ; la segmentation d'images pour trouver des régions d'intérêts et le recalage d'images afin de déterminer les correspondances entres régions d'intérêts. Les méthodes actuelles de constructions d'atlas peuvent nécessiter des interventions manuelles, souvent fastidieuses, variables, et sont en outre limitées par leurs mécanismes internes. Principalement, le recalage d'images dépend d'une déformation incrémentales d'images sujettes à des minimums locaux. Le recalage n'est ainsi pas optimal lors de grandes déformations et ces limitations requièrent la nécessité de proposer de nouvelles approches pour la construction d'atlas.

Les questions de recherche de cette thèse se concentrent donc sur l'automatisation des méthodes actuelles ainsi que sur la capture de déformations complexes de structures anatomiques, en particulier sur le cœur et le cerveau. La méthodologie adoptée a conduit à trois objectifs de recherche spécifiques. Le premier prévoit un nouveau cadre de construction automatisé d'atlas afin de créer le premier atlas humain de l'architecture de fibres cardiaques. Le deuxième vise à explorer une nouvelle approche basée sur la correspondance spectrale, nommée *FOCUSR*, afin de capturer une grande variabilité de formes sur des maillages. Le troisième aboutit finalement à développer une approche fondamentalement différente pour le recalage d'images à fortes déformations, nommée les *démons spectraux*.

Le premier objectif vise plus particulièrement à construire un atlas statistique de l'architecture des fibres cardiaques à partir de 10 cœurs *ex vivo* humains. Le système développé a

mené à deux contributions techniques et une médicale, soit l'amélioration de la segmentation de structures cardiaques et l'automatisation du calcul de forme moyenne, ainsi que notamment la première étude chez l'homme de la variabilité de l'architecture des fibres cardiaques. Pour résumer les principales conclusions, les fibres du cœur humain moyen varient de $\pm 12°$, l'angle d'hélix s'étend entre $-41°(\pm 26°)$ sur l'épicarde à $+66°(\pm 15°)$ sur l'endocarde, tandis que l'angle transverse varie entre $+9°(\pm 12°)$ et $+34°(\pm 29°)$ à travers le myocarde. Ces résultats sont importants car ces fibres jouent un rôle clef dans diverses fonctions mécaniques et électro-physiologiques du cœur.

Le deuxième objectif cherche à capturer une grande variabilité de formes entre structures anatomiques complexes, plus particulièrement entre cortex cérébraux à cause de l'extrême variabilité de ces surfaces et de leur intérêt pour l'étude de fonctions cognitives. La nouvelle méthode de correspondance surfacique, nommée FOCUSR, exploite des représentations spectrales car l'appariement devient plus facile et rapide dans le domaine spectral plutôt que dans l'espace Euclidien classique. Dans sa forme la plus simple, FOCUSR améliore les méthodes spectrales actuelles par un recalage non rigide des représentations spectrales, toutefois, son plein potentiel est atteint en exploitant des données supplémentaires lors de la mise en correspondance. Par exemple, les résultats ont montré que la profondeur des sillons et de la courbure du cortex cérébral améliore significativement la correspondance de surfaces de cerveaux.

Enfin, le troisième objectif vise à améliorer le recalage d'images d'organes ayant des fortes variabilités entre individus ou subis de fortes déformations, telles que celles créées par le mouvement cardiaque. La méthodologie amenée par la correspondance spectrale permet d'améliorer les approches conventionnelles de recalage d'images. En effet, les représentations spectrales, capturant des similitudes géométriques globales entre différentes formes, permettent de surmonter les limitations actuelles des méthodes de recalage qui restent guidées par des forces locales. Le nouvel algorithme, nommé *démons spectraux*, peut ainsi supporter de très grandes déformations locales et complexes entre images, et peut être tout autant adapté à d'autres approches, telle que dans un cadre de recalage conjoint d'images. Il en résulte un cadre complet de construction d'atlas, nommé *démons spectraux multijoints*, où la forme moyenne est calculée directement lors du processus de recalage plutôt qu'avec une approche séquentielle de recalage et de moyennage.

La réalisation de ces trois objectifs spécifiques a permis des avancées dans l'état de l'art

au niveau des méthodes de correspondance spectrales et de construction d'atlas, en permettant l'utilisation d'organes présentant une forte variabilité de formes. Dans l'ensemble, les différentes stratégies fournissent de nouvelles contributions sur la façon de trouver et d'exploiter des descripteurs globaux d'images et de surfaces. D'un point de vue global, le développement des objectifs spécifiques établit un lien entre : a) la première série d'outils, mettant en évidence les défis à recaler des images à fortes déformations, b) la deuxième série d'outils, servant à capturer de fortes déformations entre surfaces mais qui ne reste pas directement applicable à des images, et c) la troisième série d'outils, faisant un retour sur le traitement d'images en permettant la construction d'atlas à partir d'images ayant subies de fortes déformations. Il y a cependant plusieurs limitations générales qui méritent d'être investiguées, par exemple, les données partielles (tronquées ou occluses) ne sont pas actuellement prises en charge les nouveaux outils, ou encore, les stratégies algorithmiques utilisées laissent toujours place à l'amélioration.

Cette thèse donne de nouvelles perspectives dans les domaines de l'imagerie cardiaque et de la neuroimagerie, toutefois, les nouveaux outils développés sont assez génériques pour être appliqués à tout recalage d'images ou de surfaces. Les recommandations portent sur des recherches supplémentaires qui établissent des liens avec la segmentation à base de graphes, pouvant conduire à un cadre complet de construction d'atlas où la segmentation, le recalage, et le moyennage de formes seraient tous interdépendants. Il est également recommandé de poursuivre la recherche sur la construction de meilleurs modèles éléctro-mécaniques cardiaques à partir des résultats de cette thèse. En somme, les nouveaux outils offrent de nouvelles bases de recherche et développement pour la normalisation de formes, ce qui peut potentiellement avoir un impact sur le diagnostic, ainsi que la planification et la pratique d'interventions médicales.

ABSTRACT

Research on human anatomy, in particular on the heart and the brain, is a primary concern for society since their related diseases are among top killers across the globe and have exploding associated costs. Fortunately, recent advances in medical imaging offer new possibilities for diagnostics and treatments. On the other hand, the growth in data produced by these relatively new technologies necessitates the development of efficient tools for processing data. The focus of this thesis is to provide a set of tools for normalizing measurements across individuals in order to study complex anatomical characteristics.

The normalization of measurements consists of bringing a collection of images into a common reference, also known as atlas construction, in order to combine measurements made on different individuals. The process of constructing an atlas involves the topics of segmentation, which finds regions of interest in the data (e.g., an organ, a structure), and registration, which finds correspondences between regions of interest. Current frameworks may require tedious and hardly reproducible user interactions, and are additionally limited by their computational schemes, which rely on slow iterative deformations of images, prone to local minima. Image registration is, therefore, not optimal with large deformations. Such limitations indicate the need to research new approaches for atlas construction.

The research questions are consequently addressing the problems of automating current frameworks and capturing global and complex deformations between anatomical structures, in particular between human hearts and brains. More precisely, the methodology adopted in the thesis led to three specific research objectives. Briefly, the first step aims at developing a new automated framework for atlas construction in order to build the first human atlas of the cardiac fiber architecture. The second step intends to explore a new approach based on spectral correspondence, named FOCUSR, in order to precisely capture large shape variability. The third step leads, finally, to a fundamentally new approach for image registration with large deformations, named the *Spectral Demons* algorithm.

The first objective aims more specifically at constructing a statistical atlas of the cardiac fiber architecture from a unique human dataset of 10 *ex vivo* hearts. The developed framework made two technical, and one medical, contributions, that are the improvement of the segmentation of cardiac structures, the automation of the shape averaging process, and more importantly, the first human study on the variability of the cardiac fiber architecture. To

summarize the main finding, the fiber orientations in human hearts has been found to vary with about $\pm 12°$, the range of the helix angle spans from $-41°(\pm 26°)$ on the epicardium to $+66°(\pm 15°)$ on the endocardium, while, the range of the transverse angle spans from $+9°(\pm 12°)$ to $+34°(\pm 29°)$ across the myocardial wall. These findings are significant in cardiology since the fiber architecture plays a key role in cardiac mechanical functions and in electrophysiology.

The second objective intends to capture large shape variability between complex anatomical structures, in particular between cerebral cortices due to their highly convoluted surfaces and their high anatomical and functional variability across individuals. The new method for surface correspondence, named FOCUSR, exploits spectral representations since matching is easier in the spectral domain rather than in the conventional Euclidean space. In its simplest form, FOCUSR improves current spectral approaches by refining spectral representations with a nonrigid alignment; however, its full power is demonstrated when using additional features during matching. For instance, the results showed that sulcal depth and cortical curvature improve significantly the accuracy of cortical surface matching.

Finally, the third objective is to improve image registration for organs with a high inter-subject variability or undergoing very large deformations, such as the heart. The new approach brought by the spectral matching technique allows the improvement of conventional image registration methods. Indeed, spectral representations, which capture global geometric similarities and large deformations between different shapes, may be used to overcome a major limitation of current registration methods, which are in fact guided by local forces and restrained to small deformations. The new algorithm, named *Spectral Demons*, can capture very large and complex deformations between images, and can additionally be adapted to other approaches, such as in a groupwise configuration. This results in a complete framework for atlas construction, named *Groupwise Spectral Demons*, where the average shape is computed during the registration process rather than in sequential steps.

The achievements of these three specific objectives permitted advances in the state-of-the-art of spectral matching methods and of atlas construction, enabling the registration of organs with significant shape variability. Overall, the investigation of these different strategies provides new contributions on how to find and exploit global descriptions of images and surfaces. From a global perspective, these objectives establish a link between: a) the first set of tools, that highlights the challenges in registering images with very large deformations,

b) the second set of tools, that captures very large deformations between surfaces but are not applicable to images, and c) the third set of tools, that comes back on processing images and allows a natural construction of atlases from images with very large deformations. There are, however, several general remaining limitations, for instance, partial data (truncated or occluded) is currently not supported by the new tools, or also, the strategy for computing and using spectral representations still leaves room for improvement.

This thesis gives new perspectives in cardiac and neuroimaging, yet at the same time, the new tools remain general enough for virtually any application that uses surface or image registration. It is recommended to research additional links with graph-based segmentation methods, which may lead to a complete framework for atlas construction where segmentation, registration and shape averaging are all interlinked. It is also recommended to pursue research on building better cardiac electromechanical models from the findings of this thesis. Nevertheless, the new tools provide new grounds for research and application of shape normalization, which may potentially impact diagnostic, as well as planning and performance of medical interventions.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF APPENDICES

INTRODUCTION

Better Understanding of Human Anatomy

The human anatomy has fascinated mankind since the very early ages. In a particular manner, the initial quest for the center of the mind and soul was a recurrent theme in many civilizations and led to intermingled studies on the human heart and brain. Starting with the Egyptians, who, in the *Book of Hearts* (*circa* 1550BC) (Wilkins, 1992), gave the heart a central role in mental and physical illnesses; with the Chinese, who, in the *Book of Songs* (*circa* 700BC) (Schwartz, 1985), had similar beliefs in making the heart the source of intellect and sentiments (the words *heart* and *mind* share even the same Chinese character); or with the Greeks, who, with Aristotle (384BC–322BC), also viewed the heart as the primary organ for the mind and emotions, the comprehension of anatomy changed throughout the years with careful examination of human bodies. For instance, Alcmaeon (*circa* 540BC–500BC) already started to differentiate the role of the heart and brain by considering the latter as the seat of intelligence (Codellas, 1932), later, Galen (129AD–*circa* 200AD) also challenged the central role of the heart and gave equal importance to the brain and liver (Siegel, 1976). Sadly, progress in Science slowed down during the Dark Ages, although, notorious advances were still made in the Islamic world, where, perhaps most importantly, Ibn al-Nafis (1213–1288) formulated the modern concepts of the circulatory system (West, 2008). Later during the Renaissance, interests in human anatomy came back in the Western world, with, for instance, the well-known anatomical drawings of da Vinci (1452–1519, Fig. 1) (Keele, 1964), the seminal anatomical atlas *De humani corpois fabrica* (Vesalius, 1543) of Vesalius (1514-1564), or the study of Harvey (1578–1657) on cardiac mechanical functions (Harvey, 1628). Among Harvey's disciples was also Willis (1621–1675) who coined the term *Neurology* and pioneered research on brain anatomy and its functions (Willis, 1664). Since that time, the Industrial

Figure 1: Leonardo da Vinci (1452–1519) was among the firsts to study human anatomy by blending art and science with surprisingly accurate sketches of the heart and brain *(possible self-portrait, original kept at the Royal Library of Turin, Italy)*.

Age saw a flurry of scientific advances, including the influential work of *"Le règne animal"* (Cuvier, 1817) by Cuvier (1769–1832) and *"On the Origin of Species"* (Darwin, 1859) by Darwin (1809–1882), that made, both, a great progress in the field of comparative anatomy. These findings were collected in highly detailed anatomical atlases, however, the variability of anatomy, its functions, and associated diseases remain yet to be fully understood.

Modern medicine is still searching for a better understanding of human anatomy and pathology, particularly in cardiology and neurology, since these fields continue to address major concerns for society. For instance, nowadays, 71 million Americans (more than a fifth of the population) live with a form or another of cardiovascular disease, costing more than $258 billion and claiming more than 930,000 lives annually (by far the number one killer) (Thom *et al.*, 2006). The Worldwide Health Organization estimates that, worldwide, 17 millions of deaths are attributable to heart-related diseases (30% of deaths). Not far in the rank of top causes of death is also dementia, and particularly Alzheimer's disease (70% of all cases of dementia) that claims lives for more than 100,000 persons in the US each year (400,000 worldwide) and with related costs rising to $100 billion annually (Meek *et al.*, 1998).

The better understanding of complex anatomical and functional characteristics of the heart and brain is essential to a better diagnosis and treatment for any related diseases. Among the possible tools available today, recent advances in medical imaging provide a panoply of images on various organs and multiply the possibilities for research and comprehension on many pathologies. Unfortunately, the growth in data produced by these relatively new technologies necessitates the development of new tools in order to adequately process these images. The purpose of this thesis is to precisely develop new methods for constructing atlases in order to analyze the variability of complex anatomical and functional characteristics. The research presented in the following chapters was conducted with images of the heart and brain. The next section provides, therefore, background information on these organs that is necessary for reading this thesis. Following this summary are the main contributions and an overview of the organization of the manuscript.

Background Information

Two organs are studied across the thesis: the human heart and brain. As previously mentioned, their anatomy and functions have always been intriguing to many civilizations. Here,

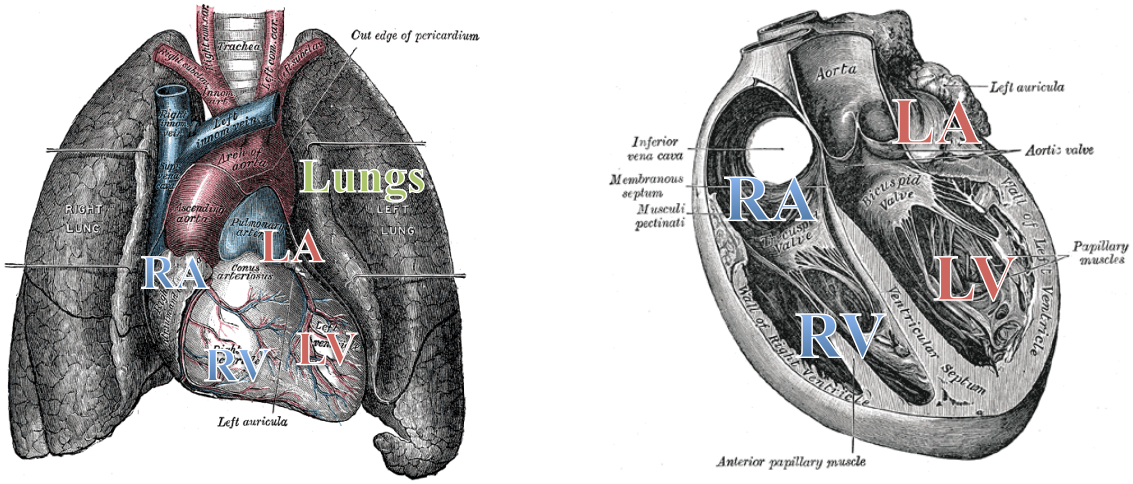

The heart pumps blood through the lungs and body

Left and right side of the heart

Figure 2: A human heart has 4 chambers: Left Atrium (LA), Left Ventricle (LV), Right Atrium (RA), Right Ventricle (RV). Blood 1) enters the heart in the RA, 2) transfers to the RV; 3) RV contraction pushes blood into the lungs; 4) blood returns in the heart in the LA, 5) transfers to the LV; 6) LV contraction ejects blood throughout the body *(Images adapted from Gray's anatomy, plate 490 and 498).*

they are demystified by explaining the general concepts on, *first*, the human heart and, *second*, the human brain, that are necessary for reading the following chapters.

The Heart

The Merriam-Webster dictionary defines the heart as *"a hollow muscular organ that acts as a force pump maintaining the circulation of blood"*. It has about the size of a fist and is located in the chest behind the breastbone. It is at the heart of the cardiovascular system and its role is to ensure blood flows through the body. The heart may be one of the most fascinating organs in mammals since it unceasingly pumps blood throughout life and undergoes an extreme deformation during its beating. The mechanics underlying such motion may be attributable to a complex cardiac electrophysiology and anatomical architecture, described as following.

Cardiac Anatomy and Physiology

The heart of mammal species is divided into left and right sides, and consists of 4 cardiac chambers, an atrium and a ventricle on each side. An atrium receives blood before transferring it to the ventricles, and each ventricle acts as a powerful pump ejecting blood out of the

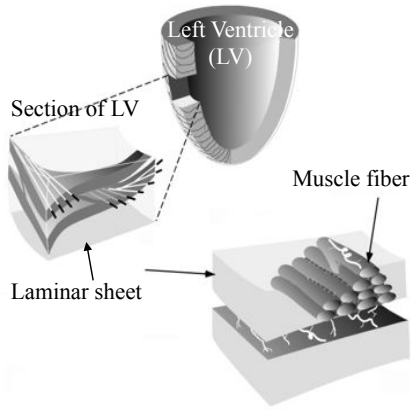

a) Organization of cardiac fibers as laminar sheets

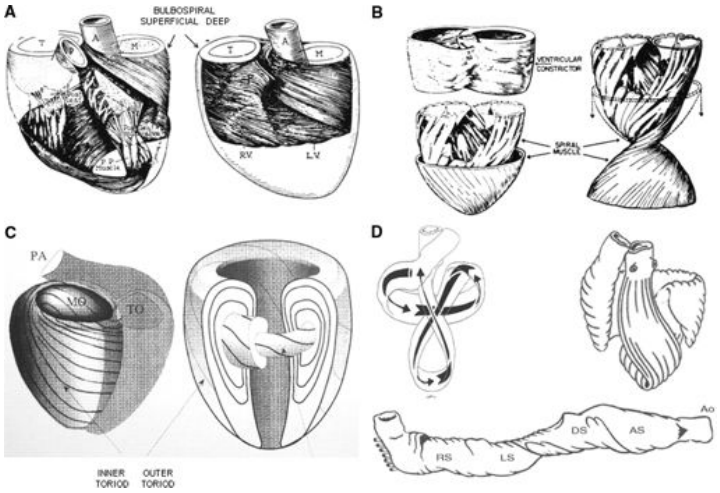

b) Possible structures of the cardiac fiber architecture

Figure 3: *a)* Illustration of the fiber organization in a segment of the left ventricle (LV). The cardiac fibers are organized as laminar sheets. The torsion of the LV is related to the fiber orientation on the inner and outer surface of the ventricle, moreover, electrical conductivity is greatest along the fiber direction. *b)* Four possible architecture. The cardiac muscle may be formed by: A, 4 myocardial bundles (Mall, 1911), B, by a spiral muscle sandwiched onto itself (Rushmer *et al.*, 1953), C, by toroidal surfaces resembling doughnuts (Streeter, 1979), or D, by a unique band folded around itself (Torrent-Guasp *et al.*, 2005) *(Images adapted from (LeGrice et al., 1995) and (Buckberg et al., 2008)).*

heart. The cardiac cycle is summarized in Fig. 2 and can be decomposed in 6 main steps:

1. *Right Atrium Diastole* — During the relaxation of the heart, deoxygenated and low-pressure blood enter the right atrium (RA) from the inferior and superior *venae cavae* (IVC, SVC).

2. *Right Atrium Systole* — The contraction of the heart pushes blood through the tricuspid valve, into the right ventricle (RV).

3. *Right Ventricle Systole* — The contraction of the RV expels blood through the pulmonary valve, into the pulmonary artery (PA) where it is oxygenated in the lungs.

4. *Left Atrium Diastole* — During the relaxation of the left atrium (LA), the oxygenated blood comes back from the lungs, through 4 pulmonary veins (PV) and into the LA.

5. *Left Atrium Systole* — The contraction of the LA and the ongoing expansion of the left ventricle (LV) cause blood to be transferred through the mitral valve, into the LV.

6. *Left Ventricle Systole* — The strong contraction of the LV, which has the thickest muscle of all chambers, ejects the newly oxygenated blood though the aortic valve, into the aorta, where it distributed across the whole body.

In effect, the left and right sides function simultaneously, producing the familiar sound of the cardiac beating (the blood stopped by a valve produces a beat). In a cycle, two successive beats (dubbed the *"lub-dub"*) are produced by the closure of, first, the tricuspid and mitral valves, and second, the pulmonary and aortic valves. A perturbation in this sound may be a sign of a heart defect.

A healthy heart has a good synchronization of cardiac muscles during contraction (systole) and relaxation (diastole). Systole first occurs in the upper chambers (LA and RA), and subsequently, in the two lower chambers (LV and RV). These contractions are caused by an electrical activation that propagates throughout the heart.

Cardiac Fiber Architecture

The heart, likewise any muscle, is mostly composed by muscular fibers. They are locally organized as laminar sheets (Streeter *et al.*, 1969; LeGrice *et al.*, 1995) and have different orientations on the inner surface (endocardium) and outer surface (endocardium) of the cardiac muscle (illustrated in Fig. 3 a). This complex fiber architecture determines various cardiac mechanical functions (Costa *et al.*, 2001) and cardiac electrophysiology patterns (Hooks *et al.*, 2002) since the torsion of the LV is directly related to the orientation of these cardiac fibers (Russel *et al.*, 2009) and electrical conductivity is greatest along these fibers (Hooks *et al.*, 2007).

The exact description of this fiber architecture is not fully known in humans. Its understanding has been mostly limited by tedious histological studies and is largely speculated from observations in other species. Different theories exist and are surveyed in (Buckberg *et al.*, 2008) (illustrated in Fig. 3 b). The hypothesis is that the cardiac muscle may be formed by 4 myocardial bundles (Mall, 1911), by a spiral muscle sandwiched onto itself (Rushmer *et al.*, 1953), by toroidal surfaces resembling doughbuts (Streeter *et al.*, 1969; Torrent-Guasp *et al.*, 2005), or by a unique band folded around itself (Torrent-Guasp *et al.*, 2005).

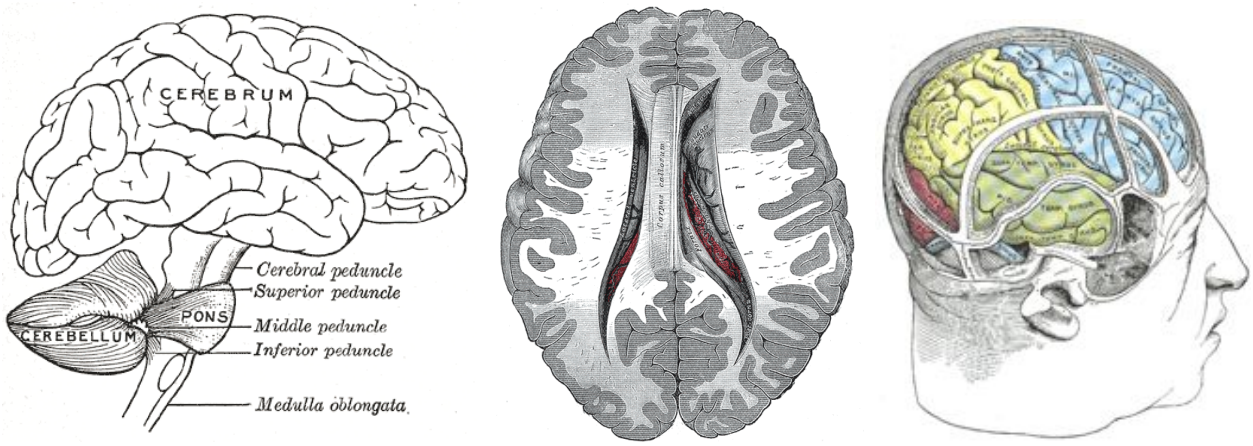

a) Cerebrum and cerebellum *b)* Grey and white matter *c)* Four lobes of the cortex

Figure 4: The brain is composed of many parts. The most prominent is the cerebrum, composed of grey and white matter. The cortical surface is divided into four lobes and is important since it is thought to be responsible for perception and cognition *(Images from Gray's anatomy, plate 677, 737 and 1197).*

The Brain

The brain is the center of the nervous system and is associated with perceptive and cognitive functions. It is contained within the skull and has a general structure common to most mammal species. Because of its functions, neuroscience invests a large amount of efforts in studying the brain, yet, so little is known on the mechanisms of cognition. This trend is, however, changing since recent advances in neuroimaging enable the visualization of cerebral activity. The need to map these functional images onto anatomy is consequently important. In the following, the general anatomy of the brain and its surface is described.

Anatomy

The brain consists of several distinctive parts and is attached to the spinal cord with the brainstem (Fig. 4 a). Functions of the brainstem involve primary needs (breathing, digestion, sense of danger). Attached to the brainstem is the cerebellum, a separate structure underneath the cerebral hemispheres. Its functions may be linked to motor control. The area that connects the brainstem to the hemispheres forms the limbic system. It comprises many substructures, such as the hippocampus linked to memory and navigation, and is surrounded by other important brain parts, such as the thalamus or hypothalamus, responsible for metabolic processes for relaying information within inner structures.

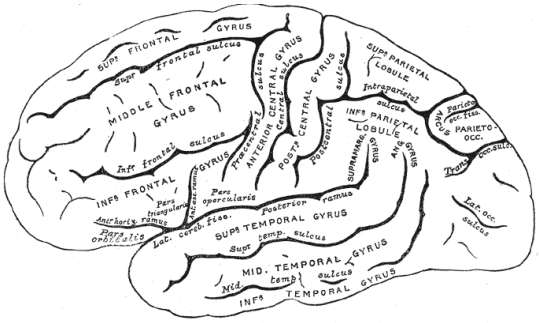

a) Lateral view of the main cortical gyri and sulci

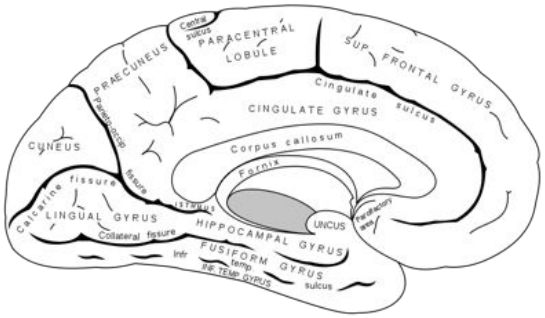

b) Medial view of the main cortical gyri and sulci

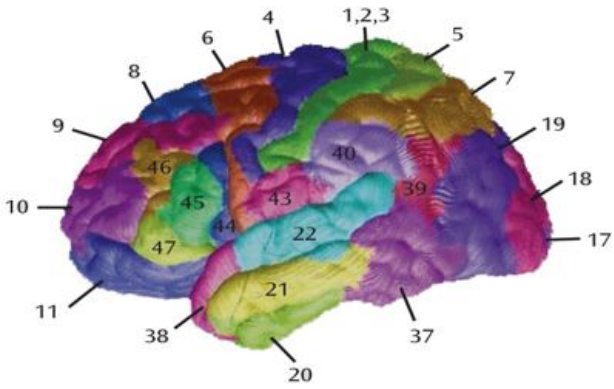

c) Brodmann areas on lateral view

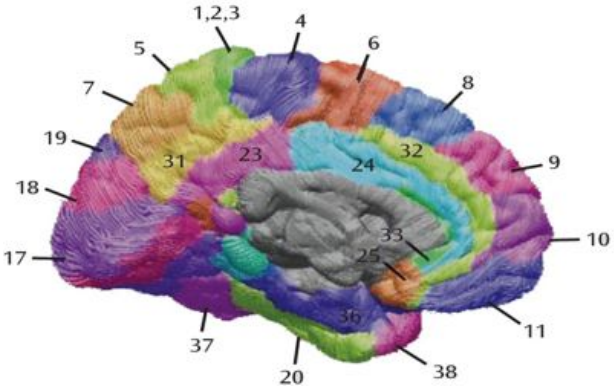

d) Brodmann areas on medial view

Figure 5: The brain surface, named the cortex, is composed of ridges and furrows, named gyri and sulci. These main structures are shown on *a)* a lateral view and *b)* a medial view. *c,d)* The architecture of neuronal cells may be used to divide the cortical surface into 51 Brodmann areas, shown in *(Images from Gray's anatomy, plate 726 and 727, and Brodmann areas adapted from (Foland-Ross et al., 2011))*.

The most prominent brain structure is the *cerebrum*, divided into two hemispheres that are connected by the *corpus callosum*, a bundle of neural fibers that facilitates interhemispheric communication. The *cerebrum* is composed of gray and white matter, and voids between hemispheres compose the ventricular system (Fig. 4 b). More precisely, the grey matter is distributed across the brain surface and consists of a layered architecture of neuronal cell bodies, while the white matter is essentially composed of axons, or fibers, passing information between different areas. The brain surface, referred to as the cortex, has the particularity of being highly convoluted in humans (Toro and Burnod, 2005). The ridges and furrows (or bumps and valleys) of the cortical surface are respectively named *gyri* and *sulci* (Fig. 5).

The cortex in early development is initially smooth and starts to fold in relatively stable

primary convolutions at midgestation (e.g., the calcarine and central sulcus), soon followed by a secondary and tertiary convolutions (folds on folds), which are much more variable across a population (Ono *et al.*, 1990; Welker, 1990). In fact, while a newborn has a relatively smooth brain surface (Dubois *et al.*, 2008), the tertiary convolutions continue throughout the first year of life (Chi *et al.*, 1977; Armstrong *et al.*, 1995; Garel *et al.*, 2001).

Cortical Areas

The center of neural activity is thought to reside on the cortical surface, which is divided into four lobes (Fig. 4 c). The frontal lobe may be involved in planning, problem solving and high cognitive functions such as behavior and emotions. The occipital lobe, in the back of the brain, is responsible for processing visual information, including shape and color recognition. The parietal lobe, on the top of the brain, is responsive to sensation such as touch and pressure. The temporal lobe, near the ears, is responsive to smell and sound, and may be related to memory.

Interestingly, the neuronal architecture changes across these cortical areas. It is thought that distinctive cytoarchitecture, differing by its organization of neuronal layers, may be related to dedicated cerebral functions. This idea has been popularized by Brodmann (Brodmann, 1909), which divided the cortex into 51 areas, each associated with a specific cytoarchitecture. For instance, the primary visual cortex is associated with Brodmann area 17 (Fig. 5 c,d). Although Brodmann areas remain the reference, other partitions exist, for instance, in 14 areas (Campbell, 1903), in 50 areas (Smith, 1907), or in 107 areas (von Economo and Koskinas, 1925, 2007).

Contributions

Research across the centuries has established key knowledge in human anatomy, however, as described above, so little is sometimes known in specific aspects of cardiology and neurology. The purpose of this thesis is to develop a set of tools that address precisely a few of these knowledge gaps. In particular, the construction of atlases, which provides a better understanding of human anatomy, necessitates improved methods for registering images and averaging shapes. As a continuation of the research presented earlier, the human brain and heart will have a particular interest throughout the thesis.

Firstly, the research will focus on the human brain, and more particularly, on the cor-

tical surface. As mentioned previously, neuroscience studies that try to find correlations between anatomical and functional areas will provide a better understanding on the mechanics of perception and cognition. Such studies depend on efficient and accurate methods for matching cortical surfaces across a population. Contributions will be made in this particular topic. Indeed, current methods for matching cortical surfaces are often prohibitively expensive in computational time. A new approach based on spectral correspondence will be developed and will yield the original FOCUSR algorithm. The results will show accurate matching computed at much greater speed than current methods. A few experiments will be presented in order to show the applicability of FOCUSR in real neuroscience studies.

Secondly, the research will focus on the human heart, and more precisely, on the cardiac fiber architecture. Contributions will be made by developing a framework to construct the first statistical atlas of the human cardiac fibers. The automation of the atlas construction requires original segmentation methods and robust schemes for averaging shapes. Besides the development of new methods, the variability study of the different fiber structures across the heart will provide key findings relevant to a better understanding of various cardiac mechanical functions and electrophysiology.

Additionally, the registration of cardiac images that are subject to very large deformations will be addressed by developing a fundamentally new approach for image registration. Inspired by the findings of FOCUSR, a new algorithm, named *Spectral Demons*, will be developed. The generic algorithm will be shown to be applicable to any conventional method for registration with examples that extend both, the classical Demons algorithm and a new groupwise registration framework. The experiments will show that this fundamentally new algorithm can naturally capture very large and complex deformations and can demonstrate substantial improvements over conventional approaches. *Spectral Demons* will be a key method for any application that involves general image registration.

Manuscript Overview

This introductory chapter gave the general context of this thesis. The manuscript is composed of three main parts. *Firstly*, the introduction reviews the current state-of-the-art for constructing atlases, and provides the basis for the rationale behind the research questions and their objectives, with a particular focus on the human heart and brain. *Secondly*, the methods and results are presented with a study on the complexity of the human cardiac fiber

architecture, which required the development of a new framework for atlas construction, and with additionally, a study on the high shape variability of the heart and the brain, which required the development of new spectral matching methods. *Lastly*, the final chapters provide a general discussion that highlights the links between different findings and determines the general implications, advances and limits of the developed tools and results. Each part is composed in more details as following:

Part I – INTRODUCTION AND BACKGROUND

- **Introduction** – This introductory chapter provides the general context, background information necessary for reading the manuscript (a global description of the human heart and brain anatomy), as well as the main contributions of the thesis.

- **Chapter 1** – The current state-of-the-art for atlas construction is established in this chapter by reviewing each building block. More specifically, this survey focuses on the current limitations found in the literature on image segmentation, image registration and shape averaging.

- **Chapter 2** – The rationale of the general methodology is explained in this chapter by presenting the problem statement and the proposed research questions. The research objectives are consequently established and aim at characterizing the human cardiac fiber architecture and at capturing large deformations between surfaces and images.

Part II – METHOD AND RESULTS

- **Chapter 3** – The first journal article is presented here. It provides a new automated framework for constructing the first human atlas of the cardiac fiber architecture. The fiber variability across the myocardium is studied for the first time in humans. Appendix A provides an additional variability study on the laminar sheet, and Appendix B provides a preliminary comparison between healthy and abnormal hearts.

- **Chapter 4** – The second journal article is presented here. It presents FOCUSR, an improved spectral correspondence method for matching generic meshes as well

as brain surfaces. It is based on an original direct feature matching technique that can use virtually any additional information capable of helping correspondence.

- **Chapter 5** – This additional chapter completes the theme initiated by Chapter 4. In its first sections, FOCUSR is used to construct an atlas of cerebral cortical surfaces. The experiments focus on the principal modes of shape variations as well as the shape of the primary cortical folding. In the last sections of the chapter, the new improved spectral matching approach is used for image registration. The new method, named *Spectral Demons*, is additionally extended in a groupwise framework in order to construct atlases, and it demonstrated substantial improvements over conventional approaches.

Part III – DISCUSSIONS AND CONCLUSIONS

- **Chapter 6** – A general discussion establishes the links between previous work and the new findings with a focus on their implications, advances, and limitations.
- **Conclusion** – This final chapter concludes the thesis by summarizing the key contributions and gives recommendations and perspectives for future work.

Guidelines for Reading the Thesis

This section provides guidelines for reading the thesis and for finding specific information. As introduced in this chapter, the general purpose of research is to develop new tools that facilitate measurements of complex characteristics across a population. This manuscript may, therefore, be read in several ways depending on the reader's interests:

Image registration and surface matching – The reader interested in shape normalizing methods (i.e., image registration and surface matching) may focus particularly on Sec. 3.2.2, 4.2 and 5.4. He will find all three different approaches developed in this thesis. The first approach is based on the conventional Demons algorithm and uses simplified images, the second approach is based on spectral correspondence and can be used with cortical surfaces as well as with generic meshes, and the third approach is, in a similar manner, based on spectral matching and is capable of capturing very large deformations between images.

Shape averaging – The reader interested in computing an average shape may focus on Sec. 3.2.2, 5.2.1, and 5.4.2. The first framework is based on Guimond's shape averaging (Guimond *et al.*, 2000), the second approach averages corresponding surfaces, and the third framework is a groupwise approach where the average shape is computed in parallel with image registration.

Statistical framework – The reader interested in tools for statistical analysis may focus on Sec. 3.2.3 and 5.2.1. The first approach uses the Log-Euclidean metric (Arsigny *et al.*, 2006b), and the second approach finds the principal modes of shape variations among a set of surfaces.

Application in computational anatomy – The reader interested in the applicability of the developed tools may focus on Sec. 3.3, A.3, 5.2.2 and 5.3.2. The first two mentioned sections are the results of variability studies on the human atlas of the cardiac fiber architecture. He will find the average angles of the fiber and laminar sheet structures across several myocardial segments, as well as their normal angular variations. The last two mentioned sections are applications of FOCUSR where the principal variations of cortical surfaces are presented along an exploration of the primary cortical folding pattern.

PART I

BACKGROUND

CHAPTER 1 LITERATURE REVIEW

Medical imaging plays a key role in our society and routinely save lives; yet, image processing becomes increasingly essential in many medical applications. For instance, it provides precise measurements, digests information, and assists interventions and planning of surgeries. The focus of this thesis lies on image processing tools, summarized in Fig. 1.1, that are necessary for building computational atlases of human organs. This construction requires, firstly, image segmentation, which extracts the region of interest, and secondly, image registration, which aligns regions of interest into a common space. This chapter will, therefore, review the literature on these topics with the following order: image segmentation, image registration, and lastly, atlas construction.

1.1 Segmentation

Image segmentation is the process of localizing regions of interest in an image (objects, boundaries, or landmarks), and the challenge for this task remains in attaining accuracy and precision for medical applications (mislabeled pixels could lead to wrong diagnostics, worse, misplaced dose delivery). The segmentation quality is also crucial when building atlases since it ensures that the right data is used during construction.

Early segmentation methods can be categorized into edge-based and region-based approaches. Edge-based methods find boundaries in images (e.g., convolution of filters such as the so-called Mexican hat (Marr and Hildreth, 1980) or more advanced filters (Canny, 1986; Deriche, 1987)), while region-based methods find regions instead of boundary and have the advantage of using the whole image in a global manner, i.e., regions has a broader scope than boundaries. This region-based approach may be tackled with a top-down scheme, where the image is successively split (e.g., thresholding (Otsu, 1979; Sezgin and Sankur, 2004)), or with a bottom-up scheme, where pixels are merged to form regions (e.g., regions are grown from "phagocyte" seeds (Brice and Fennema, 1970; Adams and Bischof, 1994)). These two schemes may be unified with the use of a structural tree (Horowitz and Pavlidis, 1976) that represents pixels with its leaves and the whole image with its root (e.g., the watershed method (Vincent and Soille, 1991) builds such tree by merging catching basins during flooding of an image space).

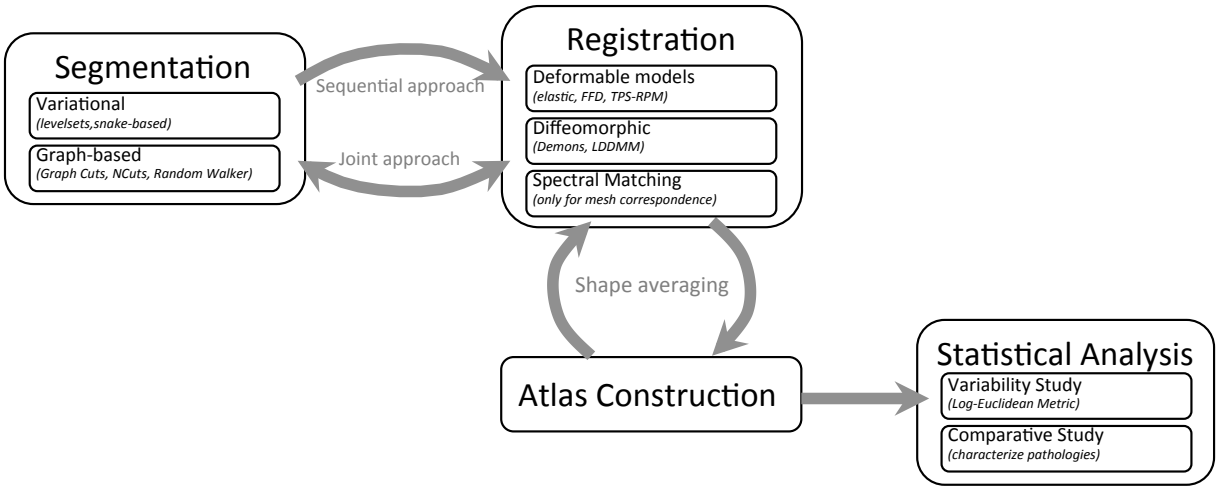

Figure 1.1: Fundamental elements necessary to construct an atlas. **Segmentation**, to extract regions of interest. **Registration**, to align the regions of interest in a common reference. **Shape averaging**, to find the average structure in a population. **Statistical Analysis**, to extract meaningful information from a population.

Jean-Michel Morel (Morel and Solimini, 1995) had the feeling that *"most segmentation algorithms try to minimize* [...] *one and the same segmentation energy"*. On the same note, Stuart Geman and his brother Donald Geman (Geman and Geman, 1984) already introduced in 1984 a Bayesian framework (Cooper *et al.*, 1981; Marroquin, 1984) for image analysis which marks a milestone in image processing, while David Mumford and Jyant Shah proposed in their seminal paper (Mumford and Shah, 1985, 1989; Mumford, 1994)) **to formulate segmentation as an energy minimization problem**. From this anecdotal feeling, modern segmentation methods are now, for a vast majority of them, formulated as an energy minimization problem, often based on, or inspired by the piecewise smooth functional of Mumford and Shah (i.e, they often have a similarity criterion and smoothness constraints). The state-of-the-art for image segmentation (surveyed in (Pal and Pal, 1993; Pham *et al.*, 2000; Noble and Boukerroui, 2006; Zhang *et al.*, 2008a; Heimann and Meinzer, 2009; Petitjean and Dacher, 2011)) may be categorized into variational approaches and graph-based approaches. A third category may also be set from methods relying on an atlas, which register a pre-segmented atlas onto an image, however, this would rather fall as being a registration approach.

1.1.1 Variational Approaches

The segmentation of an image can be generated with a slow, tractable, transformation of an initial solution (e.g., a circle in an image is morphed to fit an object of interest). Such a variational approach formulates segmentation within a Euler-Lagrangian framework (i.e., variables are updated with infinitesimal steps). These update schemes may be categorized as either explicit or implicit.

The explicit formulation was tackled by Michael Kass *et al.* (Kass *et al.*, 1988) with a deformable contour, known as *snake*, that moves *explicitly* toward an object boundary, pulled by external forces computed from the image. These forces contain high-level information, such as contour curvature (keeping the snake smooth), and low level information, such as image gradient direction (pulling the snake toward the object boundary). In (Xu and Prince, 1997), forces derived from the gradient field are diffused in order to extend these pulling forces to a greater distance far away from the object boundary. These extended forces, named gradient vector flow, improve dramatically the convergence of snakes toward boundaries. However, the main limitation of these active contour formulations is that they do not naturally handle changes in topology (the snake cannot split and merge).

The implicit formulation, proposed by Stanley Osher and James A. Sethian (Osher and Sethian, 1988), transposes segmentation into a higher dimension (i.e., by deforming a surface instead of a contour). Their core idea is that the intersection of a 3D surface with a 2D plane (the image) defines an *implicit* contour (the zero level set, which represents an object boundary). As the surface evolves, its level sets may undergo splits and merges implicitly, and therefore, changes in topology are naturally handled (tracking of splits and merges is no longer necessary). Two types of forces are typically used: Internal forces, that keep the surface smooth, and external forces, derived from image gradients (Malladi *et al.*, 1995) that pull the surface toward object boundaries. This level set method is a versatile framework that is capable with a broad range of applications and energies (even the snake energy (Caselles *et al.*, 1995) and gradient vector flow (Paragios *et al.*, 2004)). One popular energy, proposed by Tony Chan and Luminita Vese (Chan and Vese, 1999), focuses on finding the right balance between image regions rather than optimizing primarily a boundary cost (multiple regions are handled in (Yezzi *et al.*, 1999)). Additionally, simplifications and computationally efficient versions of the level set method exist (Chopp, 1993; Adalsteinsson and Sethian, 1995; Li *et al.*, 2005; Mitiche and Ayed, 2010), notably the Fast Marching algorithm (Sethian, 1996,

1999), a special case where the surface is always expanding or collapsing.

1.1.2 Graph Based Approaches

In the 1920s Max Wertheimer fathered the Gestalt theory movement (Wertheimer, 1925, 1938; Desolneux *et al.*, 2007). He suggested that cognition is a whole and cannot be explained by summing each parallel operation alone. Applied to visual perception, this may be interpreted as that global properties, such as image partitions, cannot be found by combining separate low-level cues (e.g., pixel intensities) and requires an organization and grouping of these cues (e.g., pixels alone means little, while their relations may be meaningful). Using a graph to segment an image is an elegant implementation of this theory (the graph organizes low-level cues). For instance, graph leaves may correspond to image pixels and global properties (e.g., an energy) may be modeled with the structural organization of graph edges and nodes. Regions are subsequently found by partitioning the graph. A few approaches exist; for instance, minimum spanning trees (Zahn, 1971; Chen and Pavlidis, 1990; Felzenszwalb and Huttenlocher, 2004) organize the graph as a tree and merge regions into branches. Moreover, Zhenyu Wu and Richard Leahy (Wu and Leahy, 1990, 1993) were among the first to introduce a global optimization approach. The proposed minimum cut (in terms of its total edge weights) optimizes a boundary cost function by finding an optimal partition of the graph (Greig *et al.*, 1989; Wu and Leahy, 1993). This approach can be computationally efficient (Boykov and Kolmogorov, 2004, 2001) but suffers from a bias toward short boundaries (i.e., small boundaries cost less). The ratio cut (Wei and Cheng, 1989; Hagen and Kahng, 1992; Cox *et al.*, 1996; Chan *et al.*, 1993) (also known as average cut (Sarkar and Soundararajan, 2000; Soundararajan and Sarkar, 2001, 2003)) and normalized cut criteria, although computationally less efficient, overcome this shrinkage bias toward short boundaries. Furthermore, these approaches based on eigendecomposition have strong links (Perona and Freeman, 1998; Amir and Lindenbaum, 1998; Williams and Thornber, 1998; Weiss, 1999) with spectral theory. In the following, the widely used (minimum cut) Graph Cuts algorithm and graph spectral methods will be reviewed.

Graph Cuts

The minimum cut criterion (Greig *et al.*, 1989; Wu and Leahy, 1993) favors a cut through the graph whose sum of severed edge weights is minimal (i.e., finding the cut with the minimal

cost). Fortunately, polynomial time algorithms (Karger and Stein, 1996) can be used to solve the equivalent problem of maximum flow in a graph (Ford and Fulkerson, 1962). Yuri Boykov and Vladimir Kolmogorov (Boykov and Kolmogorov, 2004, 2001) introduced the seminal **Graph Cuts algorithm** and showed that it **guarantees a global optimal solution for binary variables**. Moreover, image segmentation between foreground and background (Boykov and Jolly, 2000; Boykov and Funka-Lea, 2006) can be performed interactively and there are improved variants of the Graph Cuts algorithm[1] when time (Juan and Boykov, 2006; Delong and Boykov, 2008) or memory (Lombaert *et al.*, 2005; Xu *et al.*, 2003; Lerm *et al.*, 2010) becomes an issue. The segmentation of large 4D datasets (3D+time) is also possible with the use of temporal links in the graph (Wolz *et al.*, 2010; Lombaert and Cheriet, 2010b; Lombaert *et al.*, 2011f).

Multi-labeled variables may be additionally optimized in a Graph Cuts based framework (Boykov *et al.*, 1998, 2001; Kolmogorov and Zabih, 2002, 2004; Freedman and Drineas, 2005; Veksler, 1999) that guarantees a solution (for metric energies) bounded within a known factor of the global optimum. Recently, Label Cost (Delong *et al.*, 2011, 2010a,b) adds further regularization in the framework. The Graph Cuts algorithm attained a vast popularity in computer vision and finds applications beyond image segmentation with, in particular, few recent attempts in image registration (Kim *et al.*, 2003; Tang and Chung, 2007; So and Chung, 2009, 2010; So *et al.*, 2011; Bhat *et al.*, 2006; So and Chung, 2011; Lombaert *et al.*, 2007; Lombaert and Cheriet, 2012)[2].

Graph-Spectral Methods

The ratio cut criterion (Wei and Cheng, 1989; Hagen and Kahng, 1992; Chan *et al.*, 1993; Sarkar and Soundararajan, 2000; Soundararajan and Sarkar, 2001, 2003) was proposed to find optimal partitions in electronic circuit designs. The minimum cut was indeed not adequate as it favors small cuts in a graph (i.e., the cut isolating one single graph node is very small). To avoid such shrinkage bias, the optimal ratio cut creates partitions with equal sizes (same number of nodes in each partition). The minimal ratio cut separates two regions with similar sizes.

The normalized cut criterion (Shi and Malik, 1997, 2000) goes along the same princi-

[1] First hierarchical approach (Lombaert *et al.*, 2005) proposed by the candidate

[2] Extension using landmarks (Lombaert *et al.*, 2007) and joint image denoising/registration (Lombaert and Cheriet, 2012) proposed by the candidate

ple (balancing the size of each partitions), however, the size of a partition is defined with its connectivity degree (strong if nodes have many connections within a region, weak if nodes are loosely connected). More precisely, the minimal normalized cut separates regions with similar areas that are adjusted for their node densities.

The Laplacian matrix (Fig. 1.2) represents the connectivity between nodes in a graph (possible applications surveyed in (Mohar, 1997; Spielman, 2010)). Interestingly, finding its smallest (non zero) eigenvector (also called the Fiedler vector (Chung, 1997b)) actually gives the minimal ratio cut. The minimal normalized cut can be found similarly by decomposing the normalized Laplacian matrix. Binary image segmentation (Shi and Malik, 1997, 2000) is, therefore, performed by simply thresholding the Fiedler vector of a graph (each component of this vector gives the probability of a pixel to belong to a region). The Fiedler vector plays, in fact, an important role in spectral graph theory (Chung, 1997b; Luxburg, 2007) and gives a probabilistic foundation for graph-based segmentation approaches (Meila and Shi, 2000, 2001; Robles-Kelly, 2005).

When seed points are available, the minimization of the ratio cut and of the normalized cut can be solved with linear systems of equations (Grady and Schwartz, 2005; Grady, 2006a; Grady and Schwartz, 2006; Daneshgar and Javadi, 2012). On the same line, Leo Grady introduced the Random Walker algorithm (Grady *et al.*, 2005; Grady, 2005, 2006b) that generates a segmentation satisfying the Laplace equation with Dirichlet boundary conditions (which are set with seed points). Initialization with pre-computed eigenvectors of the weighted Laplacian matrix allows faster Random Walker and can be interpreted as a seeded Normalized Cuts algorithm (Grady and Sinop, 2008). The random walker paradigm can also explain the Graph Cuts algorithm (Sinop and Grady, 2007), and recently, the Power Watershed algorithm (Couprie *et al.*, 2009, 2011) proposed to unify the Random Walker, Graph Cuts, and the Minimum Spanning Tree. **The graph Laplacian operator plays an important role in graph-based segmentation methods** (Fig. 1.3).

1.2 Registration

Image registration is the process of matching images into a common space in order to find their correspondences. In a medical context, aligning images aids clinical interpretation by highlighting differences in regions of interest between individuals (**inter-patient** registration) or in a single patient (**intra-patient** registration) due to **structural** (e.g., morpholog-

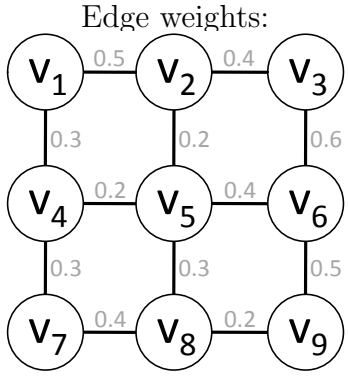

Laplacian matrix of the corresponding graph:
$L = D - W$

$$\begin{pmatrix} 0.8 & -0.5 & & -0.3 & & & & & \\ -0.5 & 1.4 & -0.4 & & -0.2 & & & & \\ & -0.4 & 1.0 & & & -0.6 & & & \\ -0.3 & & & 0.8 & -0.2 & & -0.3 & & \\ & -0.2 & & -0.2 & 0.8 & -0.4 & & -0.3 & \\ & & -0.6 & & 0.4 & 1.5 & & & -0.5 \\ & & & -0.3 & & & 0.7 & -0.4 & \\ & -0.3 & & & & & -0.4 & 0.9 & -0.2 \\ & & & & & -0.5 & & -0.2 & 0.7 \end{pmatrix}$$

W corresponds to the adjacency matrix (entries in black above, giving interrelationships between nodes), D corresponds to the degree matrix (entries in gray above, giving the the total connections per node)

Figure 1.2: Example of a Laplacian matrix associated with a graph (here a 3×3 lattice). The matrix $L = D - W$ represents the graph connectivity. The edge weight $w_{i,j}$ between nodes v_i and v_j is reported in $L_{i,j}$ and $L_{j,i}$. The node degree d_i (the sum of all edge weights from this node) is reported in the diagonal element L_i (in gray). The sum of entries in each column is zero.

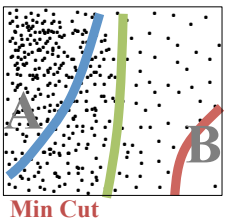

	Min Cut	Ratio Cut	Normalized Cut
Cost	$\text{cut}(A, B)$	$\frac{\text{cut}(A,B)}{\lvert A\rvert} + \frac{\text{cut}(A,B)}{\lvert B\rvert}$	$\frac{\text{cut}(A,B)}{\deg(A)} + \frac{\text{cut}(A,B)}{\deg(B)}$
Energy	$x^T L x$	$\frac{x^T L x}{x^T x}$	$\frac{x^T L x}{x^T D x}$
Description	find minimal cut	find minimal cut with equal size $\lvert A\rvert$ and $\lvert B\rvert$	find minimal cut with equal $\deg(A)$ and $\deg(B)$
Solution	eigenvector of L in null space	smallest non trivial eigenvector of L	smallest non trivial eigenvector of $D^{-\frac{1}{2}} L D^{-\frac{1}{2}}$

Figure 1.3: Comparison of Graph Based Segmentation methods. **Left:** Examples of cuts in an extreme configuration where nodes are all concentrated on the left, sparser on the right, and with a lone node on the bottom right. **Right:** Min Cut finds the cut between region A and B with minimal cost. Ratio Cut adds equipartition (A and B have similar size). Normalized Cut favors similar connectivity within regions (A and B have similar connectivity degree). L is the Laplacian matrix that incorporates edge weighting of the graph, and D is the node degree matrix of the graph.

ical growth) or **functional** changes (e.g., organ mechanics). Moreover, registration becomes particularly challenging when an organ undergoes severe deformation (e.g., a beating heart) or high variation in shape (e.g., brain surface folding patterns).

Historically, registration methods were categorized, firstly, as *rigid* (simple translation and rotations), and secondly, as *nonrigid* (different local transformations). For instance, simple translations could be used to automate digitized angiography (Venot and Leclerc, 1984). Jean Talairach and Pierre Tournoux (Talairach and Tournoux, 1988) proposed in 1967 to study brains in a common stereotactic space (known as the Talairach coordinate system) by warping brains with affine transformations (to correct for translation, rotation and scaling). Furthermore, ambiguous local minima can be avoided with block-matching optimization (Ourselin *et al.*, 2000; Ourselin, 2002). Affine transformations might be sufficient to a register brain image with small changes between scans, however, in most medical applications, complex transformations often exist between images of soft-tissues. A broad variety of algorithms exist with wide surveys in, for instance, (Brown, 1992; van den Elsen *et al.*, 1993; Mcinerney and Terzopoulos, 1996; Maintz and Viergever, 1998; Fitzpatrick *et al.*, 2000), and more recently, in (Zitova, 2003; Crum *et al.*, 2004) with a focus on cardiac imaging (Mäkelä *et al.*, 2002) and brain imaging (Gholipour *et al.*, 2007). They can be categorized in numerous ways (based on their transformation type, type of features (Rohr *et al.*, 2003; Hellier and Barillot, 2003; Azar *et al.*, 2006), similarity measure, optimization scheme (Glocker *et al.*, 2008)). This thesis focuses on the generic approach where each pixel can move freely with a displacement field.

1.2.1 From Deformable Models to Vector Fields

Demetri Terzopoulos *et al.* (Terzopoulos *et al.*, 1987) unified shape and motion description in an elastic deformation model, which is governed by dynamics that resist to stretching, bending, and twisting. In the context of image matching, deformable models are transformed until they match images. The underlying transforming forces are derived from differences between images (Terzopoulos, 1980; Mcinerney and Terzopoulos, 1996). As an aside, this framework already suggests links between image segmentation and image registration since a deformable surface may be used to either segment an image (similarly to the snake mentioned earlier) or register images (by transforming itself from one image to a second image). Rather than using a deformable surface with elastic constraints, Gary Christensen *et al.* (Christensen

et al., 1996) proposed to apply such elastic constraints on a deformation field (i.e., each pixel moves with an associated vector), where external forces guide the matching of images while preserving smoothness of the deformation. Two physical models are presented: an elastic model should be preferred for small displacements (Davatzikos and Prince, 1994) while a viscous fluid model should be considered for large displacement (Christensen *et al.*, 2002; D'Agostino *et al.*, 2003).

Vector fields can also be computed by solving the optical flow equation that assumes a point keeps a similar intensity over time:

$$I_0(x, y) = I_1(x + \delta x, y + \delta y).$$

Bruce Lucas and Takeo Kanade (Lucas and Kanade, 1981a,b) assume a constant point displacement within a neighborhood (all neighboring points move coherently). Such displacement field can be computed locally with the least squares approach. Berthold Horn and Brian Schunk (Horn and Schunck, 1981, 1992) use an iterative approach by minimizing a global energy with a smoothness constraint on the whole displacement field. Additionally, mass conservation constraints can be applied on the displacement field (Song and Leahy, 2002). This is relevant for intra-patient registration of incompressible organs, such as muscles and hearts. Rather than using image intensity to drive the optimization of the displacement field, statistical shape information may also be used (Wang and Staib, 1998, 2000).

Interpolation

In 1917, Sir D'Arcy Wentworth Thompson suggested in his book *On Growth and Form* (Thompson, 1917) that shape differences of living organisms, due to homologies of organs, can be modeled with simple transformations. He often used a grid with few control points to model such transformations. Similarly, the computation of image transformations could benefit from being interpolated with only a few control points in a spline-based transformation model such as the Free-Form Deformations (Sederberg and Parry, 1986; Lee *et al.*, 1995, 1997; Bardinet *et al.*, 1996; Rueckert *et al.*, 1999) or the Thin-Plate Splines (Duchon, 1976; Meinguet, 1979; Bookstein, 1989).

The Free-Form Deformation (FFD) (Sederberg and Parry, 1986) is an extension of a spline-based interpolation model. The underlying model interpolates any position on a

curve from control points that influence the shape of the curve from their positions and a set of weighted basis functions (often B-splines). The interpolation of positions can be similarly extended on surfaces and generalized to any dimension with hyper surface patches. The FFD model is a lattice composed of such hyper patches where each node of the lattice is a control point. Put differently, an object embedded in the lattice is deformed when lattice nodes are displaced. Moreover, the influence of control points in B-splines is local and thus, suitable for localized deformations, however, the direct manipulation is not natural with FFD and may require complex computation of inverse deformations (Hsu *et al.*, 1992). In the medical field, Eric Bardinet *et al.* (Bardinet *et al.*, 1996) used FFD to track motion in cardiac sequences and Daniel Rueckert *et al.* (Rueckert *et al.*, 1999) successfully used FFD for nonrigid registration of medical images.

The Thin-Plate Spline (TPS) (Bookstein, 1989) models the deformation of a plate (e.g., a metallic surface) that passes through several anchor points and has minimal bending. These anchor points directly control the plate. The role of TPS is to interpolate the height of such surface from the positions of the displaced anchor points. Jean Duchon (Duchon, 1976) and later Jean Meinguet (Meinguet, 1979) were the firsts to propose the thin-plate equation as an interpolation model. Fred Bookstein (Bookstein, 1989) proposed to use a relaxation term affecting the influence of all control points over the surface, while Grace Wahba (Wahba, 1990) studied approximation schemes by adding localized uncertainty to each control point (anisotropic case (Rohr *et al.*, 2001)). Furthermore, the generalization to dimensions higher than 3D requires the use of radial basis functions (Sprengel *et al.*, 1996). In a concrete application where a set of features (e.g., salient points, lines) needs to be registered, the correspondences between features are first established before finding the complete transformation with TPS. This is a basis of the Robust Point Matching algorithm (TPS-RPM) (Rangarajan *et al.*, 1997; Chui *et al.*, 1999; Chui and Rangarajan, 2000b,a) where registration is performed by alternatively updating correspondence and transformation. TPS may also be used for segmentation (Lombaert and Cheriet, 2010a)[3].

1.2.2 Diffeomorphic Registration

Modeling the transformation between two images with free displacements at each pixel does not guarantee diffeomorphism (one-to-one mapping). For instance, space can be wrongly

[3]First use of TPS for image segmentation (Lombaert and Cheriet, 2010a) proposed by the candidate

folded, which would break smoothness and reversibility of the transformation. One approach is to use an interpolation model that guarantees smoothness and diffeomorphism, for instance, by using FFD (Rueckert *et al.*, 2006; Craene *et al.*, 2009) or TPS (Marsland and Twining, 2002). Diffeomorphism might be guaranteed by constraining the registration to small displacements (i.e., the shape is slowly deformed with small and tractable updates). This may be, therefore, a long and costly process. Following this idea, a temporal dimension may be added to the displacement field $\phi(x,t)$ which becomes related to a velocity field $v(x,t)$ with:

$$\frac{d}{dt}\phi(x,t) = v(x+\phi(x,t),t). \tag{1.1}$$

Large Deformation Diffeomorphic Metric Mapping (LDDMM) (Miller *et al.*, 2002; Beg *et al.*, 2005) uses non-stationary velocity fields (velocities vary through time) to compute geodesic paths on a tangent space of a Riemannian manifold (i.e., in a high dimensional space that represents all possible transformations, one path describes the deformation of an image from one configuration to another). Multiscale kernels for LDDMM (Risser *et al.*, 2010a,b, 2011; Sommer *et al.*, 2011a; Sommer, 2011a; Sommer *et al.*, 2011b) decouple the contribution of each smoothing scale (highly localized versus global deformation) and multi-scale characteristics can evolve simultaneously for possibly faster convergence and multi-scale statistics. These approaches typically consist of elegant mathematical frameworks but are, however, computationally expensive. Stationary velocity fields (fixed velocities) simplify the computation of LDDMM (using exponential maps) but still remain costly to compute (Ashburner, 2007; Hernandez *et al.*, 2007). A solution to Eq. (1.1) may indeed be:

$$\phi(x) = \exp(v(x)) \tag{1.2}$$

The Demons algorithm (Thirion, 1998; Vercauteren *et al.*, 2007, 2008, 2009a) considers the velocity field as stationary. It shows similar results than stationary LDDMM, however, it uses a much more efficient algorithm (Hernandez *et al.*, 2008) that is suitable for fast image registration. The original (non-diffeomorphic) algorithm (Thirion, 1998) decoupled the optimization of the similarity term and the regularization term by introducing a hidden variable (i.e., demons forces) that transfers correspondence updates between alternations (i.e., demons residing in all pixels warp the displacement field from their positions). The

minimized energy (Cachier *et al.*, 2003) in the Demons algorithm is:

$$E(I_0, I_1, c, \phi) = \alpha_i^2 \text{Sim}(I_0, I_1 \circ c) + \alpha_x^2 \text{dist}(c, \phi)^2 + \alpha_T^2 \text{Reg}(\phi). \tag{1.3}$$

In a first step, that is controlled with weight α_i, the transformation ϕ is fixed, and the similarity term is optimized (e.g., $\text{Sim} = (I_0(x) - I_1 \circ c(x))^2$), which yields the updates $u = \delta c$ of the correspondences between I_0 and $I_1 \circ \phi$. In a second step, that is controlled with weight α_T, the correspondences c are fixed, and the regularization takes place (e.g., $\text{Reg} = \nabla ||\phi||^2$) by smoothing the transformation $\phi = K \star (c + u)$ with a convolution kernel K. Ideally, correspondences c must be close to the transformation ϕ, e.g., $\text{dist}(c, \phi) = ||c - \phi||$, and α_x controls how they may differ (i.e., the maximal update step is α_x^{-1}).

Additionally, the Demons algorithm can be defined to enforce incompressibility of the deformation (Mansi *et al.*, 2009, 2010; Mansi, 2010), which is satisfied by directly enforcing the Jacobian to one (Rohlfing *et al.*, 2003; Haber and Modersitzki, 2004) or using a divergence-free constraint (Bistoquet *et al.*, 2008; Saddi *et al.*, 2007). Other variants include the multichannel Demons, which allows spatio-temporal registration (Peyrat *et al.*, 2008, 2010; Forsberg *et al.*, 2011), or the Spherical harmonics, which enables the registration of spherical meshes, such as brain surfaces (Yeo *et al.*, 2010a). We now provide a walk-through of the evolution of the classical demons (Thirion, 1998) in order to understand its symmetric diffeomorphic version (Vercauteren *et al.*, 2008, 2009a).

Classic Demons The classical Demons (Thirion, 1998) decouples, as mentioned earlier, the minimization of Eq. (1.3) by computing and regularizing alternatively an update field described in Alg. (1). The smoothing of the updates with a convolution kernel K_{fluid} (typically Gaussian) acts as a fluid-like regularization, while the smoothing of the correspondences with K_{diff} acts as a diffusion-like regularization. Moreover, additions of displacement fields, such as $\phi + u$, can be replaced with compositions, $\phi \circ u$ (i.e., $(\phi \circ u)(x) = \phi(u(x))$).

Log Demons The classical Demons algorithm (Thirion, 1998) is not diffeomorphic in nature. The transformation can create folding in space and there is no guarantee that the transformation is invertible. To this effect, the Diffeomorphic Demons (Vercauteren *et al.*, 2007) optimizes Eq. (1.3) over the space of diffeomorphism (i.e., the space of smooth and invertible transformations such that $\phi^{-1}\phi = I$) through the use of exponentials of velocity

Algorithm 1 Classical Demons

Input: Images F, M, and initial transformation ϕ.
Output: Transformation ϕ mapping F to M.

repeat

- Find updates u mapping F to $M \circ \phi$
  $u(p) = -\frac{\delta(p)}{||J^p||^2 + \frac{\alpha_x^2}{\alpha_i^2}|\delta(p)|} J^{pT}$,
  with $\delta(p) = F(p) - M \circ \phi(p)$ and $J^p = -\nabla_p^T (M \circ \phi)$
- Smooth updates: $u \leftarrow K_{\text{fluid}} \star u$.
- Update transformation: $\phi \leftarrow \phi + u$.
- Smooth transformation: $\phi \leftarrow K_{\text{diff}} \star \phi$.

until convergence

Algorithm 2 Log Demons

Input: Images F, M, and initial velocity field v.
Output: Transformation $\phi = \exp(v)$ mapping F to M
ϕ is diffeomorphic and invertible ($\phi^{-1} = \exp(-v)$)

repeat

- Find updates u mapping F to $M \circ \exp(v)$.
- Smooth updates: $u \leftarrow K_{\text{fluid}} \star u$.
- Update velocity field: $v \leftarrow \log(\exp(v) \circ \exp(u))$
  (BCH approximation: $v \leftarrow v + u$).
- Smooth velocity field: $v \leftarrow K_{\text{diff}} \star v$.

until convergence

Algorithm 3 Symmetric Log Demons

Input: Images F, M, and initial velocity field v.
Output: Transformation $\phi = \exp(v)$ mapping F to M
ϕ is diffeomorphic, invertible, and symmetric.

repeat

- Find updates $u_{F \to M}$ mapping F to $M \circ \exp(v)$.
- Find updates $u_{M \to F}$ mapping M to $F \circ \exp(-v)$.
- Average updates: $u \leftarrow \frac{1}{2}(u_{F \to M} - u_{M \to F})$.
- Smooth updates: $u \leftarrow K_{\text{fluid}} \star u$.
- Update velocity field: $v \leftarrow \log(\exp(v) \circ \exp(u))$
  (approximated with $v \leftarrow v + u$).
- Smooth velocity field: $v \leftarrow K_{\text{diff}} \star v$.

until convergence

Algorithm 4 Exponential $\phi = \exp(v)$ of a velocity field v

Input: Velocity field v.
Output: Diffeomorphic exponential map $\phi = \exp(v)$.

- Choose N such that $2^{-N}v$ is close enough to 0,
  e.g., such that $\max ||2^{-N}v|| \leq 0.5$ pixels.
- Scale velocity field $\phi \leftarrow 2^{-N}v$.

for N times **do**

- Square $\phi \leftarrow \phi \circ \phi$.

end for

Figure 1.4: Walk-trough of the evolution of the classical demons toward its symmetric diffeomorphic version (differences between algorithms are colored).

fields that are estimated with the scaling and squaring method described in Alg. (4). The Log-Demons algorithm computes updates on velocity fields rather than on displacement fields (Alg. (2)). Consequently, the transformation ϕ is simply a byproduct defined by $\phi = \exp(v)$. Such an exponential map guarantees diffeomorphism of the transformation (i.e., the exponential is diffeomorphic by nature, $\exp(v)^{-1} \exp(v) = \exp(-v + v) = I$). Additionally, the inverse of a transformation comes at no extra cost with $\phi^{-1} = exp(-v)$. Moreover, the composition of velocity fields can be estimated with the Baker-Campbell-Hausforff (BCH) formula (Hausdorff, 1906; Campbell, 1897; Poincaré, 1899; Baker, 1902), where, for instance (Vercauteren *et al.*, 2008; Bossa *et al.*, 2007) shows that a satisfactory approximation could be $\exp(v) \circ \exp(u) \approx \exp(v + u)$.

Symmetric Demons In the Symmetric Demons algorithm (Vercauteren *et al.*, 2008, 2009a), the computation of the transformation from I_0 to I_1, regardless to their order, yields the same symmetric transformations, i.e., the algorithm is robust to the order of input images: $\phi_{I_0 \to I_1} = \phi^{-1}_{I_1 \to I_0}$. This is achieved by averaging symmetric velocity fields in Alg. (3). Overall, this algorithm presents, therefore, a framework with diffeomorphism, reversibility, and symmetry, which are properties that will be retained during the development of the thesis.

Alternatively, symmetry was also proposed in earlier frameworks, for instance, (Christensen and Johnson, 2001; Johnson and Christensen, 2002; Kumar *et al.*, 2003) use an inverse consistency of transformations. On the same line, symmetry has been extended to 3 images, which yields the notion of transition where $\phi_{1 \to 3} = \phi_{1 \to 2} \circ \phi_{2 \to 3}$ (Christensen and Johnson, 2003; Geng *et al.*, 2005; Skrinjar and Tagare, 2004).

1.2.3 Note on Similarity Measures

In earlier methods, registration was performed by manually selecting feature points (Maintz and Viergever, 1998). The automation of the registration required the use of similarity measure between images, and in particular between images of different modalities (e.g., CT and MRI). To this effect, intensity based approaches were introduced in (Woods *et al.*, 1992) where the variance of the intensity ratios between images is minimized. Soon after, Derek Hill *et al.* (Hill *et al.*, 1993) introduced the use of a joint histogram for multimodal registration. In the same group, Petra Van den Elsen *et al.* (van den Elsen *et al.*, 1995) introduced a method based on the correlation of the gradient.

Mutual Information André Collignon *et al.* (Collignon *et al.*, 1995b) and Colin Studholme *et al.* (Studholme *et al.*, 1995) proposed to use entropy as a similarity measure. While many definitions exist for the entropy, including the one from Claude Shannon (Shannon, 1948), they all measure the information contained in a message or an image (these concepts are in fact derived from the machine learning and information theory fields (Shannon, 1948)). Later, Paul Viola and Sandy Wells (Viola and Wells, 1995) and André Collignon *et al.* (Collignon *et al.*, 1995a; Maes *et al.*, 1997) found, in the same year, a new approach based on the mutual information between two images, which uses the joint intensity histogram between these two images.

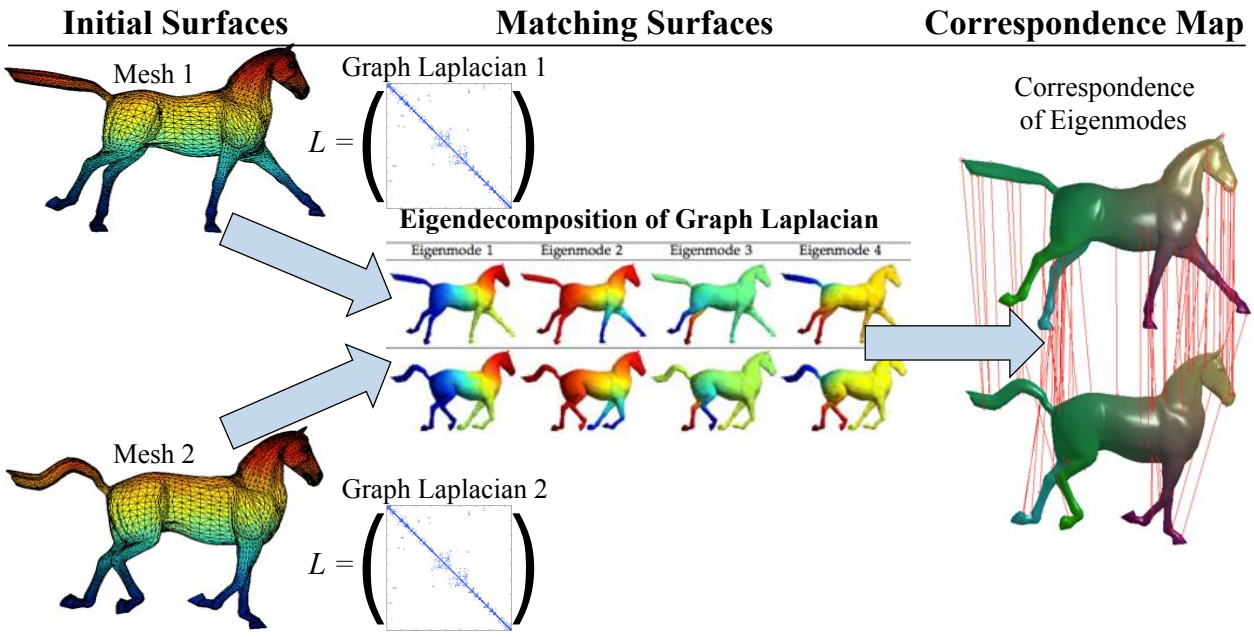

Figure 1.5: Spectral matching consists of matching spectral signatures (eigenvectors of the Laplacian matrix associated with each mesh).

This similarity measure may, however, vary depending on the overlap between images. If their overlap is small, the mutual information between both images is going to be small, and *vice versa*. To correct this bias, the normalized mutual information was introduced in (Studholme, 1999). Recent methods based on mutual information are surveyed in (Pluim *et al.*, 2003). It is in fact known to be one of the best suitable similarity measures for multimodal registration.

Kullback-Leibler distance Another important similarity measure, also derived from the information theory, is the Kullback-Leibler distance (Kullback, 1959), which can be used for registration (Chung *et al.*, 2002; Guetter *et al.*, 2005). It measures the distance from an expected distribution (a joint histogram) and an observed distribution. This similarity measure proved to be more robust in methods using *a priori* distributions in a multiresolution approach (Gan *et al.*, 2004). The same authors of the later work also proposed to use joint class histograms (Chen *et al.*, 2003), which use labels instead of intensities.

1.2.4 Spectral Matching

Spectral matching (Chung, 1997b) utilizes a graph *spectrum* to produce a vertex correspondence between two graphs (or meshes), and essentially consists of matching the spectral decomposition of matrices representing graphs. The decomposition of a graph Laplacian (Grady and Polimeni, 2010) into a set of Laplacian eigenvalues and eigenvectors is illustrated in Fig. 1.5 and reveals the shape harmonics, which provide a spectral signature for each point of a mesh. Matching vertices are consequently found with similarities in these spectral signatures.

Spectral methods have been used in many fields, such as in the segmentation and registration of shapes in images (Reuter, 2009; Leordeanu and Hebert, 2005), in the indexing of structures from structural trees (Shokoufandeh *et al.*, 2005), in shape modeling (Chung *et al.*, 2011; Seo *et al.*, 2011), graph partitioning (Chan *et al.*, 1995), or in clustering of data (Shi and Malik, 2000; Bengio *et al.*, 2004; Bach and Jordan, 2004). Their use in shape matching is often limited to hierarchical matching (e.g., matchings limbs, or large surface areas, in body models). Few medical applications of spectral methods exist and are mostly targeted to brain studies, for instance, on the geometric patterns of the cortical folding (Niethammer *et al.*, 2007; Reuter *et al.*, 2009; Shi *et al.*, 2009, 2010), or on the smoothing of brain surfaces (Anqi *et al.*, 2006).

Shinji Umeyama (Umeyama, 1988) and later Guy Scott and Christopher Longuet-Higgins (Scott and Longuet-Higgins, 1991) pioneered the use of spectral methods for the correspondence problem. On the same ideas, Larry Shapiro and Michael Brady (Shapiro and Brady, 1992) compared ordered eigenvectors of a proximity matrix to find correspondences. Their work served as a basis for future spectral correspondence methods. For instance, variants include the use of different proximity matrices using different kernels, the use of the adjacency matrix, shock graphs (Pelillo *et al.*, 1999; Siddiqi *et al.*, 1999), different normalized Laplacian matrices, or the use of Multi-Dimensional Scaling (Schwartz *et al.*, 1989; Elad and Kimmel, 2003; Wuhrer *et al.*, 2009; Bronstein *et al.*, 2006, 2007). Recent surveys covering the use of spectral methods in the past fifteen years are available in (van Kaick *et al.*, 2011; Zhang *et al.*, 2010). Among these methods, Diana Mateus *et al.* (Mateus *et al.*, 2008) proposed an original unsupervised spectral method with an alternative approach to spectral reordering that is based on the use of eigenvector histograms and the alignment of eigenvectors with a probabilistic point matching (Carcassoni and Hancock, 2003). Varun Jain and Richard

Zhang (Jain and Zhang, 2006) also addressed the alignment of eigenvectors with a nonrigid deformation where they used the Thin Plate Splines.

1.2.5 Joint Segmentation and Registration

As segmentation and registration can both be solved through energy minimization, they may be combined into the same optimization framework. This approach has been little explored. Typically, feature-based registration uses points (Arun *et al.*, 1987), edges or surfaces (Borgefors, 1986, 1988; Zhang, 1994) found in a preliminary segmentation step. The registration depends, therefore, on this feature localization. Rather than seeing both steps as separate processes, they can be solved simultaneously. For instance, in (Bansal *et al.*, 1998), segmentation and registration are solved in a unified framework; however, two steps are still used in this algorithm, where segmentation and registration are solved individually. In (Yezzi *et al.*, 2003, 2001), both processes are expressed in the same formulation within a variational framework. Nonrigid transformations are also addressed in (Unal and Slabaugh, 2005). Furthermore, solving simultaneously segmentation and registration within the same framework can answer the registration of partial data as proposed in (Periaswamy and Farid, 2003, 2006). More precisely, the registration uses a joint segmentation approach that treats differently the existing and missing data between two images.

The levelset method also offers the possibility to simultaneously segment and register two images (Leventon *et al.*, 2002; Tsai *et al.*, 2003). Other unified approaches are based on the Expectation Maximization algorithm (Dempster *et al.*, 1977), where pixels are simultaneously labeled and matched to an atlas (Pohl *et al.*, 2005, 2006). Alternatively, a probabilistic framework is presented in (Ashburner and Friston, 2005), where image registration, tissue classification, as well as bias field correction, are computed with a mixture of Gaussian models that incorporates intensity variation, nonrigid displacements, and tissue probability maps.

These joint approaches provide elegant unifying frameworks. In fact, **solving simultaneously what may be one unique problem** could become a leading approach for both segmentation and registration. Moreover, the earlier review of spectral graph methods showed strong links with the decomposition of the Laplacian matrix (Shi and Malik, 1997, 2000; Meila and Shi, 2000). This decomposition was shown in the last section to be at the core of spectral matching methods and suggests that further links exist between segmentation

and registration via the decomposition of the Laplacian matrix.

1.3 Atlas Construction

A medical atlas compounds statistics from a population by providing an average data with its global or local variability. It enables numerous studies, for instance, in neuroimaging, such as the analysis of Alzheimer's disease, asymmetry in brain hemispheres, classification of population, or characterizing specific syndromes ((Toga and Thompson, 2001) provides a survey in neuroimaging applications). Emerging methods (Grenander and Miller, 1998) in segmentation and registration, reviewed above, plays a crucial role in ameliorating the accuracy and precision of atlases. Homology, which is the basis for comparative biology (studying similar structures or functions), brings the philosophical debate (Fitch, 2000) on how to construct an atlas (Crum *et al.*, 2003). Indeed, the same anatomical structure (e.g., one brain area) may have different functions between individuals, while on the contrary, one same function may be performed by different structures between individuals (e.g., the variability of cognitive areas over the brain surface (Maudgil *et al.*, 1998)). One arising question is: should we register functional areas (e.g., Brodmann areas), or should we register structural areas (e.g., sulcal patterns) in order to perform relevant statistical analysis? **The emphasis on structural alignment (i.e., averaging shapes)** is favored as it allows the important transfer of functional data to structural data (Miller and Qiu, 2009). This transfer is at the core of neuroimaging studies on functional variability since it determines the locations of neural functions and establishes links between them.

Among the active groups constructing reference atlases, there are, in neuroimaging, the Montreal Neurological Institute (MNI) (Evans *et al.*, 1993), the International Consortium for Brain Mapping (ICBM which targets the use of 7,000 subjects) (Mazziotta *et al.*, 2001a,b), or in cardiac imaging, the euHeart initiative (Smith *et al.*, 2011) the Cardiac Atlas Project (CAP, targeting 3,000 subjects) (Backhaus *et al.*, 2010; Fonseca *et al.*, 2011; Young and Frangi, 2009), as well as other cardiac physiome projects (Smith *et al.*, 2009; Coveney *et al.*, 2011). Additionally, from unique datasets of cardiac fiber imaging, a few statistical atlases have been constructed with canine hearts (Peyrat *et al.*, 2007; Peyrat, 2009; Helm, 2005; Helm *et al.*, 2005a,b; Peyrat *et al.*, 2007; Gilbert *et al.*, 2007) and human hearts (Lombaert *et al.*, 2011b,e, 2012b)[4].

[4]First human atlas (Lombaert *et al.*, 2011b,e, 2012b) proposed by the candidate

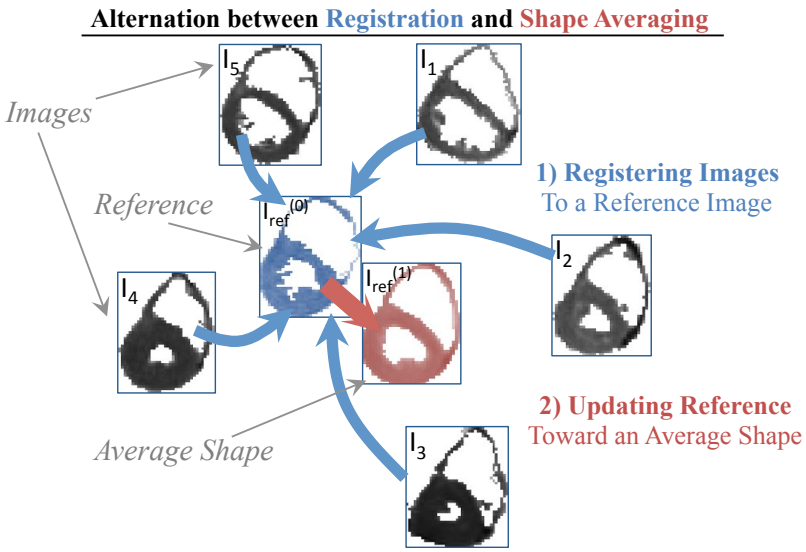

Figure 1.6: Averaging the shape of all images, based on Guimond *et al.* (Guimond, 2000). Firstly, all images are registered to an arbitrary reference image (blue arrows). Secondly, the reference is transformed with the average of all inverse transforms (ideally, this average should be zero, meaning that the reference image is at equal distance to all images) (red arrow). Lastly, the averaging process is iterated until convergence.

The increasing number of possible subjects and the advances foreseen in medical imaging, notably with better resolutions, create a demand for more efficient methods in the construction of atlases. Most atlases are currently built using sequential registration and segmentation. For instance, a model can be adapted to find classifications of tissues (Cocosco *et al.*, 2003; van Leemput *et al.*, 1999; Fischl *et al.*, 2004; Marroquin *et al.*, 2003) with the use of probabilistic maps (Cocosco *et al.*, 2003) or Expectation-Maximization approaches (van Leemput *et al.*, 1999). Joint approaches also exist to build atlases with a simultaneous segmentation and alignment of images (Leventon *et al.*, 2002; Tsai *et al.*, 2003; Ribbens *et al.*, 2009, 2010). In more details, the following sections will review the major methods for the computation of the average shape as well as for performing statistical analysis on atlases.

1.3.1 Average Shape

The underlying idea of building an atlas is to find the average shape of all images, and to register all these images to their average. The registered set of images enables consequently statistics within a common reference. More precisely, the average shape is computed iteratively by deforming an arbitrary reference until it reaches an average shape that best

describes all other images (illustrated on Fig. 1.6).

Deformation may be controlled with, for instance, manual placement of control points, mostly using thin-plate splines (Park *et al.*, 2003; Peyrat *et al.*, 2007). Images can also be registered automatically and jointly (group-wise) to their average shape. For instance, the sum of all displacements fields, transforming images to the reference, can be forced to be zero and to have minimal joint entropy between images (Studholme and Cardenas, 2004). A simultaneous approach registering all images concurrently is also possible (Bhatia *et al.*, 2004). Moreover, the modes of shape variations may be optionally provided (Gotland and Sabuncu, 2011). As an alternative approach, the average shape may be computed from a multi-dimensional scaling on a distance matrix, where the notion of distance between images is expressed as a bending energy (in terms of thin-plate splines (Park *et al.*, 2005) or polyharmonic splines (Marsland *et al.*, 2003; Marsland and Twining, 2004)).

The concept of geodesic averaging instead of a Euclidean averaging is at the core of diffeomorphic methods. It allows unbiased construction of atlases (Avants and Gee, 2004a; Joshi *et al.*, 2004) (i.e., the average shape is independent from the choice of initial reference). This initial reference evolves toward the shape with the least difference with the rest of images (Joshi *et al.*, 2004), where differences are expressed in the space of diffeomorphism. More precisely, the transformation of the reference shape follows a geodesic path on a Riemannian manifold (the space of diffeomorphic transformations) defined by the images and converges to a point representing the optimal transformation (the geodesic average on the manifold) from the reference to the rest of images (Avants and Gee, 2004b,a). The LDDMM framework provides an elegant mathematical framework for averaging shapes (Beg and Khan, 2006; Bossa *et al.*, 2007), with for instance, probabilistic segmentations (Lorenzen *et al.*, 2006), or with the automatic placement of landmarks for the construction of a ventricle atlas (Beg *et al.*, 2004). When a large number of images is available, Lilla Zöllei *et al.* (Zollei *et al.*, 2005) propose a method also based on this central tendency and minimizes the total pixel entropy using a multiresolution gradient descent.

More advanced techniques (Allassonnière *et al.*, 2007, 2010) recently surfaced with a forward scheme approach (i.e., formulating the registration functional with transformations of the atlas toward all images, instead all images toward the atlas) and handle photometric (intensity), and geometric (shape) variations in a Bayesian framework where geometric variability is modeled as a hidden variable. The reference image may also be generated from a

hyper-template (Ma *et al.*, 2008). Moreover, forward schemes may alternatively be modeled with *currents* (Durrleman *et al.*, 2009b, 2008, 2011) to represent average and modes of geometric shapes. This use of *currents* has been successfully applied to build an atlas of white matter fibers (Durrleman *et al.*, 2009a). However, the complexity of these approaches may be perceived as prohibitive.

Alexandre Guimond *et al.* (Guimond, 2000) proposed an efficient iterative algorithm to construct atlases. They use a groupwise diffeomorphic registration, where an initial reference $I_{\text{ref}}^{(0)}$ evolves toward a geodesic average $I_{\text{atlas}} = \text{atlas}(I_1, ..., I_n)$. More precisely, at each iteration $(t+1)$, the previous reference image, $I_{\text{ref}}^{(t)}$, is updated with the average of all inverse transforms:

$$I_{\text{ref}}^{(t+1)} = \frac{1}{N} \sum_{i=1}^{N} I_i \circ \phi_{I_i \to I_{\text{ref}}^{(t)}} \circ \phi_{I_{\text{ref}}^{(t)} \to I_{\text{ref}}^{(t+1)}}. \tag{1.4}$$

The initial reference image may be initialized arbitrarily with any image (e.g., $I_{\text{ref}}^{(0)} = I_1$). The transformation maps and velocity fields from all hearts to the reference image (all $\phi_{I_i \to I_{\text{ref}}^{(t)}}$'s and $v_{I_i \to I_{\text{ref}}^{(t)}}$'s) are computed using separate diffeomorphic registrations. It is assumed that using the average of all transformations (all $\phi_{I_{\text{ref}}^{(t)} \to I_i}$'s) will bring the reference image closer to the average shape. Recalling that the transformation is, for stationary velocity fields, the exponential map of the velocity field $\phi^{-1} = exp(-v)$, hence, the transformation update is:

$$\begin{aligned}
\phi_{I_{\text{ref}}^{(t)} \to I_{\text{ref}}^{(t+1)}} &= \phi^{-1}_{I_{\text{ref}}^{(t+1)} \to I_{\text{ref}}^{(t)}} \\
&= \exp\left(-v_{I_{\text{ref}}^{(t+1)} \to I_{\text{ref}}^{(t)}}\right) \\
&= \exp\left(-\frac{1}{N} \sum_{i=1}^{N} v_{I_i \to I_{\text{ref}}^{(t)}}\right),
\end{aligned}$$

In practice, Guimond's algorithm converges to an average shape with about five iterations; however, it necessitates, in each iteration, complete registrations of all images to the reference image.

1.3.2 Statistical Analysis

Statistics on atlases may be performed on the analysis of the average shape, but may also be extended to second-order measures. For instance, the complex variations (Chung, 2001) of

the underlying deformations between shapes or images may be analyzed through a Principal Component Analysis (PCA) on point coordinates (Cootes and Taylor, 1994; Cootes *et al.*, 1995). The variations can be studied on a point level (how the position of a point varies in space across a dataset), or more interestingly, on a global level (how a point varies with respect to the other points). Indeed, it must be remembered that the variations of point positions are partially correlated across the surface or image. In fact, if the position of a point varies along a certain direction, there is a high chance that the positions of the neighboring points will vary in a similar manner. A study on the global shape variations will reveal the correlations between all points. The principal modes of variations have many applications, including classification (Timoner *et al.*, 2002), recognition (Turk and Pentland, 1991; Golland *et al.*, 2001), or can also guide image registration (Ashburner *et al.*, 1997) by favoring deformations modeled with the strongest principal modes of variations.

However, the study of more complex characteristics, such as diffusion tensor fields (Jones *et al.*, 2002), may not be adequate with the standard Euclidean metric. To this effect, non-linear structures may be modeled with Riemannian geometry, for instance, to model deformations (Rumpf and Wirth, 2009, 2011) in a shape space (Kendall, 1984; Durrleman *et al.*, 2009b; Durrleman, 2010) or to model the space of diffusion tensors (symmetric positive definite matrices, or the SO3 Lie group) (Fletcher and Joshi, 2007, 2004; Pennec *et al.*, 2006). Statistics such as the mean, variance, or PCA can be computed directly on the manifold yielding tools such as geodesic PCA (Fletcher and Joshi, 2004; Fletcher *et al.*, 2004; Basser and Pajevic, 2000; Sommer, 2011b; Huckemann and Hotz, 2009; Huckemann *et al.*, 2010; Huckemann, 2011). Intrinsic manifold distances are used instead of standard Euclidean distances in order to capture the complexity of the underlying Riemannian structure.

Statistics can also be performed on a tangent space of a Riemannian manifold (Pennec *et al.*, 2006). For instance, statistics on diffusion tensor imaging can be computed on an exponential map with the Log-Euclidean metric (Arsigny *et al.*, 2006a,b). Including registration uncertainties can also improve tissue classification (Simpson *et al.*, 2011). In fact, the Log-Euclidean metric is preferred for its simple and fast framework. Moreover, the mean is computed using simple Euclidean metric on logarithm of diffusion matrices with the Karcher mean (Karcher, 1977) (also known as the Fréchet mean).

1.4 Summary

The fields mentioned in the literature review are summarized in Fig. 1.1. The problem of constructing atlases has been long studied and involves many related fields of research such as image segmentation and image registration. All these topics are interconnected and it is possible to foresee new approaches bridging gaps between image segmentation, registration, and shape averaging. This thesis proposes new tools that provide improvements in the reviewed areas. The next chapter gives the principal research questions with the main objectives, as well as the rationale behind the methodology.

CHAPTER 2 GENERAL METHODOLOGY

2.1 Problem Statement

Every clinician's goal is to, firstly, detect abnormality in an individual, secondly, judge if it poses risk to life and, ultimately, cure it. The question of how to define the normality of complex characteristics naturally arises in this context (what is normal and how it varies). However, such **statistics on anatomical and functional variability requires the normalization of measurements across subjects**. Such normalization (the construction of atlases) maps many subjects, structures, and measurements into a common space and highlights the variability of complex characteristics within a population, for instance, the variability of the fiber architecture, brain surfaces or locations of cerebral activities. Analyzing the variability of such complex characteristics depends on the quality of image acquisition, the segmentation of these images (extracting the regions of interests), and the registration of all images into a reference space (finding the correspondences between images). Notwithstanding the challenges of each step, image registration becomes particularly complex (still unsolved) with organs undergoing a large shape variation due to inter-patient or intra-patient differences in morphology (e.g., growth) or function (e.g., a beating heart). This thesis addresses the general question: **How to produce tools for analyzing the variability of complex anatomical characteristics?**

2.2 Research Questions

The specific research questions are focused on the development of shape normalizing methods demonstrating a sufficient accuracy that enables new advances in the understanding of the human anatomy.

2.2.1 How to Characterize the Human Cardiac Fiber Architecture?

Question 1: *How to construct an atlas from DTMRI that characterizes the average cardiac fiber architecture and its variability in the human population?* ■

The cardiac fiber architecture is still not fully known in humans, and yet, plays an important role in electro-physiology, in mechanical functions, and in remodeling processes. Recently, DTMRI can provide measures of diffusion within the fiber structures. The construction of a statistical atlas from DTMRI of the cardiac fiber architecture (Chapter 3), (Lombaert *et al.*, 2011b, 2012b) will allow a better understanding of its complete nature in humans along with its normal variability (Lombaert *et al.*, 2012b, 2011d,e). Such an atlas will also be relevant to new methods detecting abnormalities derived from the fiber architecture (Lombaert *et al.*, 2011c).

2.2.2 How to Capture Large Anatomical Shape Variabilities?

Question 2: *How can new spectral methods find correspondences between anatomical structures exhibiting large shape variability?* ■

The matching of anatomical structures is currently limited in the presence of a significant variability in shapes (due to morphology or function). Two well-known examples are: the human brain surface which is highly convoluted, and the heart shape which undergoes a complex deformation. A new spectral method for fast and accurate surface matching (Chapter 4) (Lombaert *et al.*, 2011a, 2012a) will be relevant to neuroimaging studies of structural variability and functional variability (Chapter 5). Additionally to surface matching, a new spectral approach for image registration will allow the capture of large cardiac deformations due to the inter-subject and intra-subject variability (Chapter 5).

2.3 Research Objectives

The preceding questions suggest the main research objective: **To produce tools for the normalization of complex anatomical structures**, as well as to illustrate the use of these tools in challenging and clinically relevant variability studies. The main target is fulfilled in the following three specific objectives.

2.3.1 Build a Human Atlas of the Cardiac Fiber Architecture

Objective 1: *To construct a human atlas of the cardiac fiber architecture that permits a variability study within a population and an investigation on the heart abnormality.* ■

The study of the cardiac fiber architecture has been limited until recently to tedious histological slices. Diffusion imaging can now be used to measure the orientation of the fiber structures in three dimensions. In order to construct an atlas of diffusion images, **an automated framework** will be developed (Chapter 3) (Lombaert *et al.*, 2011b, 2012b). The construction of the first human atlas of the cardiac fiber architecture will allow the study on the **variability of the fiber structures** within a population and across several segments of the myocardium (Lombaert *et al.*, 2012b, 2011e). The laminar organization of the fibers can also be analyzed in order to verify the hypothetic **presence of two populations of laminar sheets** (Lombaert *et al.*, 2011d). Furthermore, the comparison of hearts, classified as abnormal, against the atlas of healthy hearts provides answers on whether **differences in the fiber architecture might explain abnormality** (Lombaert *et al.*, 2011c).

2.3.2 Improve Spectral Correspondence

Objective 2: *To improve spectral correspondence with high accuracy and fast computation time in order to measure large anatomical and functional variability on human brain surfaces.* ■

The matching of surfaces with large shape variability is often a compromise between accuracy and fast computation time. An encouraging path is the use of spectral representations that have the benefit of being largely invariant to changes in poses. This property allows, therefore, the matching of shapes with high variability between configurations, or undergoing large deformations. One fit example of such shapes is the human cerebral cortex because its surface is highly convoluted and varies significantly across individuals. The matching of human brain surfaces constitutes, therefore, an excellent application for benchmarking robustness and performance of **a new fast and accurate spectral correspondence** method (Chapter 4) (Lombaert *et al.*, 2011a, 2012a). Furthermore, the spatial normalization of brain surfaces will allow a **variability study on the location of several sulcal regions** (Lombaert *et al.*, 2011a, 2012a) and of **several cognitive areas** (Chapter 5), as well as an **exploration of the shape of the primary cortical folding** (Chapter 5).

2.3.3 Unify Spectral Correspondence and Demons

Objective 3: *To unify spectral correspondence with Demons in order to create a joint*

image segmentation and image registration method capable of handling large deformations of shapes. ■

Image registration methods currently depend on local iterative approaches driven by update forces derived from intensity metrics. They are limited when strong disparities exist between images and results in the incapacity to compute local coherent updates between images. This is often palliated with a rigid alignment of images prior to their registration in order to have maximal local coherence, however, this does not solve the underlying problem. It is proposed to create a **new registration method based on the Demons algorithm that performs spectral correspondence** instead of using image gradients (Chapter 5). The new method combines the advantages from both approaches, namely, diffeomorphism, transformation reversibility, symmetry, and the natural handling of large deformations invariant to isometry. Furthermore, this unification effectively creates a **joint approach for image segmentation and image registration**.

2.4 General Approach

Methods for shape-normalization, which are essential for the construction of atlases, have been developed and enabled new variability studies on complex characteristics that lead to a better understanding of the human anatomical and functional variations of the heart and the brain. The formulated hypotheses have been verified by the achievement of the research objectives. This is demonstrated in two ways, firstly by developing new methods for atlas construction, and secondly by studying their variability.

2.4.1 Atlas Construction

Three new shape normalizing methods relevant to the construction of atlases have been developed. They permitted the construction of an atlas of cardiac fibers, the spatial normalization of brain surfaces, and settled new limits for the registration of images with large deformations.

The first objective investigated the practical value of using fast and simple methods to construct an atlas. It was experimented with cardiac images of diffusion tensors. The average heart was computed iteratively by evolving a reference heart toward the shape that best describes all provided hearts. In each iteration, the registrations of all hearts to the evolving

reference were facilitated by the use of simplified images with diffeomorphic Demons. These simplified images were generated with a semi-automatic method based on Graph Cuts using topological constraints. The diffusion images associated with each heart were warped into a common space to create the first human atlas of the cardiac fiber architecture (Chapter 3) (Lombaert *et al.*, 2011b, 2012b).

The second objective investigated a new approach to match surfaces of an anatomy having a large disparity, with an application on highly convoluted human brain surfaces. The matching of such surfaces provides, indeed, a well-suited benchmarking environment for the robustness, the accuracy and the performance of the new spectral approach. The developed method relies on a spectral representation of a graph derived from the surface that is invariant to isometry and naturally handles a large disparity between surfaces. Features helping alignment can be additionally incorporated in a graph in terms of weighting on graph edges (as differences of feature values) and also on graph nodes (as pointwise information). The correspondence between surfaces are generated with the nearest neighboring points (a fast approach) between their associated spectral representations. This new spectral correspondence method improves accuracy while maintaining a fast computation time. Spatial normalization of several brain surfaces with a variety of shapes was possible and revealed a minimal error (Chapter 4) (Lombaert *et al.*, 2011a, 2012a). Moreover, the method also performed successfully with more generic surface meshes, improving, in a broader scope, the challenge of general mesh correspondence (Lombaert *et al.*, 2012a).

The third objective aimed at introducing a fast image registration method that naturally handles large deformations. The new method uses diffeomorphic Demons where updates are computed with spectral correspondence. Properties of both approaches are effectively combined. Furthermore, the links between spectral methods and segmentation algorithms bring the intriguing fact that the new, so-called *Spectral Demons*, method is actually a joint approach bridging a gap between image segmentation and image registration. The experiments showed that images of hearts with a large shape disparity could be registered and segmented correctly and jointly with this new method (Chapter 5).

2.4.2 Variability Study

The construction of atlases and the development of new shape normalization methods enabled the possibility for new studies on anatomical and functional variabilities of hearts and

brain surfaces.

The atlas of the human cardiac fiber architecture is an important contribution to a better understanding and to a more accurate modeling of the cardiac functions derived from the fiber architecture. The variability of the fiber architecture is studied across all hearts and within specific ventricular segments (Chapter 3) (Lombaert *et al.*, 2012b, 2011b,e,d). The statistical study has been performed using the Log-Euclidean framework, which considers the space of possible computations on diffusion tensor images. The average fiber angles, analyzed across several ventricular segments, with their normal variations provide a new reference in cardiology by advancing specific knowledge on the cardiac fiber architecture in humans. A study on the laminar organization of the fibers also revealed a complex variable structure and the possible presence of two populations of laminar sheets. This has been observed in several myocardial segments of the left ventricle (Lombaert *et al.*, 2011d). Additionally, preliminary results on a comparative study between a population of healthy hearts and abnormal hearts suggest that a higher variability of the fiber structure directions could possibly characterize abnormality of a heart (Lombaert *et al.*, 2011c). These variability studies and their conclusions demonstrate the relevance of the proposed framework.

The creation of a new fast and accurate method for surface matching based on an improved spectral correspondence, lead to study the variability of several sulci locations on the brain surface (Chapter 4) (Lombaert *et al.*, 2011a, 2012a). Such study, previously computationally expensive, showed results concurring with a ground truth data and gave confidence that the method can be used for further relevant neuroimaging studies. The matching of brain surfaces with different smoothing levels, representing the course of the brain surface convolution across the ages, enabled the exploration of the shape of the primary cortical folding. The results showed that several Brodmann areas (known to be related to specific cognition) align better on surfaces smoothed with an optimal level, which suggests that the primary folding and these specific smoothed surfaces have a similar shape around the Brodmann areas of interest (Chapter 5). This experiment demonstrates that the new spectral correspondence method enables new clinically relevant studies even in the presence of large disparity in the anatomy. Moreover, it is expected that the unified method using spectral correspondence within the diffeomorphic Demons framework can be similarly used for anatomically relevant studies since it showed successful registration of images with large shape variability.

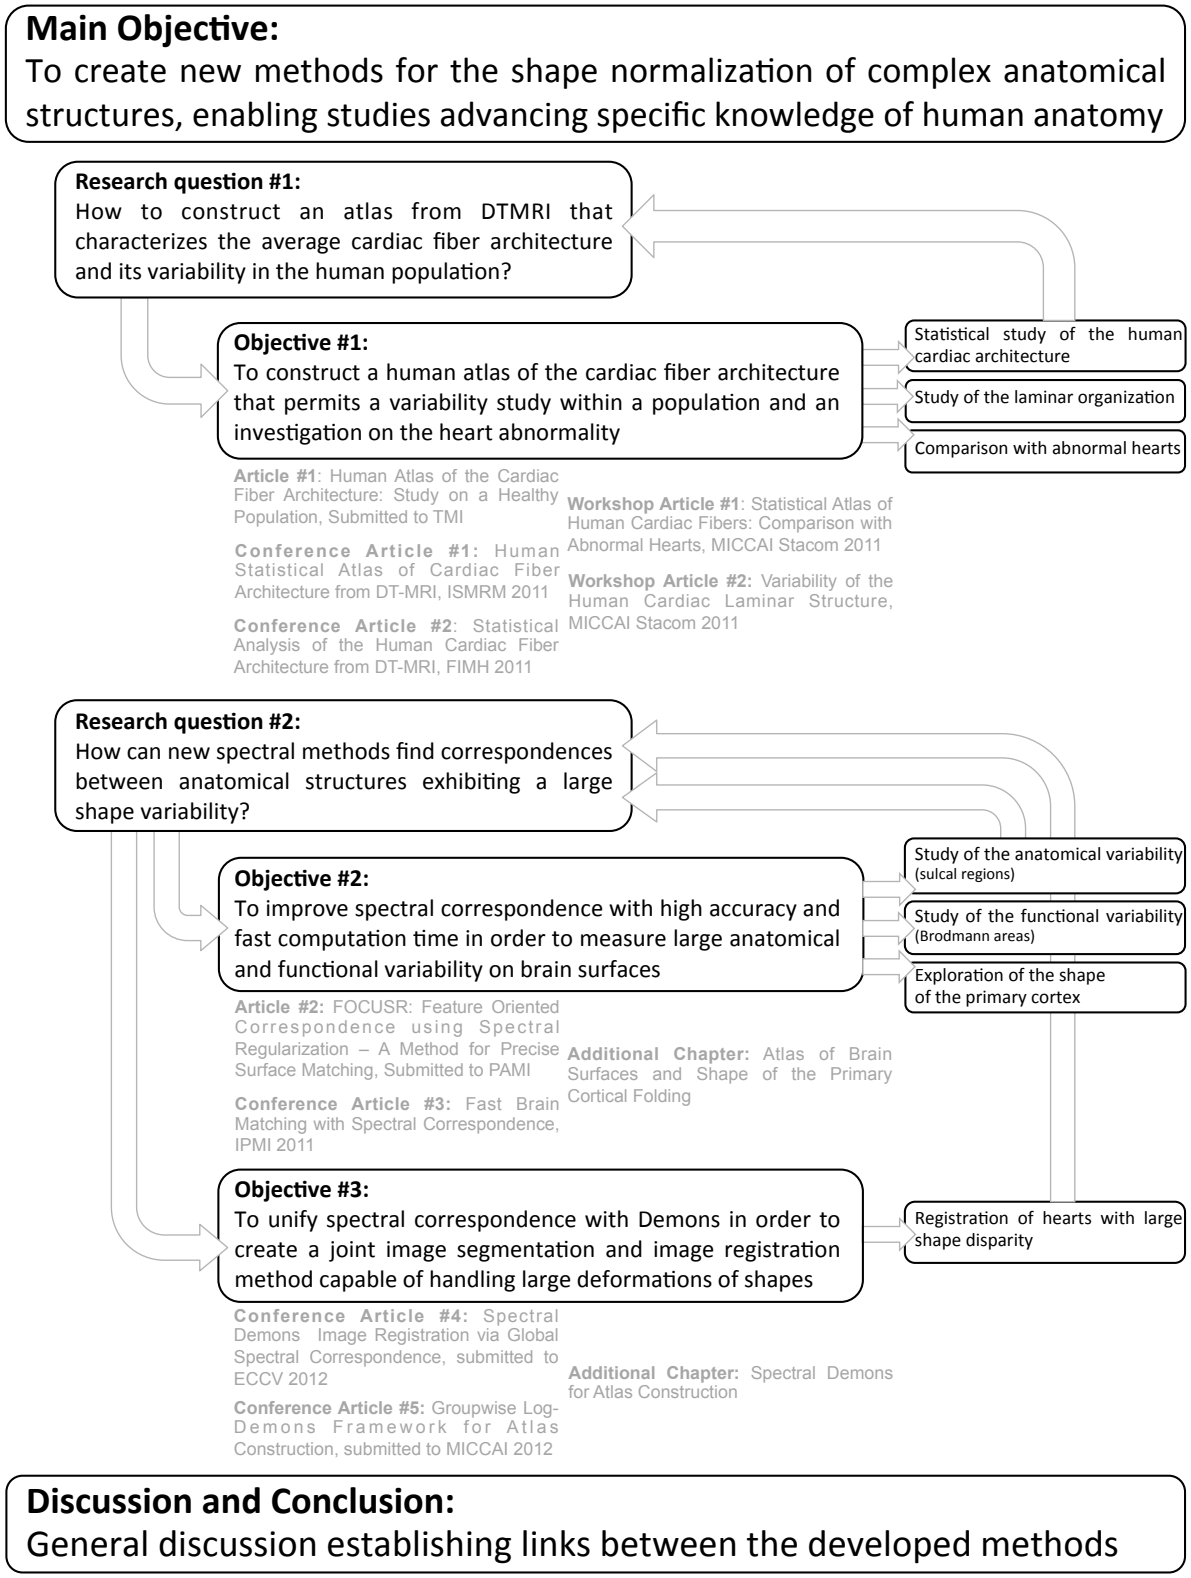

Figure 2.1: Structure of the methodology with corresponding articles

PART II

METHODS AND RESULTS

CHAPTER 3 HUMAN ATLAS OF THE CARDIAC FIBER ARCHITECTURE: STUDY ON A HEALTHY POPULATION

Herve Lombaert[1,4], Jean-Marc Peyrat[1], Pierre Croisille[2], Stanislas Rapacchi[2], Laurent Fanton[2], Farida Cheriet[4], Patrick Clarysse[2], Isabelle Magnin[2], Hervé Delingette[1], Nicholas Ayache[1]

Presentation

This chapter presents the article *"Human Atlas of the Cardiac Fiber Architecture: Study on a Healthy Population"* (Lombaert *et al.*, 2012b) submitted to **IEEE TMI** *(Transactions on Medical Imaging)*, sent on June 27th, 2011, revised on July 1st, 2011, and accepted for publication on March 23rd, 2012. Two initial articles (Lombaert *et al.*, 2011e,b) have been published in the conference **ISMRM** *(International Society for Magnetic Resonance in Medicine)* held in Montreal, QC, in May 2011, as well as in the conference **FIMH** *(Functional Imaging and Modeling of the Heart)* held in New York, NY, in May 2011. The FIMH paper has received the **Best Paper Award**. Additionally, further studies on the cardiac atlas (Lombaert *et al.*, 2011c,d) have been published in the STACOM Workshop (Statistical Atlas and Computational Models of the Heart) of the conference **MICCAI** *(Medical Image Computing and Computer Assisted Intervention)* held in Toronto, ON, in September 2011. The objective of this article is to develop a new automated framework in order to construct the *first* human atlas of the cardiac fiber architecture. This is a joint work between [1]**INRIA, Sophia Antipolis**, *France*; [2]**CREATIS**, *Lyon, France*; [3]**Siemens Molecular**, *Lyon, France*; and [4]**École Polytechnique**, *Montreal, Canada*.

Abstract

Cardiac fibers, as well as their local arrangement in laminar sheets, have a complex spatial variation of their orientation that has an important role in mechanical and electrical cardiac functions. In this paper, a statistical atlas of this cardiac fiber architecture is built for the first time using human datasets. This atlas provides an average description of the human cardiac fiber architecture along with its variability within the population. In this study, the population is composed of 10 healthy human hearts whose cardiac fiber architecture is imaged *ex vivo* with DT-MRI acquisitions. The atlas construction is based on a computational

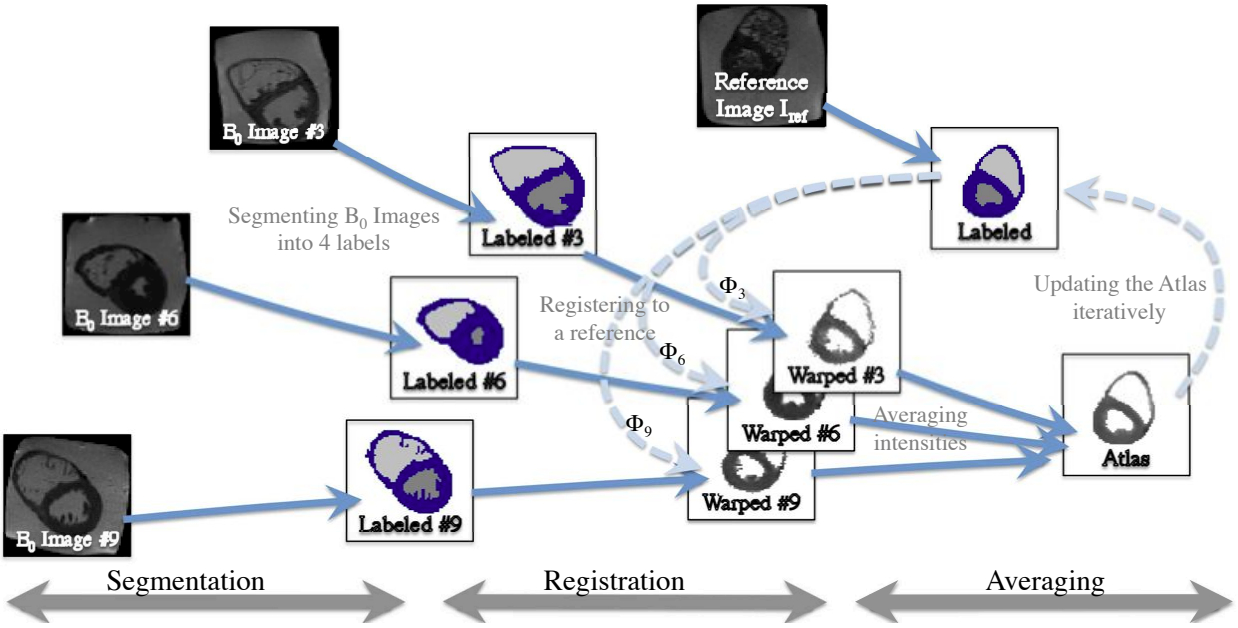

Figure 3.1: *Construction of the morphological atlas*: From the acquired images, the myocardium and the blood masses are segmented as described in Section 3.2.2. The images are then aligned and deformed toward a reference image. The registration is performed using Non-rigid Symmetric Diffeomorphic Log-Demons as explained in Section 3.2.2. The atlas is constructed iteratively by averaging acquired images in the average shape of the atlas as explained in Section 3.2.2.

framework that minimizes user interactions and combines most recent advances in image analysis: Graph Cuts for segmentation, Symmetric Log-domain Diffeomorphic Demons for registration, and Log-Euclidean metric for diffusion tensor processing and statistical analysis. Results show that the helix angle of the average fiber orientation is highly correlated to the transmural depth and ranges from $-41°$ on the epicardium to $+66°$ on the endocardium. Moreover, we find that the fiber orientation dispersion across the population ($\pm 13°$) is lower than for the laminar sheets ($\pm 31°$). This study, based on human hearts, extends previous studies on other mammals with concurring conclusions and provides a description of the cardiac fiber architecture more specific to human and better suited for clinical applications. Indeed, this statistical atlas can help to improve the computational models used for radiofrequency ablation (RFA), cardiac resynchronization therapy (CRT), surgical ventricular restoration, or diagnosis and follow-ups of heart diseases due to fiber architecture anomalies.

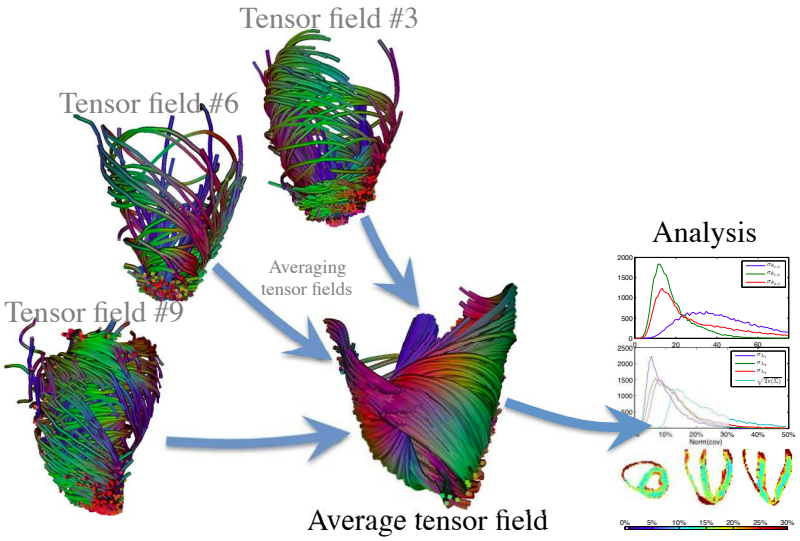

Figure 3.2: *Construction of the statistical atlas* of the human cardiac fiber architecture. All acquired tensors are warped using the transformation found during the construction of the morphological atlas. With the average tensor and its covariance matrix, the statistical analysis is carried on. The tensor field illustrations show tractographies generated with MedINRIA.

3.1 Introduction

Cardiovascular disease is a leading cause of death in developing countries. The understanding of the cardiac muscle structure and functions is essential for the diagnosis and the treatment of many heart pathologies. The cardiac fiber architecture, a complex organization of the myocardium fibers (Streeter *et al.*, 1969; LeGrice *et al.*, 1995), determines various cardiac mechanical functions (Costa *et al.*, 2001), cardiac electrophysiology patterns (Hooks *et al.*, 2002), and remodeling processes (Wu *et al.*, 2006). The assessment of structural variability is important for a better understanding of the human heart physiology (Cirillo, 2009; Cirillo *et al.*, 2010) and for a more accurate modeling of the heart. Such models with a complete description of the fiber architecture can be used for clinical applications. For instance, the planning of radiofrequency ablation (RFA) (Relan *et al.*, 2011) or cardiac resynchronization therapy (CRT) (Sermesant *et al.*, 2011), with respectively the localization of the zone to ablate and the positioning of the pacing electrodes, can be optimized using electro-mechanical models of the heart whose fiber structure is a key information for more accurate simulations and predictions.

The fibrous nature of the heart has been known for centuries, Harvey reported as early

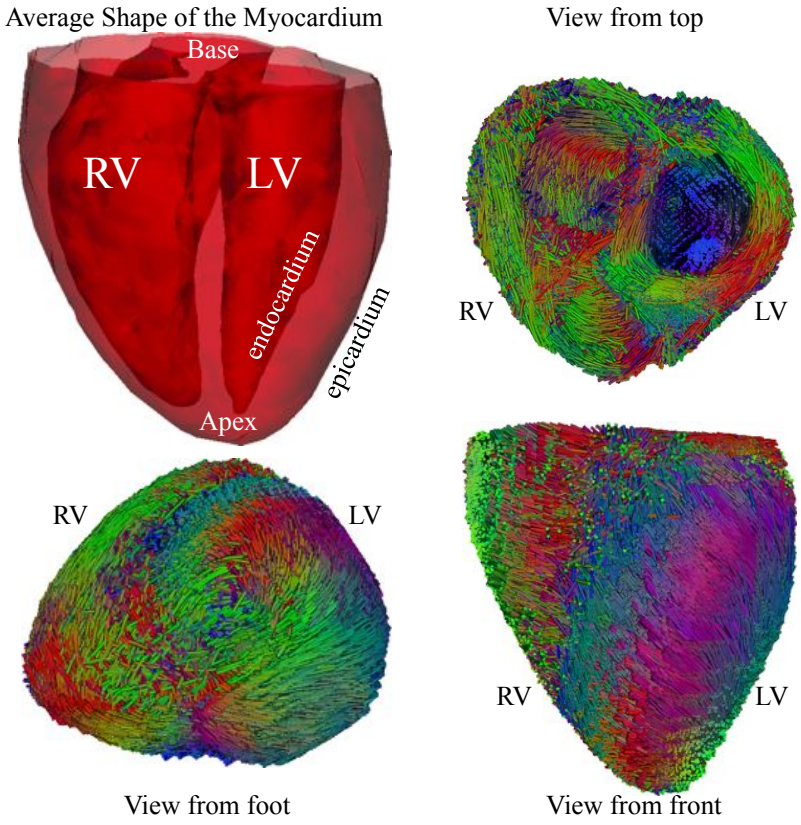

Figure 3.3: *Average geometry and fiber orientations* with *(upper left)* the average shape of the myocardium, *(upper right)* short axial view from top, *(lower left)* apical view from foot, *(lower right)* long axial view from front. The *r,g,b* color of the tracked fibers indicates the *x,y,z* components of the local orientation of the fiber.

as 1628 in *De Motu Cordis* (Harvey, 1628) the functional role of the cardiac fibers. In the modern era, Streeter (Streeter *et al.*, 1969) suggested that the cardiac fibers are organized as nested surfaces. Further studies on the cardiac fiber architecture also revealed common features among species (Bronzino, 1995). The fibers, locally organized as laminar sheets, appear to be consistently structured as two counter wound spirals wrapping around the heart clockwise on the epicardial surface and counterclockwise on the endocardial surface (Streeter *et al.*, 1969; LeGrice *et al.*, 1995). The helical myocardial band of Torrent-Guasp (Torrent-Guasp *et al.*, 2005) is another model (triggering controversial discussions (von Segesser, 2005)) that hypothesizes the existence of a single muscular band folding onto itself to form the whole heart. Nevertheless, cardiac fiber studies mainly focus on the local orientation of fibers with their angle on the tangent plane and on the horizontally normal plane of the epicardium called respectively the helix and the transverse angles. In humans, the

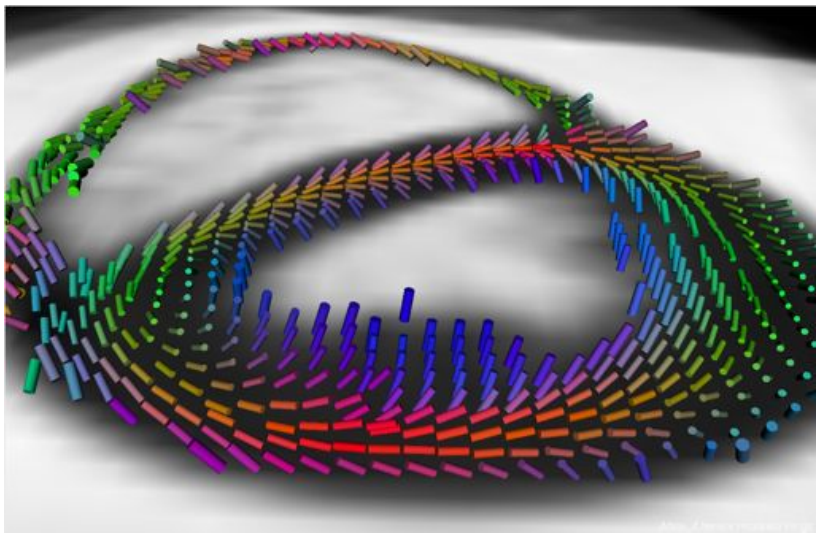

Figure 3.4: *Average fiber orientation* seen on a short axis slice. The left ventricle is on the foreground and the right ventricle on the background. The cylinders indicates the direction of the first eigenvector in the average diffusion tensor field. Color indicates the local orientation of the fiber. The transmural variation is visible in the left ventricle and in the midwall. The vertical papillary muscles are visible in darker blue.

variability of the cardiac fiber architecture was studied using tedious work on histological slices (Greenbaum *et al.*, 1981), and thus hard to assess completely in 3D. So far, more detailed cardiac fiber architecture has been mainly speculated from studies on other species (Nielsen *et al.*, 1991).

The cardiac fiber architecture of single human hearts have already been visualized and studied using DT-MRI (Zhukov and Barr, 2003; Rohmer *et al.*, 2007; Seemann *et al.*, 2006). However, due to an extreme rarity of healthy human hearts (they are rather transplanted than used for research), no statistical study has yet been performed on humans. With an access to a unique post-mortem human dataset (Frindel, 2009; Frindel *et al.*, 2009; Rapacchi *et al.*, 2010), the presented work aimed at building, for the first time, a statistical atlas of the human cardiac fiber architecture from 10 healthy *ex vivo* hearts imaged with DT-MRI. This work is an extended version of preliminary results (Lombaert *et al.*, 2011b). Here, the method is fully explained and we provide an extensive statistical study on the fiber variability as well as a refined analysis of the transmural distribution of the fiber helix and transverse angles.

The atlas construction and the population variability study are performed using a refined

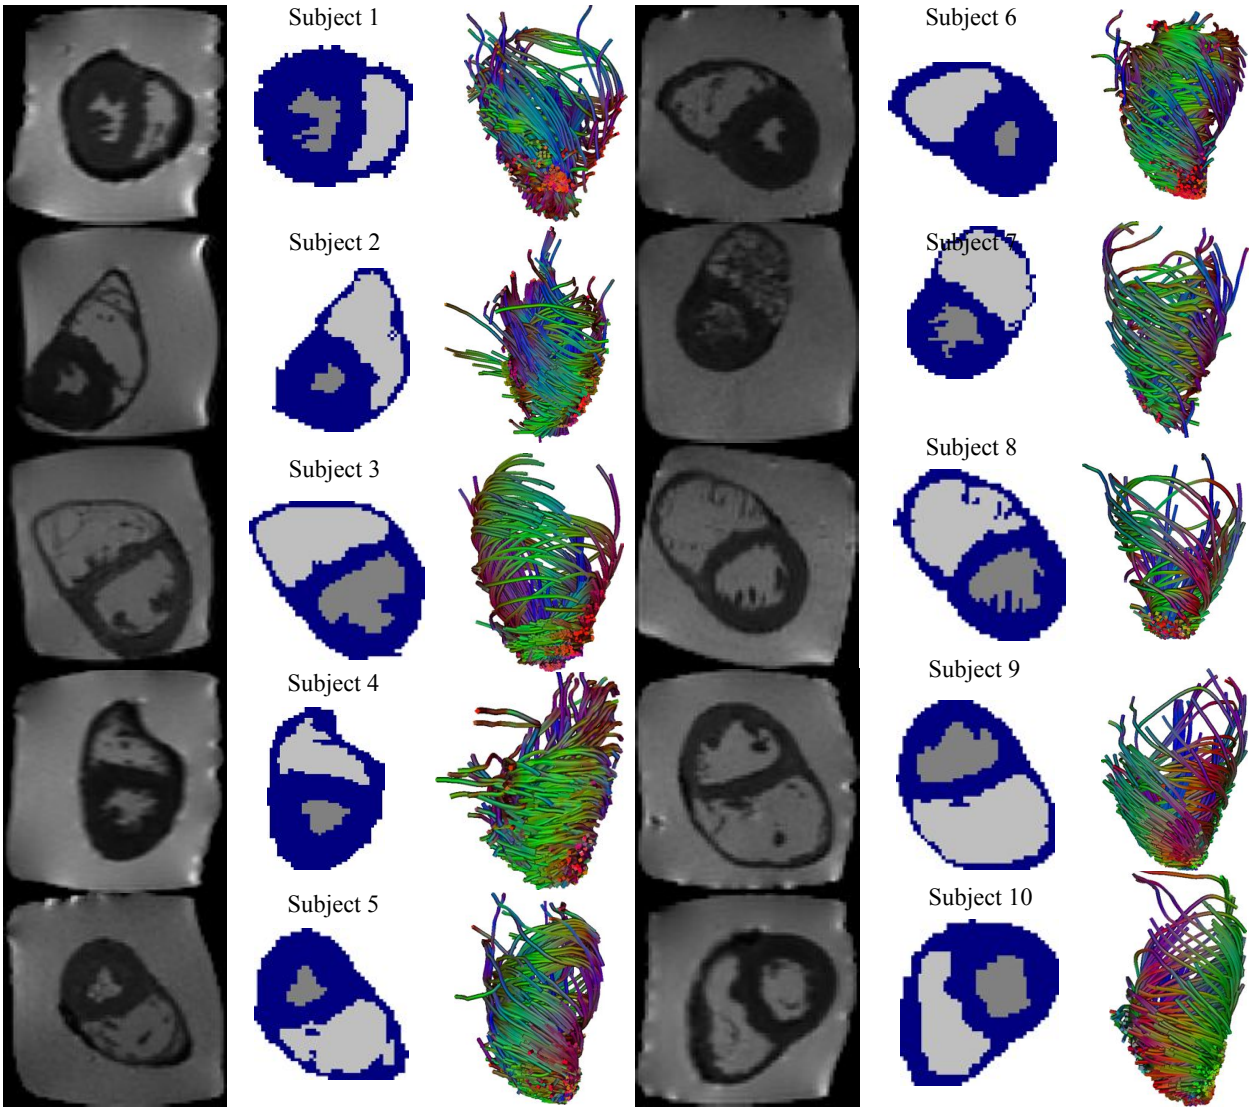

Figure 3.5: *Dataset* consisting of 10 healthy *ex vivo* human hearts. On the left, the b=0 images acquired in the DTI studies. In the middle, segmented myocardium (in blue) with left (in dark gray) and right (in light gray) ventricular chambers. On the right, tractographies of a few fibers that illustrates the fiber orientation (computed with MedINRIA, where coloring indicates the local fiber orientation).

version of the computational framework proposed by Peyrat *et al.* (Peyrat *et al.*, 2007). The method has been automated and adapted to process the available images, which are at a lower resolution than in (Peyrat *et al.*, 2007). It is described in details for reproducibility and can easily be reused for studies that includes a large number of hearts. Firstly, we present the strategy for the construction of the atlas eased by the use of a more recent registration method, the Symmetric Log-demons (Vercauteren *et al.*, 2008; Mansi *et al.*, 2010). The registration is also guided by a segmentation step involving minimal user interaction using Graph Cuts (Boykov and Jolly, 2000) to cope with topological issues that could occur with a basic thresholding on images with lower resolution than in (Peyrat *et al.*, 2007). Secondly, we use the same statistical framework on diffusion tensors as in (Peyrat *et al.*, 2007) to analyze the variability of the fiber architecture. Thirdly, we present the main contribution: the provision and analysis of a human statistical atlas, performed both globally over the whole myocardium and locally within several myocardial segments. Common features of the cardiac fiber architecture along their variabilities are shown for the first time in a population of healthy human subjects. Finally, we discuss possible improvements of our method and give perspectives on our results.

3.2 Material and Methods

The dataset and its acquisition protocol, consisting of anatomical $b = 0$ images and their associated DT-MRI tensor fields, is first presented. The transverse anisotropy in each tensor field provides confidence for the presence of a secondary fiber structure (i.e., the laminar sheet). The method constructing the DT-MRI atlas is described with the semi-automatic segmentation of the myocardium followed by the fully automatic groupwise registration of the anatomical images (Fig. 3.1), and with thereafter the warping of the tensor fields (Fig. 3.2). Finally, the statistical framework on diffusion tensors is explained. It provides an average representation (Fig. 3.3 and 3.4) of the cardiac fiber architecture as well as its variability in terms of fiber and laminar sheet orientations.

3.2.1 Data Acquisition

The hearts were explanted within 24 hours after death and *ex vivo* imaging was performed without any additional delay. The hearts did not exhibit any patent contracture compared to what is commonly observed in organs of more than a few days after death. Special care was

Table 3.1: Characteristics of the human dataset of 10 healthy hearts (Subject weight, height and age, with myocardial weight, max allowed weight, and septal thickness.).

Subject	Weight	Height	Age	Heart Weight (limit)	Sex	Septal thick.
1	60 Kg	158 cm	74	385 g (420 g)	F	12 mm
2	74 Kg	166 cm	47	385 g (444 g)	M	14 mm
3	102 Kg	192 cm	17	435 g (521 g)	M	10 mm
4	74 Kg	180 cm	47	430 g (444 g)	M	11 mm
5	77 Kg	165 cm	27	350 g (425 g)	F	11 mm
6	85 Kg	189 cm	21	365 g (473 g)	M	14 mm
7	63 Kg	175 cm	20	320 g (406 g)	M	12 mm
8	97 Kg	187 cm	21	410 g (506 g)	M	11 mm
9	77 Kg	179 cm	21	380 g (450 g)	M	12 mm
10	84 Kg	177 cm	50	460 g (473 g)	M	10 mm

taken to carefully remove any remaining blood or mural thrombi by flushing the ventricles with an isotonic solution. The ventricles were then filled with an hydrophilic and isotonic gel to remove any air and to restore a normal anatomic configuration. Finally, the hearts were placed in a plastic container filled with the same hydrophilic gel. The posterior wall of the atria were partially removed to give access to the atria-ventricular valves. The preparation and handling of the hearts were established with forensic specialists to avoid any additional delay and any potential changes to the common forensic procedure that was performed after imaging on the fresh non-fixed entire hearts.

Each heart with its plastic container was placed in a 16 element head coil and was imaged with a 1.5T MR scanner (Avanto Siemens). The main MRI parameters of the diffusion weighted EPI sequence (bipolar scheme) are: TE/TR = 69/6500 ms, 6 excitations (with polarity alternation), voxels of $2 \times 2 \times 2$ mm^3, parallel imaging (GRAPPA with acceleration factor 2), partial Fourier (6/8), base resolution matrix of 128, BW = 1628 Hz/Px, echo spacing of 0.7 ms, 12 directions, and $b = 1000$ s.mm^2 with a PSNR = 16.44 dB. The acquisition protocol is detailed in (Frindel, 2009; Frindel *et al.*, 2009). The b=0 images along their myocardial masks and illustrative fiber tractographies (generated using a spin glass model (Fillard *et al.*, 2009)) are shown on Fig. 3.5.

DT-MRI measures the local spatial diffusion within one voxel. The maximal local direction of diffusion, revealed by the primary eigenvector $\mathbf{v}_1$ of the diffusion tensor matrix D, occurs along the main structure, i.e., the fiber, while the secondary eigenvector, $\mathbf{v}_2$, is assumed to lay within the laminar sheet. The tertiary eigenvector, $\mathbf{v}_3$, would thus be related

to the normal of the laminar sheet (illustrated in Fig. 3.6 (a)). Although strong evidences exist toward this assumption (Helm *et al.*, 2005a,b; Scollan *et al.*, 2000; Tseng *et al.*, 2003), we first need to ensure that pairs of eigenvectors and eigenvalues are distinguishable and hold meaningful structural information. The corresponding null hypothesis would imply transverse isotropy (i.e., equal secondary and tertiary eigenvalues $d_2 = d_3$). Rejecting it, by observing a transverse anisotropy, would give an indication that a secondary structure exists in the transverse plane of the primary eigenvector. We do so by showing that the eigenvalues are not equal and that the eigenvectors are locally structured.

Since $d_2 > d_3 > 0$, the ratio d_3/d_2 is defined between 0 and 1 and the mode of the ratio distribution should be close to one in case of transverse isotropy. The figure 3.7 shows that even though the distributions of the eigenvalues are very close, the peak of the ratio distribution d_3/d_2 is at 0.86 and suggests that the secondary and the tertiary eigenvalues are not equal. Furthermore, Helm *et al.* (Helm *et al.*, 2005b; Helm, 2005) explain that in case of transverse isotropy the secondary and tertiary eigenvectors should rotate randomly around the principal axis of diffusion. They reject this null hypothesis with a Kolmogorov-Smirnov (K-S) test where the local distribution of the angle $\alpha = \cos^{-1}(|\mathbf{v}_2 \cdot \mathbf{v}_n|)$, between the secondary eigenvector $\mathbf{v}_2$ and a local reference vector $\mathbf{v}_n$ defined in each myocardial segment, is not uniform (i.e., testing if α is randomly distributed, or not, on the transverse plane). The local reference, $\mathbf{v}_n$, is fixed and is chosen to be the best orthogonal vector to the set of first eigenvectors within each AHA segment. Furthermore, the maximal distance between the cumulative distribution $F_\alpha(\theta)$ of the angle α and the cumulative distribution θ/π of the uniform distribution gives a measure of confidence with the R-value: $R = \max_{\theta=[0,2\pi]}\{|\theta/\pi - F_\alpha(\theta)|\}$ (see (Helm *et al.*, 2005b; Helm, 2005) for more details). The critical R-value $R_c = 0.06$ is chosen (Helm, 2005) using the average number of pixels in each AHA segment (1247 pixels), i.e., if $R > R_c$, the distribution of the angle α can be considered non-uniform. The table 3.2 shows for all AHA segments of all hearts R-values with a significance level of 1% consistently well above $R_c = 0.06$ (away from the uniform distribution) which give another indication that the diffusion is transversely anisotropic, i.e., there is a secondary structure in the human fiber architecture.

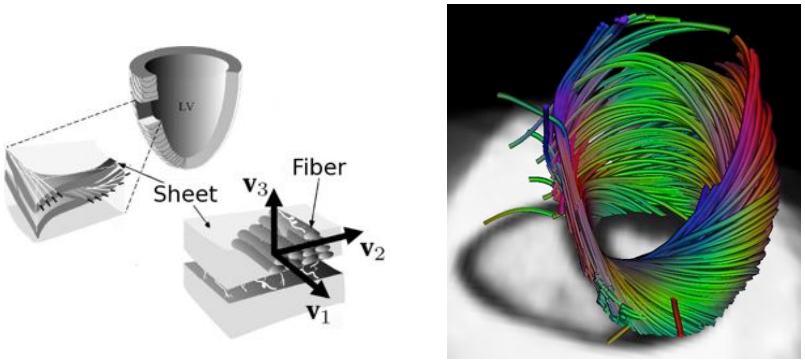

(a) Tensor eigenvectors with respect to the fiber and laminar sheet

(b) Atlas fiber tractography of the left ventricle and a short axis slice

Figure 3.6: *Fiber architecture*: *(a)* Illustration of the Tensor eigenvectors, $\mathbf{v}_{1,2,3}$, with respect to the fiber and the laminar sheet (adapted from (LeGrice *et al.*, 1995)). *(b)* The transmural variation of the fiber orientations in the left ventricle is visible in the fiber tractography (computed with MedINRIA) of the average tensor field.

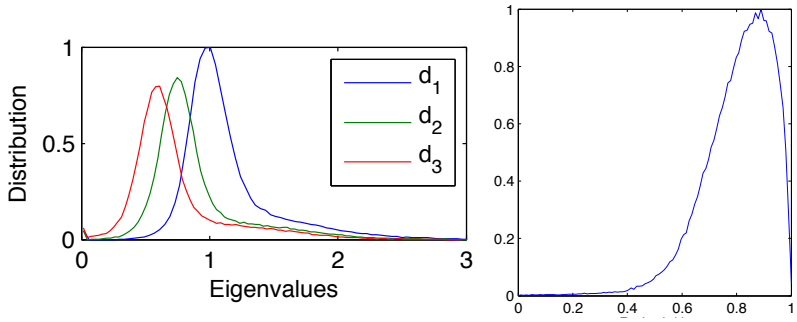

(a) Distribution of eigenvalues

(b) Transverse anisotropy

Figure 3.7: *Transverse anisotropy*. *(a)* Distribution of the eigenvalues for all subjects. *(b)* Distribution of the ratio between the third and second eigenvalues, d_3/d_2, showing a transverse anisotropy for all subjects (i.e., the ratio peak at 0.86).

Table 3.2: R value, confidence that the angular distribution of the 2nd eigenvector around the 1st eigenvector is not uniform (i.e., transverse anisotropy). Mean value for each AHA segment with its standard deviation.

AHA	R Value	AHA	R Value	AHA	R Value
1	0.18 (± 0.10)	7	0.24 (± 0.06)	13	0.25 (± 0.13)
2	0.21 (± 0.10)	8	0.19 (± 0.07)	14	0.18 (± 0.10)
3	0.23 (± 0.09)	9	0.21 (± 0.06)	15	0.23 (± 0.09)
4	0.21 (± 0.10)	10	0.18 (± 0.08)	16	0.22 (± 0.10)
5	0.18 (± 0.12)	11	0.20 (± 0.06)	17	0.23 (± 0.09)
6	0.24 (± 0.06)	12	0.19 (± 0.05)		

3.2.2 Atlas Construction

The statistical atlas is constructed using four steps, illustrated in Fig. 3.1. First, the myocardium of each heart is segmented out on the b=0 image of the DT-MRI acquisition. They are similar to T2-weighted images providing anatomical information with the advantage of being acquired and corrected for motion together with diffusion-weighted images (DWI) and thus aligned with the estimated DT-MRI. Information on the fiber architecture (i.e., any directional data such as DT-MRI or corresponding DWI) is purposely omitted from the registration process in order to avoid introducing any bias in the study of the fiber variability. Second, each myocardium is registered to a reference image using solely the b=0 images and the myocardium masks. Third, the reference image is deformed toward the morphological average of all hearts by repeating the second and this third step iteratively until convergence. Fourth, and last, the resulting deformation fields computed from the registration process are used to warp all tensor fields to the morphological atlas (Fig. 3.2).

Myocardium Segmentation

Due to the noisy nature of the available images, the boundary between the myocardium and the surrounding gel can not be captured with a simple thresholding segmentation. This is in contrast with the canine dataset used in (Helm *et al.*, 2005a,b; Helm, 2005; Sundar *et al.*, 2006; Peyrat *et al.*, 2007; Peyrat, 2009; Peyrat *et al.*, 2007) where the myocardium was clearly outstanding from its background. A semi-automatic method has thus been designed with minimal user interaction in mind. The Graph Cut algorithm (Boykov and Jolly, 2000) has been chosen for its efficiency and ease of use. The following details are provided for reproducibility of the results. An underlying 3D graph is constructed from a 3D image and a global optimal solution partitions the graph into an *object* and a *background*. The graph edges between neighboring pixels, say p and q with intensities i_p and i_q, are weighted with $w_{p,q} = \exp(-(i_p - i_q)^2/2\sigma^2)$, where σ is a penalizing term (e.g., the standard deviation of intensity differences). By setting different weights on both directions, boundaries from dark objects to bright backgrounds (e.g., dark myocardium within bright gel) can be favored (or vice versa, for bright objects on dark backgrounds). A directional edge from a dark pixel p to a bright pixel q (i.e., $i_p < i_q$) is privileged when its weight is $w_{p\to q} = \lambda w_{p,q}$ with $\lambda < 1$, while $w_{q\to p} = w_{p,q}$ remains the same. Two special nodes are added in the graph, a source representing the *object* and a sink representing the *background*. The algorithm needs only

a few *seed* points, connected to the two special nodes, in the *object* and in the *background*. The Graph Cut algorithm finds a global optimal cut separating the nodes connected to the source from those connected to the sink, thus segmenting the *object* from its *background*. Fast interactions are also possible to correct any missed segmentation. Further details are available in (Boykov and Jolly, 2000). In our workflow, the user marks on any 3D slice of a b=0 image a few points in the left ventricular and in the right ventricular blood pools. From these marks, three automatic 3D binary segmentations are performed:

- *Finding the Heart*: A graph is constructed from the whole image and the heart, including the blood pool, is isolated from its background using the marked points as *object seeds* and using the points (found automatically) lying on a sphere surrounding the initial marked points (a sphere large enough to be outside the heart) as *background seeds*. A boundary from a dark object (the dark myocardium) to a bright object (bright gel) is privileged with $\lambda = 0.95$ when $i_p < i_q$.

- *Finding the Myocardium*: A smaller graph is constructed from the previous heart mask and the myocardium is outlined from the blood pool using the initial marked points as *background seeds* and using the voxels lying on the inner boundary of the heart mask as *object seeds*. Here again, edges from dark to bright are privileged with $\lambda = 0.95$.

- *Differentiating the Ventricles*: A graph is constructed from the myocardium mask and the blood pool is partitioned into the left and the right sides using the initial marked points in the left and in the right ventricles as *object* and *background seeds*. As both ventricular fillings have similar intensities, directional edges are identical (i.e., $\lambda = 1$).

This strategy ensures a topologically valid myocardium mask. Indeed, the first two graph cuts correspond to the epicardium and the endocardium and isolate the myocardium with no unconnected pieces while guaranteeing a mask of at least one voxel thick. Holes are thus not possible across the myocardial wall which is really thin in the right ventricle. From the generated binary masks, a 3D segmented image is created with four labels: background, left and right ventricles, and myocardium. The labeled images are shown on Fig. 3.5.

Heart Registration

The pairwise registrations of N hearts, $\{I_i\}_{i=1,\dots,N}$, to a reference image, I_{ref}, give the diffeomorphic transformations, $\{\phi_{I_i \to I_{\text{ref}}}\}_{i=1,\dots,N}$, required for the warping of the tensor fields.

Fig. 3.1 shows three hearts before and after registration. To obtain a better registration and a fast convergence, the registration is performed in three steps:

- *Ventricle Rigid Alignment*: First, the labeled images are reoriented by aligning the centers of mass of each ventricle. A block matching algorithm (Ourselin *et al.*, 2000) refines this rigid registration.

- *Mask Non-rigid Registration*: Second, the aligned labeled images are transformed using a nonrigid registration (Vercauteren *et al.*, 2008; Mansi *et al.*, 2010).

- b=0 *Non-rigid Registration Refinement*: Third, and last, the previous transformation initializes the nonrigid registration of the masked b=0 images where only the myocardial regions are considered.

The pairwise registrations are performed with the Symmetric Log-domain Diffeomorphic Demons (Vercauteren *et al.*, 2008; Mansi *et al.*, 2010). In this extension of the Diffeomorphic Demons (Vercauteren *et al.*, 2009b), the transformation ϕ is constrained to be in the one parameter subgroup of diffeomorphisms with stationary velocity field by parameterizing ϕ with the velocity field v, such that $\phi = \exp(v)$. Thus, the Log-domain Diffeomorphic Demons algorithm can be formulated as an alternate optimization scheme of the following energy with respect to the velocity fields v_c (hidden variable) and v:

$$E(I_{\text{ref}}, I_{\text{flo}}, v, v_c) = \tfrac{1}{\sigma_i^2} \parallel I_{\text{ref}} - I_{\text{flo}} \circ \exp(v_c) \parallel^2$$
$$+\tfrac{1}{\sigma_x^2} \parallel \log(exp(-v) \circ \exp(v_c) \parallel^2 + \tfrac{1}{\sigma_T^2} \parallel \nabla v \parallel^2,$$

where $\sigma_{i,x,T}$ are penalizing terms, I_{ref}, a reference image, I_{flo}, a floating image, v, the velocity field such that $I_{\text{flo}} \circ \exp(v) \equiv I_{\text{ref}}$, and the velocity field v_c, a hidden variable whose exponential is called the correspondence. The optimization of this energy is implemented in (Vercauteren *et al.*, 2009b) with two consecutive smoothing steps which use two parameters. Both are widths of smoothing kernels, one for the smoothing of the update field, $\sigma_{K_{\text{fluid}}} = 1.4$, and has a fluid behavior, and the second is for the smoothing of the displacement field, $\sigma_{K_{\text{diff}}} = 1.0$, and has an elastic behavior.

Since the Demons algorithm is based on a sum of squared differences (SSD) and applied to MR images, whose intensities are not standardized, a histogram matching with the reference

image is performed before each non-linear registration of b=0 images.

We use here the symmetric version of the Log-demons (Vercauteren *et al.*, 2008) where the registration is independent from the choice of the reference image, i.e., $\phi^{-1}_{I_{\text{flo}}\to I_{\text{ref}}} = \phi_{I_{\text{ref}}\to I_{\text{flo}}}$. The forward and backward velocity fields, $v_{I_{\text{flo}}\to I_{\text{ref}}}$ and $v_{I_{\text{ref}}\to I_{\text{flo}}}$, are computed independently as detailed previously and then averaged.

The Log-domain Diffeomorphic Demons has the advantage of providing stationary velocity fields that can be averaged or negated to respectively average and invert corresponding diffeomorphic transformations (Arsigny *et al.*, 2006a). This property is particularly useful and efficient in the iterative steps of atlas construction detailed in the following subsection 3.2.2.

Construction of the Morphological Atlas

The construction of the atlas follows Guimond's *et al.* method (Guimond *et al.*, 2000) where the reference heart converges iteratively toward an average heart shape. In practice, only 5 iterations are required and the process is independent to the initial reference heart. At each iteration $(t+1)$, the reference heart from previous iteration (t), $I^{(t)}_{\text{ref}}$, is recomputed using the average of all inverse transforms, i.e., $I^{(t+1)}_{\text{ref}} = \frac{1}{N}\sum_{i=1}^{N} I_i \circ \phi_{I_i\to I^{(t)}_{\text{ref}}} \circ \phi_{I^{(t)}_{\text{ref}}\to I^{(t+1)}_{\text{ref}}}$ to converge to an average heart shape and intensities. The transformation updating the reference image is $\phi_{I^{(t)}_{\text{ref}}\to I^{(t+1)}_{\text{ref}}} = \exp\left(\frac{1}{N}\sum_{i=1}^{N} v_{I^{(t)}_{\text{ref}}\to I_i}\right)$. This can be computed in the Log-domain using the negated velocity fields: $v_{I^{(t)}_{\text{ref}}\to I_i} = -v_{I_i\to I^{(t)}_{\text{ref}}}$, where all $v_{I_i\to I^{(t)}_{\text{ref}}}$'s were previously computed during the pairwise registration of all hearts to the reference image, thus: $\phi_{I^{(t)}_{\text{ref}}\to I^{(t+1)}_{\text{ref}}} = \exp\left(-\frac{1}{N}\sum_{i=1}^{N} v_{I_i\to I^{(t)}_{\text{ref}}}\right)$. All steps and iterations of the atlas construction, the N pairwise registrations (times the number of iterations), are fully automated with no user interaction required. This is again in contrast with the method used in (Peyrat *et al.*, 2007) which required numerous user interaction with manual landmarks positioning.

Tensor Warping

The original tensor fields, $\{\tilde{D}^{(i)}\}_{i=1,\ldots,N}$, associated with each $b=0$ image, are warped to the converged average heart shape using the final transformations $\{\phi_{I_i\to I_{\text{ref}}}\}_{i=1,\ldots,N}$. Since diffusion tensors hold directional information, the warping of diffusion tensor fields includes a tensor reorientation step. Among two most common reorientation strategies (Alexander *et al.*, 2001), the Finite Strain strategy is preferred to the Preservation of the Principal

Direction for its preservation of geometric features (Peyrat *et al.*, 2007). With the Finite Strain strategy, the reorientation of diffusion tensors is defined as the rotational component $R = (AA^t)^{-1/2}A$ of the local linear approximation $A = \nabla\phi^{-1}$ of the first derivative of the inverse transformation ϕ at each voxel. In the next sections, the warped tensor fields are referred as $\{D^{(i)}\}_{i=1,\dots,N}$ such that $D^{(i)} = R \cdot \tilde{D}^{(i)} \circ \phi_{I_i \to I_{\text{ref}}} \cdot R^t$.

3.2.3 Statistics on Tensor Fields

Statistics on diffusion tensor fields is not straightforward due to the nature of the diffusion tensor space. Indeed, the diffusion tensor space of symmetric positive definite matrices does not have a vector space structure with the standard Euclidean metric (Jones *et al.*, 2002). Among valid metrics proposed in the literature (Pennec *et al.*, 2006), the Log-Euclidean metric (Arsigny *et al.*, 2006b) is preferred for its simple and fast framework with a closed form solution. Usual statistics, such as the mean and the covariance, are thus computed with:

$$\overline{D} = \exp\left(\frac{1}{N}\sum_{i=1}^{N}\log(D^{(i)})\right) \tag{3.1}$$

$$\Sigma = \frac{1}{N-1}\sum_{i=1}^{N}\text{vec}(\Delta D^{(i)}) \cdot \text{vec}(\Delta D^{(i)})^t, \tag{3.2}$$

where $\Delta D^{(i)} = \log(D^{(i)}) - \log(\overline{D})$ and $\text{vec}(D) = (D_{11},\ \sqrt{2}D_{12},\ D_{22},\ \sqrt{2}D_{31},\ \sqrt{2}D_{32},\ D_{33})^t$ is the condensed form of a diffusion tensor (redundant terms of the symmetric matrix $(D_{ij})_{i,j=1,2,3}$ are gathered). The mean diffusion tensor field, $\overline{D}$, is computed from N warped tensor fields $\{D^{(i)}\}_{i=1,\dots,N}$.

Fig. 3.5 shows the tractography of a few fibers from six subjects. Fig. 3.6 (b) shows the tractography of the average diffusion tensor field. The variability of the diffusion tensors, for a particular voxel, is embedded in the covariance matrix, Σ.

The measure of the global variability can be assessed with the trace of the covariance matrix:

$$\sqrt{\text{Trace}(\Sigma)} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}\|\text{vec}(\Delta D^{(i)})\|^2} \tag{3.3}$$

In the Log-space, $\Delta D^{(i)} = \log\left(D^{(i)}\right) - \log\left(\overline{D}\right)$ is equivalent to a ratio. The square root $\sqrt{\text{Trace}(\Sigma)}$ expresses, therefore, a relative ratio of all diffusion tensors, $\{D^{(i)}\}_{i=1,\ldots,N}$, to the mean $\overline{D}$.

The measure of local variability of the cardiac fiber structures relies on the analysis of the eigenvectors of the mean diffusion tensor, $\mathbf{v}_{1,2,3}$ extracted from $\overline{D}$. The standard deviations of orientation differences, $\{\sigma_{\theta_{i,j}}\}_{i,j=1,2,3}$, between the coupled $(\mathbf{v}_i, \mathbf{v}_j)$ axes around $\mathbf{v}_k$ (i.e., the variability of how the orthogonal axes $(\mathbf{v}_i, \mathbf{v}_j)$ rotates around $\mathbf{v}_k$) are computed by projecting the covariance matrix Σ onto the orthonormal bases $\{W_{i,j}\}_{i,j=1,2,3}$:

$$\sigma_{\theta_{i,j}} = \arctan\left(\frac{1}{2(\lambda_i - \lambda_j)^2}[\text{vec}(W_{i,j})^t \cdot \Sigma \cdot \text{vec}(W_{i,j})]\right)^{\frac{1}{2}} \tag{3.4}$$

where $\lambda_{1,2,3}$ are the tensor eigenvalues in the Log-space, and the orthonormal bases are defined with:

$$\begin{aligned} W_{2,3} &= \left(\mathbf{v_3} \cdot \mathbf{v_2}^t + \mathbf{v_2} \cdot \mathbf{v_3}^t\right)/\sqrt{2} \\ W_{1,3} &= \left(\mathbf{v_3} \cdot \mathbf{v_1}^t + \mathbf{v_1} \cdot \mathbf{v_3}^t\right)/\sqrt{2} \\ W_{1,2} &= \left(\mathbf{v_2} \cdot \mathbf{v_1}^t + \mathbf{v_1} \cdot \mathbf{v_2}^t\right)/\sqrt{2}. \end{aligned}$$

The statistics are performed directly on tensors, i.e., eigendecomposition is only necessary on the average diffusion tensor field. There is thus no need to account for the sign ambiguity, which is inherent from the eigenvector extraction (Basser and Pajevic, 2000), when computing statistics.

The variability of the tensor eigenvalues measures the coherence of the diffusion along the three principal directions. The eigenvalues of all the tensors are expressed in the Log-space, $\{\lambda^{(i)}_{1,2,3}\}_{i=1,\ldots,N}$ such that $\lambda = \log(d)$ where d is the eigenvalue in the Euclidean-space. Their standard deviations σ_{λ_1}, σ_{λ_2}, and σ_{λ_3}, are the projections of the covariance matrix Σ onto the orthonormal bases $W_{1,1}$, $W_{2,2}$, and $W_{3,3}$:

$$\sigma_{\lambda_i} = \sqrt{\text{vec}(W_i)^t \cdot \Sigma \cdot \text{vec}(W_i)}, \tag{3.5}$$

where the orthonormal bases are defined with:

$$W_1 = \mathbf{v}_1 \cdot \mathbf{v}_1^t$$
$$W_2 = \mathbf{v}_2 \cdot \mathbf{v}_2^t$$
$$W_3 = \mathbf{v}_3 \cdot \mathbf{v}_3^t.$$

In the Log-space, the standard deviation σ_λ is a ratio of all eigenvalues expressed in the Euclidean space, $\{d^{(i)}\}_{i=1,\dots,N}$, relative to their mean $\overline{d}$.

3.3 Experimental Results

The statistical study of the human cardiac fiber architecture is presented here from a global point of view, showing the variability within a population of the whole diffusion tensor, to a more detailed approach, showing the variability of the principal directions of diffusion.

Due to many external factors, such as the temperature during acquisition (known to affect the diffusion values), the tensor fields have different scalings among all hearts. The histograms of the traces of the tensor matrices show significant disparities among a population. Therefore, before using them for statistical analysis, all acquired tensor fields, $\{\tilde{D}^{(i)}\}_{i=1,\dots,N}$, are normalized using the modes of the histograms of the tensor traces for each patient. For the i^{th} patient:

$$\tilde{D}^{(i)} \leftarrow \tilde{D}^{(i)} \frac{\frac{1}{N}\sum_{j=1}^{N} \text{mode}\left(\text{Trace}(\tilde{D}^{(j)})\right)}{\text{mode}\left(\text{Trace}(\tilde{D}^{(i)})\right)}$$

The eigenvectors of the tensors fields remain unchanged with this normalization, and the eigenvalues are equally scaled in order to be able to compare all tensor fields on a common reference.

3.3.1 Fiber Variability

The global variability of the tensor field is measured with the square root of the trace of the covariance matrix (Eq. 3.3). Its histogram shows a peak of variability at 13.2% (Fig. 3.8 (a,b)). It is important to note that this variability also includes acquisition and registration errors. While certain areas, such as around the apex, might show a higher variability due to

a less organized fiber structure, other areas, such as the right ventricle, show an increased variability maybe due to registration errors. Indeed, large deformations in low resolution images make the registration challenging.

The distributions of the standard deviations of the three eigenvalues are presented on Fig. 3.8 (b). The primary eigenvalue is less variable than the two other eigenvalues. The mode of the standard deviations (Eq. 3.5) of the primary eigenvalue from the one of the statistical mean, expressed in percentage, is $\sigma_{\lambda_1} = 5.35\%$ compared to $\sigma_{\lambda_2} = 6.35\%$ and $\sigma_{\lambda_3} = 8.69\%$.

The orientation of the diffusion in human cardiac fibers has a different variability in each principal direction. The variability around $\mathbf{v}_1$, $\mathbf{v}_2$, and $\mathbf{v}_3$, is measured with the rotation of the tensor transverse $(\mathbf{v}_2, \mathbf{v}_3)$, longitudinal $(\mathbf{v}_1, \mathbf{v}_3)$, and equatorial $(\mathbf{v}_1, \mathbf{v}_2)$ planes around their normals. The standard deviations are formulated by Eq. 3.4. These angular variabilities are presented in Fig. 3.9. The dispersion of the fiber orientation, $\mathbf{v}_1$, has a standard deviation of rotation in the tensor longitudinal plane with a mode of $\sigma_{\theta_{1,3}} = 11.5°$, and in the tensor equatorial plane, $\sigma_{\theta_{1,2}} = 13.0°$. The variability of the laminar sheet orientation is described with the rotation of the tensor transverse plane around the fiber. Its standard deviation is $\sigma_{\theta_{2,3}} = 31.1°$. The laminar sheet structure is, thus, in humans, much more variable than the fiber structure. This concurs with previous canine studies (Helm *et al.*, 2005a,b; Peyrat *et al.*, 2007; Gilbert *et al.*, 2007) where the fiber orientation (with $\sigma_{\theta_{1,3}} = 7.9°$ and $\sigma_{\theta_{1,2}} = 7.7°$) is more stable than the laminar sheet orientation (with $\sigma_{\theta_{2,3}} = 22.7°$).

3.3.2 Transmural Distribution

The variability of the fiber orientation is also analyzed by measuring the distribution of the helix angle (vertical angle, or helical pitch, of the fiber, Fig. 3.10 (a)) and the transverse angle (horizontal angle of the fiber, Fig. 3.10 (b)) along the transmural depth. These angles are defined to range from $-90°$ to $+90°$. For instance, an helix angle of $-90°$ is a fiber pointing toward the base of the heart, an helix angle of $+90°$ is a fiber pointing toward the apex, and a transverse angle of $-90°$ is a fiber pointing toward the ventricular cavity, a transverse angle of $+90°$ is a fiber pointing outward the ventricular cavity.

(i.e., fiber pointing toward the base of the heart) to $+90°$ (i.e., fiber pointing toward the apex). A prolate ellipsoidal model of the heart (Nielsen *et al.*, 1991) is fitted to the morphology of the statistical atlas. The transmural variation of the fiber orientation is,

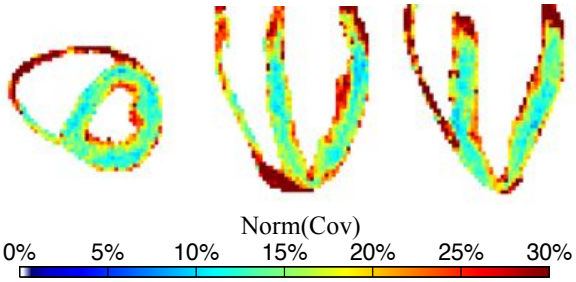

(a) Global variability of the whole tensor field

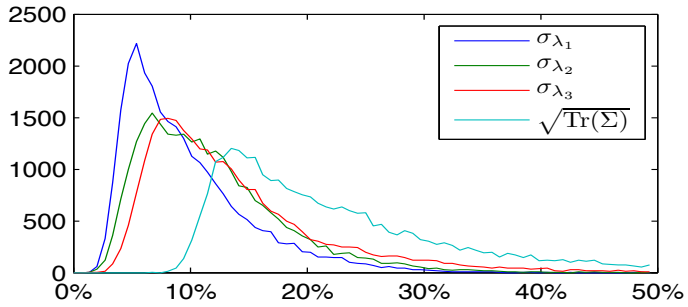

(b) Histogram of the eigenvalue variability and the global variability

Figure 3.8: *Global variability*: (a) Variability of the eigenvalues in the statistical atlas of the tensor fields. The distribution of the primary eigenvalue standard deviations (expressed in percentage) of all hearts have a mode of $\sigma_{\lambda_1} = 5.35\%$ from the statistical mean. The second eigenvalues have a mode of $\sigma_{\lambda_2} = 6.35\%$. The tertiary eigenvalues have a mode of $\sigma_{\lambda_3} = 8.69\%$. (b) The histogram of the global variability ($\sqrt{\mathrm{Tr}(\Sigma)}$, Eq. 3.3) (expressed in percentage) of the whole diffusion tensor shows a mode of 13.2%.

thus, easier to measure with the mapped prolate ellipsoidal coordinates. The distributions are presented in a joint histogram (Fig. 3.10 (c,d)) where the angle distribution, on the vertical axis, is plotted against all transmural distances, on the horizontal axis.

The helix angle varies globally from $-41°(\pm 26°)$ (plus or minus its standard deviation) on the epicardium to $+66°(\pm 15°)$ on the endocardium. The mode of the helix angle is $-48°$ on the epicardium and $+64°$ on the endocardium. The range of the helix angle in humans, 107°, concurs with the results of a canine study (Peyrat, 2009) that showed a range of 110°. It appears to be consistent across the myocardium, as seen on Fig. 3.11, only the apex presents a drastic difference. The helix angle is strongly correlated to the transmural distance with a correlation factor of $\rho = 0.831$, i.e., this suggests a linear dependence between the helix angle and the transmural location. This correlation also concurs with the findings of the histological study (Greenbaum *et al.*, 1981) of *postmortem* human hearts. Their sparse distribution of helix angles varies from about $-40°$ on the epicardium to about $+40°$ on the endocardium.

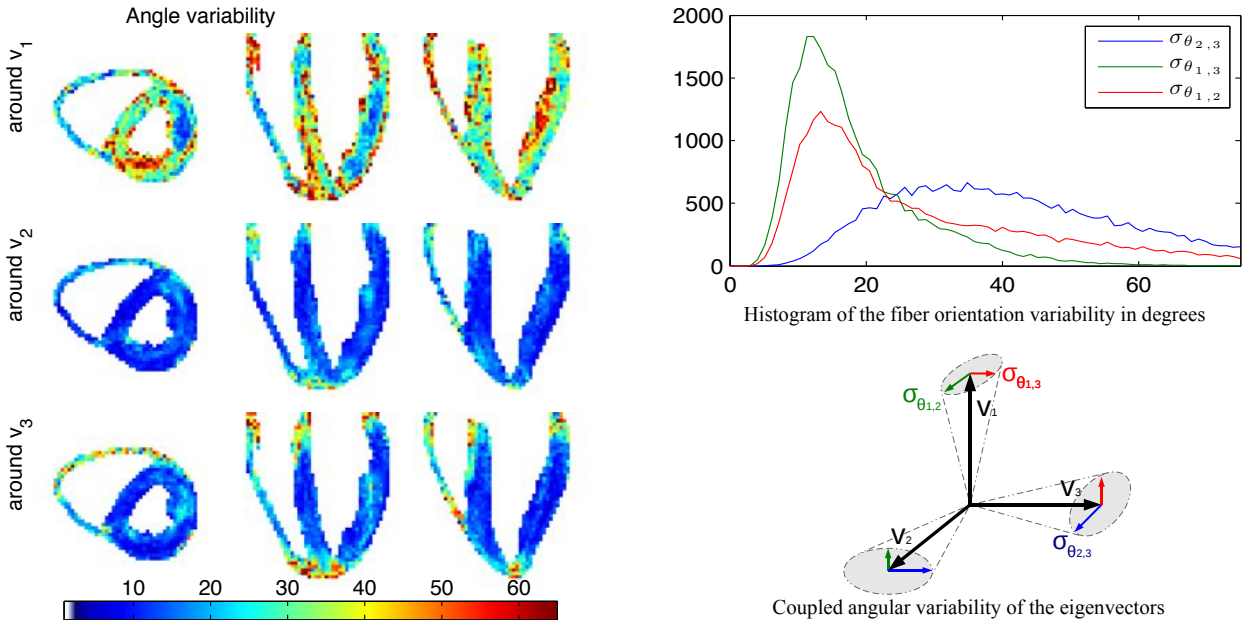

Figure 3.9: Standard deviation of the fiber orientation (expressed in degree). *Top row of the images and in the blue histogram*: Variability of the laminar sheet orientation, $\sigma_{\theta_{2,3}}$ around $\mathbf{v}_1$, with a mode at 31.1°. *Middle row and in the green histogram*: Variability of the fiber orientation, $\sigma_{\theta_{1,3}}$ around $\mathbf{v}_2$, with a mode at 11.5°. *Bottom row and in the red histogram*: Variability of the fiber orientation, $\sigma_{\theta_{1,2}}$ around $\mathbf{v}_3$, with a mode at 13.0°.

The main difference between our results and the ones in (Greenbaum *et al.*, 1981) (human data) and (Peyrat, 2009) (canine data) is the absence of inflection points in the shape of the helix angle distribution. This might be due to the low resolution of our dataset. On the other hand, the transverse angle is less correlated with $\rho = 0.286$, i.e., it depends less on the fiber location. The average transverse angles suggest indeed that the fibers are relatively parallel to the epicardium with an average angle of $+7°(\pm 31°)$. The angles appear to be more stable in the midwall with an average transverse angle of $+9°(\pm 12°)$. The transverse angles are higher on the endocardium with $+34°(\pm 29°)$.

The delineation of the myocardium into 17 AHA segments (by the American Heart Association (Cerqueira *et al.*, 2002)) in the left ventricle gives more details in each of these regions. The angular variability is summarized in Fig. 3.11. The correlation factor of the helix angle appears to be high across all myocardial segments. The apex (segment 17), however, shows a low correlation factor of $\rho = 0.355$. A canine study (Peyrat, 2009) showed similar results with a coherent helix angle, also with small dissimilarities across the segments.

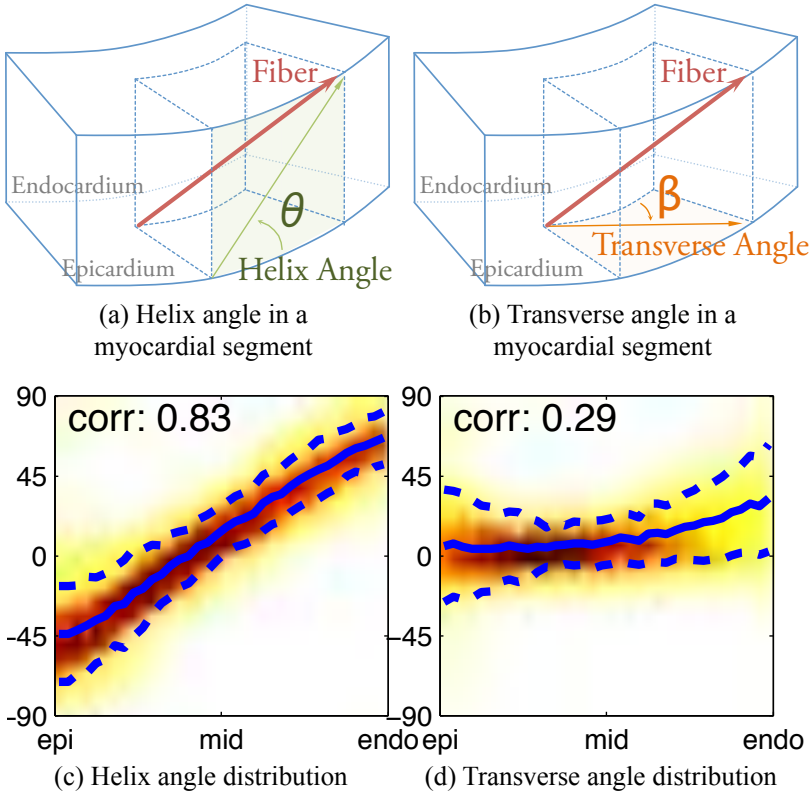

(a) Helix angle in a myocardial segment

(b) Transverse angle in a myocardial segment

(c) Helix angle distribution

(d) Transverse angle distribution

Figure 3.10: *Transmural variability*: *(a)* Illustration of the helix angle in a myocardial segment (angle on the tangent plane of the epicardium, in green), *(b)* of the transverse angle (angle on the horizontal normal plane, in orange), *(c)* joint histogram showing the distribution of the helix angle, varying from $-41°$ to $+66°$, and *(d)* the transverse angle, which shows fibers more parallel to the epicardium. The x axis is the transmural distance from epicardium *(left)* to endocardium *(right)*. The y axis is the helix or transverse angle. The color is the probability distribution of an angle for a particular transmural distance. The thick blues lines are the average angles across the wall, and the dashed lines are the one-standard-deviation envelope. The correlation factor between the angle and the transmural distance is given.

3.4 Limitations

The statistical analysis brings out a challenging task, still unsolved: distinguishing the true variability of the fiber structure from errors due to acquisition and registration. The statistical analysis is likely to suffer from several limitations. Starting with the imaging of the cardiac fibers, the choice of the acquisition protocol affects the image resolution and noise. It was shown in (Frindel *et al.*, 2007) that among a panel of acquisition protocols applied specifically to human heart *ex vivo* DWI with a comparison of different sets of directions

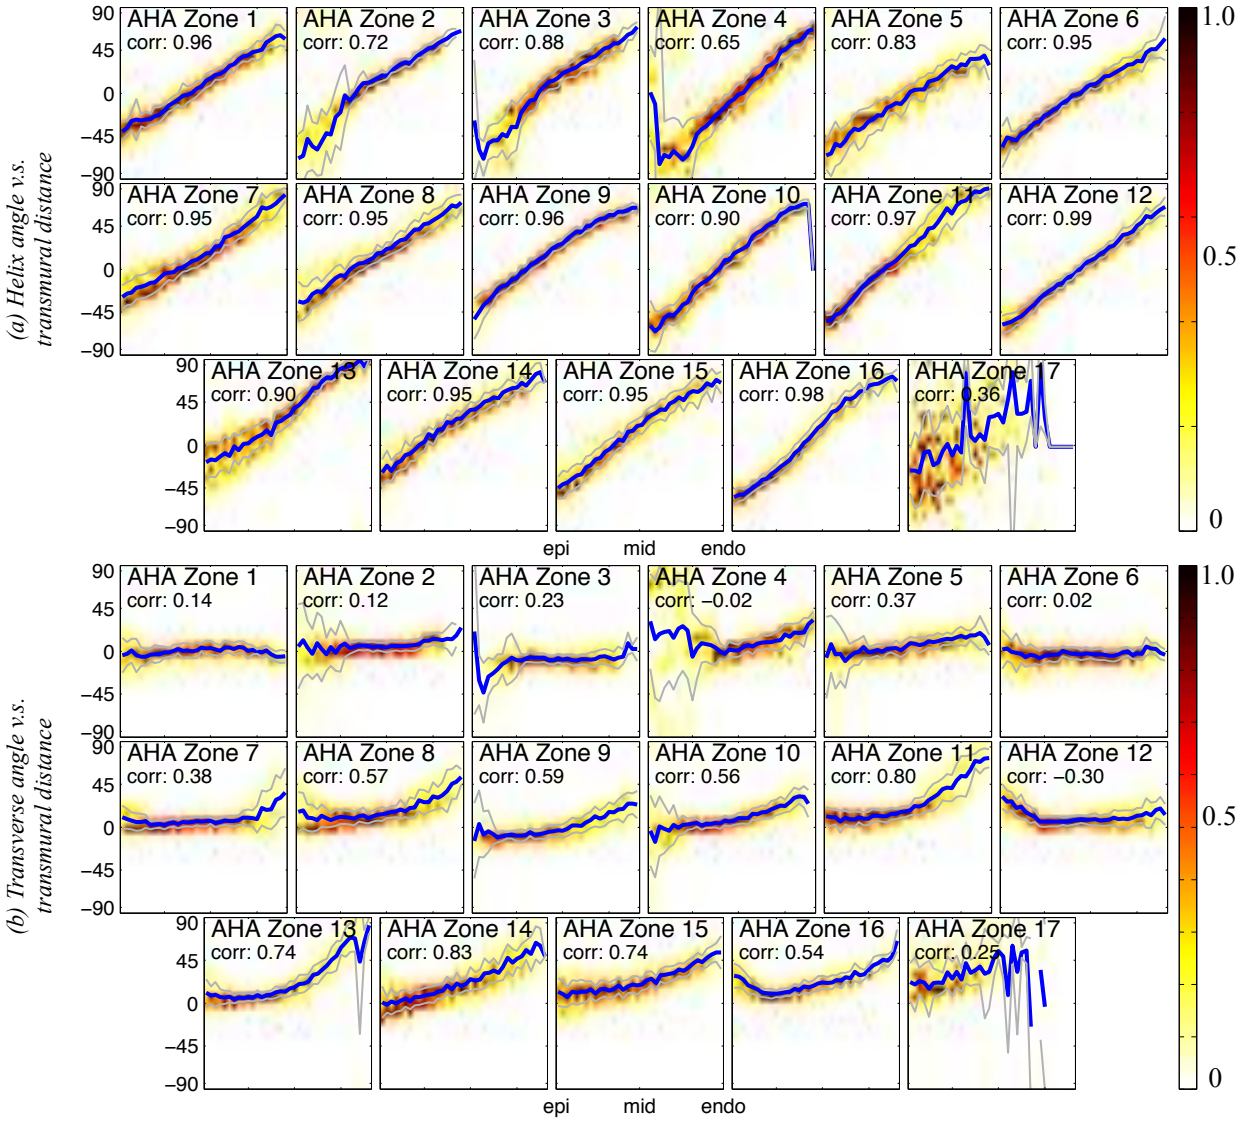

Figure 3.11: Joint histograms showing the distribution along the transmural depth of *(a)* the helix angle and *(b)* the transverse angle of the human cardiac fibers, from the epicardium (*left side of histogram*) to the endocardium (*right side of histogram*), in the 17 AHA segments of the left ventricle. The thick blue lines are the average angles across the wall, and the light gray lines are the one-standard-deviation envelope. The correlation factor between the angle and the transmural distance is given for each segment.

and repetitions, the use of 12 directions with 4 repetitions gives for *ex vivo* the best results in terms of fiber direction coherence, while only 6 directions are preferred in (Hasan *et al.*, 2001). Using 48 unique directions might for instance provide a better angular resolution in addition to a better signal to noise ratio. The imaging protocol described in (Frindel, 2009; Frindel *et al.*, 2009) shows that the peak signal-to-noise ratio during acquisition of the dataset is 16.44 dB. The variability due to acquisition can be assessed more precisely by imaging several times the same heart, or with bootstrapping resampling methods (Jones, 2008; Hasan *et al.*, 2992; Pajevic and Basser, 2003) when the number of acquisitions is limited. Our results include these inherent uncertainties and can be further refined with more precise acquisition schemes. The variability study can also be limited by a partial volume effect which is enhanced in the presence of larger voxel sizes. For instance, the myocardium in our dataset (which has a resolution of 2mm^3) appears as 6 or 7 voxel wide. The spatial resolution was mostly limited due to time constraints. Indeed, the hearts had to be imaged soon after death within a limited amount of time before being given back for autopsy in its original state. Additionally, the lack of fixative or any solutions to prevent contracture may have influenced the fiber architecture. Nonetheless, our framework is independent of the acquisition protocol and our results can be refined in the future with any improved diffusion tensor imaging protocol.

The method used to construct the atlas also has limitations. For instance, the segmentation of the myocardium has not been designed to remove the few voxels representing the fat lying on the myocardium. This is for instance observed in subject 1 (Fig. 3.5) where darker voxels surrounds part of the myocardium. Another limitation comes from the large variability in the heart shapes among a population, which makes the registration a challenging step with the presence of large deformations. This is for example particularly true in the right ventricular and apical regions. The Diffeomorphic Demons, guaranteeing smooth direct and inverse deformation fields, were chosen for this reason. Although no ground truth is available to assess the registration accuracy, the Dice metrics between the myocardium masks of all subjects and the atlas (defined as the ratio of intersection to addition set: $2\left(|S| \cap |T|\right) / \left(|S| + |T|\right)$, with $|S|$ and $|T|$ being mask volumes), which range from 0.86 to 0.88, can provide a certain confidence that there is a good boundary overlap between all registered hearts and the atlas. Moreover, the statistical study is currently limited (Young, 1978) by the number of available hearts. Since human hearts classified as *healthy* are ex-

tremely difficult to obtain, we chose to use the maximal number of hearts available to us. Our analysis on the variability of the fiber architecture is thus limited by the inherent errors due to acquisition and registration.

3.5 Discussion and Conclusions

The work reported in this paper provides the first statistical atlas of the human cardiac fiber architecture that enabled a quantitative analysis of the fiber variability among a healthy population. The statistical atlas has been constructed from 10 *ex vivo* hearts with a method minimizing user interaction. The hearts were nonrigidly registered by purposely omitting any micro structure information (e.g., fiber orientations) in order to avoid a bias in the statistical study of the fiber architecture (e.g., similarly to studying the variability of fingerprint lines, we would not include information on these line orientations, as this is exactly what we want to measure; we want to see how variable the line orientations are in different regions and across a population; registering directly these fingerprint lines would in fact remove, or lessen, any variability in their orientations). The transmural walls were all registered in a similar manner, all based on the myocardium shapes. The spatial variability of the fiber architecture were thus preserved in the registered tensor fields. The deformation fields, obtained while constructing the morphological atlas of the heart, were used to warp all tensor fields, with the finite strain reorientation strategy, to the atlas reference. Another approach would have been to warp directly the diffusion-weighted images and reestimate the diffusion tensor. This approach would have required a strategy for the local deformation of the gradient orientations which is also not necessarily trivial. A mean tensor field was computed using the Log-Euclidean metric. The variability of the mean tensor field is expressed in a covariance matrix. Global statistics on the whole diffusion tensor highlights myocardial regions of high variability. The diffusion tensors in the compact left ventricular myocardium remain stable with a global variability of 13.2%. The statistical framework provides means to study the variability of the eigenvectors in specific directions, where, the fibers are shown to vary with $\pm 11.5^\circ$ in the $\mathbf{v}_{1,3}$ plane and with $\pm 13.0^\circ$ in the $\mathbf{v}_{1,2}$ plane. Their variability is coherent across the whole myocardium. The secondary and tertiary eigenvectors, assumed to be associated with the laminar sheet normal, are, however, less stable compared to the fiber structure with a variability of $\pm 31.1^\circ$. This concurs with a previous canine study (Peyrat *et al.*, 2007). Such high variability either shows that the laminar sheet structure is

less organized than the fiber structure (e.g. the presence of two populations of orthogonal laminar sheets randomly distributed over a given heart (Helm, 2005; Lombaert *et al.*, 2011d)), or not be present everywhere. Additionally, the distribution of the actual angular values of the fiber orientation has been studied across several myocardial segments. The helix angle spans from $-41°(\pm 26°)$ on the epicardium to $+66°(\pm 15°)$ on the endocardium. The angular transmural distribution shows an helix angle highly correlated with the transmural depth. The histological study of *postmortem* human hearts (Greenbaum *et al.*, 1981) similarly observe a transmural correlation of the helix angle. Differences in the shape and in the variance of the helix angle distribution between our results and the ones in (Greenbaum *et al.*, 1981; Peyrat *et al.*, 2007) might be due to our coarser resolution. The average transverse angle shows that the cardiac fibers are relatively parallel to the myocardium epicardium. The small deviations from the average transverse angle might also be due to an underlying spiral architecture of the fibers, where fibers tracked initially on the epicardium create a spiral around the ventricle and find themselves at last to be on the endocardium. This change might occur in the apical sections where the deviations from zero are more apparent than in the basal sections.

The statistical study of the fiber architecture which plays a key role in mechanical and electrical cardiac functions, gives a better understanding of the human heart by providing the spatial distribution of fiber angles with their variations. These numbers will eventually be refined with the availability of more *ex vivo* human hearts and with DT-MRI acquisitions at higher resolutions. Moreover, the computation of the statistical atlas is simplified with the developed semi-automatic method. The human statistical atlas will thus be easily refined with future acquisitions of healthy hearts. Larger multi-population studies involving many hearts will likewise benefit from our method. For instance, cardiac pathologies could be characterized by comparing the fiber architecture against this atlas of healthy hearts (Lombaert *et al.*, 2011c). A more appropriate comprehension of the human cardiac fiber architecture is also relevant to the creation of more elaborate computational models that could for instance be used for the planning of radiofrequency ablation (RFA) and cardiac resynchronization therapy (CRT), or to fiber-based surgical treatments (Cirillo, 2009; Cirillo *et al.*, 2010), which offer a promising perspective to the restoration of failing ventricles. The use of this human statistical atlas could also ideally improve the diagnosis and the follow-up of cardiac diseases related to fiber structural defects. With ongoing research in *in vivo* DT-MRI or in

Shear Wave Imaging (Lee *et al.*, 2012), extrapolation of sparse *in vivo* acquisitions with an accurate human atlas could pave the way for more personalized *in vivo* imaging (Toussaint *et al.*, 2010) and cardiac modeling.

Acknowledgement

The authors wish to acknowledge helpful comments from Leon Axel, constructive discussions with Xavier Pennec, Maxime Sermesant and Marco Lorenzi from the Asclepios Team, as well as the anonymous reviewers for their valuable comments and suggestions.

CHAPTER 4 FOCUSR: FEATURE ORIENTED CORRESPONDENCE USING SPECTRAL REGULARIZATION – A METHOD FOR ACCURATE SURFACE MATCHING

Herve Lombaert[1,3], Leo Grady[1], Jonathan R. Polimeni[2], Farida Cheriet[3]

Presentation

This chapter presents the article ‘*“FOCUSR: Feature Oriented Correspondence using Spectral Regularization – A Method for Accurate Surface Matching”* (Lombaert *et al.*, 2012a) submitted to **IEEE PAMI** *(Transactions on Pattern Analysis and Machine Intelligence)*, sent on January 6th, 2012, and currently under revision. An initial article (Lombaert *et al.*, 2011a) was published in the conference **IPMI** *(Information Processing in Medical Imaging)* held in Irsee, Germany, in July 2011. The objective of this article is to develop a new approach based on spectral correspondence for brain surface matching. This is a joint work between [1]**Siemens Corporate Research**, *Princeton, NJ, USA*; [2]**Martinos Center at Harvard Medical School**, *Cambridge, MA, USA*; and [3]**École Polytechnique**, *Montreal, Canada*.

Abstract

Existing methods for surface matching are limited by the trade-off between precision and computational efficiency. Here we present an improved algorithm for dense vertex-to-vertex correspondence that uses direct matching of features defined on a surface and improves it by using spectral correspondence as a *regularization*. This algorithm has the speed of both feature matching and spectral matching while exhibiting greatly improved precision (distance errors of 1.4%). The method, FOCUSR, incorporates implicitly such additional features to calculate the correspondence and relies on the smoothness of the lowest-frequency harmonics of a graph Laplacian to spatially regularize the features. In its simplest form, FOCUSR is an improved spectral correspondence method that nonrigidly deforms spectral embeddings. We provide here a full realization of spectral correspondence where virtually *any* feature can be used as additional information using weights on graph edges, but also on graph nodes and as extra embedded coordinates. The full power of FOCUSR is presented in a real case

scenario with the challenging task of brain surface matching across several individuals. Our results show that combining features and regularizing them in a spectral embedding greatly improves the matching precision (to a sub-millimeter level) while performing at much greater speed than existing methods.

4.1 Introduction

Mesh correspondence is a key step in many applications of computer vision and whose precision and speed are crucial. It is at the core of studies on shape variability and on object and motion analysis. In the medical field, precision is essential and a fast method enables investigations on large studies between organs or individuals. The challenge of shape matching is to find the dense correspondences mapping all points on one surface to their equivalent points on a second surface. This task becomes particularly arduous when the matching involves highly convoluted surfaces or two surfaces representing different poses of an articulated object. Early solutions (Besl, 1988) to this problem, aligning surface models, were limited to rigid transformations (Rusinkiewicz and Levoy, 2001), or relied on fiducial markers placed on the surfaces (Audette *et al.*, 2000; Schreiner *et al.*, 2004). Methods based on deformable surfaces (Shelton, 2000; Hahnel *et al.*, 2003; Zhang *et al.*, 2008b) could find nonrigid transformations. However, to keep these approaches tractable, prior knowledge on the underlying deformation between surfaces (Lin, 1999; Blanz and Vetter, 1999), or the use of control points (Allen *et al.*, 2003; Sumner and Popović, 2004), is often required to restrain the search domain. Rather than optimizing for a deformation, other approaches would directly solve for the correspondence map (Anguelov *et al.*, 2004), avoiding iterative deformations of the surfaces. Moreover, these surfaces may be correlated with measurable *features* other than their explicit mesh geometry. For example, the method used in FreeSurfer (Fischl *et al.*, 1999), a leading tool for brain surface reconstruction and matching, uses geometric features such as cortical curvature and sulcal depth (the depth in the cortical folding pattern) (Lohmann *et al.*, 2008) to drive the warping of one brain surface into another surface. However, despite its precision, FreeSurfer suffers from a substantial computational burden, taking hours to compute a correspondence map between typical brain surface models consisting of hundreds of thousands of vertices. Needless to say, the incorporation of additional features is a convenient aspect for a matching algorithm.

A direct method of matching two surfaces based on features (e.g., the geometry of the

cerebral cortex in brain matching, or texture intensities for articulated object matching) is to treat the available features as characteristic *signatures* which can be used to identify each vertex within the surface mesh. With these signatures, a vertex on one surface could be mapped to the vertex on a second surface which most closely resembles the same characteristic features (e.g., by computing a Euclidean distance between the feature vectors). This feature matching technique would have the merits of being fast (e.g., computable within Voronoi cells) and flexible enough to allow any set of features to drive the matching. Unfortunately, this feature matching technique would completely ignore the spatial organization of the surface vertices and result in a highly non-smooth mapping between the surfaces. Our approach to the matching problem seeks to preserve the speed and flexibility of direct feature matching and address the problem of smooth mapping by using an improved spectral correspondence as a *regularization*.

Spectral correspondence (Chung, 1997b) utilizes a graph (mesh) *spectrum*, which is the set of Laplacian eigenvalues and eigenvectors (illustrated on Fig. 4.1), to produce a vertex correspondence between two graphs (meshes). The key utility of spectral correspondence in our context is to provide a spatial regularization on the correspondence map. This regularization is enabled by the fact that the low-frequency eigenvectors (those corresponding to small eigenvalues) are spatially smooth, as they represent low-frequency harmonics (Grady and Polimeni, 2010). Put differently, all neighboring nodes will have a small change in the values of these harmonics, meaning that a correspondence driven by these harmonics will map neighboring nodes to neighboring locations in the range space. The value of the harmonics at each vertex are known as the *spectral coordinates* of the vertex. At its core, our technique for spectral regularization is to supplement the direct feature matching technique described above by extending the vertex *signature* to additionally include the *spectral coordinates* of each vertex. Matching vertices are revealed with similarities in such extended signatures. Fortunately, including the spectral coordinates in our matching maintains the speed and flexibility of the simple technique. We call our method FOCUSR for Feature Oriented Correspondence Using Spectral Regularization.

Spectral methods have been used in many fields, such as in the segmentation and registration of shapes in images (Reuter, 2009), in the indexing of structures (Shokoufandeh *et al.*, 2005), or in the clustering of data (Shi and Malik, 2000; Bengio *et al.*, 2004; Bach and Jordan, 2004). Their use in shape matching is often limited to hierarchical matching

(e.g., matchings of limbs in body models, or of large surface areas). Few medical applications of spectral methods exist and are targeted to brain studies in order to study the geometrical patterns of the anatomical structures of the brain such as the cortical folds (Niethammer *et al.*, 2007; Reuter *et al.*, 2009; Shi *et al.*, 2009, 2010) and with the smoothing of cortical surfaces (Anqi *et al.*, 2006). Spectral coordinates have also been used directly for graph partitioning (Chan *et al.*, 1995). Umeyama (Umeyama, 1988) and later Scott and Longuet-Higgins (Scott and Longuet-Higgins, 1991) pioneered the use of spectral methods for the correspondence problem. Shapiro and Brady (Shapiro and Brady, 1992) compared ordered eigenvectors of a proximity matrix to find correspondences. Their work served as a basis for future spectral correspondence methods. Variants include the use of different proximity matrices using different kernels, the use of the adjacency matrix, shock graphs (Pelillo *et al.*, 1999; Siddiqi *et al.*, 1999), different normalized Laplacian matrices, or the use of Multi-Dimensional Scaling (Schwartz *et al.*, 1989; Elad and Kimmel, 2003; Wuhrer *et al.*, 2009; Bronstein *et al.*, 2006, 2007). Recent surveys covering the use of spectral methods in the past fifteen years are available in (van Kaick *et al.*, 2011; Zhang *et al.*, 2010). Mateus *et al.* (Mateus *et al.*, 2008) proposed an original unsupervised spectral method with an alternative to eigenvalue ordering based on eigenvector histograms and refining the eigenvectors alignment with a probabilistic point matching within the framework of the EM algorithm (Carcassoni and Hancock, 2003). Jain and Zhang (Jain and Zhang, 2006) approach the eigenvector alignment problem with a nonrigid deformation based on Thin Plate Splines.

Spectral correspondence has presented several difficulties that act as a barrier to its widespread adoption. Specifically, when computing the eigenvectors for two surfaces, the signs of the eigenvectors need to be aligned (the eigenvectors are ambiguous to sign), the eigenvectors sometimes require reordering (due to near algebraic multiplicity of the eigenvalues causing ordering changes of the spectral coordinates). Additionally, spectral matching methods typically start with a rigid alignment of the eigenvectors to account for translation and scaling of the spectral coordinates. Small variations however exist in the spectral coordinates (due to non perfect shape isometry, e.g., local expansion and compression within meshes). There is therefore a need for robust nonrigid point correspondence between spectral coordinates. Furthermore, the use of vertex features has not been fully realized in previous work on spectral correspondence, which have incorporated these features only to produce edge weights (measuring changes between neighboring features) rather than as node

weights (using the features themselves). We address and improve all these aspects of spectral correspondence while additionally using the spectral coordinates to provide a smooth regularization of the simple feature matching technique. This work makes several contributions to dense surface matching:

- Extending simple surface feature matching with spectral regularization. The integration of spectral components in extended signatures alongside feature characteristics provides a natural means of regularization (i.e., matching extended signatures reveal matching vertices).

- Nonrigid alignment of the multidimensional embeddings (i.e., of the extended signatures, rather than only the spectral coordinates).

- The weighting of *nodes* in the graph Laplacian, which controls the influence of each node during correspondence.

- A global approach to handle automatically the sign ambiguity and the rearrangement of the graph Laplacian eigenvectors.

After detailing FOCUSR in the next section, we show in controlled experiments that it outperforms both direct feature matching and conventional spectral correspondence. Firstly, we demonstrate that nonrigid alignment of the spectral coordinates improves the direct matching method, while showing at the same time that FOCUSR can be used with a variety of generic meshes. We chose deformed meshes of galloping animals and changing facial expressions matched with their references. The comparisons of the correspondence maps with their exact ground truth reveal minimal errors. Secondly, we expose the full power of FOCUSR with the use of additional features and assess its precision with the challenging task of brain surface matching. Indeed, while the sulcal and gyral folding pattern of the human cerebral cortex are somewhat stable across individuals, some geometric variability does exist (Hinds *et al.*, 2008), making the direct use of the folding geometry unsuitable for surface matching. This application to the problem of brain matching provides a platform for FOCUSR where the use of additional features available in the brain data—such as cortical Gaussian curvature, sulcal depth, and cortical thickness—can improve the matching precision. We show that FOCUSR produces results in a fraction of the time required by FreeSurfer while maintaining the same level of precision. We believe that this large gain in processing speed would

make possible new brain studies that were previously limited by computational burden, or, more generally, studies on meshes that wish to use non-standard features for driving the correspondence.

4.2 Methods

We begin our exposition of FOCUSR by detailing a simple technique for feature matching that does not preserve smoothness of the mapping between surfaces (Fig. 4.2 a). We then describe how spectral coordinates can be used to regularize feature matching. We re-examine and improve each step in the spectral correspondence process to overcome previous limitations with spectral correspondence. The algorithm is summarized in Fig. 4.3. Code implementation in Matlab is available at `http://step.polymtl.ca/~rv101/focusr`[1].

4.2.1 Direct Feature Matching to Provide Vertex Correspondence

Assume that we have two graphs, $\mathscr{G}_1 = \{\mathscr{V}_1, \mathscr{E}_1\}$ and $\mathscr{G}_2 = \{\mathscr{V}_2, \mathscr{E}_2\}$ (with vertices and edges) such that a correspondence $\phi : v_i \in \mathscr{V}_1 \rightarrow v_j \in \mathscr{V}_2$ is desired. Note that we do not require that $|\mathscr{V}_1| = |\mathscr{V}_2|$ or $|\mathscr{E}_1| = |\mathscr{E}_2|$ (i.e., meshes can have different sizes and structures). Consequently, there is no guarantee that the mapping is one-to-one and may not be invertible. We will use the terms *node*, *vertex* and *point* interchangeably to describe a member of $\mathscr{V}_1$ or $\mathscr{V}_2$. Given a set of K features $\mathbf{X}_i$ at every node $v_i \in \mathscr{V}_1$, and a set of K features $\mathbf{Y}_j$ at every node $v_j \in \mathscr{V}_2$, our goal is to use these features to produce a correspondence ϕ.

A direct feature matching approach to producing this correspondence would be to set

$$\phi(v_i) = \min_{v_j \in \mathscr{V}_2} ||\mathbf{X}_i - \mathbf{Y}_j||, \tag{4.1}$$

which could be computed quickly by precomputing a Voronoi tessellation of the range space. Unfortunately, this simple technique has several inadequacies. Specifically, the technique based on the Voronoi tesselation does not properly account for global changes in the feature space (e.g., due to a global scaling or translation), nor does it utilize the neighborhood structure provided by the edge sets in any way (i.e., there is no spatial regularity to the mapping in the sense that neighbors in the domain are unlikely to remain neighbors in the range).

[1] This private link is going public after publication

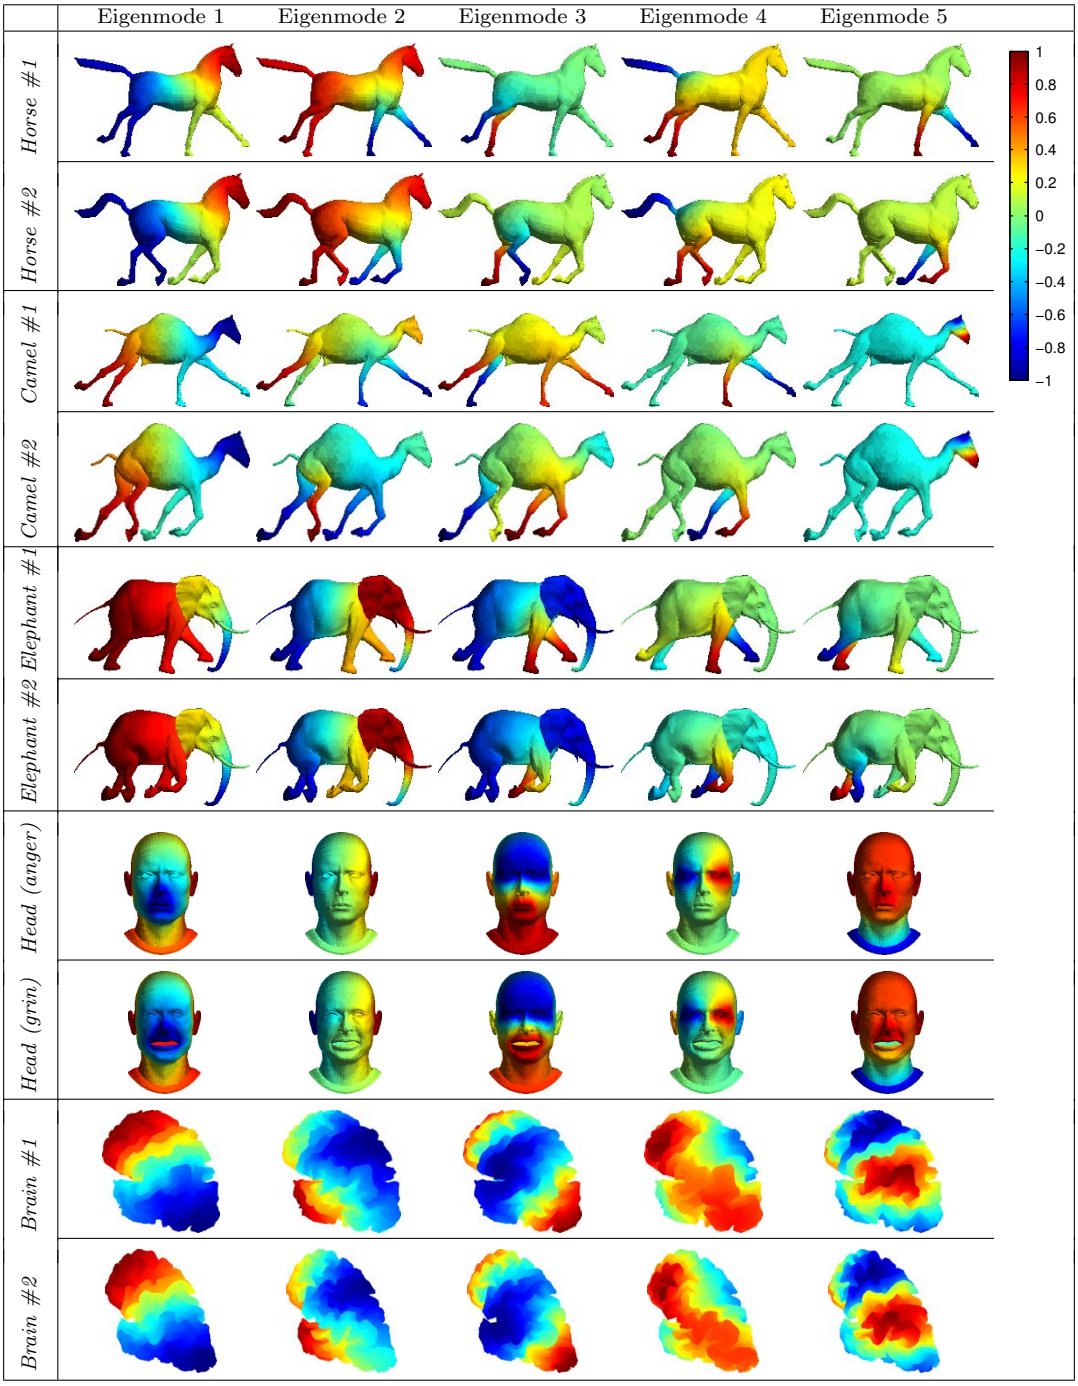

Figure 4.1: Two frames from a sequence of a galloping horse, galloping camel, galloping elephant, head poses, and two human brain surfaces. Each row shows the first five spectral components (eigenmodes) of a model, given by the eigendecomposition of the graph Laplacian of the model (eigenmodes have been reordered and their sign adjusted, so both sets are equivalent). The coloring indicates the value of a spectral component (the eigenmodal value between -1 and $+1$) for each vertex over the surface, i.e., the set of color values within a row corresponds to the 5D spectral coordinates of all vertices.

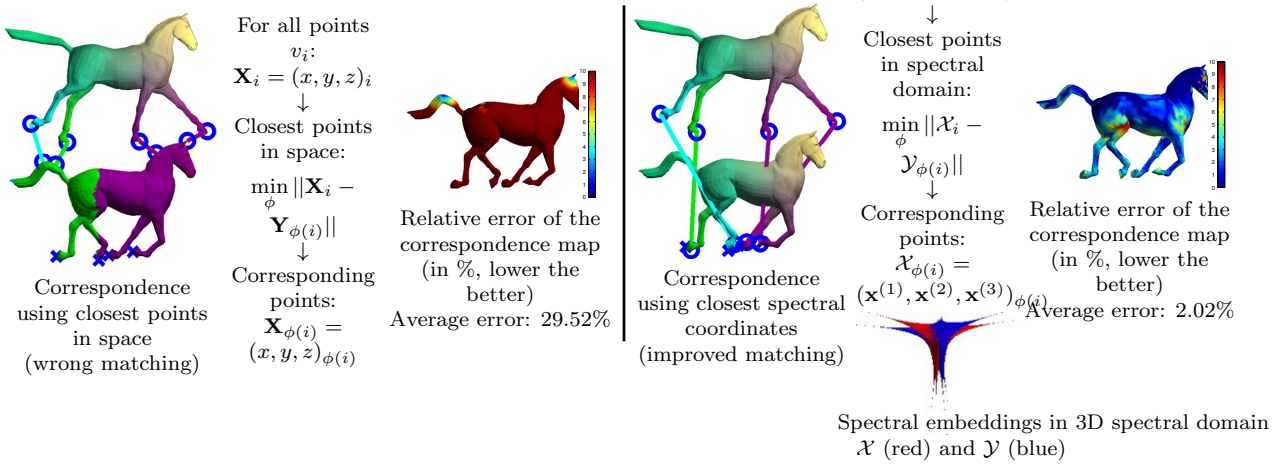

(a) Finding closest points in *space* (missed legs)

(b) Finding closest points in the *spectral domain* (matched legs)

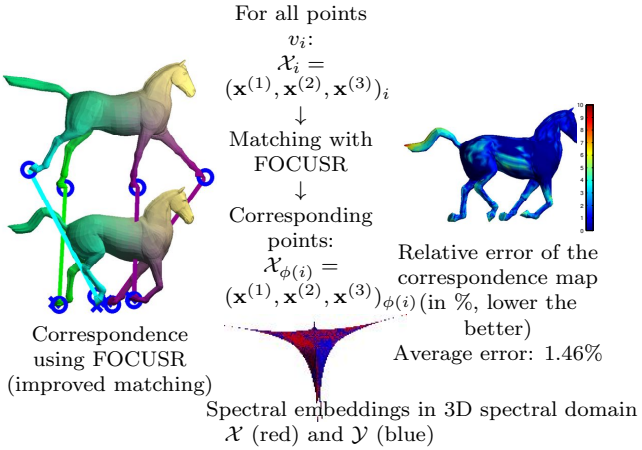

(c) Using FOCUSR in its simplest setting(nonrigidly deformated embeddings)

Figure 4.2: *Direct matching* (coloring indicates correspondence, and links and circles indicate matching of leg extremities, crosses indicate ground truth) : *(a)* Finding closest points in *space*: this naive correspondence map is computed by finding for each point of model X its closest point in *space* of model Y (match X with Y). This strategy generates an inconsistent correspondence map. *(b)* Finding closest points in the *spectral domain*: the correspondence map is computed by finding for each point of model X its closest *spectral* equivalent in model Y (match $\mathcal{X}$ with $\mathcal{Y}$ instead of X with Y). Even though the meshes are not aligned in space (they are translated), their spectral embeddings (red is $\mathcal{X}$, blue is $\mathcal{Y}$, both use three eigenmodes for 3D visualization) are almost perfectly superimposed in the spectral domain. *(c) FOCUSR in its simplest setting* (with no additional features) : Our method performs matching in the *spectral domain* (with lower error over the surface) and improves the alignment of the spectral embeddings.

Global changes in the feature space can be accounted for by using a more sophisticated point correspondence than what is described in Eq. (4.1). Robust Point Matching (Chui, 2003) with a Thin Plate Spline-based transformation is often used for 2D or 3D registration. However, with this approach the final registration depends on the number and choice of the control points. A more recent approach to the point correspondence problem is the Coherent Point Drift (CPD) method (Myronenko and Song, 2009) which is fast and demonstrates excellent performance. To summarize this method, the registration is treated as a Maximum Likelihood problem where Gaussian Mixture Model centroids are fit into a point set. There is no assumption on the global transformation between point sets. Instead, the evolution of the transformation is constrained with a motion coherence (Myronenko and Song, 2009). The CPD algorithm offers the possibility to perform matching on a subset of the points (for increased speed) while computing the transformation in the continuous domain (i.e., the continuous transformation, found with only a subsample of $\mathscr{V}_1$ and $\mathscr{V}_2$, can be applied on all points of $\mathscr{V}_2$ and thus find a dense matching between $\mathscr{V}_1$ and $\mathscr{V}_2$). Furthermore, each feature (i.e., each coordinate of $\mathbf{X}_i$ or $\mathbf{Y}_j$) can be weighted in order to accentuate or reduce its influence.

Although CPD provides a method to account for global transformation in the feature space between the two graphs, it is still necessary to incorporate spatial regularity into the mapping such that neighboring points in $\mathscr{V}_1$ map to neighboring points in $\mathscr{V}_2$. Note that a strict neighbor-to-neighbor mapping is only possible when the two graphs are isomorphic. Since we target a more general scenario, we want to account for neighborhood relationships by promoting a correspondence that maps nearby nodes in $\mathscr{V}_1$ (based on $\mathscr{E}_1$) to nearby nodes in $\mathscr{V}_2$ (based on $\mathscr{E}_2$). Our strategy for promoting spatial regularization is to supplement $\mathbf{X}_i$ and $\mathbf{Y}_j$ with the *spectral coordinates* at nodes v_i and v_j before applying the CPD point correspondence. The values of the spectral coordinates over a few sample surfaces are illustrated in Fig. 4.1. The fundamental difference between the use of $\mathbf{X}$ and $\mathbf{Y}$ as general feature vectors (illustrated with 3D coordinates (x, y, z)) and as spectral coordinates is demonstrated in Fig. 4.2 b. The low-frequency spectral coordinates are dependent on the geometry of the surface, and these coordinates are effectively more stable across articulated shapes or highly deformable shapes, i.e., normalizing these shapes in a same referential. Additionally, they are known to be spatially smooth (see below) in accordance with the low-frequency harmonics of an elastic surface (Grady and Polimeni, 2010). In the next section we will review spectral

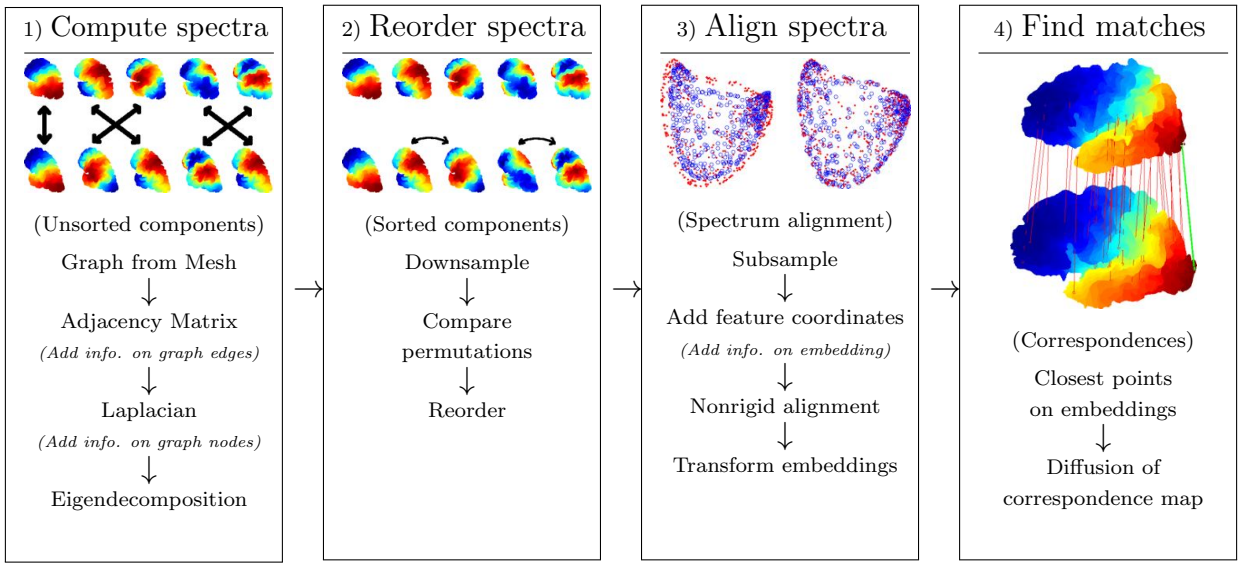

Figure 4.3: FOCUSR overview for matching a pair of surfaces. *First*, we build a graph out of each surface mesh and set the graph edge weights and graph node weights to construct the Laplacian matrix. The eigendecomposition of each graph's Laplacian matrix reveals its spectral components. *Second*, we reorder the spectral components by finding the optimal permutation of components between the pair of meshes. *Third*, regularization is performed by matching the spectral embeddings. *Finally*, corresponding points are found with closest points in both spectral embeddings, and the final correspondence map is diffused.

coordinates, and demonstrate improvements to traditional methods for solving some of the difficulties associated with comparing spectral coordinates from two graphs.

4.2.2 Spectral Coordinates

We may define the $|\mathscr{V}| \times |\mathscr{V}|$ adjacency matrix W of a graph in terms of *affinity weights* (see (Grady and Polimeni, 2010)), which are derived from a given distance metric $\mathrm{dist}(i,j)$ between two neighboring vertices (v_i, v_j). The elements of the weighted adjacency matrix are given by

$$W_{ij} = \begin{cases} 1/\operatorname{dist}(i,j) & \text{if } \exists\, e_{ij} \in \mathscr{E}, \\ 0 & \text{otherwise} \end{cases} \tag{4.2}$$

The matrix W provides a weighting on the graph edges derived from the given distance metric. The distance may be derived from the geometry via the vertex coordinates $\mathbf{x} = (x, y, z)^{\mathrm{T}}$ embedded in space (e.g., $\mathrm{dist}(i,j) = \|\mathbf{x}_i - \mathbf{x}_j\|$, the distance between nodes v_i and

v_j), from feature vectors (e.g., $\text{dist}(i,j) = \|\mathbf{F}_i - \mathbf{F}_j\|$, where $\mathbf{F} = (\mathbf{f}^{(1)}, \ldots, \mathbf{f}^{(K)})^{\mathrm{T}}$ for K features), or both. The more general edge weighting between vertices v_i and v_j uses the ℓ_2 norm between extended vectors:

$$w_{i,j} = \text{dist}(i,j) = \|(\mathbf{x}_i, \gamma\mathbf{F}_i) - (\mathbf{x}_j, \gamma\mathbf{F}_j)\|\,, \tag{4.3}$$

where $(\mathbf{x}, \gamma\mathbf{F})$ is the concatenation of the 3D coordinate values $\mathbf{x} = (x, y, z)^{\mathrm{T}}$ with the K feature values $\mathbf{F} = (\mathbf{f}^{(1)}, \ldots, \mathbf{f}^{(K)})^{\mathrm{T}}$. The $K \times K$ diagonal matrix γ contains the K weights controlling the influence of each feature. To compensate for the different scalings of the feature values, each feature vector $\mathbf{f}^{(k)}$ is normalized with respect to the range of the 3D coordinate values $\mathbf{x}$ (i.e., feature values are normalized such that $\min(\mathbf{f}^{(k)}) = \min(\mathbf{x})$ and $\max(\mathbf{f}^{(k)}) = \max(\mathbf{x})$).

The general Laplacian operator on a graph was formulated in (Grady and Polimeni, 2010) as a $|\mathscr{V}| \times |\mathscr{V}|$ matrix with the form:

$$L = G\,(D - W)\,, \tag{4.4}$$

where D, the degree matrix, is a diagonal matrix defined as $D_{ii} = \sum_j W_{ij}$, and G is the diagonal matrix of node weights. Typically in spectral correspondence G is set to identity $G = I$, or to $G = D^{-1}$. However, we propose here to replace the default assignment $G = D^{-1}$ with *any* meaningful node weighting. In particular, we propose to use a function of feature *magnitudes* to establish the (positive-valued) node weighting based on the assumption that nodes with significant features are of more interest to match precisely (i.e., nodes with large weight have a greater influence on the spectral correspondence than low-weight nodes). For example, if half of the nodes in a graph had a large weight and the other half had a small weight, the Laplacian eigenvectors would closely resemble the eigenvectors of the large-weight subgraph. The diagonal of matrix G contains the general node weights for each vertex v_i:

$$w_i = G_{ii} = \frac{1}{d_i} \cdot \frac{1}{\displaystyle\sum_{k=1}^{K} \gamma_i \rho\left(f_i^{(k)}\right)}, \tag{4.5}$$

where d_i is the node degree (i.e., D_{ii}), γ is the previously mentioned feature weights, and $\rho(\cdot)$ is a function that enforces positive values (e.g., $\rho(f) = f^2$ or $\rho(f) = \exp(f)$). The

denominator in Eq. (4.5) contains the sum of the influences of each feature on vertex v_i. We used $\rho(f) = \exp(f)$ to promote correspondence between nodes having the largest feature components (which we assume indicate greatest significance).

The right eigenvectors of the Laplacian matrix comprise the graph spectrum $\mathcal{X} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(n)})^{\mathrm{T}}$, where $n = |\mathscr{V}|$ is the number of nodes. The values over surfaces for the five lowest frequency eigenvectors are shown on Fig. 4.1, and illustrates the stability of these eigenvectors between articulated or highly deformable shapes. Each eigenvector[2] $\mathbf{x}^{(u)}$ is a column matrix with n values, and represents a different (weighted) harmonic on a mesh surface that corresponds to an inherent property of the mesh geometry. This is in comparison with extrinsic properties such as the spatial location of points (i.e., point coordinates vary when the model takes a different pose). The n values $(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(n)})$ give the *spectral coordinates* of node v_i (i.e., a coordinate in a *spectral* domain). The first eigenvector $\mathbf{x}^{(1)}$ is the trivial (uniform) eigenvector, and the eigenvectors associated with the lower non-zero eigenvalues (e.g., $\mathbf{x}^{(2)}, \mathbf{x}^{(3)}$) represent coarse (i.e., low-frequency) intrinsic geometric properties of the shape. The first of them $\mathbf{x}^{(2)}$ is called the *Fiedler vector* (Chung, 1997b), while eigenvectors associated with higher eigenvalues (e.g., $\mathbf{x}^{(n-1)}, \mathbf{x}^{(n)}$) represent fine (high-frequency) geometric properties. For example, in Fig. 4.1, the values of $\mathbf{x}^{(2)}$ increase along a virtual centerline depicting the global shape of the models (a coarse intrinsic property), while the values of $\mathbf{x}^{(5)}$ depict finer details of the models.

To illustrate why spectral coordinates corresponding to small eigenvalues transition smoothly and slowly across neighboring nodes, consider the Rayleigh quotient

$$\lambda = \frac{x^{\mathrm{T}} L x}{x^{\mathrm{T}} G^{-1} x} = \frac{\sum_{e_{ij} \in \mathscr{E}} w_{ij} (x_i - x_j)^2}{\sum_{v_i \in \mathscr{V}} \frac{1}{w_i} x_i^2}. \tag{4.6}$$

The minimum value of λ is the smallest eigenvalue for L. If the minimization of λ over x is conducted in the space orthogonal to the eigenvector corresponding to the smallest eigenvalue, then the minimum λ is the second smallest eigenvalue (the Fiedler value, in our case). Put differently, all of the eigenvectors corresponding to the smallest eigenvalues have small values of the Rayleigh quotient in Eq. (4.6). Examining the numerator of Eq. (4.6), we see that neighboring nodes must have a small change in the spectral coordinate (eigenvector)

[2] In our notation $\mathbf{x}$ represents the 3D coordinate in space (i.e., x, y, z), and the superscripted $\mathbf{x}^{(i)}$ represents the i^{th} spectral coordinate (i.e., the i^{th} eigenvector of the graph Laplacian).

x in order for the corresponding eigenvalue to be small. However, a small edge weight indicates that the change in x across that edge may be large while still maintaining a small numerator (and therefore a small eigenvalue). Consequently, the edge weights act to enforce a smoother change between similar neighbors, but the spatial regularization is more relaxed for mapping neighboring points which are dissimilar. Examining the denominator of Eq. (4.6), we see that large node weights have the effect of reducing the influence of the node in the denominator, effectively pushing the node to take a value that minimizes the numerator (i.e., the average of its neighbors). By pushing the node to minimize the numerator, the large node weight effectively promotes maximal smoothness in the spectral coordinates at that node.

Consequently, we use the node features to enforce more spatial regularity between similar neighboring nodes (large edge weight) and to enforce more spatial regularity at unremarkable nodes (nodes with small feature magnitude and small node weight). Ultimately, this use of the node features to promote variable spatial regularization is designed to enforce a stronger correspondence between key nodes (nodes with large feature magnitude) and to enforce stronger spatial regularity between key nodes. In this manner, the key nodes (which are similar in the two meshes) are matched strongly, while the remaining nodes are matched to promote maximal spatial regularity.

4.2.3 Spectrum Ordering

Each node is represented with $M \ll |\mathscr{V}|$ spectral coordinates associated with the M smallest (non-trivial) eigenvalues, i.e., the embedded representations for meshes X and Y are $\mathcal{X}^M = (\mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(M+1)})^{\mathrm{T}}$ and $\mathcal{Y}^M = (\mathbf{y}^{(2)}, \ldots, \mathbf{y}^{(M+1)})^{\mathrm{T}}$. Unfortunately, the spectral coordinates of the two meshes may not be directly comparable as a result of two phenomena. First, there exists a *sign ambiguity* when computing eigenvectors, i.e., if $Ax = \lambda x$ (the spectral decomposition of A) then $A(-x) = \lambda(-x)$, which requires checking that each corresponding eigenvector in the two meshes has the same sign. Additionally, as a result of greater algebraic multiplicity of an eigenvalue, it may be possible that the *ordering* of the lowest eigenvectors will change, e.g., if two eigenvectors correspond to the same eigenvalue in both meshes, then the solver may compute these eigenvectors in one order for the first mesh and in the opposite order for the second mesh. A graph with an eigenvalue having algebraic multiplicity greater than one indicates symmetry in the mesh. For large meshes, symmetries (and near

symmetries) is a common problem and the eigenvectors must be reordered.

Our approach to the eigenvector reordering is to compare spectral coordinates at all pairs of closest points between the two meshes. To speed up the reordering, all eigenvectors are subsampled by randomly selecting a subset of $N < |\mathscr{V}|$ nodes (we used 500 nodes or about 0.4% of the vertices in our experiments). The spectral coordinates are normalized to the range $[0, 1]$. The pairs of closest points within these subsampled points determine the correspondence map μ (i.e., vertex $v_i \in \mathscr{V}_1$, on the first mesh, is closest to point $v_{j=\mu(i)} \in \mathscr{V}_2$, on the second mesh). Then, we simply compute for all corresponding points ($v_i \in \mathscr{V}_1 \rightarrow v_{j=\mu(i)} \in \mathscr{V}_2$) the squared difference between the coordinate values $x_i^{(u)}$ and $y_{\mu(i)}^{(v)}$. The differences for all possible pairs of eigenvectors (u, v) are gathered in a $M \times M$ dissimilarity matrix C. The dissimilarity between eigenvector $\mathbf{x}^{(u)}$, on the first mesh, and eigenvector $\mathbf{y}^{(v)}$, on the second mesh, is given by:

$$C(u, v) = \sum_{i=1}^{N} \left(x_i^{(u)} - y_{\mu(i)}^{(v)} \right)^2 .$$

The Hungarian algorithm may be used to find an optimal permutation of eigenvectors $\mathbf{y}^{(v)}$ that minimizes dissimilarity. In the same step we can remove the sign ambiguity by calculating the minimal dissimilarity between all $\mathbf{x}^{(u)}$ and $\mathbf{y}^{(v)}$, as well as between all $\mathbf{x}^{(u)}$ and $-\mathbf{y}^{(v)}$. The cost matrix used in the Hungarian algorithm is thus $Q(u, v) = \min\{C(u, v), C(u, -v)\}$. After permutation π, any eigenvector $\mathbf{x}^{(u)}$ corresponds with $\mathbf{y}^{(\pi(u))}$, and its permutation cost $Q^{(u)}$ is stored for use in the spectral alignment.

To keep the notation simple, in the next sections we assume that the spectral coordinates have been appropriately reordered and signed (i.e., $\mathcal{X}^M$ and $\mathcal{Y}^{\pi \circ M}$ will simply be denoted as $\mathcal{X}^M$ and $\mathcal{Y}^M$ such that $\mathbf{x}^{(u)}$, on the first mesh, corresponds with $\mathbf{y}^{(u)}$, on the second mesh).

4.2.4 Nonrigid Spectral Alignment

Once the reordering and sign adjustment of the eigenvectors have taken place, finding the closest points in the *spectral* domain between embeddings $\mathcal{X}$ and $\mathcal{Y}$ generates a smooth correspondence map (Fig. 4.2). However, these embedded representations contain slight differences, mostly due to perturbations of the shape isometries such as small changes in distances where the surface undergoes local expansion or compression between meshes. As illustrated on Fig. 4.4, nonrigid differences in the spectral embeddings become even more

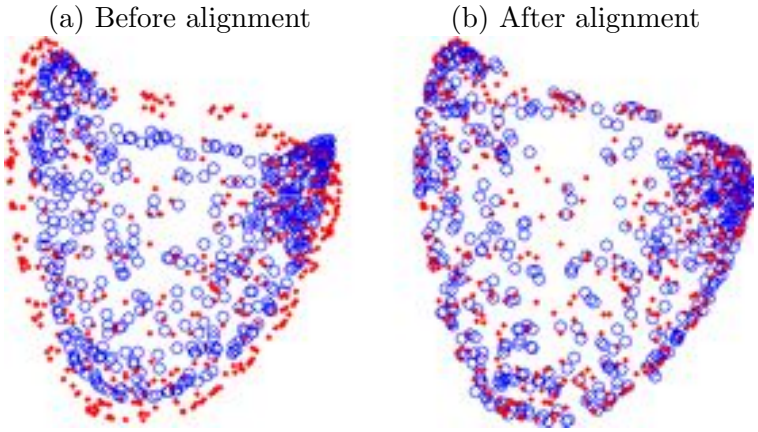

Figure 4.4: Nonrigid alignment of the two spectra corresponding to two brain surfaces. For visualization purposes, the first three eigenvectors ($X^{(2)}$, $X^{(3)}$, and $X^{(4)}$) are used as 3D coordinates (x, y, z). Red and blue are the control points used to align both spectra. Initial spectra *(a)* before and *(b)* after final alignment.

severe in highly convoluted surfaces such as brain cortices. Spectral representations need to be nonrigidly aligned.

Closest points in these nonrigidly aligned embedded representations would reveal corresponding points in both shapes (i.e., in the M-dimensional space (the *spectral* domain), if the point $v_i \in \mathscr{V}_1$ with coordinates $\mathcal{X}_i^M$, is the closest point to $v_j \in \mathscr{V}_2$ with coordinates $\mathcal{Y}_j^M$, then v_i corresponds to v_j). It is at this point where Eq. (4.1) is extended by combining the spectral coordinates, $\mathcal{X}^M$ and $\mathcal{Y}^M$, with the feature vectors, $\mathbf{F}_x = (\mathbf{f}_x^{(1)}, \ldots, \mathbf{f}_x^{(K)})^{\mathrm{T}}$ for nodes in model X, and $\mathbf{F}_y = (\mathbf{f}_y^{(1)}, \ldots, \mathbf{f}_y^{(K)})^{\mathrm{T}}$ for nodes in model Y, to enable spatial regularization in the correspondence map. The extended vectors of Eq. (4.1) becomes:

$$\mathbf{X} = (c_x \mathcal{X}^M, \beta \mathbf{F}_x), \tag{4.7}$$

$$\mathbf{Y} = (c_y \mathcal{Y}^M, \beta \mathbf{F}_y), \tag{4.8}$$

where c_x and c_y are $M \times M$ diagonal matrices that contain weights influencing each spectral coordinate, and β is a $K \times K$ diagonal matrix containing the weights for each feature (to emphasize or reflect confidence). Each feature is initially scaled, as in Eq. (4.3), to fit the values of the Fiedler vector, $\mathbf{x}^{(2)}$ (i.e., $\min(\mathbf{f}^{(k)}) = \min(\mathbf{x}^{(2)})$ and $\max(\mathbf{f}^{(k)}) = \max(\mathbf{x}^{(2)})$). The weights c of the spectral coordinates takes into account the smoothness of an eigenvector (measured by its eigenvalue $\lambda^{(u)}$) and the confidence in the reordering (measured by the

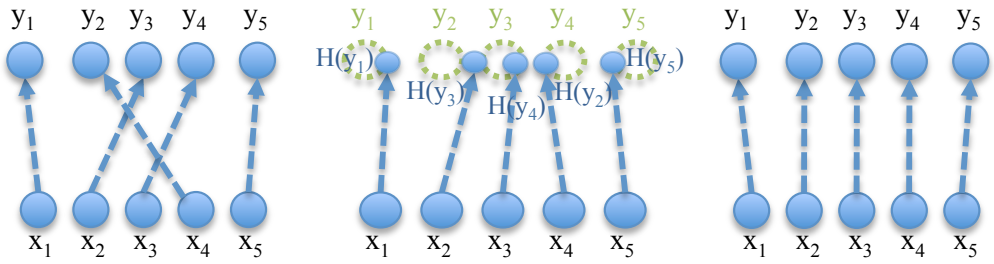

Figure 4.5: Diffusion of the correspondence map: a) The closest corresponding points on the spectral embedding might not necessarily be coherent spatially, consequently, b) their spatial coordinates are smoothed using a mean filter within the neighborhood structure, this moves the corresponding points to regularized positions ($H(y_{\phi(i)})$), and, c) the final correspondence map is reset by finding the closest nodes to these regularized positions.

permutation cost $Q^{(u)}$). Specifically, the weight, $c^{(u)}$, of the u^{th} spectral coordinate is:

$$c^{(u)} = \exp(-(Q^{(u)}\lambda^{(u)})^2/2\sigma^2), \tag{4.9}$$

where σ is a normalization factor set to

$$\sigma = \text{mean}\left\{Q^{(u)}\lambda^{(u)}\right\}_{u=1\cdots M}. \tag{4.10}$$

The alignment of these embeddings can be viewed as a nonrigid registration, $\mathbf{X} = \phi(\mathbf{Y})$. Fig. 4.4 shows the alignment challenge where the first three spectral components $(\mathbf{x}^{(2)}, \mathbf{x}^{(3)}, \mathbf{x}^{(4)})$ are used as 3D (x, y, z) coordinates for visualization purposes. The Robust Point Matching (Chui, 2003) with a Thin Plate Spline-based transformation is often used for 2D or 3D registration. However, with this approach, the final registration depends on the number and choice of the control points. We apply the recent Coherent Point Drift method (Myronenko and Song, 2009) which is scalable to N dimensions, fast, and demonstrates excellent performance in this application.

To increase speed in FOCUSR, we take advantage of the property of the Coherent Point Drift method that a continuous transformation derived from a subset of the points can be applied to all nodes of the dense embeddings (i.e., interpolation). In our case, we subsample $\mathbf{X}$ and $\mathbf{Y}$ by taking randomly a few points (in our experiments we chose 1% of the total number of vertices, roughly 1000 points).

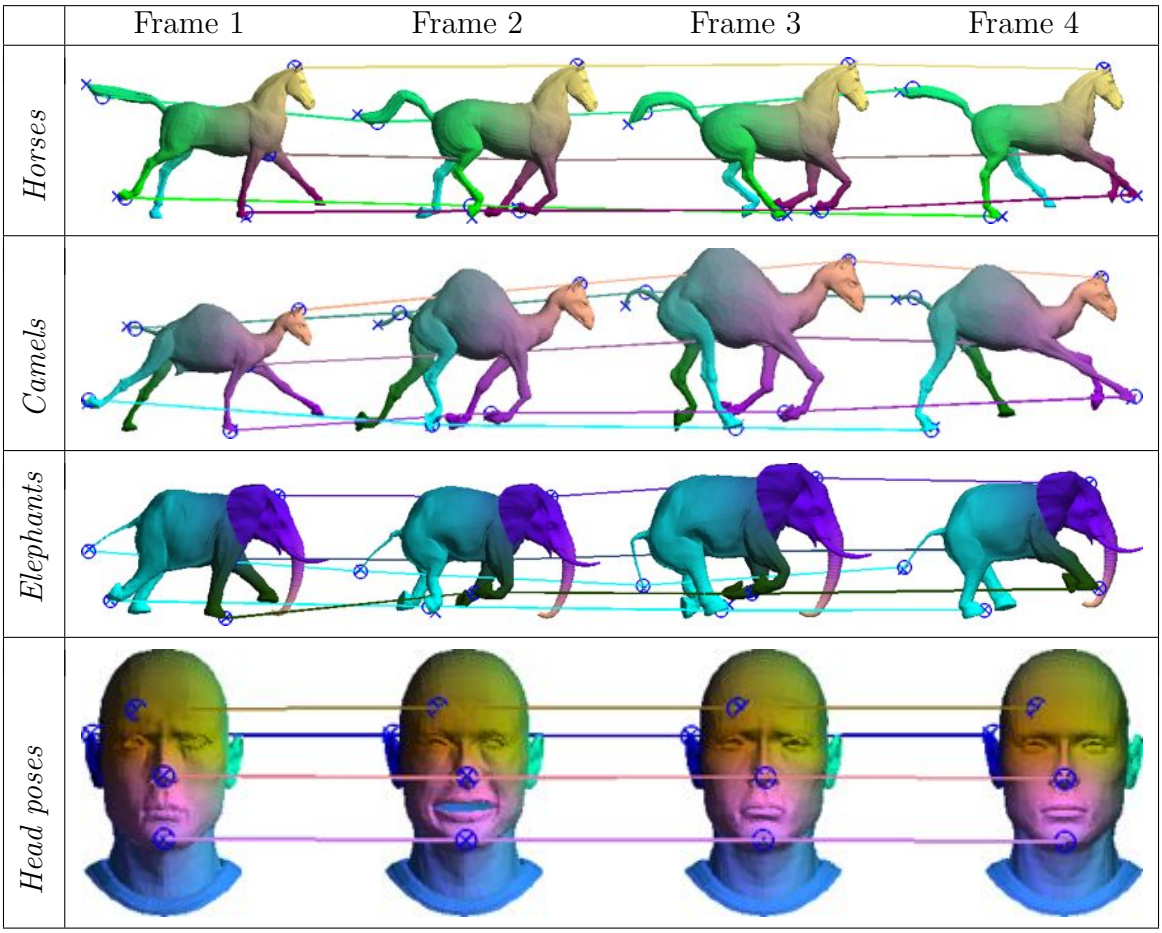

Figure 4.6: Correspondences across animated sequences for a horse gallop (average error of 1.41%(±0.57%)), a camel gallop (1.42%(±0.65%)), an elephant gallop (0.95%(±0.54%)), and facial expression changes (0.47%(±0.26%)). Models of all frames are matched with the first frame. Corresponding points have a unique color across each sequence. Five points are tracked along the sequence (colored lines) for visualization. Note that no temporal consistency was enforced (each frame was matched independently with the first frame). Blue circles show corresponding points along the sequences found with FOCUSR. Blue crosses show the true corresponding points.

4.2.5 Final Diffusion

After alignment, both embedded representations can be directly compared ($\mathbf{X} = \phi(\mathbf{Y})$), i.e. two points which are closest in the embedded representations could be treated as corresponding points in both meshes. However, the mapping is not guaranteed to be smooth, even after the CPD alignment. The spectral regularization promotes smoothness of the correspondence map, but it is possible to have irregularities in the smoothness when the features differ significantly between the two meshes. The resulting embeddings warped with the CPD, in the

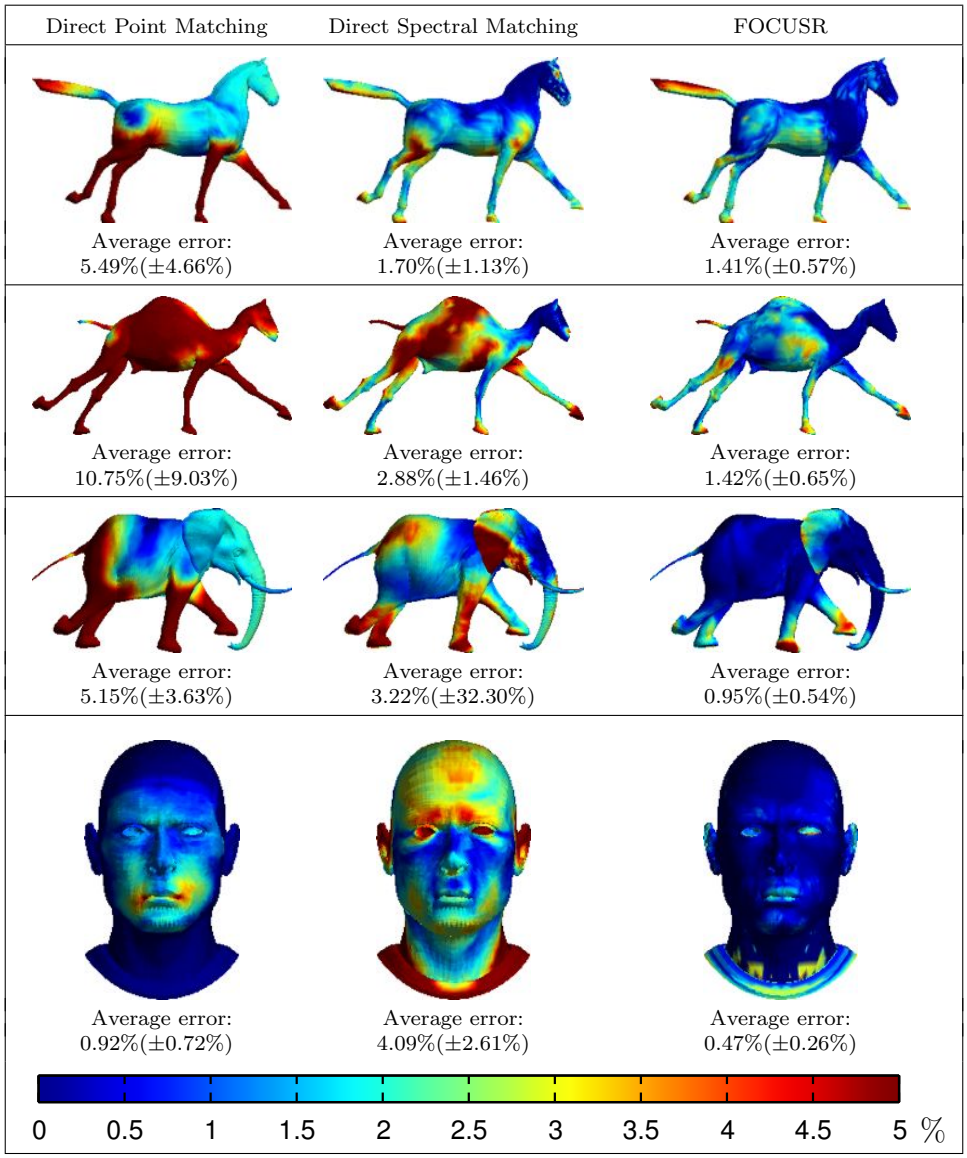

Figure 4.7: Average relative error distance from the ground truth in each animated sequence (in percentage between 0% and 5%, e.g., if a mesh is 100 mm, 5% means an average error of 5 mm, the lower percentage the better). Three settings are used: *a)* Direct matching of closest points on surfaces, *b)* direct matching of closest points on rigidly aligned spectral embeddings, and *c)* matching using nonrigidly aligned spectral embeddings in FOCUSR. While matching points in the spectral domain clearly improves correspondence, FOCUSR gets additional precision by aligning the spectral embeddings. Note that no additional features are used here in FOCUSR.

$K + M$ multidimensional space, can contain local spatial incoherence in the correspondence map (as illustrated in Fig. 4.5 a). Consequently, we include a postprocessing step to enforce additional smoothness of the correspondence map.

The correspondences obtained after CPD are used to map the second mesh vertices (target point y in Fig. 4.5 a) to the first mesh vertices (fixed points x in Fig. 4.5 a). The 3D coordinates of these mapped points on the second mesh are now treated as independent scalars and diffused on the surface of the first mesh (i.e., this moves the points of the second mesh to positions obeying the (smooth) neighborhood system of the first mesh as illustrated with points $H(y)$ in Fig. 4.5 b). We used the smoothing method in (Desbrun *et al.*, 1999) which is similar to the Laplacian smoothing, while other methods could also be used for this step. At this stage, the points x on the first mesh can be associated with either the smoothed coordinates $H(y)$ on the second mesh (i.e., vertices of the first mesh could be matched to coordinates in between the vertices of the second mesh), or with actual points on the second mesh. In our applications, we matched nodes to nodes, so the latter strategy is chosen. The correspondence map linking the first mesh to the second mesh is therefore updated by linking each point in the first mesh with the point in the second mesh which has the minimum Euclidean distance to the *diffused* geometric coordinates (shown with the new map in Fig. 4.5 c). In our experiments, 40 iterations were sufficient to diffuse the point coordinates. The fourth step in Fig. 4.3 shows a few corresponding points between two brain surfaces.

4.3 Results

To demonstrate the effectiveness of FOCUSR, we first match in a controlled experiment with a known ground truth a variety of generic meshes (3 sequences of 50 and 10 frames of models in various poses, totaling the computation of 160 matchings) and show that the use of a nonrigid alignment of spectral coordinates improves precision over a simple direct spectral matching method. We then apply FOCUSR to real data for a clinical application to brain surface matching in which matching precision is extremely important and where additional features are known to be meaningful to the accuracy of the alignment. We do so by processing and analyzing the correspondence of 264 pairs of brain surfaces using 15 different combinations of features (totaling the computation of 3,960 correspondence maps). This clinical application reveals the full power of FOCUSR where the introduction of additional

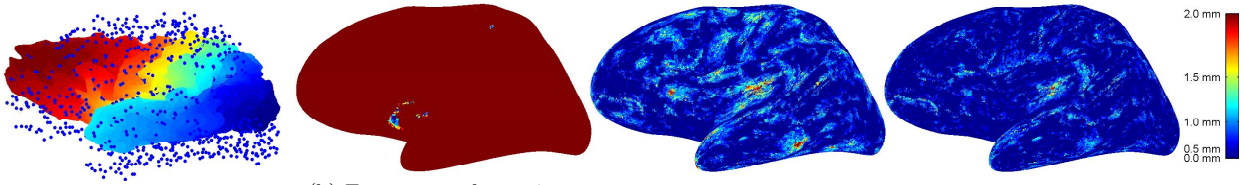

(a) deformed meshes

(b) Error map when using simple feature matching ($K = 2$) Average error: 53.11 mm.

(c) Error map when using spectral coordinates ($M = 5$) Average error: 0.35 mm.

(d) Error map for FOCUSR ($K = 2$, and $M = 5$) Average error: 0.06 mm.

Figure 4.8: Comparison with a synthetic ground truth using one brain hemisphere. *(a)*: The deformed mesh (colored with its Fielder vector) overlaid with the original vertex positions illustrated by the blue dots. *(b)*: When using simple feature matching ($K = 2, M = 0$), the mean distance error with the ground truth is as expected very high, 53.11 mm. *(c)*: When using FOCUSR with only spectral components and no additional feature ($K = 0, M = 5$), the mean distance error is 0.35 mm. *(d)*: When using FOCUSR with both spectral components and additional meaningful features ($K = 2, M = 5$), the mean distance error using this brain hemisphere is 0.06 mm. When iterating this experiment on all hemispheres, the mean distance error using FOCUSR is 0.07 mm (Note that surfaces are smoothed in *(b,c,d)* after correspondence to visualize the errors within the cortical foldings).

features improves the shape matching significantly.

4.3.1 Matching Meshes

We first begin our validation by showing that FOCUSR can find efficiently and precisely a dense correspondence between generic meshes. We use the data from (Sumner and Popović, 2004) (available publicly[3]) where animal models have been deformed in various poses. These meshes were created in (Sumner and Popović, 2004) by transferring the deformation of a sequence of source meshes to target reference meshes. We use in the dataset the sequence of a galloping animal for a horse (8,431 vertices, 50 frames), an elephant (21,887 vertices, 50 frames), and a camel (42,321 vertices, 50 frames), all illustrated in Fig. 4.6. We want to recover the deformations and assess the precision of the correspondences between all models in a sequence and the reference model. For each gallop animation, the same mesh is deformed, and all vertices across the sequence maintain a direct one-to-one correspondence with the reference mesh (i.e., node i of any mesh in the animation corresponds with node i (the same index value i) in the reference mesh). This gives a ground truth for the correspondence maps in all animations (i.e., $\phi(i) = i$) on which we can compare our method.

[3]Meshes available at `http://people.csail.mit.edu/sumner/ research/deftransfer`

We quantify precision by measuring the average distance between the locations of corresponding points found with FOCUSR and with the ground truth. That is, for all points $v_i \in \mathcal{V}_1$ in the first mesh matching the points $v_{\phi(i)} \in \mathcal{V}_2$ in the second mesh, the mean distance error is the average of the distances, $\frac{1}{N}\sum_i^N ||\mathbf{x}_i - \mathbf{x}_{\phi(i)}||$, between the real locations of the corresponding points, $\mathbf{x}_i$, and their recovered locations on the second mesh, $\mathbf{x}_{\phi(i)}$. For each gallop animation, we computed the correspondence maps of the meshes of all frames with the reference mesh. Fig. 4.7 shows the average relative distance error for all sequences when finding the closest points in space, in the spectral domain, and when using FOCUSR in its simplest setting (i.e, $K = 0$ in Eq. (4.8)). Mismatches due to nonrigid deformations (e.g., articulated limbs of the galloping animals) are the most severe when matching in the spatial domain, while these errors are attenuated when matching occurs in the spectral domain (about a 60% increase in precision). FOCUSR improves precision over the simple spectral matching by about 50%.

The relative average distance error in FOCUSR with its standard deviation (expressed in percentage of the size of a mesh) is for the whole horse gallop animation: 1.41%($\pm$0.57%) with an average computation time of 44 seconds, for the camel gallop: 1.42% ($\pm$ 0.65%) in 79 seconds, and for the elephant gallop: 0.95% ($\pm$ 0.54%) in 98 seconds (timing were performed on a 2.8 GHz Intel Pentium 4 using unoptimized Matlab code). We additionally ran the same experiment on an animation of changing facial expressions (15,941 vertices, 10 expressions) and found a relative average error of 0.47% ($\pm$ 0.26%) with on average 40 seconds of computation. All these errors remain relatively small with corresponding points found at more or less 1% of the size of the mesh from their true locations (e.g., for a mesh of 100 mm, an error of 1% is a mismatch of 1 mm). Additionally, five points of interest were tracked along each animation (between the ears, the tail tip, right rear and front paw, and on the sternum of the animals; and the right ear, left upper eyelid, nose tip, lower lip, and chin of the head).

By applying a nonrigid alignment of spectral coordinates, FOCUSR exhibits an improved level of precision (of about 1.4% error) even in the absence of using additional features. Higher errors often occurs in areas of highly nonrigid deformation, such as skin stretching (e.g., the side of the horse undergoing expansion and compression while galloping). One might also argue that displaced areas are not necessarily errors (e.g., the skin could move freely over a body by a few centimeters when galloping).

4.3.2 Brain Surface Matching

Brain surface matching is an important topic for neuroimaging studies that requires the alignment of anatomical landmarks or functional activation across a population. Specifically, cognitive function can only be studied across individuals when correspondence is made from one individual to another between activation areas of the brains. The brain morphology offers the particularity that every individual has a unique folding pattern in their cerebral cortical hemispheres while, at the same time, many large-scale similarities exist and allow correspondence between brain surfaces. Moreover, matching brain surfaces allows us to test the ability of FOCUSR to use extra features, such as the sulcal depth, the cortical Gaussian curvature, and the cortical thickness, that can potentially improve the precision of the correspondence beyond conventional spectral correspondence. We utilize the two features used by the FreeSurfer algorithm to drive alignment, which are the *sulcal depth* (Fischl *et al.*, 1999) at each point $\{s_1, s_2, \ldots, s_n\}$ (as calculated by FreeSurfer), and the surface curvature at each point, $\{\kappa_1, \kappa_2, \ldots, \kappa_n\}$. FreeSurfer outputs the mean curvature of a mesh, but in practice our method generated slightly better results when using the Gaussian curvature estimated with the method described in (Steiner and Morvan, 2003). We thus chose to test the Gaussian curvature in our feature combinations in order to avoid exploding the number of feature combinations in our experiments. In addition, FreeSurfer also supplies gray matter cortical thickness (calculated from anatomical MRI image data (Fischl and Dale, 2000)) at each point, $\{t_1, t_2, \ldots, t_n\}$, which we can additionally test as a feature to drive the alignment with FOCUSR.

To demonstrate the flexibility of FOCUSR to handle different features, different combinations of these three additional features were used in our experiment. Additionally, we independently examine the effects of using the features to define only edge weights (in Eq. (4.3)), only node weights (in Eq. (4.5)), or only as coordinates for matching (in Eq. (4.8)).

Synthetic Deformations

We begin with a synthetic experiment which is designed to demonstrate that FOCUSR profits from meaningful features to produce a precise alignment. In this experiment, we synthetically deform a brain surface such that two of the features are preserved and one feature is distorted. FOCUSR is shown to perform better when the meaningful (preserved) features are included and worse when the noise feature is included. For our experiment, we

Table 4.1: Settings used in FOCUSR for the recovery of a synthetic deformation. Checkmarks indicate whether *sulcal depth* (S), *cortical thickness* (T), or *cortical Gaussian curvature* (C) is used as features, on graph node weights, or on graph edge weights. The synthetic deformation process did not distort sulcal depth or cortical thickness, but did distort Gaussian curvature. This experiment demonstrates that FOCUSR can profit from incorporating the meaningful (undistorted) features. The reported error gives the average error distance between the matched point and the ground truth across all available hemispheres. The use of FOCUSR without any features (with only spectral components) and the simple feature matching method are provided for comparison.

Coord.			Nodes			Edges			Error
S	T	C	S	T	C	S	T	C	
✓		✓	✓		✓	✓		✓	0.14 mm
✓			✓			✓			0.13 mm
	✓			✓			✓		0.29 mm
		✓			✓			✓	0.50 mm
✓	✓		✓	✓		✓			**0.07 mm**
(with only spectral components)									0.38 mm
(simple feature matching)									53.02 mm

match one brain hemisphere with a deformed version of itself. The vertex indexing remains the same in the deformed version. Similarly as the last experiment, the true matching is thus known (i.e., $\phi(i) = i$). We severely deform one cortical surface model, where for each point (x, y, z), we apply the transformation $z' = (1 + \alpha)z$, i.e., a compression in the z-axis controlled by parameter α (we used $\alpha = 0.3$), and the transformation $x' = x + \beta r^2 / \max(r^2)$ with $r^2 = x^2 + y^2$, i.e., a radial distortion controlled by parameter β (we used $\beta = 15$). This simulates a deformation due to a drastic change in the head shape. The deformation however preserves the same mesh topology as it does not introduce any discontinuities or intersecting faces. Fig. 4.8 illustrates the position of the original hemisphere with the blue dots and the deformed hemisphere with the colored mesh. The sulcal depth and the cortical thickness are the same in both cortical meshes. The Gaussian curvature has been recomputed in the deformed mesh with the method described in (Steiner and Morvan, 2003). Therefore two of the features (sulcal depth and cortical thickness) are meaningful under this distortion and one feature (Gaussian curvature) is a distracting noise feature. The goal of this experiment is to verify if the use of additional meaningful features helps the matching precision and to measure its improvement.

If we use the simple feature-only correspondence, the error is on average across all hemi-

spheres 53.02 mm due to the fact that the correspondence map has virtually no mechanism to promote smoothness. When FOCUSR is used with only spectral components with no features (e.g., $K = 0$ and $M = 5$), we find for all hemispheres an average error distance of 0.38 mm as shown in the first error map of Fig. 4.8. Most errors appear to be located on the sulci extrema. By using FOCUSR to drive feature correspondence with spectral regularization, the error drops to 0.07 mm.

In FOCUSR, the surface features affect the correspondence by using the features as coordinates in the point matching, and/or, by using the features to set edge weights, and/or by using the features to set node weights. Now we demonstrate that the greatest precision for FOCUSR is obtained by using the features in these three ways instead of just one or two of these ways. Specifically, we iterate through all of the 512 possible combinations ($2^{3\times3}$). Table 4.1 summarizes a few combinations. We tested FOCUSR using both sulcal depth and Gaussian curvature as additional features. The average error distance across all hemispheres is in this case 0.14 mm. Adding the sulcal depth as the only additional feature yields an error of 0.13 mm; adding only the cortex Gaussian curvature yields an error of 0.50 mm; and the cortical thickness yields an error of 0.29 mm. The best combination of features for FOCUSR was obtained when using sulcal depth and cortical thickness as additional coordinates and on graph nodes, and using cortical thickness on graph edges, yielding an error of 0.07 mm. It is expected that FOCUSR should perform best with these features, since they were not changed by the synthetic deformation, but the Gaussian curvature was. The error map on a single hemisphere is shown on Fig. 4.8. The best-performing combination of features demonstrates an almost perfect matching for FOCUSR.

This experiment shows that by incorporating meaningful features FOCUSR can indeed improve the matching precision. The weighting functions used here also differs slightly from the one used in (Lombaert *et al.*, 2011a) which used the exponentials of the additional features. This experiment confirms that using stable features between two cortices (i.e., the same sulcal depth and cortical thickness) improves the cortex matching precision.

Performance Evaluation on Real Data

Cortical surface matching is a challenging problem due to the wide variability in gyral morphology and topology between individuals. There is no ground truth available for perfect brain surface matching across individuals. However, FreeSurfer (Fischl *et al.*, 1999) has been

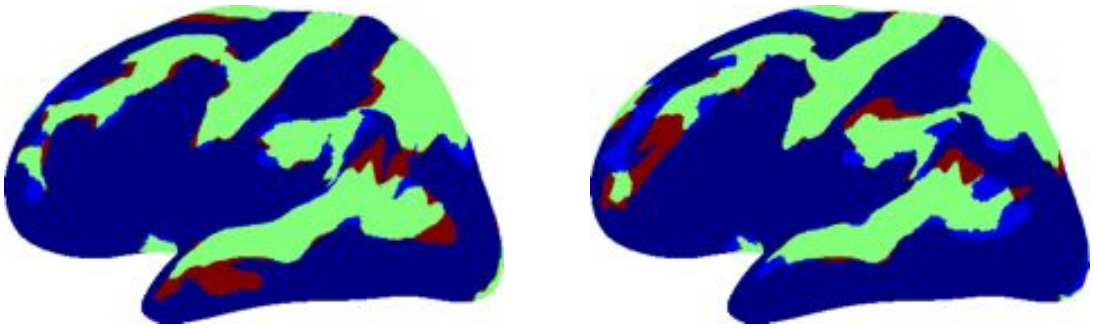

Parcel overlaps using FOCUSR
Average overlap: 71% (in 208 sec)

Parcel overlaps using FreeSurfer
Average overlap: 71% (in several hours)

Figure 4.9: In *green*, good overlap of projected sulcal regions, in *red*, wrong projection outside the sulcal regions, and in *light blue*, missed sulcal regions. *(First brain)* Correspondences computed in 208 seconds on average using FOCUSR, while *(second brain)* FreeSurfer required several hours (Note that in this visualization surfaces are smoothed to visualize the correspondence within the cortical foldings).

demonstrated to provide highly precise cortical matchings that closely align cortical areas across subjects (Hinds *et al.*, 2008) and therefore provides a reliable surrogate for our comparison. The delineations of 81 sulcal regions are available for 24 hemispheres (12 subjects). These sulcal regions were obtained using an automatic parcellation of the cortex (Fischl *et al.*, 2004) and are considered as our gold standard. Although parcellations of the cortex into named sulci and gyri are not expected to align between subjects in all cases (except for the primary folds), they do provide means to compare the two methods. We use correspondence maps generated by FreeSurfer and FOCUSR to project the parcellation areas onto different brain hemispheres and measure their overlap (illustrated on Fig. 4.9). To process a mesh of 135,000 vertices, FreeSurfer has a varying processing time which is currently on the order of several hours, while the time required by FOCUSR is on the order of 3–4 minutes. To process all our 264 possible pairs of left and right brain hemispheres, FOCUSR required on average 208 seconds (on a 2.8 GHz Intel Pentium 4 using unoptimized Matlab code). With reduced meshes of 20,000 vertices, FOCUSR performed in 19 seconds. The primary computational burden of the algorithm is the final diffusion of the correspondence map. This final step requires the smoothing of the matched mesh, which currently takes 84 seconds on average in Matlab. The total time to perform all our 264 correspondences using FOCUSR was 14 hours on a single computer, a substantial advantage compared to the several weeks required by FreeSurfer. Each overlap ratio is defined by the ratio of set intersection to set

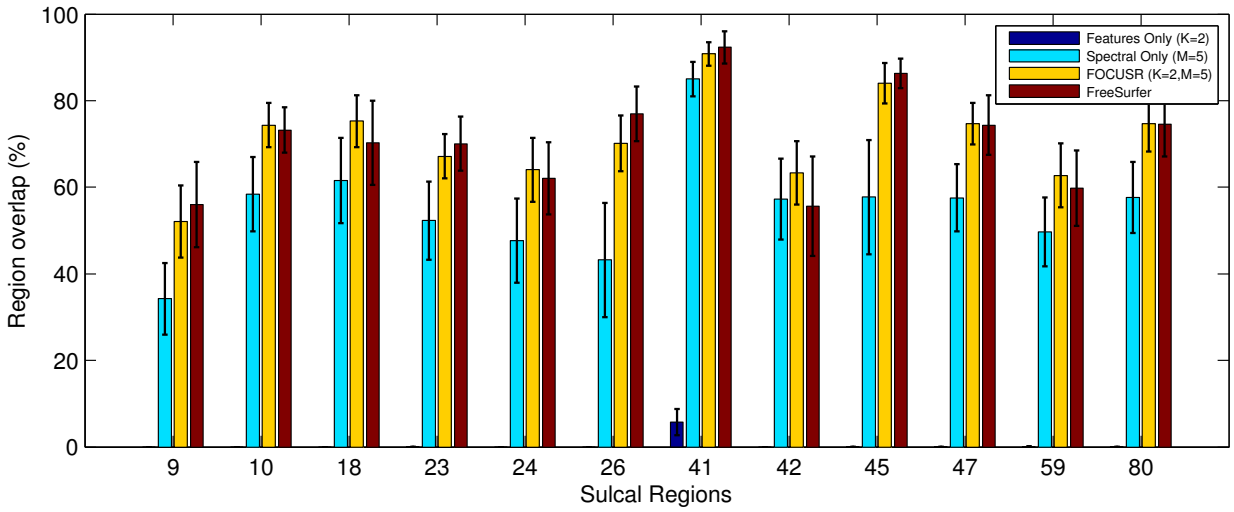

Figure 4.10: Average overlap ratios of the twelve largest sulcal regions on the right hemisphere over 264 matchings. *(Dark blue)* FOCUSR with features only, sulcal depth and cortical Gaussian curvature (i.e., $K = 2$, $M = 0$): $0.48\% \pm 0.28\%$ overlap. *(Cyan)* FOCUSR with spectral components only (i.e., $K = 0$, $M = 5$): $55.18\% \pm 9.09\%$ overlap. *(Yellow)* FOCUSR with features (sulcal depth and cortical Gaussian curvature) and spectral components (i.e., $K = 2$, $M = 5$): $71.11\% \pm 5.98\%$ overlap. *(Red)* FreeSurfer's overlap ratios (requiring weeks of computations): $70.95\% \pm 7.27\%$ overlap. FOCUSR only required 14 hours to perform all 264 matchings and is strongly correlated with FreeSurfer (correlation coefficient of $\rho = 0.897$). The error bars show the standard deviation of each overlap ratio.

union. Fig. 4.10 shows the overlap ratios for the twelve largest sulcal parcellations[1] using FOCUSR and FreeSurfer. The results of FOCUSR are correlated to FreeSurfer's overlaps with a correlation coefficient of $\rho = 0.897$.

From Fig. 4.10, we can see that FOCUSR closely matches the performance of FreeSurfer when using a similar feature set (sulcal depth and cortical curvature) to drive the correspondence (71.16% overlap for FOCUSR versus 70.95% overlap for FreeSurfer). In contrast, the pure feature matching or the use of FOCUSR with only spectral components produces results with a much lower precision (effectively null at 0.48% overlap). We now demonstrate that using features purely for edge or node weights (or purely as feature coordinates) also produces suboptimal results.

[1]Sulcal regions: 9 (*G frontal middle*), 10 (*G frontal middle*), 18 (*G occipit temp med Lingual part*), 23 (*G parietal inferior Supramarginal part*), 24 (*G parietal superior*), 26 (*G precentral*), 41 (*Medial wall*), 42 (*Pole occipital*), 45 (*S central*), 47 (*S cingulate Main part and Intracingulate*), 59 (*S intraparietal and Parietal transverse*), 80 (*S temporal superior*).

Table 4.2: Different combinations of features with spectral correspondence (three modes: as matching coordinates, to define edge weights or to define node weights) demonstrate that using all three combination modes (as we advocate in FOCUSR) gives the best performance. Additionally, different combinations of three features, *sulcal depth* (S), *cortical thickness* (T), or *Gaussian curvature* (C), demonstrate that using sulcal depth and cortical curvature in FOCUSR provides performance most similar to FreeSurfer (as expected since it relies on similar features). In comparison, FreeSurfer's overlap ratios are 72.03% in the left hemispheres, and 70.95% in the right hemispheres. The experiment was performed over all 264 pairs (from 12 brains), averaged across the twelve largest parcels

Coord.	Nodes	Edges	Overlap	Overlap
S T C	S T C	S T C	Left	Right
			55.11%	55.18%
✓ ✓ ✓			70.79%	70.57%
✓ ✓ ✓	✓ ✓ ✓		70.77%	70.60%
✓ ✓ ✓		✓ ✓ ✓	70.77%	70.42%
✓ ✓ ✓	✓ ✓	✓ ✓	70.73%	70.66%
✓ ✓ ✓	✓ ✓	✓ ✓	69.65%	70.12%
✓ ✓ ✓	✓	✓	70.67%	70.65%
✓ ✓ ✓	✓	✓	65.51%	66.84%
✓ ✓ ✓	✓	✓	70.67%	70.58%
✓ ✓ ✓	✓ ✓ ✓	✓ ✓ ✓	70.74%	70.41%
✓	✓	✓	71.10%	71.16%
✓	✓	✓	55.25%	56.77%
✓	✓	✓	55.28%	56.67%
✓ ✓	✓ ✓	✓ ✓	71.18%	71.11%
✓ ✓	✓ ✓	✓ ✓	69.64%	70.15%

Testing with Multiple Configurations

We first analyze the matching performance using different configurations of the same features used by FreeSurfer, namely sulcal depth and cortical curvature. In a second step, to demonstrate the flexibility of FOCUSR, we introduce a different feature not used by FreeSurfer and tested several combinations of features to see whether any of these combinations performs better than FreeSurfer. Additional features were incorporated in FOCUSR using Eq. (4.8), Eq. (4.3), and Eq. (4.5), with $\gamma = 1.2$ and $\beta = 0.2$ (the description of the behavior of these parameters are described in (Lombaert *et al.*, 2011a)). Overall, fifteen different combinations of additional features were used. For each combination, we ran FOCUSR on the 132 pairs ($n(n-1)$ with $n = 12$ brains) of left brain hemispheres and on the 132 pairs of right brain hemisphere (totaling 3,960 matchings, 264×15). The results are summarized in Table 4.2.

In comparison, FreeSurfer performs with an average overlap ratio for the largest parcels of 72.03%(±8.52%) in the left hemispheres (the variation is the standard deviation of all overlap ratios), and 70.95%(±7.27%) in the right hemispheres. Fig. 4.10 shows three relevant combinations.

- The *first* combination shown on Fig. 4.10 demonstrates the poor performance of the direct feature matching method where FOCUSR uses no spectral coordinates as described in Section 4.2.1 (i.e., $M = 0$, matching is a simple feature comparison using sulcal depth and cortical Gaussian curvature ($K = 2$) as used in Eq. (4.1)). The average overlap ratio on the largest parcels is effectively null at 0.38%(±0.19%) in the left hemispheres (0.48%(±0.28%) in the right hemispheres).

- The *second* combination shows FOCUSR using no features and only spectral components ($K = 0$ and $M = 5$). The average overlap ratio on the largest parcels is only 55.11%(±10.73%) in the left hemispheres (55.18%(±9.09%) in the right hemispheres).

- The *third* combination shows the full power of FOCUSR where it uses spectral components alongside sulcal depth and cortical curvature features, which are the same features driving the correspondence in FreeSurfer. The overlap ratio is as high as 71.18%(±7.63%) in the left hemispheres (71.11%(±5.98%) in the right hemispheres). This is almost a perfect match with the overlap ratios in FreeSurfer (72.03% in the left side, and 70.95% in the right side).

FOCUSR is, in the left and right cortices, equivalent with FreeSurfer's overlap ratios (71.18% vs. 72.03% in the left side, and 71.16% vs. 70.95% in the right side). It is important to note that there is no perfect combination of features to drive the correspondence. Our experiment shows that certain combinations perform better on particular parcels than on others. The best combination of extra features thus depends on which sulcal region of the brain should be matched. This finding concurs with a similar conclusion in (Yeo *et al.*, 2010b).

Dependence on the Number of Spectral Coordinates

In the previous section we demonstrated that it is optimal to use features to derive edge weights, node weights and as explicit feature coordinates. We now examine the dependence

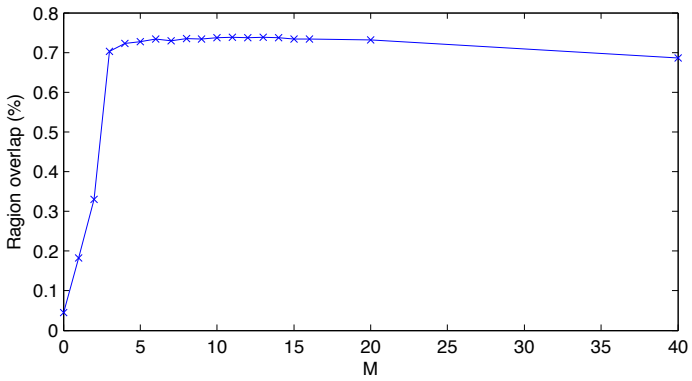

Figure 4.11: Robustness of FOCUSR to the number of eigenvectors used as spectral coordinates. When no eigenvector is used ($M = 0$, i.e., direct feature matching), the correspondence is weak, whereas the performance becomes stable after just three eigenvectors are included. The performance is measured with the overlap ratios of sulcal regions as used previously.

of the performance on the number of eigenvectors used as spectral coordinates by running the previous experiment with a varying number of eigenvectors. When no spectral regularization is used (i.e., direct feature matching with $M = 0$ eigenvectors), the algorithm relies solely on feature coordinates. As expected, the performance is weak. The plot in Fig. 4.11 shows, for the twelve largest parcels, indeed a low overlap ratio of 0.38% in the left hemispheres (0.48% in the right side) when using pure feature matching, sulcal depth and cortical Gaussian curvature, with no spectral coordinates (i.e., $M = 0$). The performance improves quickly when eigenvectors are used (i.e., $M > 0$) to spatially regularize the feature matching. These spectral coordinates provide additional means of discrimination during the optimization of the correspondence map. FOCUSR gains no further significant improvement in quality after $M > 3$. Essentially, this result demonstrates that the primary purpose of the spectral coordinates is to provide a spatial regularization, which is achieved by using only the lowest-frequency eigenvectors.

4.4 Discussion and Conclusions

This paper presents a novel method, based on spectral correspondence, for the challenging task of precise surface matching. Current methods, most of which are iterative and control local deformations of surfaces, are dependent on the extrinsic mesh geometry. They find their limitations when matching articulated or highly deformable shapes. With such challenges, additional information (e.g., texture, anatomical information, or landmark positions) can help in finding a better correspondence. For example, this is the strategy that FreeSurfer

(Fischl *et al.*, 1999) relies on to match brain surfaces. Spatial regularization becomes crucial in these methods. It is clear from our experiments that direct feature matching, with no spatial regularization (i.e., $M = 0$) exhibits very poor performance. We decided to improve this strategy by using a spectral regularization of the feature matching method and to improve spectral matching by using a nonrigid alignment. The space of regularization (i.e., the *spectral* domain) is dependent on inherent properties of the mesh geometry. This modified strategy would free our method from the limitations of matching articulated or highly deformable shapes. Spectral methods provide a natural means of regularizing solutions at speeds of several orders of magnitude faster than current methods and are independent of the mesh extrinsic geometry. Our method can implicitly incorporate additional features to drive precise correspondence and it exploits the smoothness of the lowest-frequency harmonics of a graph Laplacian to regularize the correspondence map. Present day spectral correspondence methods are not fully realized and provide matchings that are not yet reaching a clinical level of precision. Currently, only intrinsic geometry can be embedded on graph edges and no additional information can be used. We provide a full realization of spectral correspondence where virtually any feature can be used as additional information as weights in graph edges, but also on graph nodes and as extra embedded coordinates with little or no computational expense. Furthermore, rigid transformations, or older point matching methods based on Thin Plate Splines are used (Jain and Zhang, 2006) and are difficult to extend beyond 3D (Sprengel *et al.*, 1996) to multidimensional feature space.

In its simplest form, FOCUSR is an improved spectral correspondence method that utilizes nonrigid point registration. We showed in our first experiment that a nonrigid alignment of the spectral coordinates improves significantly (by about 50%) the matching precision over a direct spectral matching. Its has been demonstrated with a variety of generic models (various galloping animals and heads of varying facial expressions) that the error from a known ground truth is minimal (with 1.4% relative distance error for our matched models). The full power of FOCUSR is presented in our second set of experiments, a real-world scenario with the challenging task of brain surface matching across several individuals. We use FOCUSR with different combinations of additional features, such as sulcal depth, cortical Gaussian curvature, and cortical thickness, to improve the matching precision. The fast speed of our method allowed us to compute and analyze 3,960 correspondence maps (which is prohibitively expensive for FreeSurfer). When no regularization is used (e.g.,

$K = 2$ features and $M = 0$ spectral components), the correspondence generates a poor overlap ratio of 0.48% in the largest sulcal regions. When FOCUSR is used in its simplest form with no additional feature (e.g., $K = 0$ features and $M = 5$ spectral components), the overlap ratio is 55%. The performance of FOCUSR is improved by using additional information (e.g., $K = 2$ features and $M = 5$ spectral components), and the overlap ratio increases to a level above 71% (versus 55% when using only spectral components). Our method is effectively equivalent to FreeSurfer's level of precision (which is also around 71%) when aligning sulcal regions. However, the vast increase in speed (with a total processing time of 208 seconds on average for meshes of 135,000 vertices) and the added flexibility when using additional information gives new perspectives to previously computationally prohibitive experiments. New features (e.g., anatomical or functional features extracted from various data sources) can be quickly tested and evaluated to see if they improve cortex matching. Quick parameter sweeps can be performed to isolate the best parameter value sets. These computationally intensive experiments can help us to understand what features are consistently correlated with brain areas across individuals and what their role are during the development of the cortical folding pattern. Furthermore, the computational time could be significantly improved with a reimplementation in C++ and with parallel programming for critical sections such as the eigendecomposition (e.g., LAPACK implementations on CUDA-enabled GPUs). Approximation methods for matrix eigendecomposition such as the Nyström approximation (Fowlkes *et al.*, 2004), the Gaussian projection (Halko *et al.*, 2011), or the differentiable QR decomposition (Bach and Jordan, 2004) could be used for additional speed up in processing time.

Spectral regularization promotes the smoothness of the correspondence map, but does not guarantee it. Better relaxation schemes, such as the Relaxation Labeling used in (Zheng and Doermann, 2006), might improve the matching precision. It is also important to consider which weighting function to use and to see how generalizable the parameter values are with a larger sample set. The use of different surface metrics (Liu *et al.*, 2009) can be a promising area to investigate. The algorithm, as with other spectral methods, is also not symmetric (i.e., $\phi_{i \to j} \neq \phi_{j \to i}^{-1}$). The CPD alignment does not guarantee symmetry of the resulting transformation (i.e., the computed correspondence map matching nodes from mesh X to mesh Y might not be the same as the inverse correspondence map matching nodes from mesh Y to mesh X). Further improvement of the method will be toward achieving a better

regularization of the correspondence map and to guarantee symmetry of the correspondence map. Nevertheless, FOCUSR already presents several clear advantages over present day methods for mesh correspondence and, in particular, conventional spectral matching. It provides a fast and precise solution for general mesh correspondence that can handle articulated or highly deformable surfaces, and creates a method that can implicitly use any set of additional features to drive improved precision.

Acknowledgments

The authors would like to specially thank Bruce Fischl and Martin Reuter for their helpful comments and Gareth Funka-Lea for supporting this project, as well as the financial support from the Alexander Graham Bell Canada Graduate Scholarships of the Natural Sciences and Engineering Research Council of Canada (NSERC).

CHAPTER 5 SPECTRAL MATCHING METHODS

Presentation

The first part of this chapter (Sec. 5.2 and 5.3) is based on a recent work on the exploration of the shape of the primary cortical folding. It is based on a close collaboration between **Siemens Corporate Research**, *Princeton, NJ, USA*; **Martinos Center at Harvard Medical School**, *Cambridge, MA, USA*; and **École Polytechnique**, *Montreal, Canada*. The second part of this chapter (Sec. 5.4) is partly based on two articles recently submitted to top conferences (with a double-blind review process) in computer vision and in medical image processing. Due to their dual-submission policy, journal versions could not be produced until their publication later in the year. The objective of the submitted papers is to capture large shape variability in a novel registration method and in a novel framework for atlas construction. These submitted articles are a joint work between **INRIA, Sophia Antipolis**, *France*; **Siemens Corporate Research**, *Princeton, NJ, USA*; and **École Polytechnique**, *Montreal, Canada*.

5.1 Introduction

This chapter proposes new original investigations that are built on top of the findings from the previous two chapters. As a brief reminder, Chapter 3 presented a new framework for atlas construction based on the conventional Demons algorithm. It allowed the creation of the first human atlas of the cardiac fiber architecture and studied the variability of cardiac fibers in humans. In addition, this human atlas also served as a basis for an extra study on the cardiac laminar sheets, presented in Appendix A, as well as for a preliminary comparative study between healthy and abnormal hearts, provided in Appendix B. The proposed framework showed its limitations when the images had high shape variability. To this effect, a different strategy has been explored in order to address the limitations associated with methods using a local optimization approach. For instance, Chapter 4 presented a new approach based on spectral matching, named FOCUSR, for finding correspondences between objects and organs of high shape variability. This method prevents the optimization from being trapped in trivial local minima (a major limitation of conventional approaches based on gradient-descent) by utilizing spectral representations in a new global feature matching technique.

From these findings, new studies and new methods are now presented in order to investigate and demonstrate the relevance of using spectral matching methods for atlas construction.

In the first sections of this chapter, FOCUSR is used in practical applications. *Firstly*, the correspondences computed with FOCUSR are used to construct an atlas of human brain surfaces, and its principal modes of shape variations are briefly studied in order to characterize the shape of the human brain surface. *Secondly*, the shape of the primary cortical folding is explored by analyzing the alignment of smoothed brain surfaces. This preliminary exploration is relevant to studies on the human brain development as the human cerebral cortex undergoes an extreme convolution from being totally smooth in a fœtus to highly convoluted in the teenage years. These two additional applications (the study on the variations of the brain surface and the exploration of the primary folding) show that an improved spectral approach can be used for relevant clinical studies. On the other hand, this second developed method, FOCUSR, was designed for surfaces and is not directly applicable to images. Therefore, *finally* in the last sections of this chapter, a new approach for image registration, named the *Spectral Demons*, is presented in order to capture very large shape variations between images. The underlying Demons framework is additionally extended to perform groupwise registrations. This yields a new framework for atlas construction that, when combined with the *Spectral Demons*, is capable of using images with high shape variability. This new approach based on spectral matching addresses, therefore, the limitation of the first developed framework. In the following sections, each new topic is presented by, firstly, describing the underlying method, and secondly, presenting the results.

5.2 Atlas of Brain Surfaces

FOCUSR, developed in Chapter 4, has been shown as a fast method for brain surface matching. We now show that it can be used to compute an average shape of the brain surface (construct its atlas) and study its variability. Moreover, in order to illustrate a practical use of FOCUSR, the principal modes of shape variations are analyzed with a dataset of 16 brain surfaces. This study provides in fact the necessary tools for using FOCUSR in concrete applications for shape analysis.

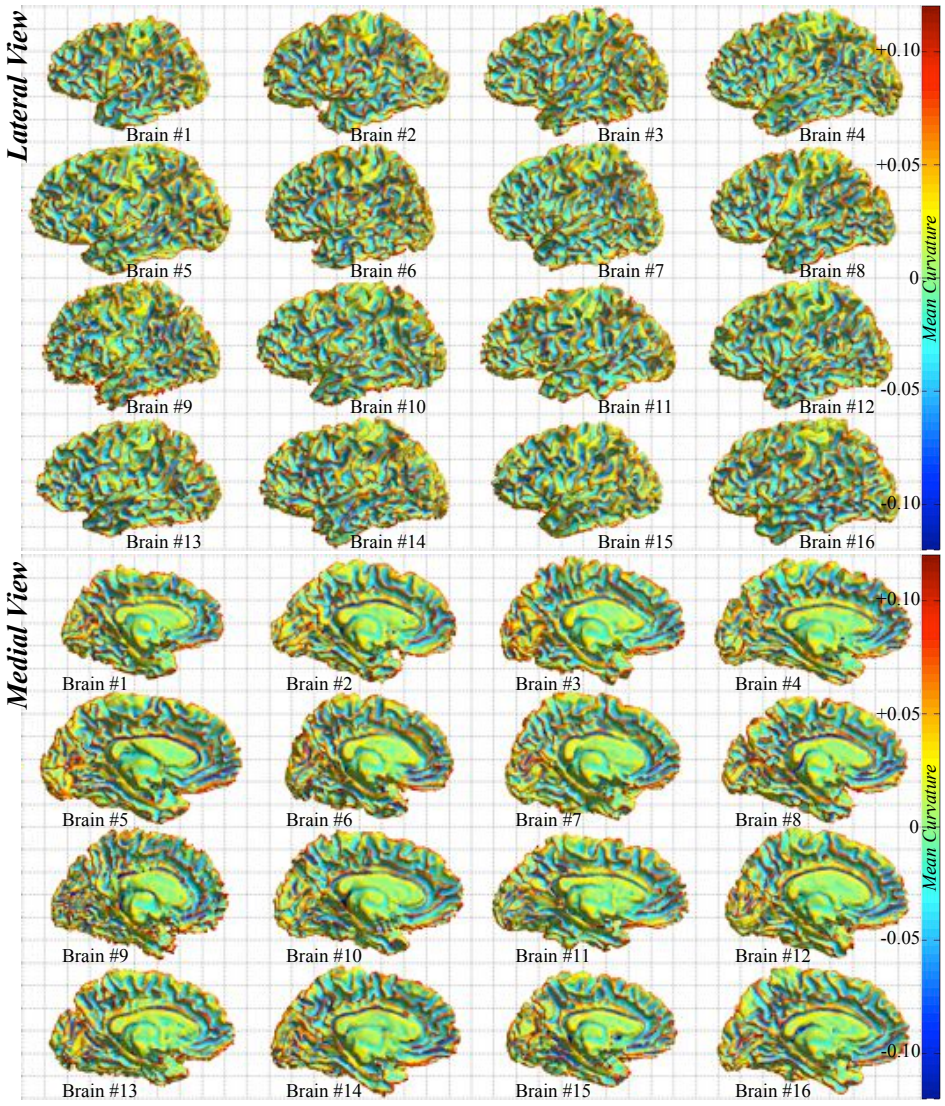

Figure 5.1: Dataset of 16 brain surfaces showing a large variability of shapes. *(top)* Lateral view, *(bottom)* Medial view. Each square in the background grid is 2 centimeters wide. The color represents the mean curvature of the surface.

5.2.1 Average and Principal Modes of Variations

Let us consider a set of N brain surfaces $\{\mathbf{s}_i\}_{i=1..N}$, where $\mathbf{s}_i$ is a column vector representing the i^{th} surface model with the coordinates of all its M vertices, $\mathbf{s}_i = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_M)^T$ (with coordinates $\mathbf{x} = (x, y, z)$, $\mathbf{s}_i$ has, therefore, the size $1 \times 3M$). In order to construct an atlas, all brain surfaces are matched with a reference brain $\mathbf{s}_{\text{ref}}$ where the correspondence map ϕ_i between one brain $\mathbf{s}_i$ and the reference $\mathbf{s}_{\text{ref}}$ is computed with FOCUSR. All points of the reference brain have, therefore, a corresponding point in each of the brain surfaces,

$\mathbf{s}_{\text{ref}} \mapsto \phi_i(\mathbf{s}_i)$. Let us now construct the $3M \times N$ matrix:

$$S = (\phi_1(\mathbf{s}_1), \phi_2(\mathbf{s}_2), \ldots, \phi_N(\mathbf{s}_N)) \,, \tag{5.1}$$

where each column represents a brain surface and where each row has the coordinates of all corresponding points across the dataset.

The average brain surface, $\bar{\mathbf{s}}$, is computed by averaging the coordinates of all corresponding points, i.e., by taking the mean of each row in S:

$$\bar{\mathbf{s}} = \frac{1}{N} \sum_{i=1}^{N} \phi_i(\mathbf{s}_i) \tag{5.2}$$

or in matrix notation: $\bar{\mathbf{s}} = \frac{1}{N}\, S\, \mathbf{1}$, where $\mathbf{1} = (1, 1, \ldots, 1)^T$ is a column vector of N ones. Note that FOCUSR directly gives a correspondence map between two surfaces without using an underlying deformation model. Statistics on the shape variations is, therefore, performed directly on the surface point positions. This is in contrast with statistical methods (Kendall, 1984; Durrleman *et al.*, 2009b; Durrleman, 2010) that analyze the deformation with complex Riemannian structures representing displacements or velocity fields.

It must be remembered that the positions of all corresponding points do not vary independently in the dataset. Indeed, the position of a point may be partially correlated with the positions of other points, for instance, if one point varies along a certain direction across the dataset, there is a high change that its neighbors will vary similarly, or that, in fact, any point located elsewhere may also have a certain correlation with its spatial variation. The study of these shape variations consists of finding the principal modes of variation, which is made with the Principal Component Analysis (PCA) of a global covariance matrix of point coordinates. The $3M \times 3M$ covariance matrix is build with:

$$C = (S - \bar{\mathbf{s}})(S - \bar{\mathbf{s}})^T. \tag{5.3}$$

The eigendecomposition of this covariance matrix $\Sigma = U \Sigma U^T$ gives the main variations $\mathbf{u}_i$ (an eigenvector is a column of U) in a $3M$ dimensional space, often referred as the shape space (Cootes *et al.*, 1995). All shapes can be represented in this shape space with a projection of $\mathbf{s}_i$ on the orthogonal space U.

However, the eigendecomposition of the $3M \times 3M$ matrix C may be prohibitively expensive when the number of points is large. The rank of the matrix C is also limited by the number of samples in the dataset, which means that at most $N-1$ non-trivial eigenvectors may be computed (there are only N linearly independent samples in the dataset). Fortunately, these first N eigenvectors can be computed by decomposing a smaller matrix:

$$L = (S - \bar{\mathbf{s}})^T (S - \bar{\mathbf{s}}) \tag{5.4}$$

The size of L is much smaller, $N \times N$, which makes its eigendecomposition much tractable, i.e., $L = V \Lambda V^T$ where each column $\mathbf{v}_i$ of V is an eigenvector of size N. It is shown in (Turk and Pentland, 1991) that $(S - \bar{\mathbf{s}})\ \mathbf{v}_i$ is also an eigenvector of C. It is, therefore, much more efficient to decompose L and use $\mathbf{u}_i = (S - \bar{\mathbf{s}})\ \mathbf{v}_i$ than trying to directly decompose C (which is often too large, hence too costly). Each eigenvector $\mathbf{u}_i$ provides the i^{th} mode of shape variation in the dataset.

5.2.2 Results and Discussion

The atlas is constructed with the dataset used in Chapter 4. It consists of 16 brain surfaces with a high shape variability and is presented on Fig. 5.1). For instance, brain #6 has a more spherical appearance than other samples, while brain #5 has a more elongated appearance than other samples.

The reference shape is arbitrarily chosen to be the first brain of the dataset. All other brains are subsequently matched against this reference by using FOCUSR and the computed correspondences across the dataset are gathered in the matrix S. Once matching is performed, the average brain surface can be computed with Eq. (5.2), while the principal modes of shape variations in the dataset are extracted by decomposing L as detailed earlier. The results are presented on Fig. 5.2 where the first three principal modes are illustrated by exaggerating each variation. More precisely, a series of synthetic brain surfaces are generated with $\bar{\mathbf{s}} + \alpha\ \mathbf{u}_i$ where α varies from -40 to +40 millimeters. From these series of synthetic surfaces, we can observe that the primary mode of variation captures a global scaling along the long axis of the brain, that the secondary mode of variation captures a global scaling on the sagittal plane, and that the tertiary mode of variation captures a more complex deformation on the brain surface. The less important modes of variations capture finer deformations on

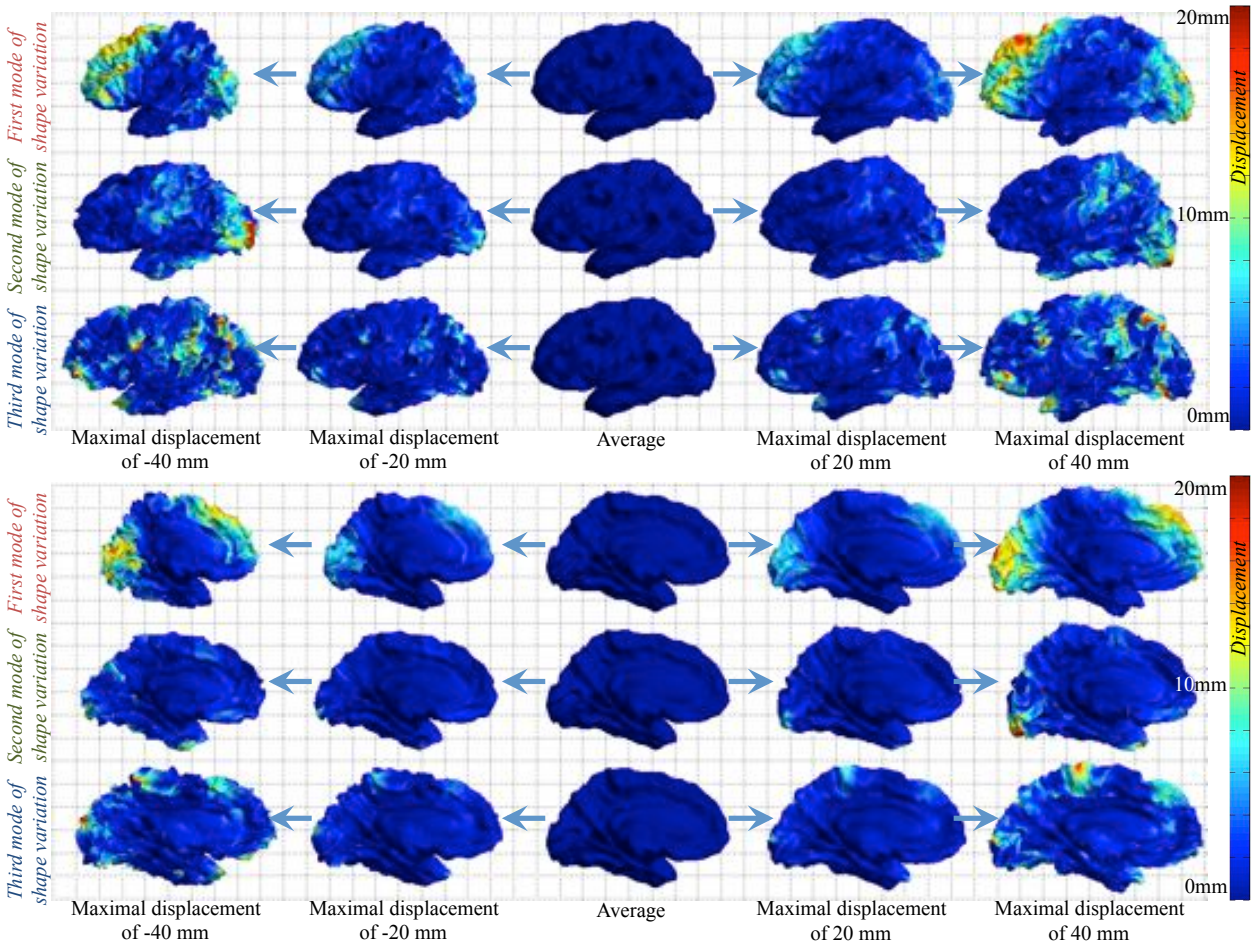

Figure 5.2: First, second, and third principal modes of variations in the dataset of brain surfaces. *(top)* Lateral view, *(bottom)* Medial view. (background squares are 2 centimeters wide). The color represents the gradual displacements of points in millimeters. The principal mode of variation appears to capture scaling along the long axis of the brain, the second mode of variation appears to capture scaling on the sagittal plane of the brain, while the tertiary mode of variation shows a more complex variation of the brain surface.

the cortical folding pattern. For instance, Fig. 5.3 shows finer details in the last non-trivial modes (the 14^{th} and 15^{th} modes).

Additionally, the projection of each brain surface onto the first and second mode of variations gives a representation of the dataset in a 2D shape space, i.e., the shape coordinates (2 scalars) of each brain surface are computed with the dot products $(\mathbf{s}_i \cdot \mathbf{u}_1)/(\|\mathbf{s}_i\|\|\mathbf{u}_1\|)$, and $(\mathbf{s}_i \cdot \mathbf{u}_2)/(\|\mathbf{s}_i\|\|\mathbf{u}_2\|)$. This 2D representation is presented in Fig. 5.4. For instance, it can be seen that brain #6 is indeed more spherical than other samples (it is at the same time the less elongated and the less flattened sample of the dataset).

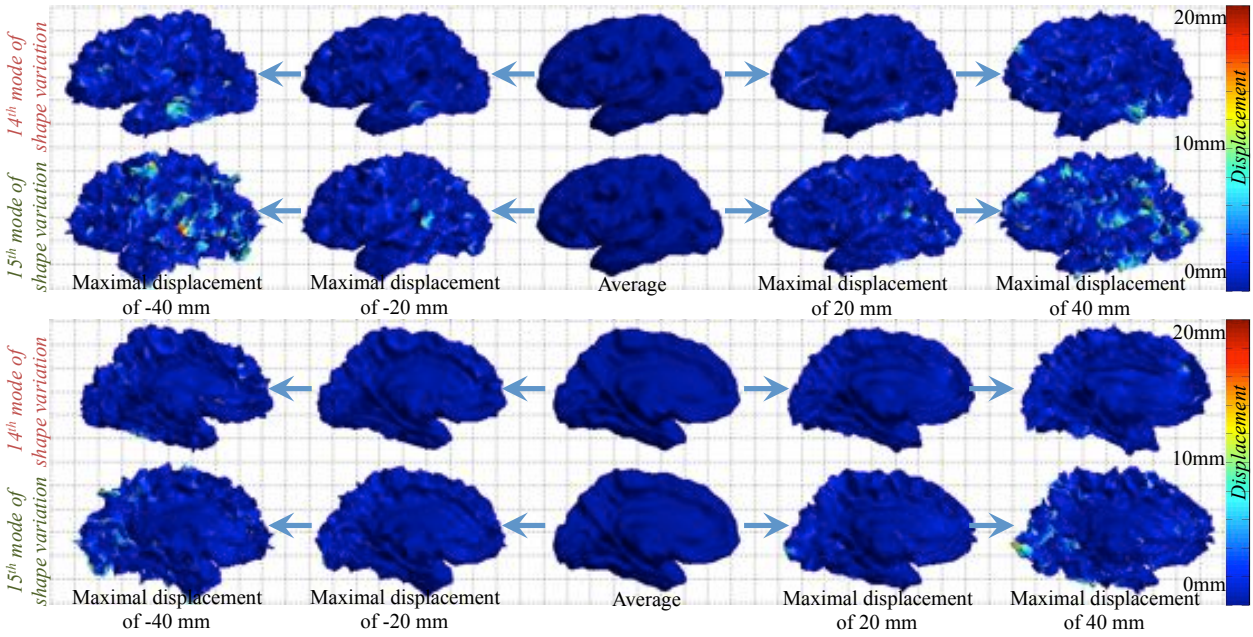

Figure 5.3: The least principal modes of variations in the brain dataset. *(top)* Lateral view, *(bottom)* Medial view The 14$^{\text{th}}$ and 15$^{\text{th}}$ modes shows complex deformations of the cortical folding, however, this analysis is not able to find the shape of the primary and secondary cortical folding.

However, the analysis of the principal modes of variations is not capable of studying the shape of the primary cortical folding pattern. In fact, none of the 16 modes revealed what may appear as the primary folding of the brain surface. It is now proposed to analyze different smoothing schemes of the brain surface, which may reveal the shape of this primary cortical folding.

5.3 Shape of the Primary Cortical Folding

The human cerebral cortex undergoes an extreme convolution in the early stage of the brain development (it is actually smooth in a human fœtus), with first the primary cortical folding (which is thought to have a stable pattern in a population), followed by the secondary and eventually the tertiary cortical folding throughout the development years (which are known to be variable across a population (Ono *et al.*, 1990; Welker, 1990)). The process of estimating the shape of the primary folding from a fully developed cortex is still unknown. To do so, the location of Brodmann areas (Brodmann, 1909), often associated with specific cognitive functions, can be studied since these areas are believed to form in different stages

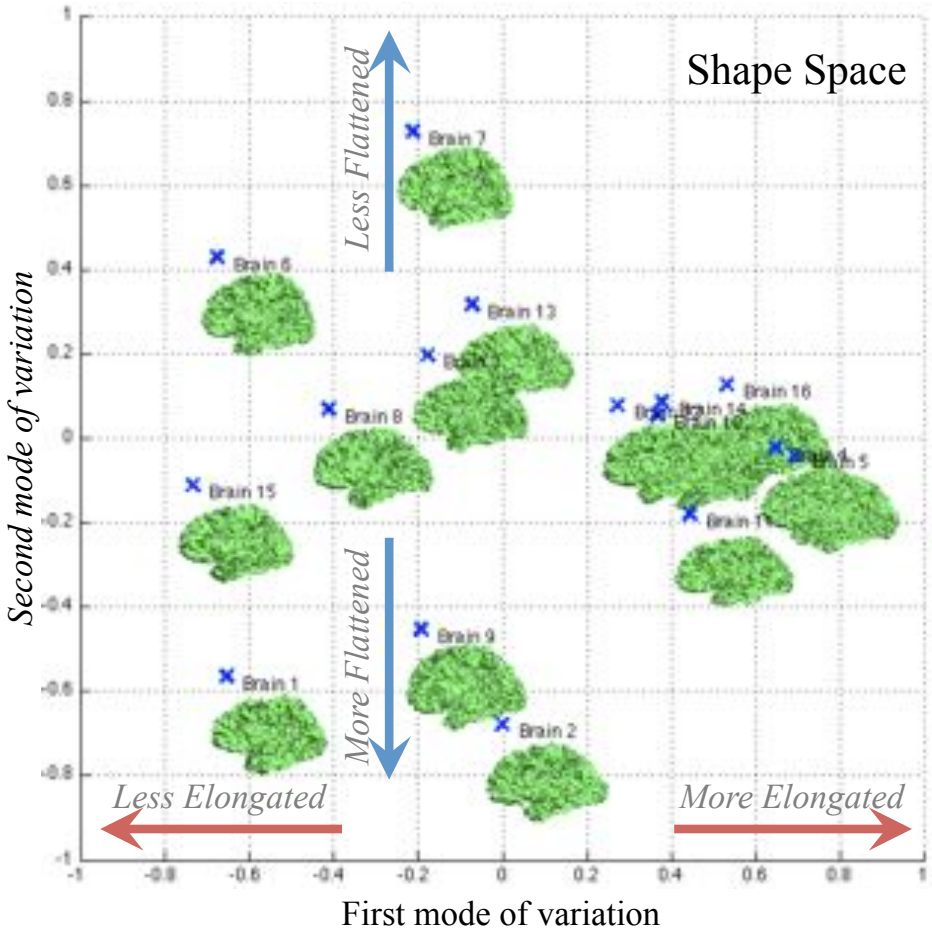

Figure 5.4: Projection of each brain surface on the first two dimensions of the shape space (each shape has 2 coordinates along the first and second principal modes of variations). It is possible to observe the distribution of brains with varying degrees of elongation (on the x-axis) and of flattening (on the y-axis).

of the brain development (Conel, 1939, 1967). In fact, the location of Brodmann areas that are formed earlier may be more stable on a surface representing the primary cortex folding. Here, we study the hypothesis that by smoothing the cortical surface, an optimal level can be found where certain Brodmann areas align better. This optimal smoothing may reveal the shape of the primary folding. In order to find this optimal level, we use FOCUSR to match brain surfaces with different types and varying degrees of smoothing. The next section briefly presents the smoothing methods, followed by the results showing the projections of Brodmann areas on the optimal smoothed cortical surface.

5.3.1 Mesh Smoothing

Three different types of smoothing method have been chosen. Each of these methods update the mesh point positions, X, with different behaviors:

Mean Filter: The mesh point coordinates are updated with the average positions of the neighboring points. The update is defined with: $X \leftarrow D^{-1}(W+I\,)X$, where I is the identity matrix, D is the degree matrix of $(W + I\,)$, and W is the adjacency matrix defined with the triangular structure of the mesh ($W_{i,j} = 1$ if points x_i and x_j are in the same triangle, 0 otherwise). This method relies on Euclidean distances and provides a very simple method of smoothing.

Heat Kernel Smoothing: The second method uses a heat kernel smoothing (Chung *et al.*, 2005). The mesh point positions are updated using the convolution of the heat kernel: $X \leftarrow K_\sigma \star X$, where σ is a diffusion time. The kernel K_σ is isotropic with respect to geodesic distances on the underlying manifold. This is achieved by using the eigenfunctions of the Laplacian representation of the mesh. The heat kernel $K_\sigma(p, q)$ is analytically given as: $K_\sigma(p, q) = \sum_{j=0}^{\infty} e^{-\lambda_j \sigma} \psi_j(p)\psi_j(q)$, where λ_j and ψ_j are respectively the j^{th} eigenvalue and the j^{th} eigenvector of the surface Laplacian matrix. This method is therefore more suited to highly folded surfaces such as the cortex.

Curvature Flow: The third method smooths the mesh using a volume preservation constraint along a curvature flow constraint (Desbrun *et al.*, 1999). Rather than using a diffusion process, a curvature flow drives the evolution of the point positions. The update is defined with: $X \leftarrow (I\ - \lambda K)^{-1}X$, where λ is a diffusion coefficient, and K is a matrix representing the curvature normals. This method preserves geometric features such as the volume and the curvature in areas of constant curvatures.

In order to translate the parameters of each method into an equivalent kernel size (how to relate the different parameters of the smoothing methods), a Dirac signal is smoothed on a flat plane (a regular triangular lattice) with all three smoothing methods (Fig. 5.5). The kernel size is estimated by measuring the diffusion of the Dirac signal (i.e., the standard deviation of the diffused signal).

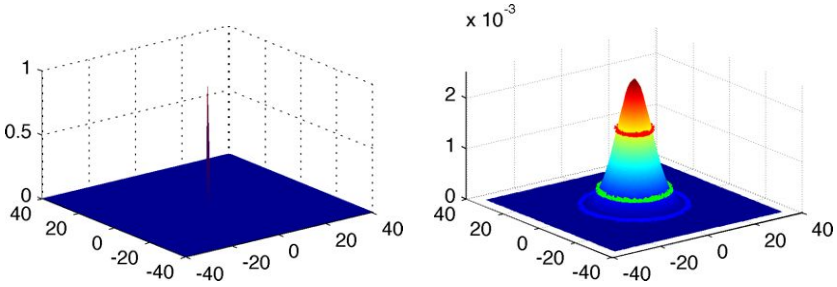

(a) Dirac signal (b) Smoothed signal

Figure 5.5: Estimating the kernel size of a smoothing method. *(a)* Original Dirac signal on a flat mesh. *(b)* Estimating the equivalent kernel size of the smoothing method with the standard deviation of the diffused signal. *Red*, *green*, and *blue* levelsets are respectively the 1σ, 2σ, and 3σ isolines.

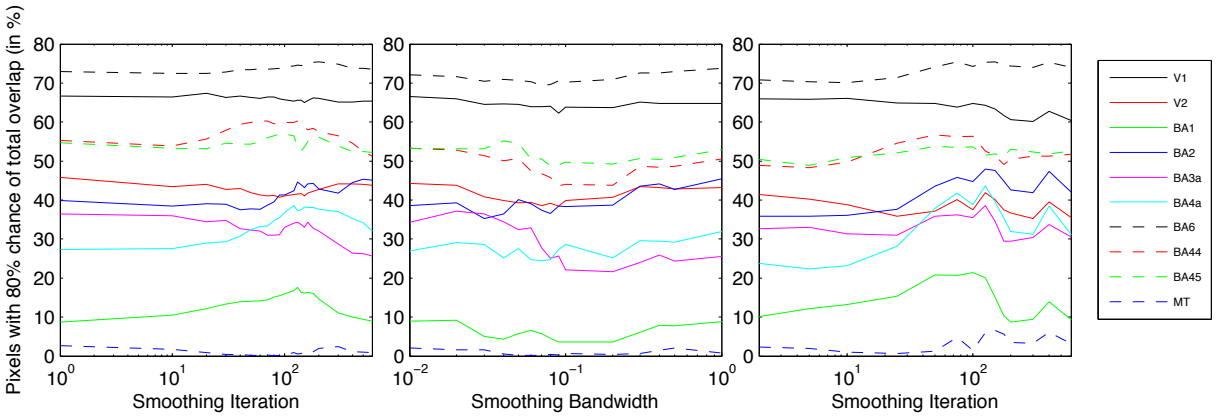

(a) Using the Mean Filter (b) Using the Heat Kernel (c) Using the Curvature Flow

Figure 5.6: Optimal level of smoothing. *(a)* Mean filter smoothing, with best overlap ratios between iterations 100 and 300 (at 160, the equivalent kernel size is $\sigma = 6.43mm$). *(b)* Heat kernel smoothing, with no clear optimality ($\sigma = 6.43mm$ at diffusion time 0.08). *(c)* Curvature flow smoothing, with best overlap ratios at iteration 125 (corresponding to $\sigma = 6.43mm$). Overlap ratios are the number of pixels (in %) of a Brodmann area that have at least 80% of chances to be in the same area in all other brains). Peaks in overlap ratios suggest that certain areas align better on smoother cortices (whose shape might be closer to shape of the primary foldings).

5.3.2 Results and Discussion

The exact localization of a cortical Brodmann area is difficult to obtain (requiring a meticulous histological study on the brain). We relied on predictions generated with FreeSurfer, which have already been shown to be accurately located for the visual cortex area V1 (Hinds *et al.*, 2009). Our dataset consists of 16 cortex surfaces with predictions on the location of areas V1, V2, BA1, BA2, BA3a, BA4a, BA6, BA44, BA45, and MT, totaling 10 Brodmann areas. For each cortex and for each smoothing method, we generated 20 meshes with an increasing level of smoothing. The correspondence maps between all pairs of cortices are generated with FOCUSR (a total of 14,400 matching, $16 \times (16 - 1) \times 20 \times 3$). Five hours were required with FOCUSR using 50 quad-core computers, while FreeSurfer would have required weeks of computation.

For each level of smoothing, the projection of Brodmann areas across all brains onto a reference brain creates probabilistic maps (i.e., the map tells if a Brodmann area aligns well at each vertex). As a measure of alignment quality, we measure the overlap ratio of a region (defined by the ratio of set intersection to set union) that finds a match on at least 80% of all other brain surfaces (i.e., the overlap ratio tells how well a Brodmann area overlaps across all brains). We iterate the process by changing the reference brain and take the average of all these overlap ratios. Fig. 5.6 shows the evolution of the overlap ratios for all three smoothing methods when the level of smoothing is increased from 0 mm to 14 mm (i.e., the size of the spread of a diffused Dirac signal, explained in sec 5.3.1, is equivalent in all methods and varies from 0 mm to 14 mm). An optimal level of smoothing seems to appear in Fig. 5.6 with parameters equivalent to using a kernel size of 6.43 mm.

Mean Filter: The alignment of Brodmann areas BA1, BA2, BA4a, BA44, BA45, BA6, and MT is better when it is performed on surfaces of smoothing levels between 100 and 300 iterations, which correspond to kernel sizes between 4.82 mm and 8.85 mm (at 160 iterations, the equivalent kernel size is 6.43 mm).

Heat Kernel Smoothing: The alignment of Brodmann areas when using this method does not show a clear improvement in the overlap ratio when smoothing the meshes (a kernel size of 6.43 mm corresponds to the algorithm diffusion time of 0.08 which shows no optimal smoothing)

Curvature Flow: Brodmann areas, in particular areas BA1, BA2, BA3a, BA4a, BA44, BA45, BA6, and MT, overlap better on meshes smoothed with fixed diffusion coefficient of

$\lambda = 1$ and 125 iterations, which also corresponds to a kernel size of 6.43 mm.

Fig. 5.7 shows the probabilistic maps (showing all Brodmann areas) at an optimal level and Fig. 5.8 shows the average shape of the cortical surfaces smoothed with all three methods using parameters equivalent to a smoothing kernel of size 6.43 mm.

Since several Brodmann areas (BA1, BA2, BA3a, BA4a, BA44, BA45, BA6, and MT) align better on surfaces smoothed with this optimal level, we might be in position to speculate that the primary cortical folding around these areas have, consequently, a shape similar to this optimally smoothed surface. When observing other areas, such as the visual cortex, V1 and V2, the lack of an optimal smoothing level could also mean that the folding convolution in these areas is more stable (i.e., a less variable process) during the brain development. For instance, a consistent relationship has already been suggested between area V1 and the surrounding cortical folding on a fully developed brain (Hinds *et al.*, 2008). A third smoothing method using a heat kernel could not lead to a clear conclusion, and shows that several smoothing methods need to be tested in future experiments.

All these findings require further investigation; in particular, even though two of our tested smoothing methods demonstrated an optimal level of smoothing, other smoothing methods (Yu *et al.*, 2007) and curvature measures (Pienaar *et al.*, 2008) need to be tested. More cortices should also be included in the study. However, we showed that FOCUSR can be an attractive solution to such study necessitating a large number of cortex matchings. We believe that FOCUSR can open new doors in neuroscience by relieving studies limited by a heavy computational burden. This experiment is one example and provides basic tools for future studies on the human brain development. Nevertheless, FOCUSR has proved itself as a fast and efficient surface matching method that is capable of capturing very large shape variability, *however*, it is not directly applicable to images. The next section introduces for the first time a spectral matching approach for non-parametric image registration.

5.4 Spectral Demons for Atlas Construction

Image registration is a building block in many applications in computer vision and medical imaging, including the construction of atlases. Current methods in the state-of-the-art uses a Euler-Lagrangian approach, i.e., iterative methods (surveyed in (Crum *et al.*, 2004)), however, their underlying update schemes rely on forces derived from image gradients and are, therefore, fundamentally limited by their *local scope* (e.g., gradients are null in texture-

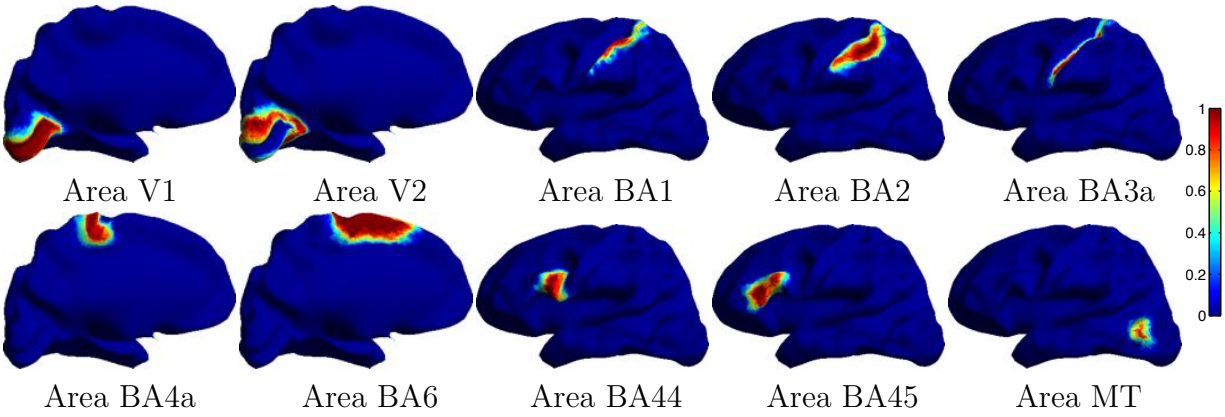

Figure 5.7: Projection of Brodmann areas V1, V2, BA1, BA2, BA3a, BA4a, BA6, BA44, BA45, and MT from 16 matched cortices onto a reference cortical mesh smoothed with curvature flow at the optimal smoothing level. Coloring is the probability of each point to belong to the Brodmann area.

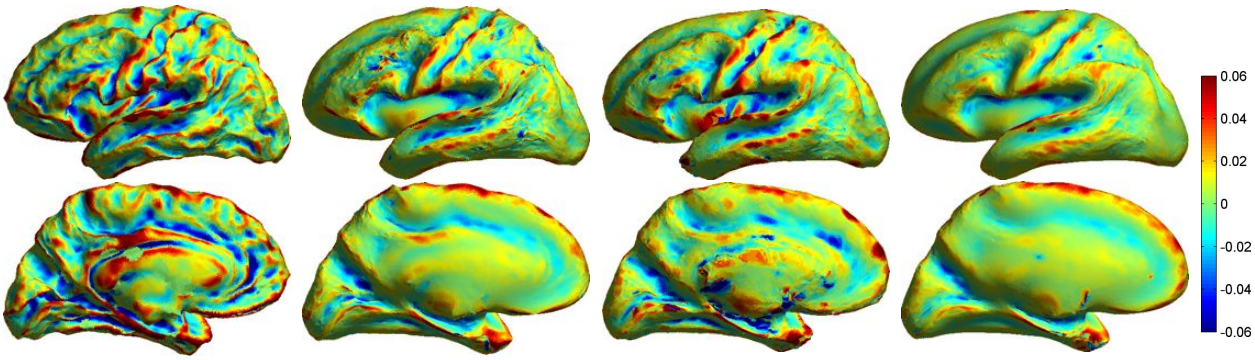

(a) Original mesh (b) Mean filter (c) Heat kernel (d) Curvature flow

Figure 5.8: Average of the cortical surfaces using the optimal smoothing levels (equivalent kernel size of $\sigma = 6.43mm$). *(a)* Original mesh (highly convoluted). Mesh smoothed: *(b)* with the mean filter (160 iterations), *(c)* with the heat kernel (diffusion time of 0.08), and *(d)* with a curvature flow (125 iterations). Coloring is the mean curvature.

less areas and the optimization is undermined by local minima). The typical response is to use a multilevel scheme to capture larger deformations in a higher resolution; however, this response does not fundamentally solve the local scope of update schemes based on image gradients.

In order to capture complex large deformations, we introduce a new approach for image registration based on a *direct feature matching* technique that has a *global scope* due to simple nearest-neighbor searches in a multi-dimensional space comprising information on *image texture* (e.g., pixel intensities), *space* (e.g., Euclidean coordinates of pixels) and

on *global image geometry*. For the latter, the graph spectral representations used in FOCUSR will enable the capture of large and complex deformations. Our new direct feature matching technique provides a geometric component with a *global scope* that can be used for image registration, i.e., the nearest-neighbor search finds correspondences with the best compromise in image similarity, spatial regularity and geometric similarity. Furthermore, our direct feature matching technique is used within a diffeomorphic framework such as the efficient symmetric *Log-Demons* algorithm (Vercauteren *et al.*, 2007, 2008). Therefore, the new method, called **Spectral Demons**, enables a symmetric and diffeomorphic registration of images undergoing *large* and *complex* deformations.

Besides the development of the *Spectral Demons* algorithm, a new framework for atlas construction is also created with the extension of the symmetric Demons to perform groupwise registration. This new simultaneous (groupwise) registration approach enables the construction of the atlas in parallel, during the registration process (rather than with a series of sequential pairwise registrations (Guimond *et al.*, 2000)). We provide *two forms* of our groupwise registration framework that we name the *Groupwise Log-Demons* (**GL-Demons**, faster and suited for local nonrigid deformations), and the *Groupwise Spectral Log-Demons* (**GSL-Demons**, slower but capable of capturing very large deformations).

In the next sections, we describe our new direct feature matching, and show how it can be used within a diffeomorphic image registration framework before presenting the groupwise extensions. The results evaluate the capability and robustness of the *Spectral Demons* to register images with highly non-local deformations as well as demonstrate how the two forms of the new groupwise framework can construct atlases from highly deformed images.

5.4.1 Spectral Demons

We begin our methodology with our simple and direct feature matching technique followed by how spectral representations are built for images and how they are used within a diffeomorphic framework for image registration.

Direct Feature Matching

Image registration warps a moving image M toward a fixed image F through a transformation ϕ that maps points from F to M (i.e., features $\mathbf{F}(\cdot)$, such as point coordinates $\mathbf{x}(\cdot) = (x, y)$ or image intensity $I(\cdot)$, match those in the transformed features $\mathbf{M}(\phi(\cdot))$, or simply $\mathbf{F} \mapsto \mathbf{M} \circ \phi$).

A *direct* approach for feature matching would find the point correspondence with a nearest-neighbor search in the feature space (e.g., with a Voronoi tessellation or a k-d tree) such that $\phi(i) = \text{argmin}_{j \in M} \|\mathbf{F}(i) - \mathbf{M}(j)\|^2$ (e.g., if point $j \in M$ has the closest intensity from the one of point $i \in F$, then $\phi(i) = j$). Obviously, matching 1D features such as image intensity would result in a highly unsmooth mapping lacking any spatial regularity (points with the closest intensities might be far apart in the images). Spatial regularity can be introduced by incorporating Euclidean coordinates in the feature space where points now have the extended coordinates $\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F)$ and $\mathbf{M} = (\alpha_i I_M, \alpha_s \mathbf{x}_M)$ (with weights α_i and α_s). A nearest-neighbor search in such extended space effectively provides similarity in pixel intensity and closeness in space between corresponding pixels and minimizes the similarity criterion:

$$\text{Sim}(F, M, \phi) = (I_F - I_{M \circ \phi})^2 + \frac{\alpha_s^2}{\alpha_i^2} \|\mathbf{x}_F - \mathbf{x}_{M \circ \phi}\|^2 \tag{5.5}$$

However, such method lacks geometric information from the images (how to naturally match objects in different poses?) and does not produce a diffeomorphic mapping (a one-to-one mapping is not guaranteed). Each issue is addressed below.

Spectral Correspondence

The spectral representation of shapes (Chung, 1997a; Grady and Polimeni, 2010; Umeyama, 1988; Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992; Jain and Zhang, 2006; Mateus *et al.*, 2008; Lombaert *et al.*, 2011a) has the strong property of being invariant to *isometry*, i.e., corresponding points between shapes in different poses would share the same *spectral coordinates* (or signature) even if they are far away in space (e.g., a point on a nose tip has a geometric description that is unique even if moved in space). By adding these spectral coordinates in our feature space, we enforce an intrinsic *geometric consistency* in our matching technique.

Spectral Graph Theory From the pixels of I_Ω (the portion of an image I bounded by a contour Ω), the connected undirected graph $\mathscr{G} = (\mathscr{V}, \mathscr{E})$ is constructed with the vertices $\mathscr{V}$ representing pixels and the edges $\mathscr{E}$ defined by the neighborhood structure of these vertices. Such graph can be represented with its adjacency matrix W in terms of affinity weights (Grady and Polimeni, 2010) where high weights are given to edges within a region of uni-

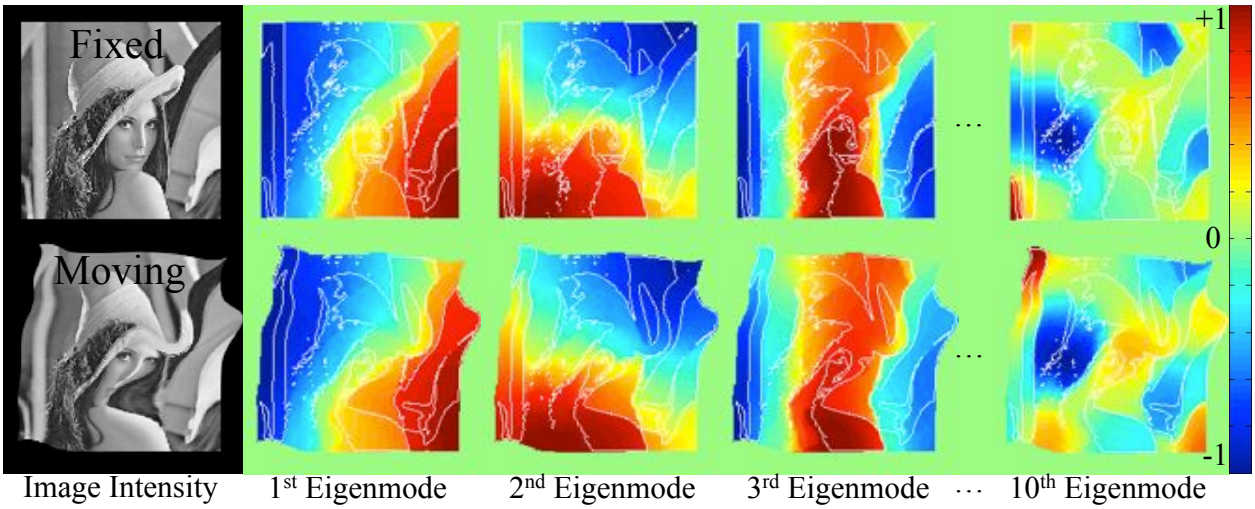

Figure 5.9: The eigenmodes of the graph Laplacian are used as geometric descriptors. They remain stable under complex deformations (i.e., invariant to isometry). The lower eigenmodes describe coarse geometric properties, while higher eigenmodes describe finer details in the images (the image contours are overlaid on the eigenmodes).

form intensities and low weights are given to edges crossing region boundaries, e.g., $W_{i,j} = \exp\left(-(I(i)-I(j))^2/2\sigma^2\right)/\|\mathbf{x}(i)-\mathbf{x}(j)\|^2$ if $(i,j) \in \mathscr{E}$ and 0 otherwise (parameter σ depends on the image noise and is set, without being limited, to $\sigma = \text{mean}\{|I(i)-I(j)|\}_{(i,j)\in\mathscr{E}}$). The (diagonal) degree matrix D provides the total weighting of all edges connected to each vertex ($D_{ii} = \sum_j W_{i,j}$) and the Laplacian matrix is defined by $L = D - W$. Here, we consider the general Laplacian operator on a graph $\mathcal{L} = G^{-1}(D-W)$ (Grady and Polimeni, 2010), i.e., a $|\mathscr{V}| \times |\mathscr{V}|$ sparse matrix where G is the (diagonal) node weighting matrix, e.g., $G = D$.

Spectral Coordinates The graph spectrum (Chung, 1997a) computed from the decomposition of the Laplacian $\mathcal{L} = \mathbf{X}^T\Lambda\mathbf{X}$ comprises the eigenvalues (in increasing order) $\Lambda = \text{diag}(\lambda_0, \lambda_1, \ldots, \lambda_{|\mathcal{V}|})$ and their associated eigenvectors $\mathbf{X} = \left(\mathscr{x}^{(0)}, \mathscr{x}^{(1)}, \ldots, \mathscr{x}^{(|\mathcal{V}|)}\right)$ ($\mathbf{X}$ is a $|\mathscr{V}| \times |\mathscr{V}|$ sparse matrix where each column $\mathscr{x}^{(\cdot)}$ is an eigenvector). The first eigenvector $\mathscr{x}^{(0)}$ is the stationary distribution (related to the expected return time of a random walker). The following eigenvectors associated with the non-zero eigenvalues are the fundamental modes of vibrations of the shape (with free ends) depicted by I_Ω (in a Riemannian sense). We thus prefer the term *eigenmode* since they are effectively functions over I_Ω, i.e., they can be visualized as images. As an example, a generic image (Lena has been chosen arbitrarily) is decomposed and each of its eigenmodes is shown in Fig. 5.9. The eigenmodes of lower modal

frequencies are harmonics depicting coarse geometric properties of I_Ω while those associated with higher eigenvalues depict finer geometric details in I_Ω. Moreover, the oscillations at a modal frequency λ occur around nodal sets (where the eigenmodal values equal 0) that reside on prominent demarcations of the shape geometry. For instance, graph-based segmentation methods (e.g., (Shi and Malik, 2000; Meila and Shi, 2000)) rely on the nodal set of $\mathscr{x}^{(1)}$, called the Fiedler vector (Chung, 1997a), to find a binary partition of an image. The number of half waves in these oscillations (or the number of extrema in the eigenmodal values) is also given by the algebraic multiplicity of their eigenvalue, n_λ.

In our approach (summarized in Alg. (5)), we consider the first k eigenmodes of lower modal frequencies $\mathscr{x}^{(1..k)}$ (i.e., the strongest intrinsic geometric descriptors). Their components represent the k-dimensional *spectral coordinates* $\mathscr{x}$ in a spectral domain[1] where each point i has the coordinates $\mathscr{x}(i) = \left(\mathscr{x}^{(1)}(i), \mathscr{x}^{(2)}(i), \ldots, \mathscr{x}^{(k)}(i)\right)$ (a truncated line of matrix $\mathbf{X}$). The spectral representation $\mathscr{x}$ has the strong property of being invariant to isometry, i.e., if F and M are images of the same object in different poses, equivalent points would share similar coordinates $\mathscr{x}_F$ and $\mathscr{x}_M$. We use this property to improve our direct feature matching by extending Eq. (5.6) with these spectral coordinates (weighted with α_g). A nearest-neighbor search between $\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F, \alpha_g \mathscr{x}_F)$ and $\mathbf{M} = (\alpha_i I_M, \alpha_s \mathbf{x}_M, \alpha_g \mathscr{x}_M)$ provides similarity in intensity, space and in intrinsic geometric characteristics, effectively minimizing:

$$\mathrm{Sim}(F, M, \phi) = (I_F - I_{M\circ\phi})^2 + \frac{\alpha_s^2}{\alpha_i^2} \left\|\mathbf{x}_F - \mathbf{x}_{M\circ\phi}\right\|^2 + \frac{\alpha_g^2}{\alpha_i^2} \left\|\mathscr{x}_F - \mathscr{x}_{M\circ\phi}\right\|^2, \tag{5.6}$$

where $\mathscr{x}_{M\circ\phi}$ are the spectral coordinates of the corresponding points in the transformed image $M \circ \phi$.

The choice of the number of spectral components $k = n_{\lambda_1}$ is motivated by the Colin de Verdière's number (Tlusty, 2010) which is in this case the multiplicity of the Fiedler vector n_{λ_1} and is also related (Tlusty, 2010) to the maximal dimension of a space in which the graph $\mathcal{G}$ can be mapped (i.e., the eigenspace of the Fiedler eigenvalue reveals the principal symmetries in I_Ω and, $n_{\lambda_1} \leq 2$ in 2D, $n_{\lambda_1} \leq 3$ in 3D). More complex symmetries in the cyclic or dihedral group could be considered with a higher k, but is not required in our method.

[1] In our notation $\mathbf{x}$ is the Euclidean coordinates (e.g., x, y, z in 3D) and superscripted $\mathscr{x}^{(u)}$ is the u^{th} component of the spectral coordinates $\mathscr{x}$

Rearrangement of the Spectra Unfortunately, the spectral coordinates $\mathscr{x}_F$ and $\mathscr{x}_M$ of points in F and M may not be directly comparable as a result of two phenomena. *Firstly*, there is a sign and scaling ambiguity between corresponding eigenmodes (if $\mathscr{x}^{(\cdot)}$ is an eigenmode of $\mathcal{L}$, so is $-\alpha\mathscr{x}^{(\cdot)}$) which requires a sign check and a scaling correction between $\mathscr{x}_F$ and $\mathscr{x}_M$. *Secondly*, the order of the eigenmodes is undefined within an eigenspace (if two eigenmodes $\mathscr{x}^{(u)}, \mathscr{x}^{(v)}$ share the same eigenvalue, their order (u, v) may differ between two images). The order is additionally perturbed with imperfections in isometry (near-symmetry creates close but not equal eigenvalues and may change order between images). We rearrange the spectral coordinates using two new simple heuristics.

The *first* issue is addressed by scaling the values of each eigenmodes in order to fit the range $[-1; +1]$. The nodal set (where $\mathscr{x}^{(\cdot)} = 0$) is thought to remain on a prominent geometric feature (an axis of symmetry in a Riemannian sense) and should not be changed. We scale thus the positive values (where $\mathscr{x}^{(\cdot)} > 0$) with $\mathscr{x}^{(\cdot)+} \leftarrow \mathscr{x}^{(\cdot)+} / \max\{\mathscr{x}^{(\cdot)+}\}$ and the negative values (where $\mathscr{x}^{(\cdot)} < 0$) with $\mathscr{x}^{(\cdot)-} \leftarrow \mathscr{x}^{(\cdot)-} / \min\{\mathscr{x}^{(\cdot)-}\}$. The *second* issue is addressed by finding the optimal permutation π such that $\mathscr{x}_F^{(\cdot)}$ and $\mathscr{x}_M^{\pi\circ(\cdot)}$ correspond with each other. The Hungarian algorithm (also used in (Mateus *et al.*, 2008; Lombaert *et al.*, 2011a)) minimizes the following 2D dissimilarity matrix:

$$C(u,v) = \sqrt{\frac{1}{|I_\Omega|}\sum_{i\in I_\Omega}\left(\mathscr{x}_F^{(u)}(i) - \mathscr{x}_M^{(v)}(i)\right)^2} + \sqrt{\sum_{i,j}\left(h_F^{\mathscr{x}_F^{(u)}}(i,j) - h_M^{\mathscr{x}_M^{(v)}}(i,j)\right)^2} \tag{5.7}$$

The first term is the difference in eigenmodal values between the images and, the second term measures the dissimilarities between the joint histograms $h(i, j)$ (a 2D matrix where the element (i, j) is the joint probability of having at the same time a pixel with intensity i and eigenmodal value $\mathscr{x}^{(\cdot)} = j$). The sign ambiguity can be removed by using, instead, the dissimilarity matrix $Q(u, v) = \min\{C(u, v), C(u, -v)\}$. To keep the notation simple, in the next sections, we assume that the spectral coordinates have been appropriately signed, scaled and reordered.

The ordered spectral coordinates provide our geometric component in our new direct feature matching. We now briefly review how diffeomorphism can be achieved for image registration.

Algorithm 5 Spectral Correspondence

Input: Images F, M.
Output: Correspondence c mapping F to M

- Compute general Laplacians $\mathcal{L}_F$, $\mathcal{L}_M$.
  $\mathcal{L} = D^{-1}(D - W)$, where
  $W_{ij} = \exp\left(-(I(i) - I(j))^2/2\sigma^2\right) / \|\mathbf{x}(i) - \mathbf{x}(j)\|^2$
  $D_{ii} = \sum_j W_{ij}$,
- Compute first k eigenmodes of Laplacians
- Reorder $\mathscr{x}_M$ with respect to $\mathscr{x}_F$ (Eq. (5.7))
- Build embeddings:
  $\mathbf{F} = (I_F, \mathbf{x}_F, \mathscr{x}_F)$; $\mathbf{M} = (I_M, \mathbf{x}_M, \mathscr{x}_M)$
- Find c mapping nearest points $\mathbf{F} \mapsto \mathbf{M}$

Diffeomorphic Registration

The minimization of Eq. (5.6) does not guarantee a one-to-one mapping between points (only closest points are assigned and undefined correspondences are possible). Such property however exists in classical methods for diffeomorphic registration such as the *Log-Demons* algorithm (Vercauteren *et al.*, 2007). Arising from the theory of Lie groups, a diffeomorphic transformation ϕ (on a Lie group structure) is related to the exponential map of a velocity field v (a Lie algebra), i.e., $\phi = \exp(v)$. In the case of stationary velocity fields, a practical and fast approximation is possible with the scaling-and-squaring method (Vercauteren *et al.*, 2007) (Alg. (6)). As an aside, the inverse of the transformation is simply $\phi^{-1} = \exp(-v)$.

The *Log-Demons* framework alternates, similarly to the Maxwell's demon, between the optimization of a similarity term, e.g., $\mathrm{Sim}(F, M \circ \exp(v)) = (I_F - I_{M \circ \exp(v)})^2$, and a regularization term, e.g., $\mathrm{Reg}(v) = \|\nabla v\|^2$, through the introduction of a hidden variable (the correspondences c) which allows a small error between alternations, e.g., $\mathrm{dist}(c, \phi) = \|c - v\|$. Moreover, invariance to the order of the input images is possible with the symmetric extension of the algorithm (Vercauteren *et al.*, 2008, 2009a). The energy of the symmetric *Log-Demons* can be written:

$$\begin{aligned} E(F, M, \exp(c), \exp(v)) = & \tfrac{1}{2}\alpha_i^2 \left(\mathrm{Sim}\left(F, M \circ \exp(c)\right) + \mathrm{Sim}\left(F \circ \exp(-c), M\right)\right) \\ & + \alpha_x^2 \mathrm{dist}(c, v)^2 + \alpha_T^2 \mathrm{Reg}(v), \end{aligned} \tag{5.8}$$

where the Euler-Lagrangian updates are computed directly on the stationary velocity field v and consist of the average of the forward and backward updates $u_{F \to M}, u_{M \to F}$ mapping F to

Algorithm 6 Exponential $\phi = \exp(v)$

Input: Velocity field v.
Output: Diffeomorphic map $\phi = \exp(v)$.
- Choose N such that $2^{-N}v$ is close to 0
  e.g., such that $\max ||2^{-N}v|| \leq 0.5$ pixels
- Scale velocity field $\phi \leftarrow 2^{-N}v$.

for N times **do**
- Square $\phi \leftarrow \phi \circ \phi$.

end for

Algorithm 7 The Log-Demons Framework

Input: Images F, M and initial velocity field v
Output: Transformation $\phi = \exp(v)$ from F to M

repeat
- Find updates $u_{F \to M}$ mapping F to $M \circ \exp(v)$
- Find updates $u_{M \to F}$ mapping M to $F \circ \exp(-v)$
- Average updates: $u \leftarrow \frac{1}{2}(u_{F \to M} - u_{M \to F})$.
- Smooth updates: $u \leftarrow K_{\text{fluid}} \star u$
- Update velocity field: $v \leftarrow \log(\exp(v) \circ \exp(u))$
  (approximated with $v \leftarrow v + u$)
- Smooth velocity field: $v \leftarrow K_{\text{diff}} \star v$.

until convergence

$M \circ \exp(c)$ and, M to $F \circ \exp(-c)$ such as (see (Vercauteren *et al.*, 2007) for more details):

$$u_{F \to M} = -\frac{I_F - I_{M \circ \phi}}{\|\nabla I_{M \circ \phi}\|^2 + \alpha_x^2 |I_F - I_{M \circ \phi}|^2} \nabla I_{M \circ \phi}. \tag{5.9}$$

In the first step, the transformation ϕ is fixed and the updates are computed (optionally smoothed with a kernel K_{fluid}, e.g., Gaussian with σ_{fluid}). In the second step, the velocity field is updated $v \leftarrow v + u$ (optionally smoothed with the a kernel $K_{\text{diffusion}}$, e.g., Gaussian with $\sigma_{\text{diffusion}}$, see (Cachier *et al.*, 2003)). The general symmetric diffeomorphic Demons framework is summarized in Alg. (7).

Spectral Demons

At this stage, we have described two methods. The first (the direct feature matching technique, Eq. (5.6)) can capture large deformations between images but does not guarantee diffeomorphism and symmetry. The second (the Demons algorithm, Eq. (5.8)) offers these properties but suffers from the local scope of the update forces derived from the image gradient (Eq. (5.9)). Both methods can benefit from the incorporation, with very little modifications, of our spectral correspondence approach in the *Log-Demons* framework.

The so-called *Spectral Demons* algorithm takes advantage of the efficient diffeomorphic framework offered by Alg. (7) and finds the correspondences between images F and M using our *global* spectral approach in place of using *local* gradient-based updates. To be more precise, the first two steps of Alg. (7) (originally computing updates with Eq. (5.9)) now perform spectral correspondence between images F and $M \circ \exp(v)$ using Alg. (5) (and respectively between M and $F \circ \exp(-v)$). This modification enables large jumps in each iteration where points are moving toward their isometric equivalent even if they are far away in space. This virtually enables the capture of very large deformations (with invariance to

Figure 5.10: Implicit image segmentation by partitioning the images with the nodal sets of the first and second eigenmodes. Image contours are overlaid for visualization.

isometry) between images as well as a faster convergence of the algorithm. The underlying energy being minimized has the form of Eq. (5.8) where its similarity term is expressed with Eq. (5.6).

Multilevel Scheme A multilevel scheme is also possible with the *Spectral Demons* where large deformations are assumed to be related with coarse geometric information. Spectral updates can thus be safely used in a lower resolution level, while finer details and local deformations remain computed in the higher levels of resolutions using the classical update forces based on the image gradient. This multilevel scheme keeps the computation of the eigenmodes tractable. On the same note, the computation of the eigenmodes can be used with the efficient Lanczos method (used by Matlab) (Lehoucq *et al.*, 1997) which has a running time of $O(n\sqrt{n}) + O(n^2)$ (Shi and Malik, 2000), where n is the number of pixels in I_Ω, while spectral matching can be performed with a k-d tree which is built in $O(n \log^2 n)$ and queried in $O(\log n)$.

Note on Image Segmentation The Laplacian eigenmodes have been demonstrated to have important properties in the field of spectral graph theory (Chung, 1997a; Shi and Malik, 2000; Luxburg, 2007; Grady and Polimeni, 2010) by providing a probabilistic foundation (Meila and Shi, 2000; Robles-Kelly, 2005) for graph-based segmentation methods. In particular, the Normalized Cut problem (Shi and Malik, 2000) finds a segmentation x by minimizing $\frac{x^T L x}{x^T D x}$ (revealed by the Fiedler vector of the normalized Laplacian $D^{-\frac{1}{2}} L D^{-\frac{1}{2}}$ (Meila and Shi, 2000)). *Spectral Demons* considers the more general Laplacian operator $\mathcal{L}$ and effectively exploits for registration the same global image description used by Normal-

ized Cuts for segmentation, i.e., eigenmodes of the (general) Laplacian operator are used for image registration. Since the Fiedler vector is an inherent part of our algorithm, binary segmentations of images come at no extra cost by taking either the positive or negative values of the Fiedler vector, $\mathscr{X}^{(1)+}$ or $\mathscr{X}^{(1)-}$. For example, the implicit segmentation of Lena's image, shown on Fig. 5.10 (with overlaid image contours), was obtained with the nodal sets of the first and second eigenmodes (positive values of the Fiedler vector gave warmer colors, negative values gave cooler colors, see also 1st and 2nd eigenmodes on Fig. 5.9). Nodal sets of higher frequency eigenmodes may additionally reveal important geometric features for meaningful segmentation, however, an exhaustive experimental study on the segmentation aspect of our registration method goes beyond the scope of this thesis.

5.4.2 Groupwise Demons

The atlas is defined as mentioned earlier, with a set of N images $\{I_i\}_{i=1..N}$ nonrigidly aligned to their average shape $\tilde{I}$. Our new shape averaging framework extends the symmetric Log-Demons algorithm (Vercauteren *et al.*, 2008) and can use classical gradient-based updates (*GL-Demons*) or an improved spectral matching for groupwise registration (*GSL-Demons*).

The *Log-Demons* algorithm (Alg. (7)) is slightly modified such that convergence is toward an average shape instead of the fixed image. Eq. (5.8) is modified accordingly to become:

$$E(F, M, c, v) = \alpha_i^2 \mathrm{Sim}(F', M') + \alpha_x^2 \mathrm{dist}(c, v)^2 + \alpha_T^2 \mathrm{Reg}(v), \text{ where} \tag{5.10}$$
$$\mathrm{Sim}(F', M') = (F' - M')^2,\ \mathrm{dist}(c, v) = \|c - v\|,\ \text{and } \mathrm{Reg}(v) = \|\nabla v\|^2$$

The similarity term incorporates diffeomorphism and symmetry with $F' = F \circ \exp(-c)$ and $M' = M \circ \exp(+c)$. Both images F' and M' effectively converge toward an average shape $\tilde{I} = F \circ \phi^{-1} + M \circ \phi$ (similar to the approaches in (Avants and Gee, 2004a; Bossa *et al.*, 2007)).

Our groupwise framework is based on Guimond's *et al.* approach (Guimond *et al.*, 2000) where they construct the average image $\tilde{I}$ *sequentially* by alternating between pairwise registrations (i.e., fixing a reference image) and updates of the average image (i.e., transforming the reference image). Our novelty is to directly compute $\tilde{I}$ *in parallel* with simultaneous (groupwise) registrations (illustrated in Fig. 5.11). To do so, Eq. (5.10) is extended to incorporate N velocity fields that warp all images $\{I_i \circ \exp(c_i)\}$ toward the average image $\tilde{I}$.

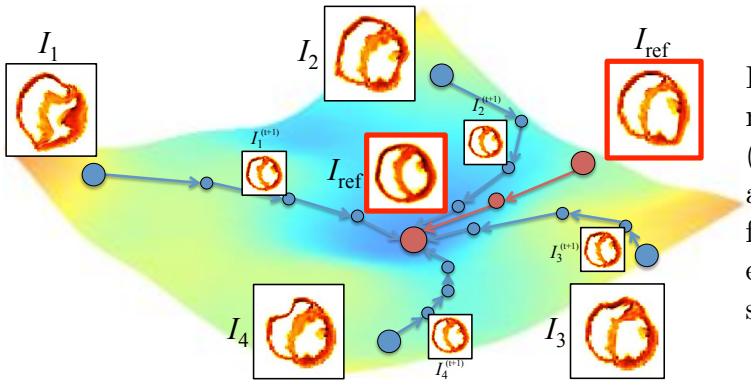

Figure 5.11: Groupwise Demons: Simultaneous registration of 4 images (blue circles) toward a reference image that evolves in the space of diffeomorphisms (manifold). The reference image converges to the average shape (middle red circle).

The new groupwise framework is summarized in Alg. (8) and the underlying energy is:

$$E(\tilde{I}, \{I_i, c_i, v_i\}) = \frac{1}{N} \sum_{i=1}^{N} \left(\alpha_i^2 \text{Sim}(\tilde{I}, I_i \circ \exp(c_i)) + \alpha_x^2 \text{dist}(c_i, v_i)^2 + \alpha_T^2 \text{Reg}(v_i) \right) \tag{5.11}$$

The reference image can be optionally generated with weighted contributions from all images (e.g., weights different than $1/N$ in order to remove outliers). The minimization of all similarity terms, $\{\text{Sim}(\tilde{I}, I_i')\}$, causes all warped images to become similar to the reference image and the sum of all velocity fields is brought to a minimal value at convergence. Similar to the convergence of (Guimond *et al.*, 2000), the Groupwise Demons framework effectively brings the reference image toward the barycenter of all images. The average image is simply generated with $\tilde{I} = \frac{1}{N} \sum_{i=1}^{N} I_i \circ \exp(c_i)$.

Groupwise Spectral Log-Demons

The update schemes based on image gradients and on spectral correspondence can be used in the Groupwise Demons framework. The *Groupwise Log-Demons* (GL-Demons) algorithm uses update forces derived from image gradients and is well suited for images with local nonrigid deformations, while the *Groupwise Spectral Log-Demons* (GSL-Demons) algorithm uses spectral correspondences as update forces (i.e., u is found with Alg. (5)) and is better suited for large and highly nonlocal deformations. GSL-Demons enables large jumps during the construction of the atlas where points move toward their isometric equivalents even if they are far away in space. The atlas construction can handle very large deformations and convergences in fewer iterations (typically 5 iterations are sufficient). The energy has the same form of Eq. (5.11) and uses the similarity term of Eq. (5.6).

Algorithm 8 Groupwise Demons Framework

Input: N images with initial reference (e.g., $\tilde{I} = I_1$)
Output: Transformations $\phi_i = \exp(v_i)$ mapping $\tilde{I}$ to I_i
Average shape is $\tilde{I} = \frac{1}{N} \sum_{i=1}^{N} I_i \circ \exp(v_i)$

repeat
for $i = 1 \rightarrow N$ **do**
- Find updates $u_i \leftarrow \text{mapping}(\tilde{I}, I_i \circ \exp(v_i))$.
  (mapping() differs in GL and GSL-Demons)
- Smooth updates: $u_i \leftarrow K_{\text{fluid}} \star u_i$.
  (convolution of a Gaussian kernel on u_i)
- Update velocity fields: $v_i \leftarrow \log(\exp(v_i) \circ \exp(u_i))$
  (approximated with $v_i \leftarrow v_i + u_i$).
- Smooth velocity fields: $v_i \leftarrow K_{\text{diff}} \star v_i$.

end for
- Get reference update: $u_{\text{ref}} = -\frac{1}{N} \sum_{i=1}^{N} v_i$
- Update velocity fields: $v_i \leftarrow v_i + u_{\text{ref}}$.
- Update reference: $\tilde{I} \leftarrow \frac{1}{N} \sum_{i=1}^{N} I_i \circ \exp(v_i)$.

until convergence

Figure 5.12: Pairs of images used in the synthetic experiments (Lena, heart, baseball player). Each left image is a fixed image, each right side is a moving image generated with a random deformation of at most 25 pixels (difficulties in red). These transformations provide our ground truth.

Multilevel Scheme

Moreover, large and complex deformations can be captured in a low resolution level with *GSL-Demons*, improving thus the processing time, while the remaining small and local deformations can be recovered with *GL-Demons* in higher resolutions. This multilevel approach keeps the computation of the eigenmodes tractable.

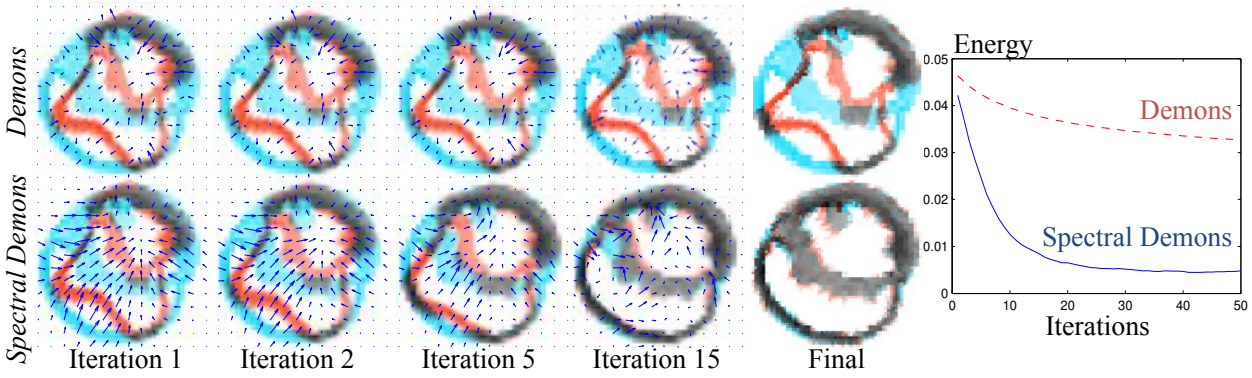

Figure 5.13: Comparison of the iterations in the *Log-Demons* and the *Spectral Demons* (our method) within the same resolution level. The update forces (indicated with the arrows and scaled for visualization) have a local scope with the *Log-Demons* and a global scope with the *Spectral Demons*. This global scope allows a faster convergence while the *Log-Demons* remains in a local minimum.

5.4.3 Results and Discussion

We now evaluate the performance of the *Log-Demons* (our benchmark) and *Spectral Demons* (our method) by registering images with large and highly nonlocal deformations. In our controlled experiments, the full power of the *Spectral Demons* can be appreciated with drastic deformations of the images. Accurate measurements, with known ground truth, are used to evaluate the improvements in registration accuracy with respect to the *Log-Demons*. We additionally provide a real application where two human brains are registered.

In a second set of experiments, *GL-Demons* and *GSL-Demons* are evaluated by constructing atlases of images with large deformations. We verify, in a controlled experiment, convergence toward an average shape, as well as in a real experiment, the construction of a 3D atlas of cardiac images.

Controlled Experiments

In the first controlled experiment, we evaluate the fundamental difference between the update schemes of the *Log-Demons* and *Spectral Demons*. To do so, we analyze the convergence rate of both algorithms and, since we are not interested here in their final performance, we compare them within the same level of resolution. The algorithms use the same parameter set ($\sigma_{\text{fluid,diffuse}} = 1, \alpha_x = 1, k = 2, \alpha_g = 0.05, \alpha_s = 0.15, \alpha_i = 0.8$). We register the images on Fig. 5.12 (Lena has size 128^2, the heart is 75^2 and the baseball player is 110×75). Each

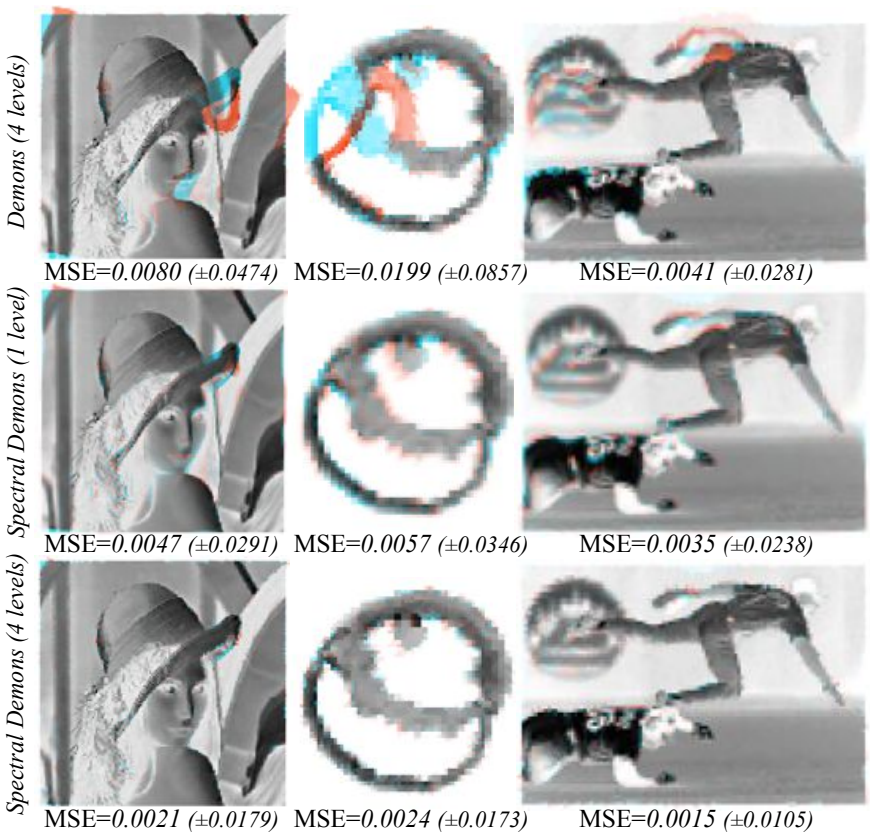

Figure 5.14: Final registrations for (*top row*) Multilevel *Log-Demons* with 4 levels, (*middle row*) *Spectral Demons* with 1 level, (*bottom row*) Multilevel *Spectral Demons* with 4 levels. The fixed image is in blue and misalignments of the registered image are in red. The mean square differences of intensities are reported along their standard deviations. *Log-Demons* is limited in areas of high deformations, while *Spectral Demons* can capture these large deformations.

moving image is generated with random diffeomorphic deformations ϕ_{truth} with displacements of at most 25 pixels, i.e., we take the exponential map of a random velocity field generated with 15 random displacements (control points randomly located) diffused across the image (Gaussian smoothing $\sigma = 10$ pixels). Notably, Lena's hat, her neck and, the player's arm are the highest registration challenges, while the cardiac image (a 2D slice of an MRI) shows a papillary muscle (red circle on Fig. 5.12) severely deformed and almost fully collapsed (the muscle forms, however, a dent in the image and provides a signature that *Spectral Demons* can understand). The iterations of the *Log-Demons* and *Spectral Demons* are compared in Fig. 5.13. It shows that within the same level of resolution, the update forces computed with spectral correspondence are coherent spatially and geometrically, i.e., points move toward

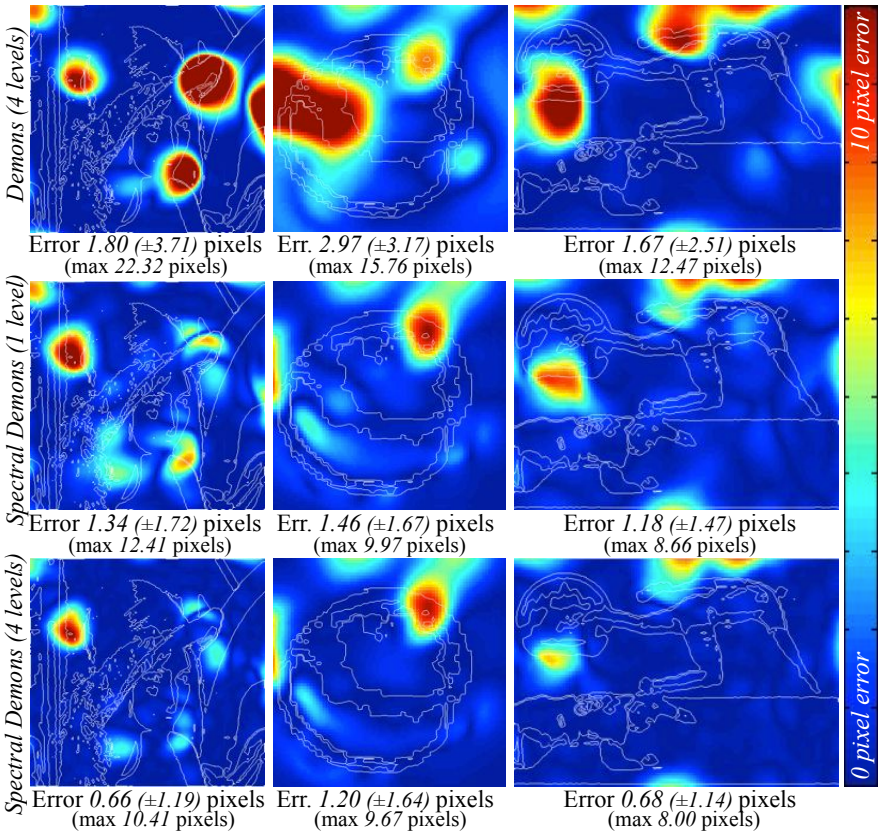

Figure 5.15: Differences (in pixels) between computed transformations and ground truth for (*top row*) Multilevel *Log-Demons*, (*middle row*) *Spectral Demons* and, (*bottom row*) Multilevel *Spectral Demons*. The fixed images are overlaid on the error maps. The Multilevel *Spectral Demons* decreases the error by 60% from the Multilevel *Log-Demons*.

their geometric equivalent, however the update forces derived from the image gradient lack any global information on the shape geometry and put the *Log-Demons* into an erroneous local minimum.

The use of a multilevel scheme allows the *Log-Demons* to capture larger deformations (Fig. 5.14) but does not change the inherent local scope of its update forces. For instance, *Log-Demons* even with 4 levels of resolution ultimately fails in recovering the extreme deformations on the anterior side of the heart, while *Spectral Demons* without a multilevel scheme can successfully register the whole myocardium with a 71% improvement in performance (mean square differences (or MSE) of intensities with ground truth from 19.9×10^{-3} to 5.7×10^{-3}). The performance is further improved when 4 levels are used (down to 2.4×10^{-3} MSE, or a 88% improvement in the heart registration). Similar results are observable with the other images. Lena's hat, her neck and, the player's arm are successfully registered using

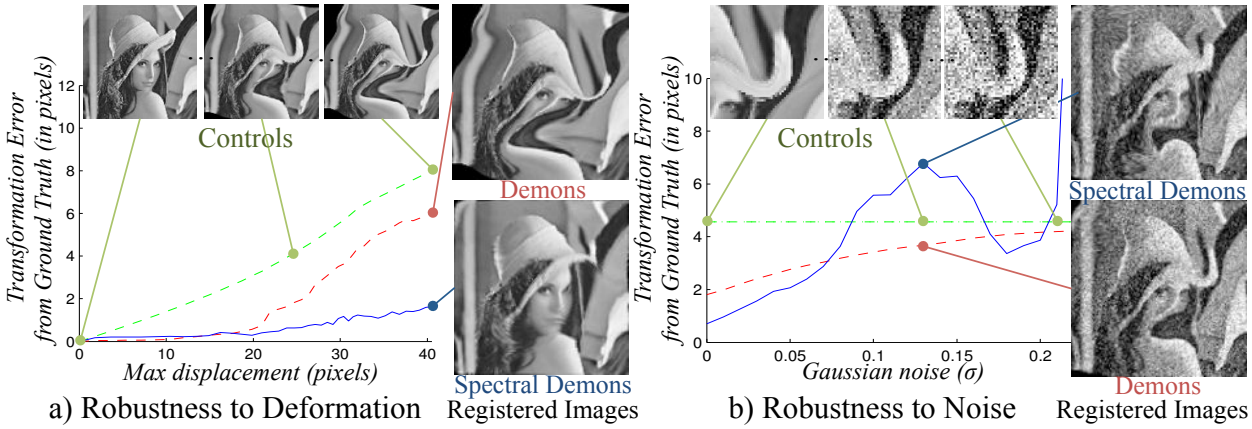

Figure 5.16: Robustness to deformation and noise. *a)* Deformations of Fig. 5.12 are amplified to a max of 40 pixels (image size is 128^2, deformed samples on top). The transformation differences with ground truth (y-axis in pixels) are smaller with *Spectral Demons* (*blue*) than with *Log-Demons* (*red*). *b)* Gaussian noise is used (samples around the hat area). *Spectral Demons* looses advantage after $\sigma > 0.075$, however even with a lower error, *Log-Demons* stops moving when increasing noise (error is similar with untransformed images (controls, in *green*)), whereas *Spectral Demons* continues to recover large deformations.

the *Spectral Demons* with respectively 73% improvement over *Log-Demons* in Lena's image and 63% in the baseball image.

Additionally, the quality of the computed registration maps ϕ is evaluated in terms of difference of displacements (in pixels) with the ground truth $\|\phi - \phi_{\text{truth}}\|$. The *Log-Demons* results in registration maps (Fig. 5.15) with larger errors in high deformation areas (e.g., Lena's hat or neck), whereas the *Spectral Demons* results in a smoother registration map with significantly less errors (62% less) in these same areas.

The cost of the global scope offered by *Spectral Demons* is increased computation time. For instance, on Lena's image, 50 iterations requires 108.49 seconds with *Log-Demons* and 201.43 seconds with *Spectral Demons*; on the heart image, 21.01 seconds with *Log-Demons* and 41.03 seconds with *Spectral Demons*; and on the baseball player's image, 42.06 seconds with *Log-Demons* and 102.17 seconds with *Spectral Demons*. We used unoptimized Matlab code on a Core 2 Duo, 2.53GHz. We now evaluate the robustness to deformation and noise of the *Spectral Demons* algorithm and show its applicability to real medical images.

Robustness to Deformation The *Spectral Demons*' robustness to deformation is evaluated on Lena's image by exaggerating the previous synthetic transformation from $\phi_0 =$

$\exp(0v)$ (zero displacements) to $\phi_2 = \exp(2v)$ (creating a maximal displacement of 40 pixels, see samples on Fig. 5.16). The performance is evaluated with the transformation differences (in pixels) from the ground truth. Both *Log-Demons* and *Spectral Demons* perform with sub-pixel accuracy on deformations below 20 pixels, however they differ with larger deformations where our method shows a greater robustness, e.g., with a deformation of 40 pixels (more than 30% of the image size) the average transformation error is 5.9 pixels with *Log-Demons* and 1.6 pixels with *Spectral Demons* (a 73% decrease).

Robustness to Noise The analysis on the robustness to noise reveals the current limitations of both algorithms. An increasing Gaussian noise is added to Lena's image, from $\sigma = 0$ to 0.25 (samples on Fig. 5.16). *Spectral Demons* performs better with noise $\sigma < 0.075$, however the comparison with control images (the unregistered noisy images) reveals that *Log-Demons* stops transforming the images when noise is increased (due to trapping in a local minimum). For instance, Fig. 5.16 shows that when using *Log-Demons* with noise $\sigma = 0.13$, the registered image is similar to its initial state, while *Spectral Demons*, even though with a larger average error (6.8 pixels versus 3.7 pixels), continues to recover large deformations (see the hat area on Fig. 5.16). With noise $\sigma > 0.2$, the corrupted images become problematic for *Spectral Demons* (graph edge weights are almost null and may need a different heuristic weighting), whereas *Log-Demons* is almost immediately trapped in a local minimum.

Registration of Medical Images

The performance of the *Spectral Demons* is evaluated in a medical application with the registration of brain MR images from the Internet Brain Segmentation Repository (IBSR, `http://www.cma.mgh.harvard.edu/ibsr`, our images are 64^3 volumes). The brain presents a wide variety of shapes and internal structures across individuals. While the cerebral cortex is particularly convoluted and is the focus of many specific surface matching techniques ((Fischl *et al.*, 1999; Yeo *et al.*, 2010a; Reuter, 2009; Lombaert *et al.*, 2011a)), the registration of internal components in the brain, such as the white or gray matter, requires a volumetric approach. We chose two individuals that have lateral ventricles with different sizes (Fig. 5.19, the moving image shows a longer ventricle). We evaluate the registration accuracy with the overlap of the provided manual segmentations of the white and gray matter (measured with the Dice metric defined as the ratio $2(A \cap B)/(|A| + |B|)$ with 1 being an optimal overlap) as

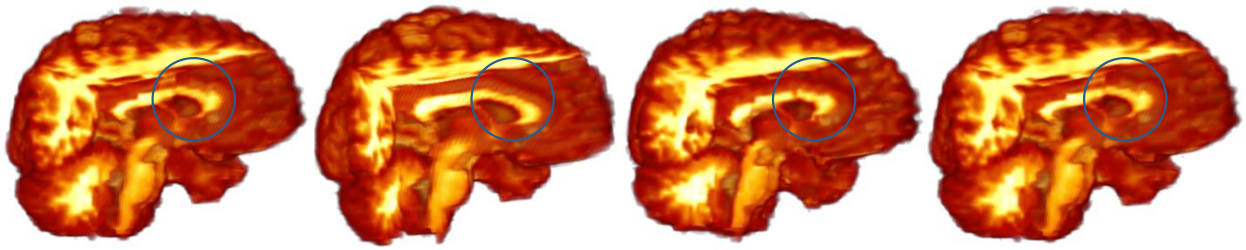

Fixed	Moving (Original)	Registered Moving (Demons)	Registered Moving (Spectral Demons)
Difference (MSE) with Fixed	*29.0 (±51.3) x10^{-3}*	*20.8 (±37.7) x10^{-3}* ↘	*19.9 (±36.2) x10^{-3}* ↘
Dice metric white matter	*0.6515*	*0.7126* ↗	*0.7166* ↗
Dice metric gray matter	*0.6931*	*0.7572* ↗	*0.7784* ↗

Figure 5.17: Registration in 3D between two brains from healthy subjects using *Log-Demons* and *Spectral Demons* (both with 4 levels). While observations show apparently similar results (ventricles are circled), the mean square differences of intensities (MSE) between the fixed and registered images reveal an increase of 4% in accuracy and precision when using *Spectral Demons*. The Dice metrics of the white and gray matter (measuring segmentation overlaps) also increase with *Spectral Demons*.

well as with the mean squared differences of pixel intensities between the fixed and registered images. The original (unregistered) setting has a Dice metric of 0.65 in the white matter and 0.69 in the gray matter. Both algorithms (using 4 levels) increase the overlap of the white and gray matter to respectively above 0.71 and 0.75, with a slight advantage to the *Spectral Demons*, however the comparison of the mean squared differences of intensities reveals a 4% improvement in accuracy and precision when using *Spectral Demons* (from 29.0×10^{-3} error in the original setting, decreasing to 20.8×10^{-3} with *Log-Demons* and to 19.9×10^{-3} with *Spectral Demons*). This experiment showed that *Spectral Demons* offers an improved performance in a real application when registering medical images.

The computation and the current implementation show again that there is room for improvements with our method. With downsampled images of size 32^3, 50 iterations require 108.49 seconds with *Log-Demons* and 201.43 seconds with *Spectral Demons*. Notably, memory becomes problematic with our unoptimized Matlab code as volumes beyond 32^3 require the decomposition of Laplacian matrices larger than $32^3 \times 32^3$ (although extremely sparse, our current code is not optimized for such large matrices).

Note on Computational Efficiency

A *fast Spectral Demons* algorithm could find a widespread use in many medical applications where speed is non negligible, notably, in interventional imaging where faster and more accurate registration and fusion of images would help the clinicians to understand what is happening inside a patient and would, therefore, help them guide their tools accordingly. The improvement of the computational efficiency of the *Spectral Demons* can be addressed through several strategies. Firstly, the algorithm could be changed to avoid a spectral decomposition in each iteration. For instance, similar to the levelset frameworks where its distance transform is reinitialized once every few iterations, the initial eigenmodes of the *Spectral Demons* could be reused and warped during a few iterations before being reinitialized. This is a valid scheme since the underlying graphs of the images in two successive iterations are assumed isometric, i.e., their Laplacian matrix should theoretically have the same set of eigenvalues and eigenmodes, and in practice, small image perturbations between iterations may slightly change the eigenmodes and could, therefore, be assumed negligible between a few iterations. Secondly, the spectral decomposition could benefit from approximation approaches. For instance, the Nyström method (Drineas and Mahoney, 2005) removes random lines and columns of the Laplacian matrix before the decomposition and approximates them later; therefore, the computation of the eigenmodes is faster due to this smaller Laplacian matrix. Moreover, the sparsity of the Laplacian matrix as well as the order of its lines and columns could be rearranged in order to yield a more efficient decomposition (Li *et al.*, 2011). Thirdly, the current implementation could greatly benefit from the use of compiled programming languages, e.g., C++, or parallel computing, e.g., GPU. The *Spectral Demons* algorithm, besides being more accurate than the conventional Demons, has, therefore, the potential of being computationally efficient.

Construction of Cardiac Atlases

After showing and discussing the fundamental differences and substantial improvements of the *Spectral Demons* algorithm with conventional *Log-Demons*, its use within the new groupwise framework is presented. More precisely, the construction of atlases is evaluated with images exhibiting large deformations. In a synthetic experiment, convergence toward an average shape is verified with both groupwise variants, the *GL-Demons* and *GSL-Demons* framework. The parameters are, therefore, similar in both variants: $\sigma_{\text{fluid,diff}} = 1, \alpha_x =$

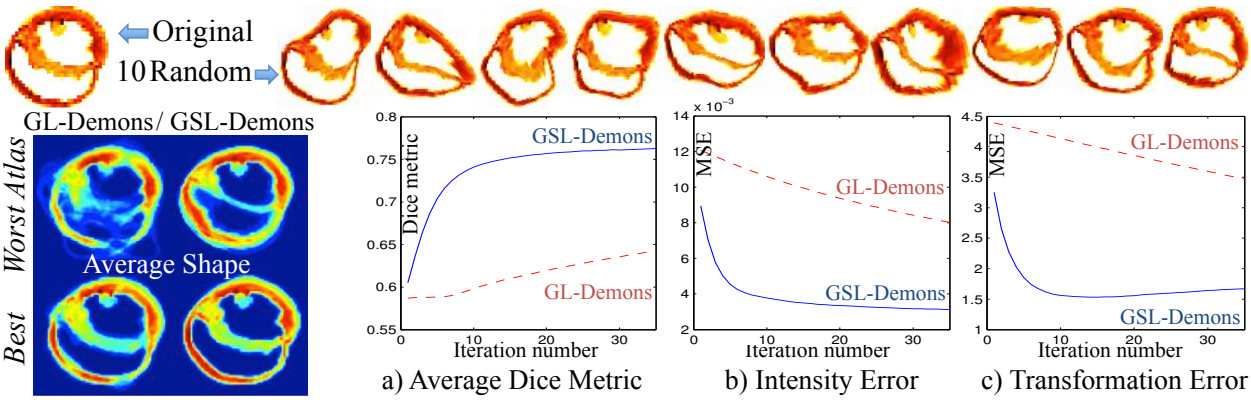

Figure 5.18: Groupwise registration of 10 images deformed randomly (100 trials, 1 sample on top row, with known ground truth) using *GL-Demons* and *GSL-Demons*, *Left)* Best and worst atlases (based on Dice metric among 100 trials) demonstrating the capability of the *GSL-Demons* to handle large deformations, *a)* Average Dice metric with ground truth, *b)* Intensity difference between average shape and ground truth, *c)* transformation error with ground truth. *GSL-Demons* converges faster toward the average shape.

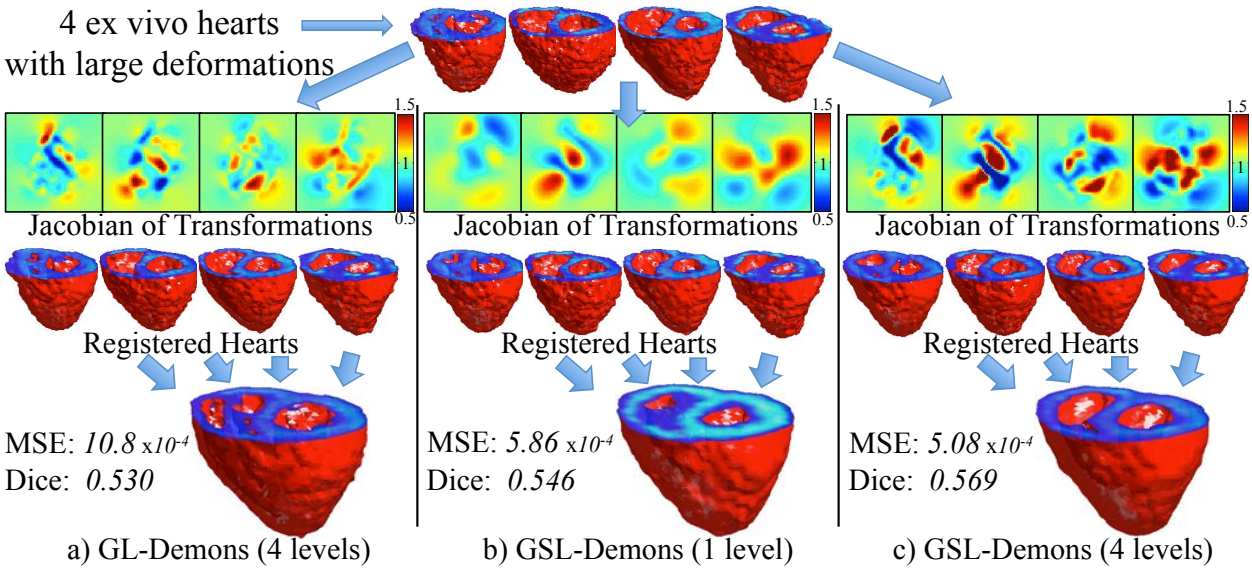

Figure 5.19: Atlas of *ex vivo* hearts (isosurfaces are shown) using *a) GL-Demons* (4 levels, showing failure in the right ventricle), *b) GSL-Demons* (1 level), *c)* and *GSL-Demons* (4 levels, with correct right ventricle). *GSL-Demons* capture successfully large deformations. Jacobian determinants (axial planes) show that spectral matching capture smooth and large deformations while gradient-based updates capture local deformations.

$1, k = 5, \alpha_g = 0.1, \alpha_s = 0.2, \alpha_i = 0.7$ in 2D). In a second experiment, both frameworks are tested with real cardiac images that exhibit a high shape variability (parameters: $\sigma_{\text{fluid,diff}} = 0.75, \alpha_x = 1, k = 5, \alpha_g = 0.25, \alpha_s = 0.35, \alpha_i = 0.4$ in 3D).

Synthetic deformations Convergence and capture of large deformations are now evaluated. $N/2$ velocity fields v are generated randomly using 15 control points with random locations in the image and random displacements of at most 15 pixels (20% of the image size) that are diffused over the image. Their forward and background transformations ($\exp(v)$ and $\exp(-v)$) are applied to an initial image I_0, holding thus the average shape to I_0 (establishing our ground truth). Since we compare the convergence and its rate, and not the final performance, the multi-level scheme (which should be used in real applications) is not applied. Fig. 5.18 shows the groupwise registrations of 10 random hearts (2D 75×75 images) through 100 trials (a total of 1000 hearts). The average Dice metric (measuring the overlap) between all computed average shapes and I_0 as well as the intensity errors (MSE) reveal that the reference shape (defined arbitrarily as one of the 10 images) evolves toward the ground truth (i.e., Dice increases and MSE decreases). Moreover, the N deformation fields become closer to the ground truth during registration. The striking difference in the convergence rates shows the full power of GSL-Demons (less than 5 iterations are required) while *GL-Demons* might not converge with such large deformations (we stopped the algorithms after 200 iterations). Time-wise, 35 iterations take 194 seconds with GSL-Demons, and 53 seconds with *GL-Demons* (using unoptimized Matlab code on a 2.53GHz Core 2 Duo). GSL-Demons shows a better performance with high deformations than *GL-Demons*.

Cardiac Atlases We now evaluate the construction of atlases with organs of high shape variability. *Ex vivo* hearts are particularly challenging to register as they present a high variability in fixture poses due to flabby ventricular walls. The human *ex vivo* DTMRI dataset (Rapacchi *et al.*, 2010; Lombaert *et al.*, 2011b) provides good candidates to evaluate our algorithms. We use four hearts ($b = 0$ images of size 64^3) that were excluded in the construction of the human atlas (Lombaert *et al.*, 2011b) due to their hypertrophy and highly deformed shapes (see Fig. 5.19). GL-Demons (with 4 resolution levels) fail in recovering the shapes of the right ventricles, while GSL-Demons successfully constructs the atlas even with 1 level of resolution (downsampled images at size 28^3). As a comparison, 35 iterations takes 40 minutes in Matlab with *GSL-Demons* and 9 minutes with *GL-Demons*. Using *GSL-*

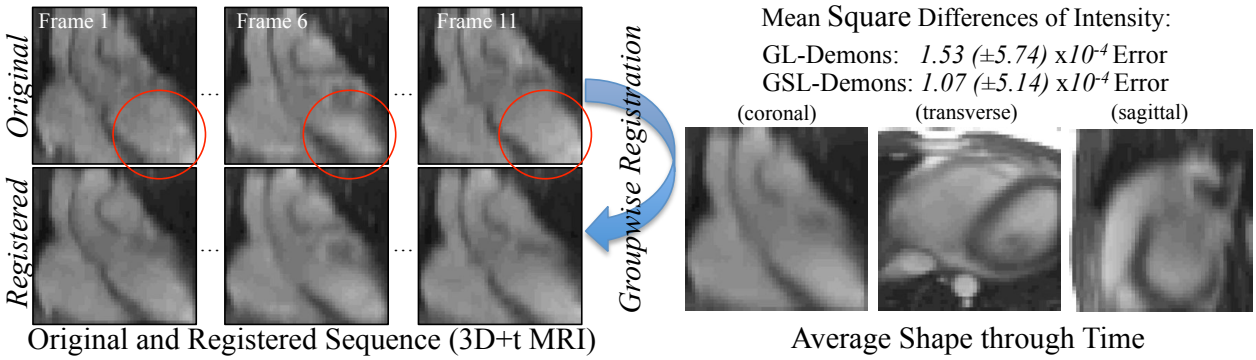

Figure 5.20: Atlas (average shape) from 11 frames of a 3D+t cardiac MRI sequence. *(Left)* Original 3D frames before and after groupwise registration (contraction/expansion is corrected). *(Right)* Computed atlas with the average shape (only the GSL-Demons version is shown). The intensity error (between all registered frames and the atlas) is reduced by 30% with GSL-Demons.

Demons with 4 resolution levels reduce the intensity error (MSE) by half (from 10.8 to 5.08). Moreover, the Jacobian determinants of the transformation fields show that the large and highly non-local deformations are successfully captured with the spectral-based update scheme (high and smooth Jacobian in Fig. 5.19 b) while local deformations are captured with the gradient-based update scheme in the higher levels of *GSL-Demons* (Fig. 5.19 c). A second cardiac atlas (Fig. 5.20) is constructed from a 3D+t MRI sequence (11 frames of 64^3 with high systolic deformations on frame 6) as an example to show that GL and GSL-Demons can be successfully used on denser images and not only on segmented hearts. The study of the *in vivo* cardiac motion is undermined with registration challenges. Both algorithms performing successfully perform a groupwise registration, even with systolic configurations (frame 6).

5.5 Conclusion

In this additional chapter, we have extended the use of spectral methods to construct atlases in two manners. *Firstly*, FOCUSR, the surface matching method developed in Chapter 4, has been shown capable of constructing atlases in concrete applications. For instance, an atlas of the human cerebral cortex has been constructed from 16 brain surfaces. The atlas reveals the average shape of the human cortex and allows a shape analysis of the principal modes of variations across the dataset. *Secondly*, in order to explore the shape of the primary cortical folding, which is relevant to studies on the brain development, FOCUSR

has been used to align a large number of cortices with varying degrees of smoothing. In fact, the speed advantage of FOCUSR allowed the matching of 14,400 pairs of brain surfaces, which would have been prohibitively expensive with current brain matching methods in the state-of-the-art. These experiments on brain surfaces provided basic tools for future studies on shape analysis as well as preliminary results on the exploration of the primary cortical folding pattern. *Lastly*, the direct feature matching technique and the spectral representations developed in FOCUSR served as basis to a fundamentally new approach for image registration. This new approach can be used within classical frameworks such as the *Log-Demons*, yielding in this case the *Spectral Demons* algorithm. It is capable of capturing very large, complex and highly non-local deformations between images, in fact, displacements of more than 30% of the image size were recovered and the registration accuracy even showed an improvement of 73% over the conventional *Log-Demons* algorithm. This substantial improvement is rendered possible by the spectral representation, which captures a *global* geometric description of the underlying Riemannian structure of images. This is in contrast with classical approaches for image registration, which rely on update schemes limited by the *local* scope of forces derived from image gradients. Additionally, the Demons algorithm has been extended to perform groupwise registration in parallel with the computation of the average shape, i.e., the atlas is constructed during the registration process. This new groupwise framework is versatile enough to be used with any variant of the Demons algorithm, including our newly developed *Spectral Demons*, which yields the *Groupwise Spectral Demons* (or *GSL-Demons*). The conducted experiments used 1000 random hearts severely deformed and showed that the computation of atlases converges toward an average shape. We additionally showed that *GSL-Demons* could construct an atlas for a challenging dataset of *ex vivo* hearts with high shape variability.

To summarize, this chapter provided additional material based on spectral correspondence. It showed that FOCUSR can be used to quickly construct atlases of brain surfaces (enabling consequently new studies previously limited by computational burden), and that spectral representations can also be used with images in order to construct atlases of organs with very large deformations. The next chapter provides a general discussion that establishes links between the tools developed in the thesis.

PART III

CONCLUSIONS

CHAPTER 6 GENERAL DISCUSSION

The previous chapters aimed at answering a few limitations of existing methods for atlas construction. Two aspects were specifically studied, *a)* how to automate the atlas construction and *b)* how to handle very large deformations in a dataset. The general methodology established three research objectives that led to several original tools and frameworks for atlas construction. *Firstly*, an automated framework was developed in order to analyze the cardiac fiber architecture, *secondly*, a new approach for finding correspondences that involves spectral matching has been investigated, and *thirdly*, these new findings allowed the creation of an original framework for accurate image registration that is capable of handling very large deformations. The development of these new tools, summarized in Fig. 6.1, are now discussed in the following sections with a global perspective on their general advances made in the state-of-the-art of surface matching and image registration, as well as with a focus on their implications and limitations. Additionally, these new methods and frameworks are shown to be more than a series of individual tools. The general discussion highlights the general trends arising from their development, and also relates them with previous work.

6.1 Surface Matching

6.1.1 Spectral Correspondence with FOCUSR

Two types of medical data have been studied during the thesis: surfaces and images. Surfaces are particularly adapted for modeling, for instance, membranes, valves or organ boundaries. These examples may, however, undergo complex deformations or shape variability that may render their matching difficult, and this is no exception to cerebral cortical surfaces studied in this thesis. Historically, surface correspondences were found by slowly deforming their models until being matched (Terzopoulos *et al.*, 1987; Terzopoulos, 1980; Mcinerney and Terzopoulos, 1996). These active models are often limited by the local scope of the forces guiding their control points. This limitation led to the development of FOCUSR in Chapter 4, which is fundamentally different than classical approaches for surface matching since it avoids the difficult deformation of models and rather aims at directly establishing correspondences between surfaces. This is in comparison with FreeSurfer (Fischl *et al.*, 1999) or Spherical Demons (Yeo *et al.*, 2010a), leading solutions for brain surface matching, which

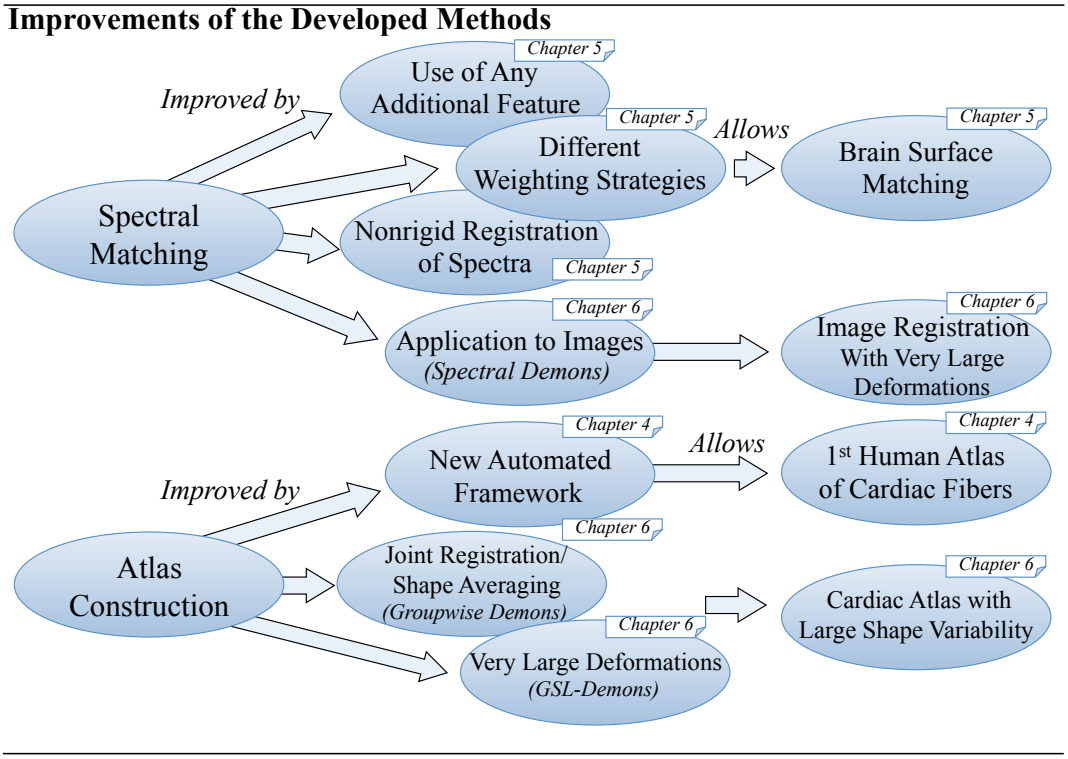

Figure 6.1: Improvements made during the thesis. Spectral Matching has been improved by introducing an original direct feature matching technique that can incorporate virtually any information capable of helping matching. It is implemented by FOCUSR (Chapter 4) for brain surface matching and by *Spectral Demons* (Chapter 5) for image registration. Both algorithms allow consequently the capture of very large shape variability. Additionally, two new frameworks for atlas construction provide new contributions in the state-of-the-art. The first framework allowed the construction of a human atlas of the cardiac fiber architecture, while the second framework enables the computation of the average shape during the registration process and can incorporates the previous *Spectral Demons* algorithm in order to build atlases with very large shape variability.

finds correspondences on inflated brain surface models. These solutions, although accurate, have a prohibiting computational cost associated with the inevitable handling of surface distortions during their inflation.

The originality of FOCUSR in brain surface matching remains in its use of global geometric descriptors captured with spectral graph representations. This underlying idea has, in fact, been investigated in the past (van Kaick *et al.*, 2011; Zhang *et al.*, 2010) but has never been successfully used for dense surface correspondence. For instance, only hierarchical structures are matched in (Reuter, 2009; Leordeanu and Hebert, 2005; Shi *et al.*, 2010), i.e., correspondences are established between large areas and not at the vertex level. Since

FOCUSR uses similar graph spectral representations, it may resemble other spectral matching methods such as (Mateus *et al.*, 2008; Carcassoni and Hancock, 2003; Jain and Zhang, 2006), however, there are several major differences. Notably, FOCUSR improves spectral correspondence with an efficient nonrigid deformation of spectral graph representations. This is, for instance, different than previous spectral approaches (Mateus *et al.*, 2008) that simply rely on a rigid transformation or (Jain and Zhang, 2006) that are constrained by their choice of nonrigid deformation model. The basic principle of FOCUSR is also different than conventional approaches since it establishes correspondences with a direct feature matching technique that can use virtually any additional information capable of helping matching. Spectral representations are, in fact, used as regularization in a multidimensional feature space. This different approach enables dense surface correspondence at the vertex level and is one explanation on why the results of Chapter 4 showed a substantial improvement in accuracy over a simple spectral matching approach. Furthermore, additional information may also be incorporated in FOCUSR as weighting on graph nodes. This original weighting strategy may be beneficial for all spectral graph matching methods since vertices can be efficiently activated or deactivated depending on the application, e.g., graph node weighting can be derived from pointwise measurements or can model vertex relevance from specific statistics.

Additionally, FOCUSR was used in Chapter 5 to construct an atlas of the human cerebral cortex. The variability of the cortical folding pattern, studied in (Ono *et al.*, 1990; Welker, 1990), and of functional areas, studied in (Fischl *et al.*, 1999; Lohmann *et al.*, 2008), are still not fully understood in humans. Their studies typically require a prohibitively large number of surface matching, however, the speed advantage of FOCUSR is beneficial for such neuroimaging studies and enables new experiments such as the ones conducted in Chapter 5 on the anatomical and functional variability of the cortical surface.

6.2 Image Registration

The tools developed during the thesis also addressed the registration of images and more precisely of cardiac MRI images. Three main approaches have been elaborated: the use of simplified images within a conventional registration framework, the application of the new spectral approach to medical images, and a joint approach for registration and shape averaging in order to construct atlases.

6.2.1 Framework for Atlas Construction

The first approach for image registration is presented in Chapter 3. It elaborates a new framework for atlas construction where image registration is facilitated with an initial labeling of major cardiac structures. These simplified images enable the automation of the atlas construction and minimize user interaction. This is in contrast with previous frameworks (Helm *et al.*, 2005b; Peyrat *et al.*, 2007) that often assume a high contrast of images, and where simple thresholding of pixel intensities is generally sufficient. The labeling of images is computed with a developed method based on Graph Cuts (Boykov and Jolly, 2000), that respects the topology of cardiac structures (i.e., the myocardial wall contains no hole). Furthermore, the registration is simplified since these labeled images prevent the alignment of ambiguous configurations of images (i.e., information is digested in images to a few labels, facilitating correspondences between images). The alignment of the internal structures within each label is performed with an additional nonrigid registration of a masked image of the myocardium. This two-level strategy for registration ensures, therefore, that the coarse and large deformations are handled by the alignment of simplified images, and that the fine deformations are handled with a final registration based on image intensity. Moreover, the new framework also improves the previous approach used with a canine dataset (Helm *et al.*, 2005b) by using a more up-to-date registration algorithm. The symmetric Demons algorithm (Vercauteren *et al.*, 2008, 2009a) was chosen, as it is well adapted for averaging shapes.

The new framework has been applied to the construction of the first human atlas of the cardiac fiber architecture, which is relevant to the study of various cardiac mechanical functions (Costa *et al.*, 2001), of cardiac electrophysiology patterns (Hooks *et al.*, 2002), and of remodeling processes (Wu *et al.*, 2006). The results showed that the variability of the fiber structures in humans concurs with studies on other species such as on dogs (Helm *et al.*, 2005a,b; Helm, 2005; Sundar *et al.*, 2006; Peyrat *et al.*, 2007; Peyrat, 2009; Gilbert *et al.*, 2007), goats (Geerts *et al.*, 2002) and rats (Bishop *et al.*, 2009). A preliminary experiment in Appendix B also shows that the fiber architecture may characterize certain cardiomyopathies. Additionally, the supplemental study in Appendix A suggests that two populations of laminar sheets may appear in several myocardial segments. These findings may be compatible with the intriguing and controversial (von Segesser, 2005) theory of the Torrent-Guasp model (Torrent-Guasp *et al.*, 2005) that considers the whole cardiac myocardium as a single band folding onto itself, i.e., this folding might explain the presence of two configurations of fiber

structures.

6.2.2 Spectral Demons

The second method, developed in Chapter 5, introduces a fundamentally new approach for image registration. It is based on the direct feature matching technique elaborated in Chapter 4 with FOCUSR. As mentioned earlier, the underlying spectral representations (Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992) of this approach enable the capture of large deformations between images, i.e., matching is facilitated because objects with significantly different configurations would still share similar spectral representations. The developed algorithm, named *Spectral Demons*, benefits from using spectral correspondence (Mateus *et al.*, 2008; Carcassoni and Hancock, 2003; Jain and Zhang, 2006) since very large deformations are naturally handled. The global geometric description provided by spectral representations actually contrasts with the local nature of the forces underlying conventional approaches for image registration. Although elegant mathematical frameworks (Beg and Khan, 2006; Bossa *et al.*, 2007; Durrleman *et al.*, 2009b, 2008, 2011) and computationally efficient frameworks (Cachier *et al.*, 2003; Vercauteren *et al.*, 2007, 2008) exist in the literature, these forces are often derived from image gradients whose local scope limits the registration to local deformations, i.e., the optimization remains local even in a multilevel scheme. The use of a spectral correspondence approach within these conventional registration frameworks provides, in fact, a missing link between image registration and the capture of very large and highly non-local deformations. For instance, by extending the Demons algorithm (Thirion, 1998; Vercauteren *et al.*, 2009a), the results showed that *Spectral Demons* substantially improves the matching accuracy over the conventional approach. The cost of such improvement is, however, an increased computational burden due to the spectral decomposition, i.e., computing eigenvectors is more costly than approximating image gradients.

6.2.3 Groupwise Registration

The third approach for image registration is presented in Chapter 5. It extends the *Log-Demons* algorithm to perform groupwise registration. In typical approaches for atlas construction (Guimond *et al.*, 2000; Helm *et al.*, 2005b), the shape averaging process is indeed limited by sequential pairwise registrations, i.e., it alternates, until convergence, between *i)*

the registration of each image with a reference and *ii)* the transformation of the reference image. The new approach, named the *Groupwise Demons* framework, computes the shape average during the registration process and removes, therefore, the need of the previous alternation scheme. A joint approach for image registration and shape averaging is also proposed in (Avants and Gee, 2004a; Bossa *et al.*, 2007), however both methods use more complex diffeomorphic frameworks (Beg *et al.*, 2005; Beg and Khan, 2006; Bossa *et al.*, 2007) that may be slow and limited to local deformations.

This new groupwise framework is declined in two variants. The first version, named the *Groupwise Log-Demons* (or *GL-Demons*), extends the symmetric *Log-Demons* to use multiple images, while the second version, named the *Groupwise Spectral Log-Demons* (or *GSL-Demons*), extends the *Spectral Demons* algorithm. The results showed that *GSL-Demons* is more robust to shape variations than *GL-Demons* and can construct cardiac atlases with only 5 iterations. This faster convergence rate compensates, therefore, for the increased computational burden inherent to the spectral decomposition. Moreover, when deformations between images become significantly large, conventional approaches based on image gradients remained often trapped in local minima (Joshi *et al.*, 2004; Avants and Gee, 2004a; Marsland *et al.*, 2003; Guimond *et al.*, 2000), whereas *GSL-Demons* correctly captured very large deformations.

6.3 Common Trends of the Developed Methods

The handling of large shape variability is a recurrent concern when matching medical data. The typical response found in the literature (Zitova, 2003; Crum *et al.*, 2004) is to use a Euler-Lagrangian approach where complex deformations of objects are handled with small and tractable updates of an iterative transformation.

All the developed methods discussed in this chapter are capable of finding correspondences between organs with significant shape variability. One interpretation is that these methods take advantage, in one form or another, of a *global* aspect in their optimization schemes. For instance, the direct feature matching technique, used in both FOCUSR and *Spectral Demons*, is in fact a *global* optimizer since it finds best nearest-neighbors in a multi-dimensional feature space. Moreover, spectral representations, used again in both FOCUSR and *Spectral Demons*, are effectively *global* geometric descriptors, since they provide unique signatures between corresponding points and render matching much less sensitive to de-

formation. Additionally, even though the framework developed in Chapter 3 is based on conventional Demons, and therefore, inherently limited by the local scope of its gradient descent optimizer, the labeling of the simplified images captures as well a *global* description of organs by providing the locations of its main components. Overall, the investigation of these different strategies provides new contributions on how to find and exploit global descriptions of images and surfaces.

From a global perspective, these new tools actually close a loop between *a)* the first set of tools, developed in Chapter 3, that highlights the challenges in registering images with very large deformations, *b)* the second set of tools, developed in Chapter 4, that captures very large deformations between surfaces but were not applicable to images, and *c)* the third set of tools, developed in Chapter 5, that comes back on treating images by allowing a natural construction of atlases from images with very large deformations.

6.4 General Limitations

Notwithstanding the advantages of each individual contribution with their identified extents, there remain general limitations that were not fully investigated. A few of them were previously identified in the thesis and are now discussed in order to find their general impact.

Complete framework: Perhaps, the most important limitation of the developed tools is that they are not proposing a completely automatic framework for atlas construction. Indeed, they still depend on a *preliminary segmentation* before constructing the atlas. This process remains a separate step that still requires user interaction. However, as noticed during the development of *Spectral Demons*, segmentation may naturally arise from spectral graph theory. Further links between spectral correspondence and segmentation were not investigated and may lead to a complete framework where automatic segmentation is an inherent part of the atlas construction.

Partial data: Additionally, there is a strong assumption in spectral methods that the graphs being matched must represent complete, or similar, objects, i.e., the topology must be equivalent across graphs. The developed tools are, therefore, currently not adapted to handle *occlusions* or *missing parts* since they may change topology and global shape. Such problems may occur when, for instance, imaging devices have a limited field of view, e.g., the left, or right part of the heart may be omitted in images or organs may not be completely visible during acquisitions.

Computational burden: From a practical point of view, the developed spectral methods are currently burdened with the inevitable *spectral decomposition*. Even though they are compensated with a fast convergence rate and the capability of handling very large deformations, this computational burden may be problematic for their widespread use in medical context where very large datasets are the norm, and where fast registration is required, notably in interventional imaging.

Multimodal data: Moreover, the developed tools are designed for specific types of data, i.e., they can match either surfaces or images, but not both. The current formalism of the tools does not allow for *multimodal data* to be registered. To this effect, the strategy of using spectral representations may lead to a fundamentally new similarity measure that would be independent of the type of used data. Spectral representations describe indeed the global geometry of data and, with further improvements, may be matched regardless of the type of data, either surfaces, images, both, or any other type of data expressible with graphs.

Spatio-temporal data: On the same line, another important general limitation remains in the fact that *spatio-temporal data* (e.g., 3D+t images) is not considered as a whole but as separate data fixed in time. The temporal aspect is, therefore, ignored in the developed tools and prevents a natural description of global characteristics over space and time. Such proper characterization may be relevant to study, for instance, the variability in growth or motion. This would require investigations on how to model adequately spatio-temporal data with graphs, and above all, how to decompose them into meaningful spectral representations.

Functional studies: On a different note, the tools developed in the thesis enabled new findings relevant to cardiology and neurology, however, the investigation did not go further by using these findings for modeling or improving comprehension of *complex cardiac or neural functions*. For instance, the complete description of the cardiac fiber architecture could be used to simulate and predict cardiac mechanics or electrophysiology, which are active fields of research and directly relevant for diagnostics, treatments, and follow-ups.

From these identified general limitations, it is clear that further work remains to be done; nevertheless, the contributions of this thesis make an important step toward a better understanding of human anatomy. This general discussion provides ground for promising future research studies, from which the most important ones are recommended in the next chapter.

CONCLUSION

This thesis has addressed the general problem of atlas construction by exploring particular problems associated with the shape variability of complex characteristics of the human anatomy. The literature, reviewed in Chapter 1, highlighted the challenges of finding correspondences between highly deformed images and also revealed the current limitations of the state-of-the-art. During the thesis, a set of tools has been developed in order to normalize measurements of complex characteristics. More precisely, the methodology in Chapter 2 led to, in a first step (Chapter 3), a new automated framework for atlas construction that enabled new studies on the complex cardiac fiber architecture, in a second step (Chapter 4), a new improved surface matching method based on spectral correspondence that enabled studies on the anatomical and functional variability of the brain surface, and in a third step (Chapter 5), a fundamentally new approach for image registration based on the findings from the earlier steps. The results and contributions from these research objectives have been discussed individually as well as from a global perspective where light is shed on their current limitations in Chapter 6. The main findings are briefly reminded in Fig. 7.1 and Fig. 7.2, where they are categorized in two groups. The first set of contributions are new methods that address the problem of constructing complex atlases, while the second set of contributions are results from variability studies on the constructed atlases. The next sections now establish the main recommendations based on the contributions found in the thesis.

Recommendations and Perspectives

The new methods developed during the thesis improve the construction of atlases by capturing larger deformations between images. Their advantages and disadvantages were discussed in Chapter 6. Here, their limitations are addressed through four major recommendations, which give the main lines and main research questions for future work.

Complete Spectral Framework for Atlas Construction

Recommendation 1: *Improve the Groupwise Spectral Log-Demons algorithm to be a complete automatic framework for atlas construction.* ■

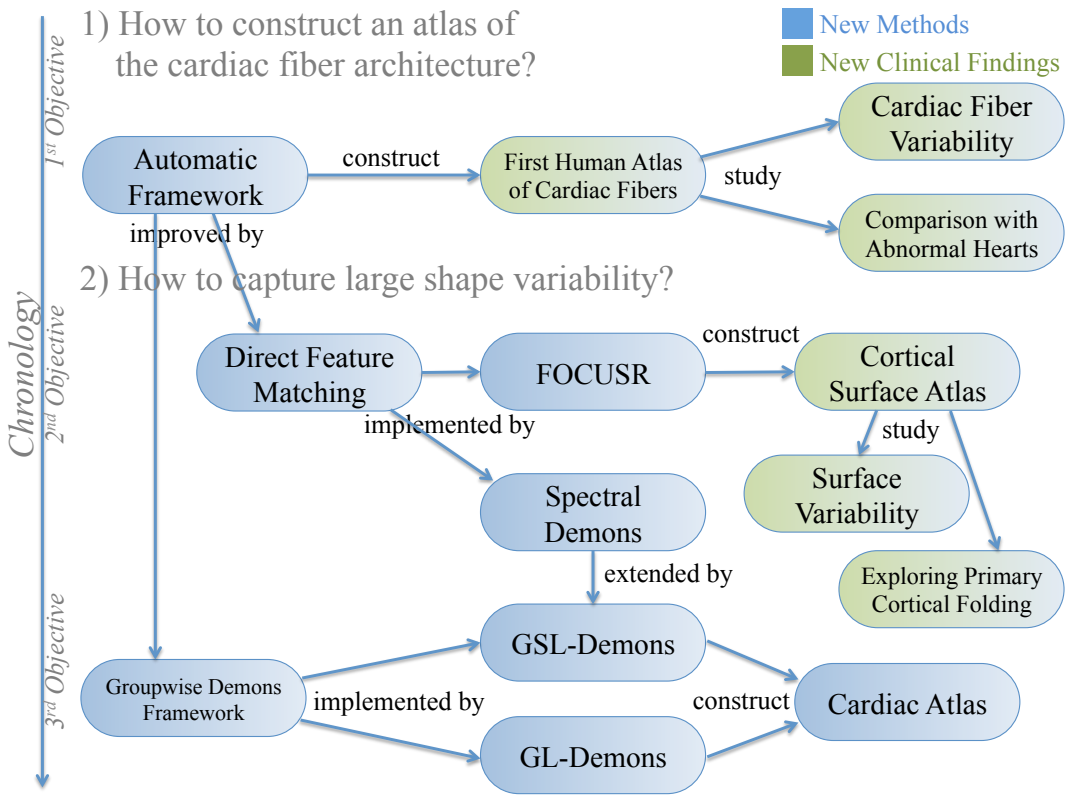

Figure 7.1: Establishing links between the main findings of the thesis (new methods in blue, new clinical findings in green). *Firstly*, an automated framework has been developed using conventional Demons and allowed the construction of the first human atlas of the cardiac fiber architecture (i.e., this answers the first research question). *Secondly*, a new approach based on a direct feature matching technique that incorporates spectral representations has been developed in order to capture a large shape variability (i.e., this answers the second research question). Two implementations has been done, FOCUSR for surfaces, and *Spectral Demons* for images, and additionally, a groupwise registration framework has been developed in order to construct more precise cardiac atlases.

The first objective led to the development of a complete framework for atlas construction, including three main steps: segmentation, registration and shape averaging. The two last objectives of the thesis led to the creation of a new framework for atlas construction where registration and shape averaging are performed within the same process, however, segmentation still need to be done beforehand. **How to create a complete framework for atlas construction based on the *Spectral Demons*?** It has been found that there are strong links between the *Spectral Demons* algorithm and graph-based image segmentation methods. In fact, their underlying energies are even similar. The segmentation aspect of the *Spectral*

New Tools for Spatial Normalization

Automated framework for atlas construction — A complete framework where hearts are segmented and registered in order to compute their average shape.

FOCUSR — A new method for matching brain surfaces as well as general meshes with high shape variability.

Spectral Demons — A fundamentally new approach for image registration where very large deformations can be captured.

Groupwise Log-Demons — A new framework for atlas construction where the average shape is computed during the registration process.

Groupwise Spectral Log-Demons — The previous framework improved with the new spectal approach in order to construct atlases with very large deformations.

New Results on Human Anatomy

First human atlas of the cardiac fiber architecture — The atlas provides the average structures of the fiber architecture with their normal variations.

Variations of the cardiac fibers — Local and global descriptions of the fiber directions.

Variations of the cardiac laminar sheets — Local and global descriptions of the laminar structures.

Comparison with healthy and abnormal hearts — The fiber architecture may be related to abnormality.

Exploration of the primary cortical folding — Smoothing the surface of a developed brain may reveal the primary cortical folding pattern.

Figure 7.2: Summary of the key contributions categorized as new methods for spatial normalization and new results on human anatomy.

Demons algorithm could be, therefore, further studied in order to yield a joint segmentation and registration method. Using this enhanced *Spectral Demons* algorithm in the groupwise demons framework could lead to a *complete* framework for atlas construction where segmentation, registration, and shape averaging are all interlinked and could potentially ensure an improved spatial consistency in the atlas.

Fast Spectral Demons

Recommendation 2: *Develop a new optimization scheme for* Spectral Demons *where spectral decomposition is efficiently utilized.* ■

The *Spectral Demons* algorithm and its groupwise extension have been shown to have a substantial improvement in accuracy when large deformations are present between registered

images. However, the main limitation is an increased computational burden due to a costly spectral decomposition. **How to overcome the computational burden of the *Spectral Demons* algorithm?** A *fast Spectral Demons* algorithm could find a widespread use in many medical applications where speed is non negligible, notably, in interventional imaging where faster and more accurate registration, as well as fusion, of images would help clinicians to understand what is happening inside a patient and would, therefore, help them guide their tools accordingly. The improvement of the computational efficiency of the *Spectral Demons* can be addressed through several strategies. *Firstly*, the algorithm could be changed to avoid a spectral decomposition in each iteration. For instance, similar to the levelset frameworks where its distance transform is reinitialized once every few iterations, the initial eigenmodes of the *Spectral Demons* could be reused and warped during a few iterations before being reinitialized. This is a valid scheme since the underlying graphs of the images in two successive iterations are assumed isometric, i.e., their Laplacian matrix should theoretically have the same set of eigenvalues and eigenmodes, and in practice, small image perturbations between iterations may slightly change the eigenmodes and could, therefore, be assumed negligible between a few iterations. *Secondly*, the spectral decomposition could benefit from approximation approaches. For instance, the Nyström method (Drineas and Mahoney, 2005) removes random lines and columns of the Laplacian matrix before the decomposition and approximates them later; therefore, the computation of the eigenmodes is faster due to this smaller Laplacian matrix. Moreover, the sparsity of the Laplacian matrix as well as the order of its lines and columns could be rearranged in order to yield a more efficient decomposition (Li *et al.*, 2011). *Thirdly*, the current implementation could greatly benefit from the use of compiled programming languages, e.g., C++, or parallel computing, e.g., GPU. The *Spectral Demons* algorithm, besides being more accurate than the conventional Demons, has, therefore, the potential of being computationally efficient.

Spectral Matching for Partial and Multimodal Data

Recommendation 3: *Improve spectral matching methods in order to handle partial and multimodal data such as truncated surface and occluded images.* ■

The development of the spectral methods in this thesis did not address the problem of partial data and multimodal data. More precisely, surfaces and images are treated sepa-

rately with FOCUSR and *Spectral Demons*, which assume no truncation in surfaces and no occlusion in images. However, in many real scenarios, computed surface models of organs may require truncation with arbitrary boundaries (e.g., surface imaging of scoliotic trunks usually shows no arms, legs or head), occlusions may appear in images (e.g., X-Ray imaging often shows occlusions from lead protective covers), and sometimes both types of data may need to be registered (e.g., matching the surface model of an organ within its related MRI volume). Furthermore, ablations may also create similar limitations when a missing area changes significantly the organ topology. **How can spectral matching methods handle partial and multimodal data?** To this effect, three strategies can be explored. *Firstly*, the registration process may minimize an energy where missing data is modeled as a separate variable. For instance, similar to (Mateus *et al.*, 2008), points with no significant correspondence could be labeled as being unmatched and a cost function could ensure spatial regularity of the missing areas. Local affine schemes (Seiler *et al.*, 2011) could also ensure additional spatial regularity. *Secondly*, user interaction may provide additional information for the registration. For instance, a user could indicate the missing areas or pairs of corresponding points. The computation of eigenmodes, or more generally of geometric descriptors, could take into consideration such user inputs. One possible approach may be inspired from the Heat Kernel Signature (Sun *et al.*, 2009; Bronstein and Kokkinos, 2010), where, for instance, the geometry could be described by heat propagation patterns generated from corresponding points selected by the user. *Thirdly*, multimodal data may be matched by developing a common similarity measure based on spectral representations of data. Indeed, spectral representations finds a global shape description regardless of the type of data as long as it is expressible with a graph, and may be consequently used to find correspondences between surfaces and images. This recommendation may, therefore, improve spectral methods to handle naturally partial and multimodal data.

Improved Cardiac Electromechanical Models

Recommendation 4: *Use the new findings on the human cardiac fiber architecture to build new accurate electromechanical models of the heart.* ■

The construction of the first human atlas of the cardiac fiber architecture enabled a variability study of the fiber structures and of the laminar structures. A complete description of

the average fiber architecture with its normal variation has been provided in three dimensions across several myocardium segments. Since the fiber architecture plays a key role in electrophysiology and in complex mechanics of the cardiac tissue, the results found in this thesis are significant contributions in the understanding of cardiac anatomy and functions in humans. From these results, a research question naturally arises: **How to create new accurate and precise models of various cardiac electromechanical functions from the human atlas found in this thesis?** Such models would advance studies on wave propagation in ventricular tissue, defibrillation, cardiac resynchronization, and in general simulations of electrophysiological interventions (Trayanova, 2011). One strategy would be to reuse the human atlas of the fiber architecture in existing finite element models of the heart (Sermesant *et al.*, 2006, 2003, 2008; Sermesant, 2003; Vadakkumpadan *et al.*, 2012). Electro-mechanical models could, therefore, benefit from an accurate description of the human cardiac fiber architecture.

Summary

The original findings of the thesis with concrete recommendations for future work are summarized on Fig. 7.3. New tools were developed to address the current limitations of shape normalization. The first research objective led to the construction of the first human atlas of the cardiac atlas, establishing concrete advances in specific knowledge of the cardiac anatomy. This atlas gives new perspectives for future electromechanical models of the heart and opens new doors for simulation and interventional tools. The subsequent research objectives led to the development of a fundamentally new registration approach where very large deformations between surfaces or images can be naturally captured. The new methods, FOCUSR and the *Spectral Demons* algorithm, were, in fact, found to be particularly relevant to neuro and cardiac imaging, but remain general enough for virtually any application that uses surface or image registration. These spectral methods could further benefit from improvements in computational efficiency and, on the theoretical point of view, from demonstrating additional links with graph-based segmentation methods. This may lead to a complete framework for atlas construction where segmentation, registration, and shape averaging are all interlinked. This thesis provided, therefore, new grounds for research and applications of shape normalization tools that may potentially have an impact on establishing diagnostics, and on planning and performing interventions.

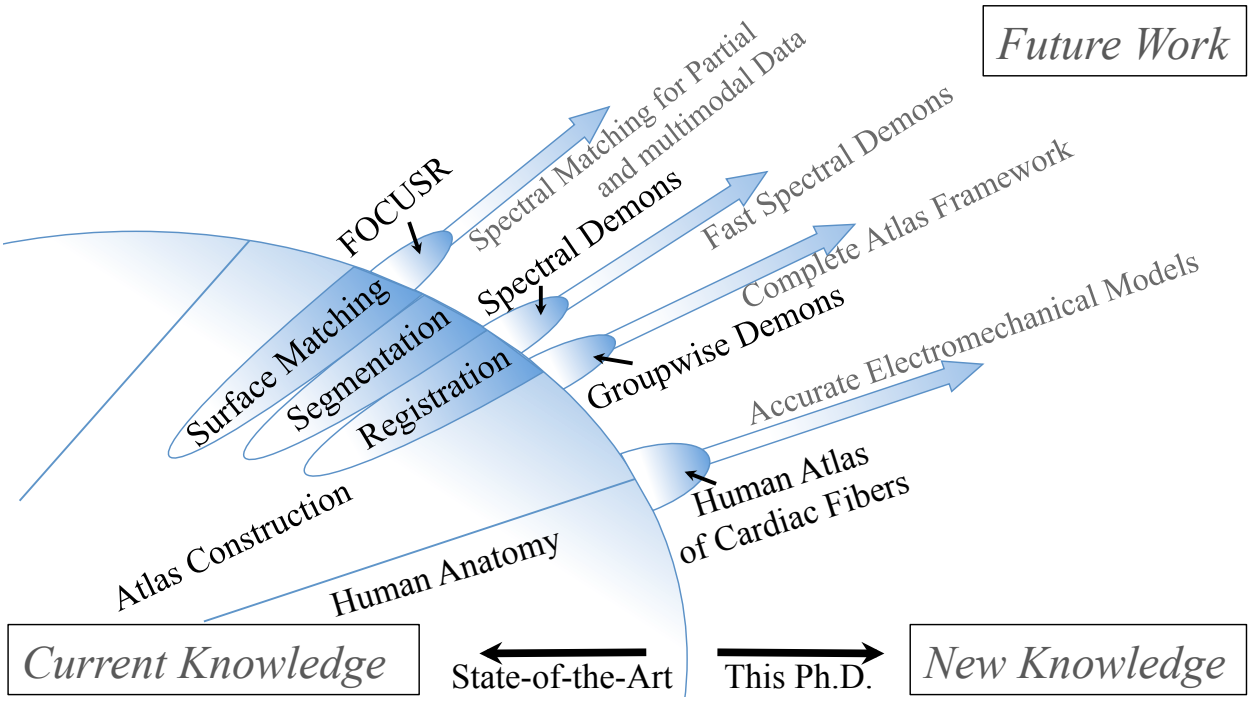

Figure 7.3: Organization of the key contributions and perspectives for future work. *a)* **FOCUSR** for surface matching, leading to an improved version handling partial data and that may be capable with multimodal data; *c)* ***Spectral Demons* algorithm** for image registration with large deformations, also leading also to a version handing partial and multimodal data, and that could further benefit for faster approaches; *c)* a **Groupwise Demons framework** for atlas construction, leading to a complete atlas framework with the development of an interlinked segmentation, registration and shape averaging method; and *d)* a **human atlas of the cardiac fiber architecture**, leading to more accurate electromechanical models of the heart.

BIBLIOGRAPHY

Adalsteinsson, D. and Sethian, J. A. (1995). A fast level set method for propagating interfaces. *J. Comput. Phys.*, 118, 269–277. 16

Adams, R. and Bischof, L. (1994). Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16, 641–647. 14

Alexander, D. C., Pierpaoli, C., Basser, P. J. and Gee, J. C. (2001). Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging*, 20, 1131–1139. 58

Allassonnière, S., Amit, Y. and Trouvé, A. (2007). Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 3–29. 33

Allassonnière, S., Kuhn, E. and Trouvé, A. (2010). Construction of bayesian deformable models via a stochastic approximation algorithm: A convergence study. *Bernoulli*, 16, 641–678. 33

Allen, B., Curless, B. and Popović, Z. (2003). The space of human body shapes: reconstruction and parameterization from range scans. *SIGGRAPH*. 587–594. 72

Amir, A. and Lindenbaum, M. (1998). A generic grouping algorithm and its quantitative analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20, 168–185. 17

Anguelov, D., Srinivasan, P., Pang, H.-C. and Koller, D. (2004). The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *Advances in Neural Information Processing Systems (NIPS)*. 33–40. 72

Anqi, Q., Bitouk, D. and Miller, M. (2006). Smooth functional and structural maps on the neocortex via orthonormal bases of the Laplace-Beltrami operator. *IEEE Transactions on Medical Imaging*, 25, 1296–1306. 29, 74

Armstrong, E., Schleicher, A., Omran, H., Curtis, M. and Zilles, K. (1995). The ontogeny of human gyrification. *Cerebral Cortex*, 5, 56–63. 8

Arsigny, V., Commowick, O., Pennec, X. and Ayache, N. (2006a). A Log-Euclidean framework for statistics on diffeomorphisms. R. Larsen, M. Nielsen and J. Sporring, editors, *MICCAI*. Berlin, Heidelberg, vol. 4190 of *Lecture Notes in Computer Science*, 924–931. 35, 58

Arsigny, V., Fillard, P., Pennec, X. and Ayache, N. (2006b). Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56, 411–421. 12, 35, 59, 204, 212

Arun, K. S., Huang, T. S. and Blostein, S. D. (1987). Least-Squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9, 698–700. 30

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, 38, 95–113. 24

Ashburner, J. and Friston, K. J. (2005). Unified segmentation. *NeuroImage*, 26, 839–851. 30

Ashburner, J., Neelin, P., Collins, D. L., Evans, A. and Friston, K. (1997). Incorporating prior knowledge into image registration. *NeuroImage*, 6, 344–352. 35

Audette, M. A., Ferrie, F. P. and Peters, T. M. (2000). An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis*, 4, 201–217. 72

Avants, B. and Gee, J. C. (2004a). Geodesic estimation for large deformation anatomical shape averaging and interpolation. *NeuroImage*, 23, S139–S150. 33, 124, 144

Avants, B. and Gee, J. C. (2004b). Shape averaging with diffeomorphic flows for atlas creation. *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*. IEEE, 595–598 Vol. 1. 33

Azar, A., Xu, C., Pennec, X. and Ayache, N. (2006). An interactive hybrid non-rigid registration framework for 3D medical images. *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*. IEEE, 824–827. 21

Bach, F. and Jordan, M. I. (2004). Learning spectral clustering. *Advances in Neural Information Processing Systems*. No. UCB/CSD-03-1249. 29, 73, 101

Backhaus, M., Chung, J. D., Cowan, B. R., Fonseca, C. G., Tao, W. and Young, A. A. (2010). The cardiac atlas project: Towards a map of the heart Patient-Specific modeling of the cardiovascular system. R. C. P. Kerckhoffs, editor, *Patient-Specific Modeling of the Cardiovascular System*, Springer New York, New York, NY, chapiter 7. 113–129. 31

Baker, H. (1902). *Proceedings of the London Mathematical Society*, vol. 1. 347–360. 26

Bansal, R., Staib, L. H., Chen, Z., Rangarajan, A., Knisely, J., Nath, R. and Duncan, J. S. (1998). A novel approach for the registration of 2D portal and 3D CT images for treatment setup verification in radiotherapy. *MICCAI*. 1075–1086. 30

Bardinet, E., Cohen, L. D. and Ayache, N. (1996). Tracking and motion analysis of the left ventricle with deformable superquadrics. *Med. Image Anal.*, 1, 129–149. 22, 23

Basser, P. J., Mattiello, J. and LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66, 259–267. 202

Basser, P. J. and Pajevic, S. (2000). Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magnetic Resonance in Medicine*, 44, 41–50. 35, 60

Becker, A. E. and Caruso, G. (1982). Myocardial disarray. a critical review. *British Heart Journal*, 47, 527–538. 210

Beg, M. F. and Khan, A. (2006). Computing an average anatomical atlas using LDDMM and geodesic shooting. *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*. IEEE, 1116–1119. 33, 143, 144

Beg, M. F., Miller, M. I., Trouvé, A. and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision*, 61, 139–157. 24, 144

Beg, M. F. F., Helm, P. A., McVeigh, E., Miller, M. I. and Winslow, R. L. (2004). Computational cardiac anatomy using MRI. *Magnetic resonance in medicine*, 52, 1167–1174. 33

Bengio, Y., Paiement, J.-F., Vincent, P., Delalleau, O., Le Roux, N. and Ouimet, M. (2004). Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. *Advances in Neural Information Processing Systems*. 29, 73

Besl, P. J. (1988). Geometric modeling and computer vision. *Proceedings of the IEEE*, 76, 936–958. 72

Bhat, P., Zheng, K. C., Snavely, N., Agarwala, A., Agrawala, M., Cohen, M. F. and Curless, B. (2006). Piecewise image registration in the presence of multiple large motions. *Proceedings of the IEEE Conference on Computer Vision and Pattern, CVPR 2006*. IEEE Computer Society, New York, NY, vol. 2, 2491–2497. 18

Bhatia, K. K., Hajnal, J. V., Puri, B. K., Edwards, A. D. and Rueckert, D. (2004). Consistent groupwise non-rigid registration for atlas construction. *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*. IEEE, 908–911 Vol. 1. 33

Bishop, M., Hales, P., Plank, G., Gavaghan, D., Scheider, J. and Grau, V. (2009). Comparison of rule-based and DTMRI-derived fibre architecture in a whole rat ventricular computational model. *Functional Imaging and Modeling of the Heart*. 87–96. 142, 202

Bistoquet, A., Oshinski, J. and Skrinjar, O. (2008). Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model. *Medical image analysis*, 12, 69–85. 25

Blanz, V. and Vetter, T. (1999). A morphable model for the synthesis of 3D faces. *SIGGRAPH*. 187–194. 72

Bookstein, F. L. (1989). Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 567–585. 22, 23

Borgefors, G. (1986). Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34, 344–371. 30

Borgefors, G. (1988). Hierarchical chamfer matching: a parametric edge matching algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10, 849–865. 30

Bossa, M., Hernandez, M. and Olmos, S. (2007). Contributions to 3D diffeomorphic atlas estimation: application to brain images. *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer-Verlag, Berlin, Heidelberg, vol. 1 of *MICCAI'07*, 667–674. 26, 33, 124, 143, 144

Boykov, Y. and Funka-Lea, G. (2006). Graph cuts and efficient n-d image segmentation. *Int. J. Comput. Vision*, 70, 109–131. 18

Boykov, Y. and Jolly, M.-P. (2000). Interactive organ segmentation using graph cuts. *MICCAI*. 276–286. 18, 51, 55, 56, 142, 204, 211

Boykov, Y. and Kolmogorov, V. (2001). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Energy Minimization Methods in Computer Vision and Pattern Recognition*. 359–374. 17, 18

Boykov, Y. and Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, 1124–1137. 17, 18

Boykov, Y., Veksler, O. and Zabih, R. (1998). Markov random fields with efficient approximations. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR 1998*. IEEE Computer Society, Santa Barbara, CA, 648–655. 18

Boykov, Y., Veksler, O. and Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, 1222–1239. 18

Brice, C. R. and Fennema, C. L. (1970). Scene analysis using regions. *Artificial Intelligence Journal*, 1, 205–226. 14

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien Dargestellt auf Grund des Zellenbaues*. Barth. 8, 109

Bronstein, A. M., Bronstein, M. I. and Kimmel, R. (2006). Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 1168–1172. 29, 74

Bronstein, A. M., Bronstein, M. l. and Kimmel, R. (2007). Calculus of nonrigid surfaces for geometry and texture manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 13, 902–913. 29, 74

Bronstein, M. M. and Kokkinos, I. (2010). Scale-invariant heat kernel signatures for non-rigid shape recognition. *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 1704–1711. 151

Bronzino, J. D. (1995). *Biomedical Engineering Handbook*, vol. 1. Springer-Verlag Berlin and Heidelberg. 48

Brown, L. G. (1992). A survey of image registration techniques. *ACM Comput. Surv.*, 24, 325–376. 21

Buckberg, G., Hoffman, J. I. E., Mahajan, A., Saleh, S. and Coghlan, C. (2008). Cardiac mechanics revisited. *Circulation*, 118, 2571–2587. 4, 5

Cachier, P., Bardinet, E., Dormont, D., Pennec, X. and Ayache, N. (2003). Iconic feature based nonrigid registration: the PASHA algorithm. *Computer Vision Image Understing*, 89, 272–298. 25, 122, 143

Campbell, A. W. (1903). Histological studies on cerebral localisation. *Proceedings of the Royal Society of London*, 72, 488–498. 8

Campbell, J. (1897). *Proceedings of the London Mathematical Society*, vol. 28. 14–32. 26

Canny, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8, 679–698. 14

Carcassoni, M. and Hancock, E. (2003). Spectral correspondence for point pattern matching. *Pattern Recognition*, 36, 193–204. 29, 74, 141, 143

Caselles, V., Kimmel, R. and Sapiro, G. (1995). Geodesic active contours. *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision*. IEEE Computer Society, Washington, DC, USA. 16

Cerqueira, M., Weissman, N., Dilsizian, V., Jacobs, A., Kaul, S., Laskey, W., Pennell, D., Rumberger, J., Ryan, T. and Verani, M. (2002). Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: a statement for healthcare professionals from the cardiac imaging committee of the council on clinical cardiology of the American Heart Association. *Circulation*, 105, 539–542. 64, 207

Chan, P. K., Schlag, M. D. F. and Zien, S. (1993). Spectral K-Way Ratio-Cut partitioning and clustering. *Design Automation, 1993. 30th Conference on*. IEEE, 749–754. 17, 18

Chan, T., Gilbert, J. and Teng, S.-H. (1995). Geometric spectral partitioning. Technical Report PARC CSL-94-15, Xerox. 29, 74

Chan, T. and Vese, L. (1999). An active contour model without edges. *SCALE-SPACE '99: Proceedings of the Second International Conference on Scale-Space Theories in Computer Vision*. Springer-Verlag, London, UK, 141–151. 16

Chen, H.-M., Chung, A. C. S., Yu, S. C. H., Norbash, A. and Wells, W. M. (2003). Multi-modal image registration by minimizing Kullback-Leibler distance between expected and observed joint class histograms. *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*. vol. 2, II–570–6 vol.2. 28

Chen, M. H. and Pavlidis, T. (1990). Image seaming for segmentation on parallel architecture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 588–594. 17

Chi, J. G., Dooling, E. C. and Gilles, F. H. (1977). Gyral development of the human brain. *Annals of Neurology*, 1, 86–93. 8

Chopp, D. L. (1993). Computing minimal surfaces via level set curvature flow. *J. Comput. Phys.*, 106, 77–91. 16

Christensen, G. E. and Johnson, H. J. (2001). Consistent image registration. *IEEE Transactions on Medical Imaging*, 20, 568–582. 27

Christensen, G. E. and Johnson, H. J. (2003). Invertibility and transitivity analysis for nonrigid image registration. *Journal of Electronic Imaging*, 12, 106–117. 27

Christensen, G. E., Miller, M. I., Vannier, M. W. and Grenander, U. (1996). Individualizing neuroanatomical atlases using a massively parallel computer. *Computer*, 29, 32–38. 21

Christensen, G. E., Rabbitt, R. D. and Miller, M. I. (2002). Deformable templates using large deformation kinematics. *Image Processing, IEEE Transactions on*, 5, 1435–1447. 22

Chui, H. (2003). A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*. vol. 89, 114–141. 79, 86

Chui, H., Rambo, J., Duncan, J. S., Schultz, R. and Rangarajan, A. (1999). Registration of cortical anatomical structures via robust 3D point matching. *Proceedings of the 16th International Conference on Information Processing in Medical Imaging*. Springer-Verlag, London, UK, IPMI '99, 168–181. 23

Chui, H. and Rangarajan, A. (2000a). A feature registration framework using mixture models. *Mathematical Methods in Biomedical Image Analysis, 2000. Proceedings. IEEE Workshop on*. IEEE, 190–197. 23

Chui, H. and Rangarajan, A. (2000b). A new algorithm for non-rigid point matching. *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*. IEEE, vol. 2, 44–51 vol.2. 23

Chung, A. C. S., Wells, W. M., Norbash, A. and Eric (2002). Multi-modal image registration by minimising Kullback-Leibler distance. *MICCAI*. Springer-Verlag, London, UK, 525–532. 28

Chung, F. (1997a). *Spectral Graph Theory*. AMS. 117, 118, 119, 123

Chung, F. R. K. (1997b). *Spectral Graph Theory*. AMS. 19, 29, 73, 82

Chung, M., Seo, S., Adluru, N. and Vorperian, H. (2011). Hot spots conjecture and its application to modeling tubular structures. K. Suzuki, F. Wang, D. Shen and P. Yan, editors, *Machine Learning in Medical Imaging*. Springer Berlin / Heidelberg, vol. 7009 of *Lecture Notes in Computer Science*, 225–232. 29

Chung, M. K. (2001). *Statistical Morphometry in Neuroanatomy*. Ph.D., McGill University, Montreal, Canada. 34

Chung, M. K., Robbins, S. M., Dalton, K. M., Davidson, R. J., Alexander, A. L. and Evans, A. C. (2005). Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage*, 25, 1256–1265. 111

Cirillo, M. (2009). A new surgical ventricular restoration technique to reset residual myocardium's fiber orientation: the "KISS" procedure. *Annals of Surgical Innovation and Research*, 3. 47, 69

Cirillo, M., Campana, M., Brunelli, F., Tomba, M. D. D., Mhagna, Z., Messina, A., Villa, E. and Troise, G. (2010). 'let's twist again': surgically induced renewal of left ventricular torsion in ischemic cardiomyopathy. *Journal of Cardiovascular Medicine*, 11, 34–39. 47, 69

Cocosco, C. A., Zijdenbos, A. P. and Evans, A. C. (2003). A fully automatic and robust brain MRI tissue classification method. *Medical image analysis*, 7, 513–527. 32

Codellas, P. S. (1932). Alcmaeon of croton: His life, work, and fragments. *Proceedings of the Royal Society of Medicine*, 25, 1041–1046. 1

Collignon, A., Maes, F., Delaere, D., Vandermeulen, P., Suetens, P. and Marchal, G. (1995a). automated multimodality image registration based on information theory. Y. Bizais, C. Barillot and R. Di Paola, editors, *Proceedings of Information Processing in Medical Imaging*. Kluwer Academic Publishers, Ile de Berder, France, vol. 3, 263–274. 27

Collignon, A., Vandermeulen, D., Suetens, P. and Marchal, G. (1995b). 3D Multi-Modality medical image registration using feature space clustering. *CVRMed '95: Proceedings of the First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine*. Springer-Verlag, London, UK, 195–204. 27

Conel, J. L. (1939). *Postnatal Development of the Human Cerebral Cortex: The Cortex of the Newborn*, vol. 1. Harvard Univ. Press. 110

Conel, J. L. (1967). *Postnatal Development of the Human Cerebral Cortex: The Cortex of the Seventy-two-Month Infant*, vol. 8. Harvard Univ. Press. 110

Cooper, D. B., Elliott, H., Cohen, F., Reiss, L. and Symosek, P. (1981). Stochastic boundary estimation and object recognition. *Image Modeling*, 63–94. 15

Cootes, T. F. and Taylor, C. J. (1994). Combining point distribution models with shape models based on finite element analysis. *Proceedings of the conference on British machine vision (vol. 2)*. BMVA Press, Surrey, UK, UK, BMVC 94, 419–428. 35

Cootes, T. F., Taylor, C. J., Cooper, D. H. and Graham, J. (1995). Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61, 38–59. 35, 106

Costa, K. D., Holmes, J. W. and Mcculloch, A. D. (2001). Modelling cardiac mechanical properties in three dimensions. *Mathematical, Physical and Engineering Sciences*, 359, 1233–1250. 5, 47, 142, 201, 210

Couprie, C., Grady, L., Najman, L. and Talbot, H. (2009). Power watersheds: A new image segmentation framework extending graph cuts, random walker and optimal spanning forest. *Computer Vision, 2009 IEEE 12th International Conference on*. IEEE, 731–738. 19

Couprie, C., Grady, L., Najman, L. and Talbot, H. (2011). Power watershed: A unifying Graph-Based optimization framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33, 1384–1399. 19

Coveney, P. V., Diaz, V., Hunter, P., Kohl, P. and Viceconti, M. (2011). The virtual physiological human. *Interface Focus*, 1, 281–285. 31

Cox, I. J., Rao, S. B. and Zhong, Y. (1996). Ratio regions: a technique for image segmentation. *Pattern Recognition, 1996., Proceedings of the 13th International Conference on*. IEEE, vol. 2, 557–564 vol.2. 17

Craene, M., Camara, O., Bijnens, B. H. and Frangi, A. F. (2009). Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences. *Proceedings of the 5th International Conference on Functional Imaging and Modeling of the Heart*. Springer-Verlag, Berlin, Heidelberg, vol. 5528 of *FIMH '09*, 437–446. 24

Crum, W. R., Griffin, L. D., Hill, D. L. and Hawkes, D. J. (2003). Zen and the art of medical image registration: correspondence, homology, and quality. *NeuroImage*, 20, 1425–1437. 31

Crum, W. R., Hartkens, T. and Hill, D. L. (2004). Non-rigid image registration: theory and practice. *The British Journal of Radiology*, 77, 140–153. 21, 114, 144

Cuvier, G. (1817). *Le Règne Animal Distribué d'Après son Organisation*. 2

D'Agostino, E., Maes, F., Vandermeulen, D. and Suetens, P. (2003). A viscous fluid model for multimodal non-rigid image registration using mutual information. *Med Image Anal*, 7, 565–575. 22

Daneshgar, A. and Javadi, R. (2012). On the complexity of isoperimetric problems on trees. *Discrete Appl. Math.*, 160, 116–131. 19

Darwin, C. (1859). *On the Origin of Species*. John Murray, London. 2

Davatzikos, C. and Prince, J. L. (1994). Brain image registration based on curve mapping. *Biomedical Image Analysis, 1994., Proceedings of the IEEE Workshop on*. 245–254. 22

Delong, A. and Boykov, Y. (2008). A scalable graph-cut algorithm for N-D grids. *IEEE International Conference on Computer Vision and Pattern Recognition*. IEEE, Los Alamitos, CA, USA, vol. 0, 1–8. 18

Delong, A., Gorelick, L., Schmidt, F. R., Veksler, O. and Boykov, Y. (2011). Interactive segmentation with super-labels. *Proceedings of the 8th international conference on Energy minimization methods in computer vision and pattern recognition*. Springer-Verlag, Berlin, Heidelberg, EMMCVPR'11, 147–162. 18

Delong, A., Osokin, A., Isack, H. N. and Boykov, Y. (2010a). Fast approximate energy minimization with label costs. *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2173–2180. 18

Delong, A., Osokin, A., Isack, H. N. and Boykov, Y. (2010b). Fast approximate energy minimization with label costs. *International Journal of Computer Vision*, 96, 1–27. 18

Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39, 1–38. 30

Deriche, R. (1987). Using canny's criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, 1, 167–187–187. 14

Desbrun, M., Meyer, M., Schröder, P. and Barr, A. (1999). Implicit fairing of irregular meshes using diffusion and curvature flow. *SIGGRAPH*. 317–324. 89, 111

Desolneux, A., Moisan, L. and Morel, J.-M. (2007). *From Gestalt Theory to Image Analysis: A Probabilistic Approach (Interdisciplinary Applied Mathematics)*. Springer, first edition. 17

Drineas, P. and Mahoney, M. W. (2005). On the Nyström method for approximating a Gram matrix for improved Kernel-Based learning. *J. Mach. Learn. Res.*, 6, 2153–2175. 133, 150

Dubois, J., Benders, M., Borradori-Tolsa, C., Cachia, A., Lazeyras, F., Ha-Vinh Leuchter, R., Sizonenko, S. V., Warfield, S. K., Mangin, J. F. and Hüppi, P. S. (2008). Primary cortical folding in the human newborn: an early marker of later functional development. *Brain : a journal of neurology*, 131, 2028–2041. 8

Duchon, J. (1976). Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Revue d'automation, informatique et recherche operationelle*, 10, 5–12. 22, 23

Durrleman, S. (2010). *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. Ph.D., Université de Nice-Sophia Antipolis. 35, 106

Durrleman, S., Fillard, P., Pennec, X., Trouvé, A. and Ayache, N. (2009a). A statistical model of white matter fiber bundles based on currents. *Information processing in medical imaging : proceedings of the conference*. vol. 21, 114–125. 34

Durrleman, S., Fillard, P., Pennec, X., Trouvé, A. and Ayache, N. (2011). Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage*, 55, 1073–1090. 34, 143

Durrleman, S., Pennec, X., Trouvé, A. and Ayache, N. (2008). A forward model to build unbiased atlases from curves and surfaces. *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy*. 34, 143

Durrleman, S., Pennec, X., Trouvé, A. and Ayache, N. (2009b). Statistical models of sets of curves and surfaces based on currents. *Medical image analysis*, 13, 793–808. 34, 35, 106, 143

Elad, A. and Kimmel, R. (2003). On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25, 1285–1295. 29, 74

Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L. and Peters, T. M. (1993). 3D statistical neuroanatomical models from 305 MRI volumes. *Nuclear Science*

Symposium and Medical Imaging Conference, 1993., 1993 IEEE Conference Record. IEEE, 1813–1817 vol.3. 31

Felzenszwalb, P. F. and Huttenlocher, D. P. (2004). Efficient Graph-Based image segmentation. *Int. J. Comput. Vision*, 59, 167–181. 17

Fillard, P., Poupon, C. and Mangin, J. F. (2009). A novel global tractography algorithm based on an adaptive spin glass model. *MICCAI*. 927–934. 52

Fischl, B. and Dale, A. M. (2000). Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 11050–5. 92

Fischl, B., Sereno, M. I., Tootell, R. B. and Dale, A. M. (1999). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8, 272–284. 72, 92, 94, 100, 131, 139, 141

Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Segonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B. and Dale, A. M. (2004). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14, 11–22. 32, 95

Fitch, W. M. (2000). Homology a personal view on some of the problems. *Trends in genetics : TIG*, 16, 227–231. 31

Fitzpatrick, J., Hill, D. and Maurer, C. (2000). *Image Registration*, SPIE, 1000 20th Street, Bellingham, WA 98227-0010 USA, vol. 2, chapiter PM80_ch8. 447–513. 21

Fletcher, P. and Joshi, S. (2007). Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, 87, 250–262. 35

Fletcher, P. T., Lu, C., Pizer, S. M. and Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23, 995–1005. 35

Fletcher, T. P. and Joshi, S. (2004). Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. Jan, editor, *Computer Vision and Mathematical Methods in*

Medical and Biomedical Image Analysis: ECCV 2004Workshops CVAMIA and MMBIA Prague, Czech Republic, May 15, 2004 Revised Selected Papers. Springer, vol. 3117. 35

Foland-Ross, L. C., Thompson, P. M., Sugar, C. A., Madsen, S. K., Shen, J. K., Penfold, C., Ahlf, K., Rasser, P. E., Fischer, J., Yang, Y., Townsend, J., Bookheimer, S. Y. and Altshuler, L. L. (2011). Investigation of cortical thickness abnormalities in lithium-free adults with bipolar i disorder using cortical pattern matching. *The American journal of psychiatry*, 168, 530–539. 7

Fonseca, C. G., Backhaus, M., Bluemke, D. A., Britten, R. D., Chung, J. D., Cowan, B. R., Dinov, I. D., Finn, J. P., Hunter, P. J., Kadish, A. H., Lee, D. C., Lima, J. A. C., Medrano-Gracia, P., Shivkumar, K., Suinesiaputra, A., Tao, W. and Young, A. A. (2011). The cardiac atlas project–an imaging database for computational modeling and statistical atlases of the heart. *Bioinformatics*, 27, 2288–2295. 31

Ford, L. R. and Fulkerson, D. R. (1962). *Flows in Networks*. Princeton University Press, Princeton, NJ. 18

Forsberg, D., Rathi, Y., Bouix, S., Wassermann, D., Knutsson, H. and Westin, C. F. (2011). Improving registration using multi-channel diffeomorphic demons combined with certainty maps. *Proceedings of the First international conference on Multimodal brain image analysis*. Springer-Verlag, Berlin, Heidelberg, MBIA'11, 19–26. 25

Fowlkes, C., Belongie, S., Chung, F. and Malik, J. (2004). Spectral grouping using the nystrom method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, 214–225. 101

Freedman, D. and Drineas, P. (2005). Energy minimization via graph cuts: Settling what is possible. 18

Frindel, C. (2009). *Imagerie par résonance magnétique du tenseur de diffusion (IRM-TD) en imagerie cardiaque humaine : traitements et premières interprétations*. Ph.D., Institut National des Sciences Appliquées de Lyon. 49, 52, 67

Frindel, C., Robini, M., Croisille, P. and Zhu, Y.-M. M. (2009). Comparison of regularization methods for human cardiac diffusion tensor MRI. *Medical Image Analysis*, 13, 405–418. 49, 52, 67, 202, 210

Frindel, C., Robini, M., Rapacchi, S., Stephant, E., Zhu, Y.-M. M. and Croisille, P. (2007). Towards in vivo diffusion tensor MRI on human heart using edge-preserving regularization. *IEEE Engineering in Medicine and Biology Society*, 6008–6011. 65

Gan, R., Wu, J., Chung, A. C. S., Yu, S. C. H. and Wells, W. M. (2004). Multiresolution image registration based on Kullback-Leibler distance. C. Barillot, D. R. Haynor and P. Hellier, editors, *MICCAI*. Springer-Verlag, Berlin Heidelberg, vol. 3216, 599–606. 28

Garel, C., Chantrel, E., Brisse, H., Elmaleh, M., Luton, D., Oury, J.-F., Sebag, G. and Hassan, M. (2001). Fetal cerebral cortex: Normal gestational landmarks identified using prenatal MR imaging. *American Journal of Neuroradiology*, 22, 184–189. 8

Geerts, L., Bovendeerd, P., Nicolay, K. and Arts, T. (2002). Characterization of the normal cardiac myofiber field in goat measured with MR-diffusion tensor imaging. *American Journal of Physiology - Heart and Circulatory Physiology*, 283, H139–H145. 142, 202

Geman, S. and Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6, 721–741. 15

Geng, X., Kumar, D. and Christensen, G. E. (2005). Transitive Inverse-Consistent manifold registration information processing in medical imaging. Springer Berlin / Heidelberg, Berlin, Heidelberg, vol. 3565 of *Lecture Notes in Computer Science*, chapiter 39. 3–30. 27

Gholipour, A., Kehtarnavaz, N., Briggs, R., Devous, M. and Gopinath, K. (2007). Brain functional localization: A survey of image registration techniques. *IEEE Transactions on Medical Imaging*, 26, 427–451. 21

Gilbert, S., Benson, A., Li, P. and Holden, A. (2007). Visualisation of dog myocardial structure from diffusion tensor magnetic resonance imaging: the paradox of uniformity and variability. *Functional Imaging and Modeling of the Heart*. 403–412. 31, 62, 142, 202

Glocker, B., Komodakis, N., Tziritas, G., Navab, N. and Paragios, N. (2008). Dense image registration through MRFs and efficient linear programming. *Medical image analysis*, 12, 731–741. 21

Golland, P., Eric, Shenton, M. E. and Kikinis, R. (2001). Deformation analysis for shape based classification. *Proceedings of the 17th International Conference on Information Processing in Medical Imaging*. Springer-Verlag, London, UK, UK, IPMI '01, 517–530. 35

Gotland, P. and Sabuncu, M. R. (2011). Modeling anatomical heterogeneity in populations. *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 5776–5779. 33

Grady, L. (2005). Multilabel random walker image segmentation using prior models. *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. IEEE, vol. 1, 763–770 vol. 1. 19

Grady, L. (2006a). Fast, quality, segmentation of large volumes - isoperimetric distance trees. *ECCV*. vol. 3, 449–462. 19

Grady, L. (2006b). Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28, 1768–1783. 19

Grady, L. and Polimeni, J. R. (2010). *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer. 29, 73, 79, 80, 81, 117, 118, 123

Grady, L., Schiwietz, T., Aharon, S. and Westermann, R. (2005). Random walks for interactive organ segmentation in two and three dimensions: implementation and validation. *MICCAI*, 8, 773–780. 19

Grady, L. and Schwartz, E. L. (2005). Isoperimetric partitioning: A new algorithm for graph partitioning. *SIAM J. Sci. Comput.*, 27, 1844–1866. 19

Grady, L. and Schwartz, E. L. (2006). Isoperimetric graph partitioning for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28, 469–475. 19

Grady, L. and Sinop, A. K. (2008). Fast approximate random walker segmentation using eigenvector precomputation. *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 1–8. 19

Grajek, S., Lesiak, M., Pyda, M., Zajac, M., Paradowski, S. and Kaczmarek, E. (1993). Hypertrophy or hyperplasia in cardiac muscle. post-mortem human morphometric study. *European Heart Journal*, 14, 40–47. 210

Greenbaum, R. A., Ho, S. Y., Gibson, D. G., Becker, A. E. and Anderson, R. H. (1981). Left ventricular fibre architecture in man. *British Heart Journal*, 45, 248–263. 49, 63, 64, 69

Greig, D., Porteous, B. and Seheult, A. (1989). Exact maximum a posteriori estimation for binary images. *Royal Journal on Statistical Society*, 51, 271–279. 17

Grenander, U. and Miller, M. I. (1998). Computational anatomy: an emerging discipline. *Q. Appl. Math.*, LVI, 617–694. 31

Guetter, C., Xu, C., Sauer, F. and Hornegger, J. (2005). Learning based non-rigid multi-modal image registration using Kullback-Leibler divergence. *MICCAI*, 8, 255–262. 28

Guimond, A. (2000). Average brain models: A convergence study. *Computer Vision and Image Understanding*, 77, 192–210. 32, 34

Guimond, A., Meunier, J. and Thirion, J. P. (2000). Average brain models: a convergence study. *Computer Vision and Image Understanding*, 192–210. 12, 58, 116, 124, 125, 143, 144, 204, 211

Haber, E. and Modersitzki, J. (2004). Numerical methods for volume preserving image registration. *Inverse Problems*, 20, 1621–1638. 25

Hagen, L. and Kahng, A. B. (1992). New spectral methods for ratio cut partitioning and clustering. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 11, 1074–1085. 17, 18

Hahnel, D., Thrun, S. and Burgard, W. (2003). An extension of the ICP algorithm for modeling nonrigid objects with mobile robots. *Proceedings of the International Joint Conference on Artificial Intelligence*. 915–920. 72

Halko, N., Martinsson, P. G. and Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM*, 53, 217–288. 101

Harvey, W. (1628). *Exercitatio anatomica de motu cordis et sanguinis in animalibus*. 1, 48

Hasan, K. M., Parker, D. L. and Alexander, A. L. (2001). Comparison of gradient encoding schemes for diffusion-tensor MRI. *Magnetic Resonance Imaging*, 13, 769–780. 67

Hasan, K. M., Parker, D. L. and Alexander, A. L. (2992). Magnetic resonance water self-diffusion tensor encoding optimization methods for full brain acquisition. *Image Analysis and Stereology*, 21, 87–96. 67

Hausdorff, F. (1906). *Die symbolische Exponentialformel in der Gruppentheorie*, Ber Verh Saechs Akad Wiss Leipzig, vol. 58. 19–49. 26

Heimann, T. and Meinzer, H.-P. (2009). Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis*, 13, 543–563. 15

Hellier, P. and Barillot, C. (2003). Coupling dense and landmark-based approaches for nonrigid registration. *IEEE Transactions on Medical Imaging*, 22, 217–227. 21

Helm, P. A. (2005). *A novel technique for quantifying variability of cardiac anatomy: Application to the dyssynchronous failing heart*. Ph.D., Johns Hopkins University. 31, 53, 55, 69, 142, 202

Helm, P. A., Beg, M. F., Miller, M. I. and Winslow, R. L. (2005a). Measuring and mapping cardiac fiber and laminar architecture using diffusion tensor MR imaging. *Annals of the New York Academy of Sciences*, 1047, 296–307. 31, 53, 55, 62, 142, 202, 205

Helm, P. A., Tseng, H.-J. J., Younes, L., McVeigh, E. R. and Winslow, R. L. (2005b). Ex vivo 3D diffusion tensor imaging and quantification of cardiac laminar structure. *Magnetic Resonance in Medicine*, 54, 850–859. 31, 53, 55, 62, 142, 143, 202, 205, 207

Hernandez, M., Bossa, M. N. and Olmos, S. (2007). Registration of anatomical images using geodesic paths of diffeomorphisms parameterized with stationary vector fields. *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*. IEEE, 1–8. 24

Hernandez, M., Olmos, S. and Pennec, X. (2008). Comparing algorithms for diffeomorphic registration: Stationary LDDMM and diffeomorphic demons. X. Pennec and S. Joshi, editors, *Mathematical Foundations of Computation Anatomy MICCAI Workshop*. 24–35. 24

Hill, D. L. G., Hawkes, D. J., Harrison, N. A. and Ruff, C. F. (1993). A strategy for automated multimodality image registration incorporating anatomical knowledge and imager characteristics. *IPMI '93: Proceedings of the 13th International Conference on Information Processing in Medical Imaging*. Springer-Verlag, London, UK, 182–196. 27

Hinds, O., Polimeni, J. R., Rajendran, N., Balasubramanian, M., Amunts, K., Zilles, K., Schwartz, E. L., Fischl, B. and Triantafyllou, C. (2009). Locating the functional and anatomical boundaries of human primary visual cortex. *Neuroimage*, 46, 915–22. 113

Hinds, O., Rajendran, N., Polimeni, J. R., Augustinack, J. C., Wiggins, G., Wald, L. L., Rosas, D. H., Potthast, A., Schwartz, E. L. and Fischl, B. (2008). Accurate prediction of V1 location from cortical folds in a surface coordinate system. *Neuroimage*, 39, 1585–99. 75, 95, 114

Hooks, D. A., Tomlinson, K. A., Marsden, S. G., LeGrice, I. J., Smaill, B. H., Pullan, A. J. and Hunter, P. J. (2002). Cardiac microstructure: implications for electrical propagation and defibrillation in the heart. *Circulation Research*, 91, 331–338. 5, 47, 142, 201, 210

Hooks, D. A., Trew, M. L., Caldwell, B. J., Sands, G. B., LeGrice, I. J. and Smaill, B. H. (2007). Laminar arrangement of ventricular myocytes influences electrical behavior of the heart. *Circulation Research*, 101, e103–e112. 5

Horn, B. K. P. and Schunck, B. G. (1981). Determining optical flow. *Artificial Intelligence*, 17, 185–203. 22

Horn, B. K. P. and Schunck, B. G. (1992). Determining optical flow. *Physics-Based Vision: Principles And Practice*, 389–407. 22

Horowitz, S. L. and Pavlidis, T. (1976). Picture segmentation by a tree traversal algorithm. *J. ACM*, 23, 368–388. 14

Hsu, E. W. and Henriquez, C. S. (2001). Myocardial fiber orientation mapping using reduced encoding diffusion tensor imaging. *Journal of Cardiovascular Magnetic Resonance*, 3, 339–347. 202

Hsu, W. M., Hughes, J. F. and Kaufman, H. (1992). Direct manipulation of free-form deformations. *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*. ACM Press, New York, NY, USA, 177–184. 23

Huckemann, S. and Hotz, T. (2009). Principal component geodesics for planar shape spaces. *Journal of Multivariate Analysis*, 100, 699–714. 35

Huckemann, S., Hotz, T. and Munk, A. (2010). Intrinsic shape analysis: Geodesic PCA for riemannian manifolds modulo isometric lie group actions. *Statistica Sinica*, 20, 1–100. 35

Huckemann, S. F. (2011). Intrinsic inference on the mean geodesic of planar shapes and tree discrimination by leaf growth. *Annals of Statistics*, 39, 1098–1124. 35

Jain, V. and Zhang, H. (2006). Robust 3D shape correspondence in the spectral domain. *IEEE International Conference on Shape Modeling and Applications*. 19. 30, 74, 100, 117, 141, 143

Johnson, H. J. and Christensen, G. E. (2002). Consistent landmark and intensity-based image registration. *IEEE Transactions on Medical Imaging*, 21, 450–461. 27

Jones, D. K. (2008). Tractography gone wild: Probabilistic fibre tracking using the wild bootstrap with diffusion tensor MRI. *IEEE Transactions on Medical Imaging*, 27, 1268–1274. 67

Jones, D. K., Griffin, L. D., Alexander, D. C., Catani, M., Horsfield, M. A., Howard, R. and Williams, S. C. R. (2002). Spatial normalization and averaging of diffusion tensor MRI data sets. *NeuroImage*, 17, 592–617. 35, 59

Joshi, S., Davis, B., Jomier, M. and Gerig, G. (2004). Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23, S151–S160. 33, 144

Juan, O. and Boykov, Y. (2006). Active graph cuts. *CVPR*. vol. 1, 1023–1029. 18

Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.*, 30, 509–541. 35

Karger, D. R. and Stein, C. (1996). A new approach to the minimum cut problem. *J. ACM*, 43, 601–640. 18

Karsner, H. T., Saphir, O. and Todd, T. W. (1925). The state of the cardiac muscle in hypertrophy and atrophy. *The American journal of pathology*, 1. 210

Kass, M., Witkin, A. and Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, V1, 321–331. 16

Keele, K. D. (1964). Leonardo Da Vinci's infludence on renaissance anatomy. *Medical history*, 8, 360–370. 1

Kendall, D. G. (1984). Shape manifolds, procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16, 81–121. 35, 106

Kim, J., Kolmogorov, V. and Zabih, R. (2003). Visual correspondence using energy minimization and mutual information. *Proceedings of the Ninth IEEE International Conference on Computer Vision, ICCV 2003*. Nice, France, vol. 2, 1033–1040. 18

Kolmogorov, V. and Zabih, R. (2002). What energy functions can be minimized via graph cuts? *ECCV (3)*. 65–81. 18

Kolmogorov, V. and Zabih, R. (2004). What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, 147–159. 18

Kullback, S. (1959). *Information theory and statistics*. John Wiley and Sons., New York. 28

Kumar, D., Geng, X., Christensen, G. E. and Vannier, M. W. (2003). Characterizing shape differences between phantom image populations via multivariate statistical analysis of inverse consistent transformations biomedical image registration. Springer Berlin / Heidelberg, Berlin, Heidelberg, vol. 2717 of *Lecture Notes in Computer Science*, chapiter 39. 367–376. 27

Lee, S., Wolberg, G. and Shin, S. Y. (1997). Scattered data interpolation with multilevel b-splines. *Visualization and Computer Graphics, IEEE Transactions on*, 3, 228–244. 22

Lee, S.-Y., Chwa, K.-Y. and Shin, S. Y. (1995). Image metamorphosis using snakes and free-form deformations. *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM Press, New York, NY, USA, 439–448. 22

Lee, W. N., Pernot, M., Couade, M., Messas, E., Bruneval, P., Bel, A., Hagege, A. A., Fink, M. and Tanter, M. (2012). Mapping myocardial fiber orientation using Echocardiography-Based shear wave imaging. *IEEE Transactions on Medical Imaging*, 31, 554–562. 70

LeGrice, I. J., Smaill, B. H., Chai, L. Z., Edgar, S. G., Gavin, J. B. and Hunter, P. J. (1995). Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *The American journal of physiology*, 269. 4, 5, 47, 48, 54, 201, 205, 209

Lehoucq, R., Sorensen, D. and Yang, C. (1997). *Arpack users' guide: Solution of large scale eigenvalue problems with implicitly restarted Arnoldi methods*. SIAM. 123

Leordeanu, M. and Hebert, M. (2005). A spectral technique for correspondence problems using pairwise constraints. *Proceedings IEEE International Conference on Computer Vision*. vol. 2. 29, 140

Lerm, N., Malgouyres, F. and Ltocart, L. (2010). Reducing graphs in graph cut segmentation. *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 3045–3048. 18

Leventon, M. E., Grimson, W. E. L. and Faugeras, O. (2002). Statistical shape influence in geodesic active contours. *Biomedical Imaging, 2002. 5th IEEE EMBS International Summer School on*. IEEE, 8 pp.+. 30, 32

Lexell, J. and Taylor, C. C. (1991). Variability in muscle fibre areas in whole human quadriceps muscle: effects of increasing age. *Journal of anatomy*, 174, 239–249. 214

Li, C., Xu, C., Gui, C. and Fox, M. D. (2005). Level set evolution without re-initialization: a new variational formulation. *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. vol. 1, 430–436. 16

Li, M., Lian, X.-C., Kwok, J. T. and Lu, B.-L. (2011). Time and space efficient spectral clustering via column sampling. *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 2297–2304. 133, 150

Lin, M. H. (1999). Tracking articulated objects in real-time range image sequences. *IEEE International Conference on Computer Vision (ICCV)*. vol. 1, 648–653 vol.1. 72

Liu, R. F., Zhang, H., Shamir, A. and Cohen-Or, D. (2009). A part-aware surface metric for shape analysis. *Eurographics*, 28. 101

Lohmann, G., von Cramon, D. Y. and Colchester, A. C. (2008). Deep sulcal landmarks provide an organizing framework for human cortical folding. *Cerebral Cortex*, 18, 1415–20. 72, 141

Lombaert, H. and Cheriet, F. (2010a). Geodesic thin plate splines for image segmentation. *Proceedings of the Int. Conf. on Pattern Recognition (ICPR)*. IEEE, 2234–2237. 23

Lombaert, H. and Cheriet, F. (2010b). Spatio-Temporal segmentation of the heartin 4D MRI images using graph cuts with motion cues. *Proceedings of the Int. Symposium on Biomedical Imaging (ISBI)*. vol. 1. 18

Lombaert, H. and Cheriet, F. (2012). Simultaneous image denoising and registration using graph cuts. *ISSPA (To Appear),*. 18

Lombaert, H., Grady, L., Polimeni, J. R. and Cheriet, F. (2011a). Spectral correspondence for brain matching. *Proceedings of the Int. Conf. on Information Processing in Medical Imagin (IPMI)*. 660–670. 38, 39, 41, 42, 71, 94, 97, 117, 120, 131

Lombaert, H., Grady, L., Polimeni, J. R. and Cheriet, F. (2012a). FOCUSR: Feature oriented correspondence using spectral regularization –a method for precise surface matching. *Submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence*. 38, 39, 41, 42, 71

Lombaert, H., Peyrat, J.-M., Croisille, P., Rapacchi, S., Fanton, L., Cheriet, F., Clarysse, P., Magnin, I., Delingette, H. and Ayache, N. (2012b). Human atlas of the cardiac fiber architecture: Study on a healthy population. *To Appear in IEEE Transactions on Medical Imaging*. 31, 38, 39, 41, 42, 45

Lombaert, H., Peyrat, J.-M., Croisille, P., Rapacchi, S., Fanton, L., Clarysse, P., Delingette, H. and Ayache, N. (2011b). Statistical analysis of the human cardiac fiber architecture from DT-MRI. *Proceedings of the Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH)*. vol. 6666, 171–179. 31, 38, 39, 41, 42, 45, 49, 135, 202, 203, 205, 210, 211, 214

Lombaert, H., Peyrat, J.-M., Fanton, L., Cheriet, F., Delingette, H., Ayache, N., Clarysse, P., Magnin, I. and Croisille, P. (2011c). Statistical atlas of human cardiac fibers: Comparison with abnormal hearts. *STACOM MICCAI*. 38, 39, 42, 45, 69, 201

Lombaert, H., Peyrat, J.-M., Fanton, L., Cheriet, F., Delingette, H., Ayache, N., Clarysse, P., Magnin, I. and Croisille, P. (2011d). Variability of the human cardiac laminar structure. *STACOM MICCAI*. 38, 39, 42, 45, 69, 209

Lombaert, H., Peyrat, J.-M., Rapacchi, S., Fanton, L., Delingette, H., Ayache, N. and Croisille, P. (2011e). Human statistical atlas of cardiac fiber architecture from DT-MRI. *Proceedings of Int. Soc. for Magnetic Resonance in Medicine (ISMRM)*. 31, 38, 39, 42, 45, 202, 210

Lombaert, H., Sun, Y. and Cheriet, F. (2007). Landmark-based non-rigid registration via graph cuts. *Proceedings of the Int. Conf. on Image Analysis and Recognition (ICIAR)*. 18

Lombaert, H., Sun, Y. and Cheriet, F. (2011f). Fast 4D segmentation of large datasets using graph cuts. *Proceedings of SPIE*, 7962, 79622H–79622H–7. 18

Lombaert, H., Sun, Y., Grady, L. and Xu, C. (2005). A multilevel banded graph cuts method for fast image segmentation. *ICCV*. IEEE Computer Society, Washington, DC, USA, 259–265. 18

Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E. and Joshi, S. (2006). Multi-modal image set registration and atlas formation. *Medical image analysis*, 10, 440–451. 33

Lucas, B. D. and Kanade, T. (1981a). An iterative image registration technique with an application to stereo vision. *Imaging Understanding Workshop*. 121–130. 22

Lucas, B. D. and Kanade, T. (1981b). An iterative image registration technique with an application to stereo vision. *Proceedings of the Seventh International Joint Conference on Artifical Intelligence, IJCAI 1981*. Vancouver, Canada, 674–679. 22

Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing*, 17, 395–416. 19, 123

Ma, J., Miller, M. I., Trouvé, A. and Younes, L. (2008). Bayesian template estimation in computational anatomy. *NeuroImage*, 42, 252–261. 34

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G. and Suetens, P. (1997). Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16, 187–198. 27

Maintz, J. B. and Viergever, M. A. (1998). A survey of medical image registration. *Medical Image Analysis*, 2, 1–36. 21, 27

Mäkelä, T., Clarysse, P., Sipilä, O., Pauna, N., Pham, Q. C., Katila, T. and Magnin, I. E. (2002). A review of cardiac image registration methods. *IEEE Trans Med Imaging*, 21, 1011–1021. 21

Mall, F. P. (1911). On the muscular architecture of the ventricles of the human heart. *Am. J. Anat.*, 11, 211–266. 4, 5

Malladi, R., Sethian, J. A. and Vemuri, B. C. (1995). Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, 158–175. 16

Mansi, T. (2010). *Image-Based Physiological and Statistical Models of the Heart, Application to Tetralogy of Fallot*. Ph.D., Ecole Nationale Supérieure des Mines de Paris. 25

Mansi, T., Pennec, X., Sermesant, M., Delingette, H. and Ayache, N. (2010). iLogDemons: A Demons-Based registration algorithm for tracking incompressible elastic biological tissues. *International Journal of Computer Vision*, 1–20. 25, 51, 57, 204, 211

Mansi, T., Peyrat, J. M., Sermesant, M., Delingette, H., Blanc, J., Boudjemline, Y. and Ayache, N. (2009). Physically-Constrained diffeomorphic demons for the estimation of 3D myocardium strain from Cine-MRI. *Proceedings of the 5th International Conference on Functional Imaging and Modeling of the Heart*. Springer-Verlag, Berlin, Heidelberg, FIMH '09, 201–210. 25

Marr, D. and Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London*. vol. 207, 187–217. 14

Marroquin, J. L. (1984). Surface reconstruction preserving discontinuities. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA. 15

Marroquin, J. L., Santana, E. A. and Botello, S. (2003). Hidden markov measure field models for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25, 1380–1387. 32

Marsland, S. and Twining, C. J. (2002). Clampled-plate splines and the optimal flow of bounded diffeomorphisms. *Statistics of Large Datasets, Proceedings of Leeds Annual Statistical Research Workshop*. 91–95. 24

Marsland, S. and Twining, C. J. (2004). Constructing diffeomorphic representations for the groupwise analysis of nonrigid registrations of medical images. *IEEE Transactions on Medical Imaging*, 23, 1006–1020. 33

Marsland, S., Twining, C. J. and Taylor, C. J. (2003). Groupwise non-rigid registration using polyharmonic Clamped-Plate splines. *MICCAI*. vol. 2879, 771–779. 33, 144

Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F. and Boyer, E. (2008). Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1–8. 29, 74, 117, 120, 141, 143, 151

Maudgil, D. D., Free, S. L., Sisodiya, S. M., LeMieux, L., Woermann, F. G., Fish, D. R. and Shorvon, S. D. (1998). Identifying homologous anatomical landmarks on reconstructed magnetic resonance images of the human cerebral cortical surface. *Journal of Anatomy*, 193, 559–571. 31

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Le Goualher, G., Feidler, J., Smith, K., Boomsma, D., Hulshoff Pol, H., Cannon, T., Kawashima, R. and Mazoyer, B. (2001a). A four-dimensional probabilistic atlas of the human brain. *Journal of the American Medical Informatics Association : JAMIA*, 8, 401–430. 31

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Le Goualher, G., Boomsma, D., Cannon, T., Kawashima, R. and Mazoyer, B. (2001b). A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (ICBM). *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 356, 1293–1322. 31

Mcinerney, T. and Terzopoulos, D. (1996). Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1, 91–108. 21, 139

Meek, P. D., McKeithan, K. and Schumock, G. T. (1998). Economic considerations in alzheimer's disease. *Pharmacotherapy*, 18. 2

Meila, M. and Shi, J. (2000). Learning segmentation by random walks. *NIPS*. 873–879. 19, 30, 119, 123

Meila, M. and Shi, J. (2001). A random walks view of spectral segmentation. T. Jaakkola and T. Richardson, editors, *Proceedings of the Eighth International Workshop on Artifical Intelligence and Statistics*. Morgan Kaufmann. 19

Meinguet, J. (1979). Multivariate interpolation at arbitrary points made simple. *Zeitschrift für Angewandte Mathematik und Physik*, 30, 292–304. 22, 23

Miller, M. I. and Qiu, A. (2009). The emerging discipline of computational functional anatomy. *NeuroImage*, 45, S16–S39. 31

Miller, M. I., Trouve, A. and Younes, L. (2002). On the metrics and euler-lagrange equations of computational anatomy. *Annual review of biomedical engineering*, 4, 375–405. 24

Mitiche, A. and Ayed, I. B. (2010). *Variational and Level Set Methods in Image Segmentation (Springer Topics in Signal Processing)*. Springer, 2011th edition. 16

Mohar, B. (1997). Some applications of laplace eigenvalues of graphs. G. Hahn and G. Sabidussi, editors, *Graph Symmetry: Algebraic Methods and Applications*, Kluwer Academic Publishers, Dordrecht, vol. 497. 227–275. 19

Morel, J. M. and Solimini, S. (1995). *Variational methods in image segmentation*. Birkhauser Boston Inc., Cambridge, MA, USA. 15

Mumford, D. (1994). *The Bayesian rationale for the variational formulation*, Ed. Dordrecht, Kluwer, The Netherlands, chapiter 5. 137–147. 15

Mumford, D. and Shah, J. (1985). Boundary detection by minimizing functionals. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. San Fransisco, 22–26. 15

Mumford, D. and Shah, J. (1989). Optimal approximation by piecewise smooth functions and associated variational problems. *Communications on Pure Applied Mathematics*, 42, 577–685. 15

Myronenko, A. and Song, X. (2009). Point-set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32, 2262–2275. 79, 86

Nielsen, P. M., Le Grice, I. J., Smaill, B. H. and Hunter, P. J. (1991). Mathematical model of geometry and fibrous structure of the heart. *American Journal of Physiology, Heart Circulatory Physiology*, 260, H1365–1378. 49, 62, 202, 205

Niethammer, M., Reuter, M., Wolter, F., Bouix, S., Peinecke, N., Koo, M. and Shenton, M. (2007). Global medical shape analysis using the Laplace-Beltrami spectrum. *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*. vol. 4791, 850–857. 29, 74

Noble, J. A. and Boukerroui, D. (2006). Ultrasound image segmentation: a survey. *IEEE Transactions on Medical Imaging*, 25, 987–1010. 15

Ono, M., Kubick, S. and Abernathey, C. (1990). *Atlas of the Cerebral Sulci*. Thieme. 8, 109, 141

Osher, S. and Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulations. *J of Comp Physics*, 79, 12–49. 16

Otsu, N. (1979). A threshold selection method from gray level histograms. *IEEE Trans. Systems, Man and Cybernetics*, 9, 62–66. 14

Ourselin, S. (2002). *Recalage d'images médicales par appariement de régions - Application à la construction d'atlas histologiques 3D*. Ph.D., Université de Nice Sophia-Antipolis. 21

Ourselin, S., Roche, A., Prima, S. and Ayache, N. (2000). Block matching: a general framework to improve robustness of rigid registration of medical images. *MICCAI*. 557–566. 21, 57

Pajevic, S. and Basser, P. J. (2003). Parametric and non-parametric statistical analysis of DT-MRI data. *Journal of Magnetic Resonance*, 161, 1–14. 67

Pal, N. R. and Pal, S. K. (1993). A review on image segmentation techniques. *Pattern Recognition*, 26, 1277–1294. 15

Paragios, N., Mellina-Gottardo, O. and Ramesh, V. (2004). Gradient vector flow fast geometric active contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, 402–407. 16

Park, H., Bland, P. H., Hero, A. O. and Meyer, C. R. (2005). Least biased target selection in probabilistic atlas construction. *MICCAI*, 8, 419–426. 33

Park, H., Bland, P. H. and Meyer, C. R. (2003). Construction of an abdominal probabilistic atlas and its application in segmentation. *IEEE Transactions on Medical Imaging*, 22, 483–492. 33

Pelillo, M., Siddiqi, K. and Zucker, S. W. (1999). Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 1105–1120. 29, 74

Pennec, X., Fillard, P. and Ayache, N. (2006). A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66, 41–66. 35, 59

Periaswamy, S. and Farid, H. (2003). Elastic registration with partial data. *Second International Workshop on Biomedical Image Registration*. Philadelphia, PA. 30

Periaswamy, S. and Farid, H. (2006). Medical image registration with partial data. *Medical Image Analysis*, 452–464. 30

Perona, P. and Freeman, W. (1998). A factorization approach to grouping. *Lecture Notes in Computer Science*, 1406, 655–?? 17

Petitjean, C. and Dacher, J.-N. N. (2011). A review of segmentation methods in short axis cardiac MR images. *Medical image analysis*, 15, 169–184. 15

Peyrat, J.-M. (2009). *Comparison of cardiac anatomy and function: statistics on fibre architecture from DT-MRI and registration of 4D CT Images*. Ph.D., Nice Sophia Antipolis University. 31, 55, 63, 64, 142, 202

Peyrat, J.-M., Sermesant, M., Pennec, X., Delingette, H., Xu, C., McVeigh, E. R. and Ayache, N. (2007). A computational framework for the statistical analysis of cardiac diffusion tensors: application to a small database of canine hearts. *IEEE Transactions on Medical Imaging*, 26, 1500–1514. 31, 33, 51, 55, 58, 59, 62, 68, 69, 142, 202, 204, 211

Peyrat, J.-M. M., Delingette, H., Sermesant, M., Pennec, X., Xu, C. and Ayache, N. (2008). Registration of 4D time-series of cardiac images with multichannel diffeomorphic demons. *MICCAI*. vol. 11, 972–979. 25

Peyrat, J.-M. M., Delingette, H., Sermesant, M., Xu, C. and Ayache, N. (2010). Registration of 4D cardiac CT sequences under trajectory constraints with multichannel diffeomorphic demons. *IEEE transactions on medical imaging*, 29, 1351–1368. 25

Pham, D. L., Xu, C. and Prince, J. L. (2000). A survey of current methods in medical image segmentation. *Annual Review of Biomedical Engineering*, vol. 2. 315–338. 15

Pienaar, R., Fischl, B., Caviness, V., Makris, N. and Grant, P. E. (2008). A methodology for analyzing curvature in the developing brain from preterm to adult. *Int. J. Imaging Syst. Technol.*, 18, 42–68. 114

Pluim, J. P., Maintz, A. B. and Viergever, M. A. (2003). Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, 22, 986–1004. 28

Pohl, K. M., Fisher, J., Grimson, W. E., Kikinis, R. and Wells, W. M. (2006). A bayesian model for joint segmentation and registration. *NeuroImage*, 31, 228–239. 30

Pohl, K. M., Fisher, J., Levitt, J. J., Shenton, M. E., Kikinis, R., Grimson, W. E. and Wells, W. M. (2005). A unifying approach to registration, segmentation, and intensity correction. *MICCAI*, 8, 310–318. 30

Poincaré, H. (1899). *Compte Rendu Academie des Sciences de Paris*, vol. 128. 1065–1069. 26

Rangarajan, A., Chui, H., Mjolsness, E., Pappu, S., Davachi, L., Goldman-Rakic, P. and Duncan, J. (1997). A robust point-matching algorithm for autoradiograph alignment. *Medical image analysis*, 1, 379–398. 23

Rapacchi, S., Croisille, P., Pai, V., Grenier, D., Viallon, M., Kellman, P., Mewton, N. and Wen, H. (2010). Reducing motion sensitivity in free breathing DWI of the heart with localized Principal Component Analysis. *International Society for Magnetic Resonance in Medicine*. 49, 135, 202, 210

Relan, J., Chinchapatnam, P., Sermesant, M., Rhode, K., Ginks, M., Delingette, H., Rinaldi, C. A., Razavi, R. and Ayache, N. (2011). Coupled personalization of cardiac electrophysiology models for prediction of ischaemic ventricular tachycardia. *Journal of the Royal Society Interface Focus*, 1, 396–407. 47

Reuter, M. (2009). Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *International Journal of Computer Vision*, 89, 287–308. 29, 73, 131, 140

Reuter, M., Wolter, F. E., Shenton, M. and Niethammer, M. (2009). Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer Aided Design*, 41, 739–755. 29, 74

Ribbens, A., Hermans, J., Maes, F., Vandermeulen, D. and Suetens, P. (2010). SPARC: Unified framework for automatic segmentation, probabilistic atlas construction, registration and clustering of brain MR images. *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 856–859. 32

Ribbens, A., Maes, F., Vandermeulen, D. and Suetens, P. (2009). An extensive validation study of non-rigid registration techniques for atlas based brain segmentation : local priors,

bias field correction and parametric intensity model. Technical report, Internal report KUL/ESAT/PSI/0904, K.U.Leuven, Leuven, Belgium. 32

Risser, L., Vialard, F., Wolz, R., Murgasova, M., Holm, D. D. and Rueckert, D. (2011). Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping. *IEEE Transactions on Medical Imaging*, 30, 1746–1759. 24

Risser, L., Vialard, F.-X., Murgasova, M., Holm, D. and Rueckert, D. (2010a). Large deformation diffeomorphic registration using fine and coarse strategies biomedical image registration. Springer Berlin / Heidelberg, Berlin, Heidelberg, vol. 6204 of *Lecture Notes in Computer Science*, chapiter 17. 186–197. 24

Risser, L., Vialard, F.-X. X., Wolz, R., Holm, D. D. and Rueckert, D. (2010b). Simultaneous fine and coarse diffeomorphic registration: application to atrophy measurement in alzheimer's disease. *MICCAI*, 13, 610–617. 24

Robles-Kelly, A. (2005). Segmentation via Graph-Spectral methods and riemannian geometry computer analysis of images and patterns. Springer Berlin / Heidelberg, Berlin, Heidelberg, vol. 3691 of *Lecture Notes in Computer Science*, chapiter 81. 661–668. 19, 123

Rohlfing, T., Maurer, C. R., Bluemke, D. A. and Jacobs, M. A. (2003). Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging*, 22, 730–741. 25

Rohmer, D., Sitek, A. and Gullberg, G. T. (2007). Reconstruction and visualization of fiber and laminar structure in the normal human heart from ex vivo diffusion tensor magnetic resonance imaging (DTMRI) data. *Investigative Radiology*, 42, 777–789. 49

Rohr, K., Fornefett, M. and Stiehl, H. S. (2003). Spline-based elastic image registration: integration of landmark errors and orientation attributes. *Comput. Vis. Image Underst.*, 90, 153–168. 21

Rohr, K., Stiehl, H. S., Sprengel, R., Buzug, T. M., Weese, J. and Kuhn, M. H. (2001). Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on Medical Imaging*, 20, 526–534. 23

Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V. and Hammers, A. (2006). Diffeomorphic registration using b-splines. *MICCAI*, 9, 702–709. 24

Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O. and Hawkes, D. J. (1999). Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18, 712–721. 22, 23

Rumpf, M. and Wirth, B. (2009). An elasticity approach to principal modes of shape variation. *Proceedings of the Second International Conference on Scale Space and Variational Methods in Computer Vision*. Springer-Verlag, Berlin, Heidelberg, SSVM '09, 709–720. 35

Rumpf, M. and Wirth, B. (2011). An Elasticity-Based covariance analysis of shapes. *International Journal of Computer Vision*, 92, 281–295. 35

Rushmer, R. F., Crystal, D. K. and Wagner, C. (1953). The functional anatomy of ventricular contraction. *Circulation Research*, 1, 162–170. 4, 5

Rusinkiewicz, S. and Levoy, M. (2001). Efficient variants of the ICP algorithm. *Proceedings on International Conference on 3-D Digital Imaging and Modeling*. 145–152. 72

Russel, I. K., Gotte, M. J. W., Bronzwaer, J. G., Knaapen, P., Paulus, W. J. and van Rossum, A. C. (2009). Left ventricular torsion: An expanding role in the analysis of myocardial dysfunction. *J Am Coll Cardiol Img*, 2, 648–655. 5

Saddi, K. A., Chefd'hotel, C. and Cheriet, F. (2007). Large deformation registration of contrast-enhanced images with volume-preserving constraint. SPIE, vol. 6512. 25

Sarkar, S. and Soundararajan, P. (2000). Supervised learning of large perceptual organization: graph spectral partitioning and learning automata. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 504–525. 17, 18

Schreiner, J., Asirvatham, A., Praun, E. and Hoppe, H. (2004). Inter-surface mapping. *SIGGRAPH*. 870–877. 72

Schwartz, B. I. (1985). *The World of Thought in Ancient China*. Belknap Press of Harvard University Press. 1

Schwartz, E. L., Shaw, A. and Wolfson, E. (1989). A numerical solution to the generalized mapmaker's problem: flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 1005–1008. 29, 74

Scollan, D. F., Holmes, A., Zhang, J. and Winslow, R. L. (2000). Reconstruction of cardiac ventricular geometry and fiber orientation using magnetic resonance imaging. *Annals of biomedical engineering*, 28, 934–944. 53

Scott, G. L. and Longuet-Higgins, H. C. (1991). An algorithm for associating the features of two images. *Proceedings of the Royal Society Biological Sciences*, 244, 21–26. 29, 74, 117, 143

Sederberg, T. W. and Parry, S. R. (1986). Free-form deformation of solid geometric models. *SIGGRAPH '86: Proceedings of the 13th annual conference on Computer graphics and interactive techniques*. ACM Press, New York, NY, USA, 151–160. 22

Seemann, G., Keller, D. U. J., Weiss, D. L. and Dossel, O. (2006). Modeling human ventricular geometry and fiber orientation based on diffusion tensor MRI. *Computers in Cardiology*. 801–804. 49

Seiler, C., Pennec, X. and Reyes, M. (2011). Geometry-aware multiscale image registration via OBBTree-based polyaffine log-demons. *MICCAI*, 14, 631–638. 151

Seo, S., Chung, M. K., Whyms, B. J. and Vorperian, H. K. (2011). Mandible shape modeling using the second eigenfunction of the Laplace-Beltrami operator. *Proceedings of SPIE*, 7962, 79620Z–79620Z–6. 29

Sermesant, M. (2003). *Modèle électromécanique du coeur pour l'analyse d'image et la simulation*. Ph.D., Université de Nice Sophia Antipolis. 152

Sermesant, M., Chabiniok, R., Chinchapatnam, P., Mansi, T., Billet, F., Moireau, P., Peyrat, J. M., Wong, K., Relan, J., Rhode, K., Ginks, M., Lambiase, P., Delingette, H., Sorine, M., Rinaldi, C. A., Chapelle, D., Razavi, R. and Ayache, N. (2011). Patient-specific electromechanical models of the heart for the prediction of pacing acute effects in CRT: A preliminary clinical validation. *Medical Image Analysis*. 47

Sermesant, M., Delingette, H. and Ayache, N. (2006). An electromechanical model of the heart for image analysis and simulation. *IEEE Trans Med Imaging*, 25, 612–625. 152

Sermesant, M., Forest, C., Pennec, X., Delingette, H. and Ayache, N. (2003). Deformable biomechanical models: Application to 4D cardiac image analysis. *Medical Image Analysis*, 7, 475–488. 152

Sermesant, M., Peyrat, J.-M. M., Chinchapatnam, P., Billet, F., Mansi, T., Rhode, K., Delingette, H., Razavi, R. and Ayache, N. (2008). Toward patient-specific myocardial models of the heart. *Heart failure clinics*, 4, 289–301. 152

Sethian, J. A. (1996). A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences of the United States of America*, 93, 1591–1595. 16

Sethian, J. A. (1999). *Level Set Methods and Fast Marching Methods*. Cambridge University Press, second edition. 17

Sezgin, M. and Sankur, B. (2004). Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic Imaging*, 13, 146–168. 14

Shannon, C. (1948). A mathematical theory of communication. Technical report, Bell System. 27

Shapiro, L. and Brady, J. (1992). Feature-based correspondence: an eigenvector approach. *Image and Vision Computing*, 10, 283–288. 29, 74, 117, 143

Shelton, C. R. (2000). Morphable surface models. *IEEE International Journal of Computer Vision*. vol. 38, 75–91. 72

Shi, J. and Malik, J. (1997). Normalized cuts and image segmentation. *International Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 731–737. 18, 19, 30

Shi, J. and Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 888–905. 18, 19, 29, 30, 73, 119, 123

Shi, Y., Dinov, I. and Toga, A. W. (2009). Cortical shape analysis in the Laplace-Beltrami feature space. *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*. 208–215. 29, 74

Shi, Y., Sun, B., Lai, R., Dinov, I. and Toga, A. W. (2010). Automated sulci identification via intrinsic modeling of cortical anatomy. *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*. 49–56. 29, 74, 140

Shokoufandeh, A., Macrini, D., Dickinson, S., Siddiqi, K. and Zucker, S. W. (2005). Indexing hierarchical structures using graph spectra. *, IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27, 1125–1140. 29, 73

Siddiqi, K., Shokoufandeh, A., Dickinson, S. J. and Zucker, S. W. (1999). Shock graphs and shape matching. *International Journal of Computer Vision*, 35, 13–32. 29, 74

Siegel, R. E. (1976). *Galen on the Affected Parts: Translation from the Greek Text With Explanatory Footnotes*. S Karger Pub. 1

Silver, M., Gotlieb, A. I. and Schoen, F. R. (2001). *Cardiovascular Pathology*. Churchill Livingstone, third edition. 203, 210

Simpson, I. J., Woolrich, M. W., Groves, A. R. and Schnabel, J. A. (2011). Longitudinal brain MRI analysis with uncertain registration. *MICCAI*, 14, 647–654. 35

Sinop, A. K. and Grady, L. (2007). A seeded image segmentation framework unifying graph cuts and random walker which yields a new algorithm. *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*. IEEE, 1–8. 19

Skrinjar, O. and Tagare, H. (2004). Symmetric, transitive, geometric deformation and intensity variation invariant nonrigid image registration. *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*. IEEE, 920–923 Vol. 1. 27

Smith, G. E. (1907). A new topographical survey of the human cerebral cortex, being an account of the distribution of the anatomically distinct cortical areas and their relationship to the cerebral sulci. *Journal of anatomy and physiology*, 41, 237–254. 8

Smith, N., de Vecchi, A., McCormick, M., Nordsletten, D., Camara, O., Frangi, A. F., Delingette, H., Sermesant, M., Relan, J., Ayache, N., Krueger, M. W., Schulze, W. H. W., Hose, R., Valverde, I., Beerbaum, P., Staicu, C., Siebes, M., Spaan, J., Hunter, P., Weese, J., Lehmann, H., Chapelle, D. and Rezavi, R. (2011). euHeart: personalized and integrated cardiac care using patient-specific cardiovascular modelling. *Interface Focus*, 1, 349–364. 31

Smith, N. P., Hunter, P. J. and Paterson, D. J. (2009). The cardiac physiome: at the heart of coupling models to measurement. *Experimental Physiology*, 94, 469–471. 31

So, R. and Chung, A. (2009). Multi-level non-rigid image registration using graph-cuts. *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*. IEEE, Washington, DC, USA, 397–400. 18

So, R. W. K. and Chung, A. C. S. (2010). Non-rigid image registration by using graph-cuts with mutual information. *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 4429–4432. 18

So, R. W. K. and Chung, A. C. S. (2011). Learning-based non-rigid image registration using prior joint intensity distributions with graph-cuts. *IEEE International Conference on Image Processing*. 709–712. 18

So, R. W. K., Tang, T. W. H. and Chung, A. C. S. (2011). Non-rigid image registration of brain magnetic resonance images using graph-cuts. *Pattern Recogn.*, 44, 2450–2467. 18

Sommer, S. (2011a). Accelerating multi-scale flows for LDDKBM diffeomorphic registration. *GPUCV workshop at ICCV*. 499–505. 24

Sommer, S. (2011b). *Anatomy in Curved Space Non-linear Modeling of Deformation and Shape for Medical Imaging*. Ph.D., Department of Computer Science, University of Copenhagen. 35

Sommer, S., Lauze, F., Nielsen, M. and Pennec, X. (2011a). Kernel bundle EPDiff: Evolution equations for Multi-Scale diffeomorphic image registration. *Scale Space and Variational Methods in Computer Vision*. 677–688. 24

Sommer, S., Nielsen, M., Lauze, F. and Pennec, X. (2011b). A multi-scale kernel bundle for LDDMM: towards sparse deformation description across space and scales. *Information processing in medical imaging : proceedings of the ... conference*, 22, 624–635. 24

Song, S. M. and Leahy, R. M. (2002). Computation of 3-D velocity fields from 3-D cine CT images of a human heart. *IEEE Transactions on Medical Imaging*, 10, 295–306. 22

Soundararajan, P. and Sarkar, S. (2001). Analysis of MinCut, average cut, and normalized cut measures. *Proc. Third Workshop Perceptual Organization in Computer Vision*. 17, 18

Soundararajan, P. and Sarkar, S. (2003). An in-depth study of graph partitioning measures for perceptual organization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25, 642–660. 17, 18

Spielman, D. A. (2010). Algorithms, graph theory, and linear equations in laplacian matrices. *PROCEEDINGS OF THE INTERNATIONAL CONGRESS OF MATHEMATICIANS 2010 (ICM 2010)*. 2698–2722. 19

Sprengel, R., Rohr, K. and Stiehl, H. S. (1996). Thin-plate spline approximation for image registration. *Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE*. vol. 3, 1190–1191. 23, 100

Steiner, D. C. and Morvan, J. M. (2003). Restricted Delaunay triangulations and normal cycle. *Proceedings of the Symposium on Computational Geometry*. 312–321. 92, 93

Streeter, D. (1979). *Gross Morphology and Fiber Geometry of the Heart*, American Physiology Society, vol. 1. 61–112. 4

Streeter, D. D., Spotnitz, H. M., Patel, D. P., Ross, J. and Sonnenblick, E. H. (1969). Fiber orientation in the canine left ventricle during diastole and systole. *Circulatory Research*, 24, 339–347. 5, 47, 48, 201, 209

Studholme, C. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32, 71–86. 28

Studholme, C. and Cardenas, V. (2004). A template free approach to volumetric spatial normalization of brain anatomy. *Pattern Recogn. Lett.*, 25, 1191–1202. 33

Studholme, C., Hill, D. L. G. and Hawkes, D. J. (1995). Multiresolution voxel similarity measures for MR-PET registration. Y. Bizais, C. Barillot and R. Di Paola, editors, *Proceedings of Information Processing in Medical Imaging*. Kluwer Academic Publishers, Ile de Berder, France, vol. 3, 287–298. 27

Sumner, R. W. and Popović, J. (2004). Deformation transfer for triangle meshes. *SIGGRAPH*. 399–405. 72, 90

Sun, J., Ovsjanikov, M. and Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. *Proceedings of the Symposium on Geometry Processing*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SGP '09, 1383–1392. 151

Sundar, H., Shen, D., Biros, G., Litt, H. and Davatzikos, C. (2006). Estimating myocardial fiber orientations by template warping. *International Symposium on Biomedical Imaging*. 73–76. 55, 142, 202

Talairach, J. and Tournoux, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain: 3-D Proportional System: An Approach to Cerebral Imaging (Thieme Classics)*. Thieme. 21

Tang, T. and Chung, A. (2007). Non-rigid image registration using graph-cuts. *Medical Image Computing and Computer-Assisted Intervention*. vol. 4791 of *Lecture Notes in Computer Science*, 916–924. 18

Terzopoulos, D. (1980). Matching deformable models to images: Direct and iterative solutions. *Optical Society of America, Topical Meeting on Machine Vision*. 164–167. 21, 139

Terzopoulos, D., Platt, J., Barr, A. and Fleischer, K. (1987). Elastically deformable models. *SIGGRAPH Comput. Graph.*, 21, 205–214. 21, 139

Tezuka, F. (1975). Muscle fiber orientation in normal and hypertrophied hearts. *The Tohoku journal of experimental medicine*, 117, 289–297. 210

Thirion, J. P. (1998). Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis*, 2, 243–260. 24, 25, 143

Thom, T., Haase, N., Rosamond, W., Howard, V. J., Rumsfeld, J., Manolio, T., Zheng, Z.-J., Flegal, K., O'donnell, C., Kittner, S., Lloyd-Jones, D., Goff, D. C. and Hong, Y. (2006). Heart disease and stroke statistics–2006 update. *Circulation*, 113, 85–151. 2, 209

Thompson, D. W. (1917). *On Growth and Form*. Courier Dover Publications. 22

Timoner, S. J., Golland, P., Kikinis, R., Shenton, M. E., Eric and Wells, W. M. (2002). Performance issues in shape classification. *Proceedings of the 5th International Conference on Medical Image Computing and Computer-Assisted Intervention-Part I*. 355–362. 35

Tlusty, T. (2010). A relation between the multiplicity of the second eigenvalue of a graph Laplacian, Courant's nodal line theorem and the substantial dimension of tight polyhedral surfaces. *Electronic Journal Of Linear Algebra*, 16, 315–324. 119

Toga, A. W. and Thompson, P. M. (2001). The role of image registration in brain mapping. *Image and vision computing*, 19, 3–24. 31

Toro, R. and Burnod, Y. (2005). A morphogenetic model for the development of cortical convolutions. *Cerebral Cortex*, 15, 1900–1913. 7

Torrent-Guasp, F., Kocica, M. J., Corno, A. F., Komeda, M., Carreras-Costa, F., Flotats, A., Cosin-Aguillar, J. and Wen, H. (2005). Towards new understanding of the heart structure and function. *European Journal of Cardio-Thoracic Surgery*, 27, 191–201. 4, 5, 48, 142

Toussaint, N., Sermesant, M., Stoeck, C. T., Kozerke, S. and Batchelor, P. G. (2010). In vivo human 3D cardiac fibre architecture: reconstruction using curvilinear interpolation of diffusion tensor images. *MICCAI*. 418–425. 70

Trayanova, N. A. (2011). Whole-Heart modeling. *Circulation Research*, 108, 113–128. 152

Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W. E. and Willsky, A. (2003). A shape-based approach to the segmentation of medical imagery using level sets. *IEEE transactions on medical imaging*, 22, 137–154. 30, 32

Tseng, W.-Y. Y., Wedeen, V. J., Reese, T. G., Smith, R. N. and Halpern, E. F. (2003). Diffusion tensor MRI of myocardial fibers and sheets: correspondence with visible cut-face texture. *Journal of Magnetic Resonance Imaging*, 17, 31–42. 53, 205

Turk, M. and Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3, 71–86. 35, 107

Umeyama, S. (1988). An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10, 695–703. 29, 74, 117

Unal, G. and Slabaugh, G. (2005). Coupled PDEs for non-rigid registration and segmentation. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*. IEEE Computer Society, vol. 1, 168–175. 30

Vadakkumpadan, F., Arevalo, H., Ceritoglu, C., Miller, M. and Trayanova, N. (2012). Image-Based estimation of ventricular fiber orientations for personalized modeling of cardiac electrophysiology. *IEEE transactions on medical imaging*. 152

van den Elsen, P. A., Maintz, J. B. A., Pol, E. J. D. and Viergever, M. A. (1995). Automatic registration of CT and MR brain images using correlation of geometrical features. *IEEE Transactions on Medical Imaging*, 14, 384–396. 27

van den Elsen, P. A., Pol, E. J. D. and Viergever, M. A. (1993). Medical image matching-a review with classification. *Engineering in Medicine and Biology Magazine, IEEE*, 12, 26–39. 21

van Kaick, O., Zhang, H., Hamarneh, G. and Cohen-Or, D. (2011). A survey on shape correspondence. *Eurographics*, 30, 1681–1707. 29, 74, 140

van Leemput, K., Maes, F., Vandermeulen, D. and Suetens, P. (1999). Automated model-based bias field correction of MR images of the brain. *IEEE transactions on medical imaging*, 18, 885–896. 32

van Leeuwenhoek, A. (1694). Correspondence with the members of the royal society in london. *Letter no 82*. 202

Veksler, O. (1999). *Efficient Graph-Based Energy Minimization Methods in Computer Vision*. Ph.D., Cornell University. 18

Venot, A. and Leclerc, V. (1984). Automated correction of patient motion and gray values prior to subtraction in digitized angiography. *IEEE Trans Med Imaging*, 3, 179–186. 21

Vercauteren, T., Pennec, X., Perchant, A. and Ayache, N. (2007). Non-parametric diffeomorphic image registration with the demons algorithm. *MICCAI*. Springer Berlin Heidelberg, Berlin, Heidelberg, vol. 4792 of *Lecture Notes in Computer Science*, 319–326. 24, 25, 116, 121, 122, 143

Vercauteren, T., Pennec, X., Perchant, A. and Ayache, N. (2008). Symmetric Log-Domain diffeomorphic registration: A Demons-Based approach. *MICCAI*. vol. 5241, 754–761. 24, 25, 26, 27, 51, 57, 58, 116, 121, 124, 142, 143, 204, 211

Vercauteren, T., Pennec, X., Perchant, A. and Ayache, N. (2009a). Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage*, 45, S61–S72. 24, 25, 27, 121, 142, 143, 211

Vercauteren, T., Pennec, X., Perchant, A. and Ayache, N. (2009b). Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage*, 45, S61–S72. 57

Vesalius, A. (1543). *De humani corporis fabrica libri septem*. Joannis Oporini, Basel. 1

Vincent, L. and Soille, P. (1991). Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13, 583–598. 14

Viola, P. and Wells, W. M. (1995). Alignment by maximization of mutual information. *Computer Vision, 1995. Proceedings, Fifth International Conference on*. 16–23. 27

von Economo, C. and Koskinas, G. N. (1925). *Die Cytoarchitektonik der Hirnrinde des erwachsenen Menschen*. Julius Springer, Vienna. 8

von Economo, C. and Koskinas, G. N. (2007). *Atlas of Cytoarchitectonics of the Adult Human Cerebral Cortex: With Full-Scale Reproductions of the Original 112 Microphotographic Plates*. S Karger Pub. 8

von Segesser, L. K. (2005). The myocardial band: fiction or fact? *European Journal of Cardio-Thoracic Surgery*, 27, 181–182. 48, 142

Wahba, G. (1990). *Spline models for observational data*, vol. 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA. 23

Wang, Y. and Staib, L. H. (1998). Elastic model based non-rigid registration incorporating statistical shape information. *MICCAI*. vol. 1496, 1162–1173. 22

Wang, Y. and Staib, L. H. (2000). Physical model-based non-rigid registration incorporating statistical shape information. *Medical Image Analysis*, 4, 7–20. 22

Wei, Y.-C. and Cheng, C.-K. (1989). Towards efficient hierarchical designs by ratio cut partitioning. *Computer-Aided Design, 1989. ICCAD-89. Digest of Technical Papers., 1989 IEEE International Conference on*. IEEE, 298–301. 17, 18

Weiss, Y. (1999). Segmentation using eigenvectors: a unifying view. *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*. IEEE, vol. 2, 975–982 vol.2. 17

Welker, W. (1990). Why does cerebral cortex fissure and fold? A review of determinants of gyri and sulci. *Cerebral cortex*, 8, 3–136. 8, 109, 141

Wertheimer, M. (1925). Über gestalttheorie. *Philosophische Zeitschrift für Forschung und Aussprache*, 1, 39–60. 17

Wertheimer, M. (1938). *Laws of Organization in Perceptual Forms*, Harcout, Brace and Company. 71–88. 17

West, J. B. (2008). Ibn al-Nafis, the pulmonary circulation, and the islamic golden age. *Journal of Applied Physiology*, 105, 1877–1880. 1

Wilkins, R. (1992). *Neurosurgical Classics*. Thieme/AANS. 1

Williams, L. R. and Thornber, K. K. (1998). A comparison of measures for detecting natural shapes in cluttered backgrounds. H. Burkhardt and B. Neumann, editors, *ECCV'98*. Springer Berlin / Heidelberg, Berlin/Heidelberg, vol. 1407 of *Lecture Notes in Computer Science*, 432–448. 17

Willis, T. (1664). *Cerebri Anatome: Cui Accessit Nervorum Descriptio Et Usus*. 1

Wolz, R., Heckemann, R. A., Aljabar, P., Hajnal, J. V., Hammers, A., Lötjönen, J., Rueckert, D. and Alzheimer's Disease Neuroimaging Initiative (2010). Measurement of hippocampal

atrophy using 4D graph-cut segmentation: application to ADNI. *NeuroImage*, 52, 109–118. 18

Woods, R. P., Cherry, S. R. and Mazziotta, J. C. (1992). Rapid automated algorithm for aligning and reslicing PET images. *J Comput Assist Tomogr*, 16, 620–633. 27

Wu, M.-T., Tseng, W.-Y. I., Su, M.-Y. M., Liu, C.-P., Chiou, K.-R., Wedeen, V. J., Reese, T. G. and Yang, C.-F. (2006). Diffusion tensor magnetic resonance imaging mapping the fiber architecture remodeling in human myocardium after infarction: correlation with viability and wall motion. *Circulation*, 114, 1036–1045. 47, 142, 210

Wu, Z. and Leahy, R. (1990). Tissue classification in MR images using hierarchical segmentation. *Nuclear Science Symposium, 1990. Conference record : Including Sessions on Nuclear Power Systems and Medical Imaging Conference, 1990 IEEE*. IEEE, 1410–1414. 17

Wu, Z. and Leahy, R. (1993). An optimal graph theoretic approach to data clustering: theory and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15, 1101–1113. 17

Wuhrer, S., Shu, C. and Bose, P. (2009). Posture invariant correspondence of triangular meshes in shape space. *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*. 1574–1581. 29, 74

Xu, C. and Prince, J. L. (1997). Gradient vector flow: A new external force for snakes. *CVPR '97: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)*. IEEE Computer Society, Washington, DC, USA. 16

Xu, N., Bansal, R. and Ahuja, N. (2003). Object segmentation using graph cuts based active contours. *CVPR*. vol. 2, 46–53. 18

Yeo, B. T. T., Sabuncu, M. R., Vercauteren, T., Ayache, N., Fischl, B. and Golland, P. (2010a). Spherical demons: Fast diffeomorphic Landmark-Free surface registration. *IEEE Trans on Med Imaging*, 29, 650–668. 25, 131, 139

Yeo, B. T. T., Sabuncu, M. R., Vercauteren, T., Holt, D. J., Amunts, K., Zilles, K., Golland, P. and Fischl, B. (2010b). Learning task-optimal registration cost functions for localiz-

ing cytoarchitecture and function in the cerebral cortex. *IEEE Transactions on Medical Imaging*, 29, 1424–41. 98

Yezzi, A., Tsai, A. and Willsky, A. (1999). A statistical approach to snakes for bimodal and trimodal imagery. *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*. IEEE Computer Society, Washington, DC, USA. 16

Yezzi, A., Zollei, L. and Kapur, T. (2001). A variational framework for joint segmentation and registration. *Mathematical Methods in Biomedical Image Analysis, 2001. MMBIA 2001. IEEE Workshop on*. 44–51. 30

Yezzi, A., Zöllei, L. and Kapur, T. (2003). A variational framework for integrating segmentation and registration through active contours. *Med Image Anal*, 7, 171–185. 30

Young, A. A. and Frangi, A. F. (2009). Computational cardiac atlases: from patient to population and back. *Experimental Physiology*, 94, 578–596. 31

Young, I. (1978). Further consideration of sample and feature size (corresp.). *IEEE Transactions on Information Theory*, 24, 773–775. 67

Yu, P., Grant, P. E., Qi, Y., Han, X., Segonne, F., Pienaar, R., Busa, E., Pacheco, J., Makris, N., Buckner, R. L., Golland, P. and Fischl, B. (2007). Cortical surface shape analysis based on spherical wavelets. *IEEE Transactions on Medical Imaging*, 26, 582–597. 114

Zahn, C. T. (1971). Graph-Theoretical methods for detecting and describing gestalt clusters. *Computers, IEEE Transactions on*, C-20, 68–86. 17

Zhang, H., Fritts, J. E. and Goldman, S. A. (2008a). Image segmentation evaluation: A survey of unsupervised methods. *Computer Vision and Image Understanding*, 110, 260–280. 15

Zhang, H., Sheffer, A., Or, D. C., Zhou, Q., van Kaick, O. and Tagliasacchi, A. (2008b). Deformation-driven shape correspondence. *Proceedings of the Symposium on Geometry Processing*. 1431–1439. 72

Zhang, H., van Kaick, O. and Dyer, R. (2010). Spectral mesh processing. *Eurographics*, 29, 1865–1894. 29, 74, 140

Zhang, Z. (1994). Iterative point matching for registration of free-form curves and surfaces. *International Journal of Computer Vision*, 13, 119–152. 30

Zheng, Y. and Doermann, D. (2006). Robust point matching for nonrigid shapes by preserving local neighborhood structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28, 643–649. 101

Zhukov, L. and Barr, A. H. (2003). Heart-muscle fiber reconstruction from diffusion tensor MRI. *Visualization*. 597–602. 49

Zitova, B. (2003). Image registration methods: a survey. *Image and Vision Computing*, 21, 977–1000. 21, 144

Zollei, L., Miller, L. E., Grimson, W. E. L. and Iii, w. (2005). Efficient population registration of 3D data. *ICCV 2005, Computer Vision for Biomedical Image Applications*. 33

APPENDIX

APPENDIX A VARIABILITY OF THE CARDIAC LAMINAR STRUCTURE

Herve Lombaert, Jean-Marc Peyrat, Laurent Fanton, Farida Cheriet, Hervé Delingette, Nicholas Ayache, Patrick Clarysse, Isabelle Magnin, Pierre Croisille

Presentation

This appendix presents the article *"Variability of the Human Cardiac Laminar Structure"* (Lombaert *et al.*, 2011c) published in the STACOM Workshop (Statistical Atlas and Computational Models of the Heart) of the conference **MICCAI** *(Medical Image Computing and Computer Assisted Intervention)* held in Toronto, ON, in September 2011. The objective of this article is to further study the variability of the laminar structure of the cardiac fibers in humans. This is a joint work between **INRIA, Sophia Antipolis**, *France*; **CREATIS**, *Lyon, France*; and **École Polytechnique**, *Montreal, Canada*.

abstract

The cardiac fiber architecture has an important role in electrophysiology, in mechanical functions of the heart, and in remodeling processes. The variability of the fibers is the focus of various studies in different species. However, the variability of the laminar sheets is still not well known especially in humans. In this paper, we present preliminary results on a quantitative study on the variability of the human cardiac laminar structure. We show that the laminar structure has a complex variability and we show the possible presence of two populations of laminar sheets. Bimodal distributions of the intersection angle of the third eigenvector of the diffusion tensor have been observed in 10 *ex vivo* healthy human hearts. Additional hearts will complete the study and further characterize the different populations of cardiac laminar sheets.

A.1 Introduction

The heart is a complex muscle that is composed with myocardial fibers organized as laminar sheets (Streeter *et al.*, 1969; LeGrice *et al.*, 1995). The cardiac fiber structures have an important role in electrophysiology (Hooks *et al.*, 2002) and in mechanical functions (Costa *et al.*, 2001) of the heart. The understanding of the cardiac fiber architecture is essential for better diagnosis and treatment of many cardiac pathologies. The fibrous nature of the heart

has been known for centuries, tracing back to as early as 1694 (van Leeuwenhoek, 1694), but has been limited to tedious histological studies (Nielsen *et al.*, 1991). The cardiac fiber structure can now be imaged with diffusion tensor magnetic resonance imaging (DT-MRI) (Basser *et al.*, 1994; Hsu and Henriquez, 2001), however the variability of the fiber structure in humans is still not well known (due to the very limited number and the value of post-mortem healthy human hearts) and is largely speculated from studies on other species (dogs (Helm *et al.*, 2005a,b; Helm, 2005; Sundar *et al.*, 2006; Peyrat *et al.*, 2007; Peyrat, 2009; Gilbert *et al.*, 2007), goats (Geerts *et al.*, 2002), and rats (Bishop *et al.*, 2009)). Recently, Lombaert *et al.* (Lombaert *et al.*, 2011b,e) constructed a statistical atlas of the human cardiac fiber architecture and assessed its variability. The fiber structure is shown to be more stable than the laminar sheet structure. They hypothesized that the higher variability of the laminar sheet could be due to the presence of two or more populations of laminar sheets (Helm, 2005). Helm *et al.* studied the variability of the cardiac laminar sheet in (Helm *et al.*, 2005b). Using 7 canine hearts, they observed a bimodal distribution of intersection angles (i.e., two populations of laminar sheet structure) in most myocardial segments of the left ventricle.

We present here the preliminary results of a study on the variability of the cardiac laminar sheet structure in humans. The methods used to construct and analyze the statistical atlas are briefly described. Next, the results show the angular variability, from the average healthy heart, of the the laminar sheet normal. The complexity of the laminar sheet structure is revealed thereafter by analyzing the distribution of the intersection angle of the laminar sheet normal. The distributions suggest the possible presence of two populations of laminar sheets in several myocardial segments of the left ventricle.

A.2 Material and Method

A.2.1 Dataset

The human dataset (Frindel *et al.*, 2009; Rapacchi *et al.*, 2010) consists of 10 healthy *ex vivo* human hearts acquired during forensic autopsies. The excised hearts were placed in a plastic container and filled with non destructive hydrophilic gel to maintain a diastolic shape. The images have been acquired on a 1.5T MR scanner (Avanto Siemens), all within 24 hours after death and prior to the examination by the pathologist, with a bipolar echo planar imaging

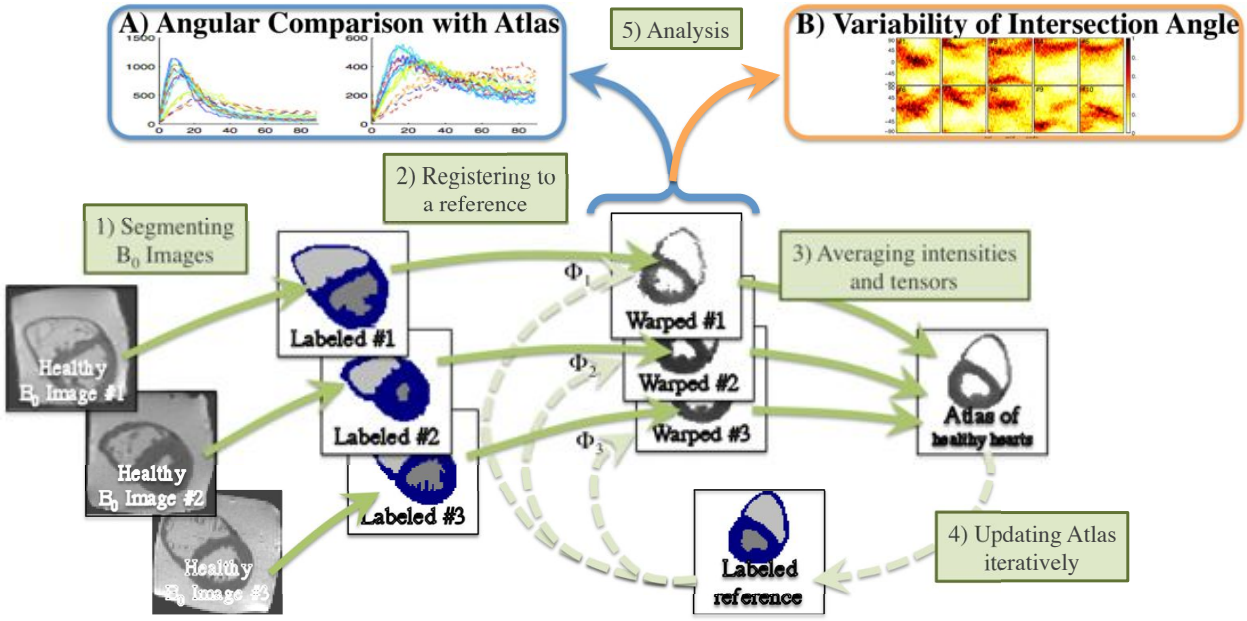

Figure A.1: *Atlas Construction*: *(1)* From the acquired images, the myocardia are segmented. *(2)* Images are then aligned and registered nonrigidly toward a reference image. *(3,4)* The atlas is constructed iteratively by averaging acquired images in the average heart shape. *Variability of the laminar sheets*: *(A)* The directions of the laminar sheet normals are compared with the atlas for each heart, and *(B)* the probability distribution of the intersection angle is analyzed.

using 4 repetitions of 12 gradient images. The diffusion-weighted images, from which are estimated the diffusion tensors, are of size 128x128x52 with an isotropic resolution of 2 mm. All cases are from extra cardiac sudden deaths, and the hearts are classified as healthy after controlling their weight, wall thickness, and subsequent pathology examination (Silver *et al.*, 2001).

A.2.2 Atlas Construction

The statistical atlas is constructed using four steps, all fully described in (Lombaert *et al.*, 2011b) and summarized here in Fig. A.1. Information on the fiber architecture (i.e., any directional data from DT-MRI) is purposely omitted from the registration process in order to avoid introducing any bias in the study of the fiber variability.

Myocardium Segmentation

— *Firstly*, the myocardium of each heart is segmented out on the B_0 image of the DT-MRI acquisition. The segmentation method is based on Graph Cuts (Boykov and Jolly, 2000).

Heart Registration

— *Secondly*, each myocardium is registered to a reference image using solely the B_0 images and the myocardial masks. The pairwise registrations are performed with the symmetric Log-domain diffeomorphic demons (Vercauteren *et al.*, 2008; Mansi *et al.*, 2010).

Construction of Healthy Atlas

— *Thirdly*, the reference image is deformed toward the morphological average of all hearts by iterating until convergence the pairwise registrations and the heart averaging steps. This atlas construction follows Guimond's *et al.* method (Guimond *et al.*, 2000).

Warping of Diffusion Tensors

— *Fourthly*, and last, the resulting deformation fields computed from the registration process are used to warp all tensor fields to the morphological atlas. The diffusion tensors are reoriented using the Finite Strain strategy since it preserves the geometric features (Peyrat *et al.*, 2007).

A.2.3 Statistical Analysis

The diffusion tensor space of symmetric positive definite matrices does not have a vector space structure with the standard Euclidean metric. The Log-Euclidean metric (Arsigny *et al.*, 2006b) provides a simple and fast framework where first order arithmetic on diffusion tensors has a closed form solution. The average diffusion tensor field, $\overline{\mathbf{D}}$, is computed from the N warped tensor fields $\{\mathbf{D}^{(i)}\}_{i=1\ldots N}$ (with $N = 10$ healthy hearts) using the Fréchet mean:

$$\overline{\mathbf{D}} = \exp\left(\frac{1}{N}\sum_{i=1}^{N}\log(\mathbf{D}^{(i)})\right) \tag{A.1}$$

The eigendecomposition of the average diffusion tensor $\overline{\mathbf{D}}$ gives the three average eigenvectors $\overline{\mathbf{v}}_{1,2,3}$. The maximal local diffusion, revealed by the primary eigenvector $\mathbf{v}_1$ occurs along

the fiber while most of the remaining diffusion occurs within the laminar sheet, where the secondary eigenvector $\mathbf{v}_2$ is thought to lay (Helm *et al.*, 2005a,b; Tseng *et al.*, 2003). The tertiary eigenvector, $\mathbf{v}_3$, corresponds to the normal of the laminar sheet.

The deviation of the cardiac laminar sheet of each heart is given with the angular difference θ from the direction of laminar sheet normal, $\mathbf{v}_3$, to the direction of the average laminar sheet normal, $\overline{\mathbf{v}}_3$. For instance, for the i^{th} heart, the angular deviation from the average heart is:

$$\theta_3^{(i)} = \arccos\left(\frac{|\mathbf{v}_3^{(i)} \cdot \overline{\mathbf{v}}_3|}{\| \mathbf{v}_3^{(i)} \| \| \overline{\mathbf{v}}_3 \|}\right) \tag{A.2}$$

The angles are defined between 0° and 90°. The absolute value of the dot product removes the inherent ambiguity in the orientation of the eigenvectors (i.e., $|a{\cdot}b| = |a{\cdot}(-b)|$). The variability of the laminar sheet can be measured with the probability distribution of the intersection angle of the third eigenvector (i.e., of the laminar sheet normal). The intersection angle (LeGrice *et al.*, 1995) is defined as the projected angle of the laminar sheet normal (in red in the right figure) onto a transverse plane (the vertical transmural plane in green in the right figure). A prolate ellipsoidal model of the heart (Nielsen *et al.*, 1991) is fitted to the morphology of the statistical atlas to ease measurements in the prolate ellipsoidal coordinates.

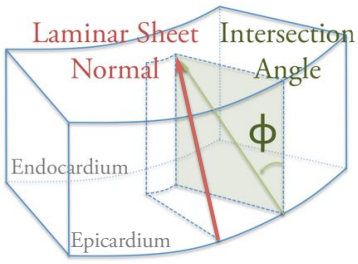

Figure A.2: Intersection angle of the 3rd eigenvector in a myocardial section.

A.3 Results

The cardiac laminar sheet was shown (Lombaert *et al.*, 2011b) to vary more than the fiber direction. In order to understand the higher variability, the distribution of the intersection angle of the laminar sheet normal is estimated in all hearts and in several myocardial segments. The distributions show the presence of possibly two populations of laminar sheets.

A.3.1 Variability of the Laminar Sheet Normal

The direction of the laminar sheet normal in each heart is compared with the ones of the average healthy heart (i.e., the atlas). The angular differences of the laminar sheet normals, given by Eq. A.2 and shown in Fig. A.4, present deviations to the average heart in several

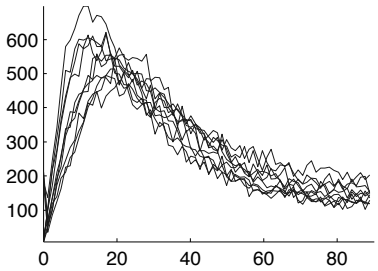

Figure A.3: Histograms of the angular deviation θ_3 (in degrees) for 10 hearts.

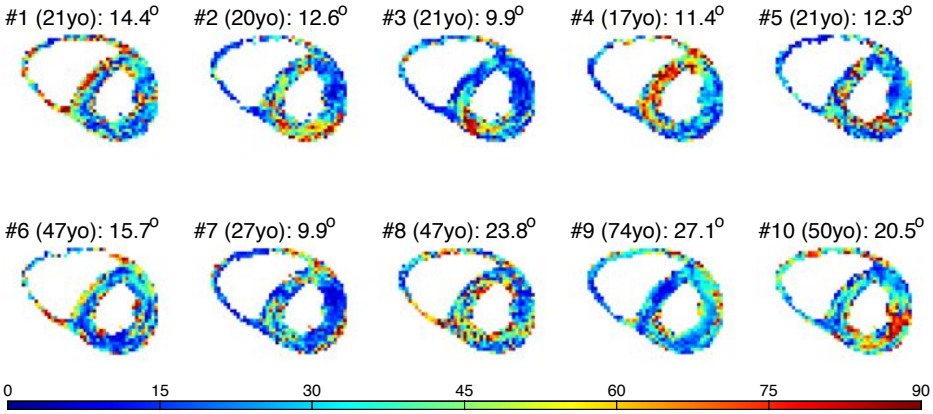

Figure A.4: Deviation, θ_3, of the laminar sheet normal of each heart to the atlas. The coloring is the angular difference in degree.

areas for each heart. The histogram of the angular differences, in Fig. A.3, shows an angular peak at $\overline{\theta}_3 = 15.77°$ (the average of the histogram modes in Fig. A.3).

A.3.2 Variability of the Intersection Angle

We now study the probability distribution of the intersection angle of the third eigenvector (i.e., the laminar sheet normal). The probability distributions are presented in a joint histograms (Fig. A.5) where the angle distribution, on the vertical axis, is plotted against all transmural distances, on the horizontal axis. Each heart appears to have a consistent distribution of laminar sheet normal directions with angles concentrated around a specific mean. Subject #3, #6, #9, and #10 appear to show two populations of laminar sheet normals (i.e., the angles are concentrated along two horizontal curves). The global joint histogram in Fig. A.5 (b) shows the probability distribution of the intersection angle (i.e., the variability of the laminar sheet normal) among all 10 hearts. Furthermore, the probability distributions

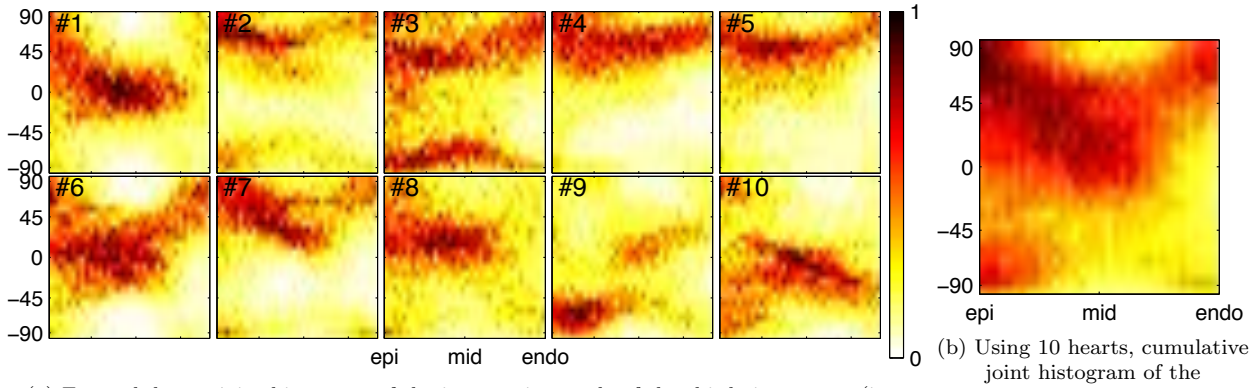

(a) For each heart, joint histogram of the intersection angle of the third eigenvector (i.e., laminar sheet variability)

(b) Using 10 hearts, cumulative joint histogram of the intersection angle.

Figure A.5: Joint histograms showing the distribution of the intersection angle from epicardium *(left side of each histogram)* to endocardium *(right side of each histogram)* for *(a)* each heart, and *(b)* all hearts combined. The x-axis is the transmural distance from epicardium. The y-axis is the distribution of the intersection angles observed at one specific transmural distance (i.e., each column is the histogram of angles for one given distance). Color is the normalized probability distribution.

in the 17 AHA segments (American Heart Association (Cerqueira *et al.*, 2002)) provide local statistics across the myocardium. More distinct clusters of laminar sheet structures are visible in Fig. A.6, in particular AHA zones 2, 3, 4, 7, 8, 9, 12, 13, and 14 show angular distributions concentrated along two horizontal curves. These curves of average angles can be estimated using Gaussian Mixture Models (i.e., for each transmural distance, the intersection angle values are clustered into two Gaussian models). This is illustrated with two blue curves in each joint histogram. Each curves indicates the estimated mean angle of one of the two Gaussian models.

A.4 Conclusion

In this paper, preliminary results of a study on variability of the human cardiac laminar structure have been presented. The cardiac fiber architecture has an important role in electrophysiology and in mechanical functions of the heart. The variability of the laminar sheets in humans is still not well known. It is thought that there are two populations of laminar sheets. Helm *et al.* (Helm *et al.*, 2005b) observed in 7 canine hearts a bimodal distribution of intersection angles of the third eigenvector (i.e., the laminar sheet normal). We similarly observed a bimodal distribution of intersection angles in human hearts. Our

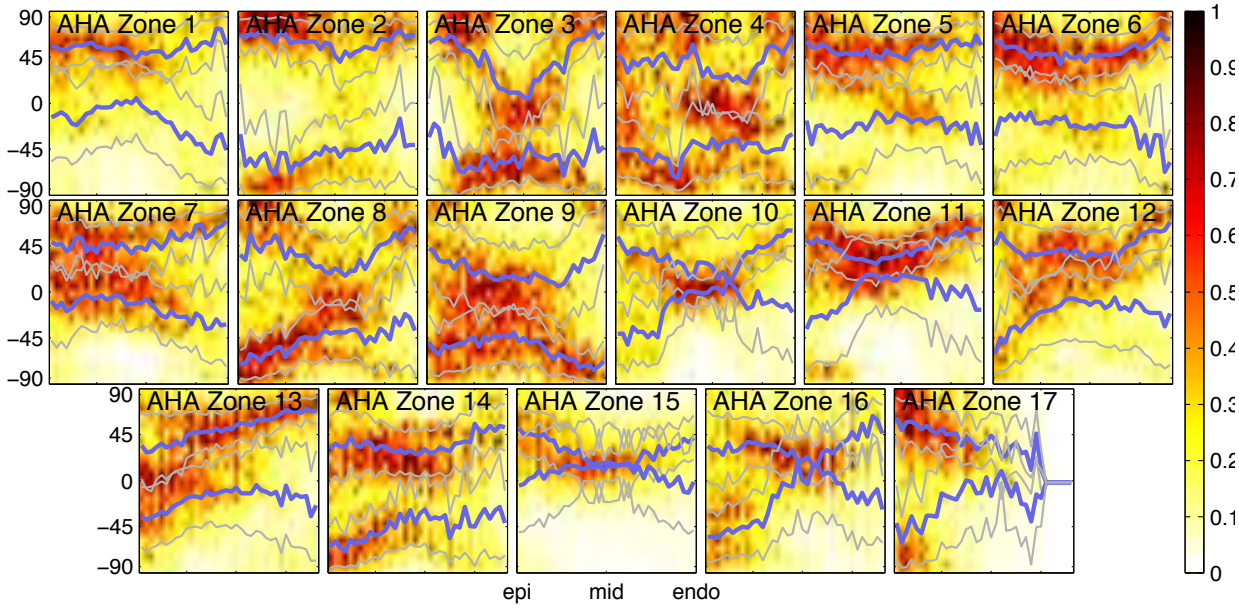

Figure A.6: Joint histograms showing of the distribution of the intersection angle in 17 AHA LV segments. Blue curves (mean angles) are found using GMM (gray lines are the one-standard-deviation envelopes). Two populations of intersection angles are visible in most segments.

preliminary results within the dataset of 10 hearts suggest the possible presence of two populations of laminar sheets. We will include additional hearts to the study and try to further characterize the different populations of cardiac laminar sheets.

Acknowledgment

The authors wish to acknowledge helpful comments from Leon Axel and the members of the Asclepios Team. The project was supported financially by the National Science and Engineering Research Council of Canada (NSERC), the Michael Smith Scholarship (CGS-MSFSS), and the EGIDE/INRIA Scholarship.

APPENDIX B COMPARISON WITH ABNORMAL HEARTS

Herve Lombaert, Jean-Marc Peyrat, Laurent Fanton, Farida Cheriet, Hervé Delingette, Nicholas Ayache, Patrick Clarysse, Isabelle Magnin, Pierre Croisille

Presentation

This appendix presents the article *"Statistical Atlas of Human Cardiac Fibers: Comparison with Abnormal Hearts"* (Lombaert *et al.*, 2011d) published in the STACOM Workshop (Statistical Atlas and Computational Models of the Heart) of the conference **MICCAI** *(Medical Image Computing and Computer Assisted Intervention)* held in Toronto, ON, in September 2011. The objective of this article is to investigate whether the cardiac fiber architecture is related to abnormality of hearts. This is a joint work between **INRIA, Sophia Antipolis**, *France*; **CREATIS**, *Lyon, France*; and **École Polytechnique**, *Montreal, Canada*.

abstract

Criteria of normality of the cardiac fibers are important in cøardiomyopathies. In this paper, we investigate the differences in the cardiac fiber structures between 10 hearts classified as healthy and 6 hearts classified as abnormal, and determine if properties of the cardiac fiber structures can be discriminants for abnormality. We compare the variability of the fiber directions from abnormal hearts to an atlas of healthy hearts. The human atlas of the cardiac fiber structures is built with an automated framework based on symmetric Log-domain diffeomorphic demons. We study the angular variability of the different fiber structures. Our preliminary results might suggest that a higher variability of the fiber structure directions could possibly characterize abnormality of a heart.

B.1 Introduction

Cardiovascular diseases are by far the number one killer in the US with over 930,000 deaths annually and 71 millions, more than a fifth of the population, live with a form of cardiovascular disease (Thom *et al.*, 2006). The characterization of the consequences led by specific cardiopathies is essential to a better diagnosis and a better treatment of these diseases. Among the possible causes, the differences in the cardiac fiber architecture could be an promising topic. The heart is composed of myocardial fibers organized in a complex laminar structure (Streeter *et al.*, 1969; LeGrice *et al.*, 1995), and the cardiac fiber structures

have an important role in electrophysiology (Hooks *et al.*, 2002), in mechanical functions (Costa *et al.*, 2001), and in remodeling (Wu *et al.*, 2006) of the heart. Changes in the fiber structures are for instance inherent in myocardial hypertrophy (Karsner *et al.*, 1925; Tezuka, 1975; Grajek *et al.*, 1993). Myocardial disarray, or disorganisation of the fibers, is also still the focus of contentious studies (Becker and Caruso, 1982). The question of normality of the cardiac fiber structures arises when trying to assess the role of myocardial disarray in cardiomyopathies. In this paper, we try to assess whether there is a difference in the cardiac fiber structures between hearts classified as normal and hearts considered as abnormal.

The directions of the fiber structures and their variability can be measured with Diffusion Tensor Imaging (DT-MRI). A human atlas of the cardiac fiber structures from DT-MRI (Lombaert *et al.*, 2011b,e) has recently been built with 10 healthy *ex vivo* hearts. We register 6 *ex vivo* hearts classified as abnormal to the atlas of healthy hearts and analyze the angular differences between the fiber structure directions of the abnormal hearts and the ones of the average healthy heart. The statistical study shows that the directions of the cardiac fiber structures vary more in abnormal hearts than in healthy hearts. The preliminary results might suggest that a higher variability of the fiber structure directions could possibly characterize abnormality.

B.2 Material and Method

B.2.1 Dataset

The human dataset (Frindel *et al.*, 2009; Rapacchi *et al.*, 2010) consists of 10 healthy and 6 abnormal *ex vivo* human hearts acquired during forensic autopsies. All cases are from extra cardiac sudden deaths. However, the true nature of deaths is not available. The images have been acquired on a 1.5T MR scanner (Avanto Siemens), all within 24 hours after death and prior to the examination by the pathologist, with a bipolar echo planar imaging using 4 repetitions of 12 gradient images. The diffusion-weighted images, from which are estimated the diffusion tensors, are of size 128x128x52 with an isotropic resolution of 2 mm. Criteria of abnormality (Silver *et al.*, 2001) are based on the heart weight (with given permitted weight limits within the 95% percentile), the septal thickness (with a maximal thickness defined at 12 mm for women and 14 mm for men), and on subsequent pathology examination.

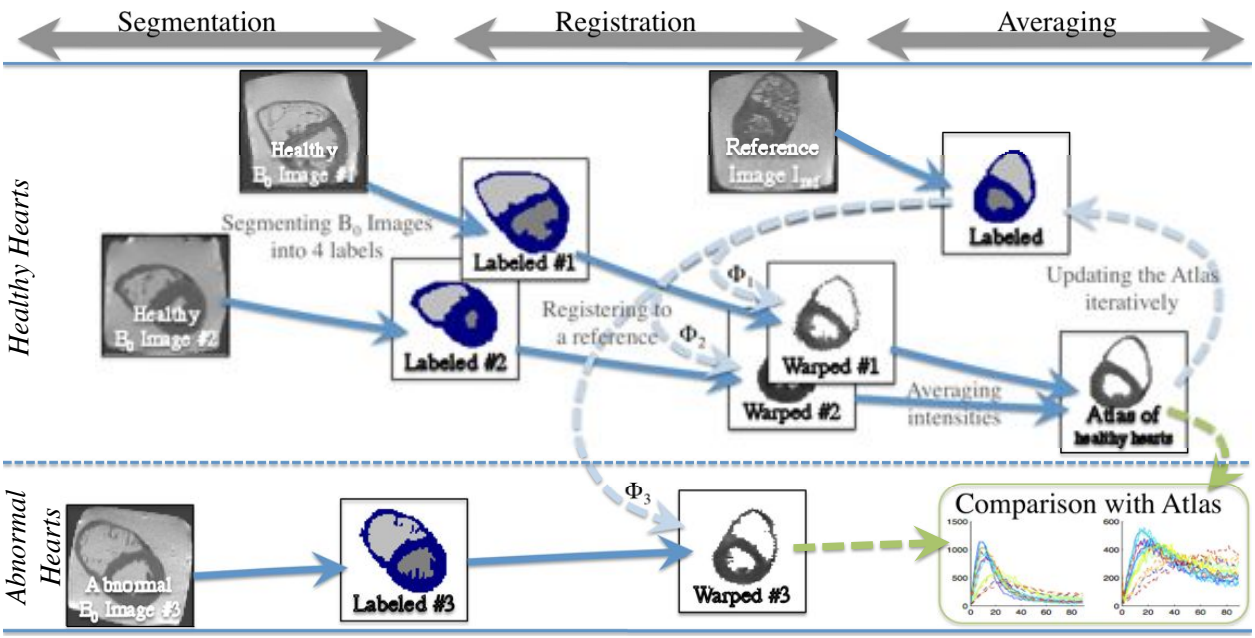

Figure B.1: *Construction of the healthy atlas*: The myocardia are segmented. Images are then aligned and registered nonrigidly toward a reference image. The atlas is constructed iteratively by averaging acquired images in the average heart shape. *Comparison with the atlas*: Abnormal hearts are registered to the average healthy heart. The cardiac fiber structures of each abnormal heart are compared with the structures of the average healthy heart.

B.2.2 Registration of Abnormal Hearts

The atlas of diffusion tensors is constructed using the automated framework described in (Lombaert *et al.*, 2011b). The method is summarized in Fig. B.1 and has four steps. The myocardium is segmented (Boykov and Jolly, 2000) and its mask is used to guide the nonrigid pairwise registration (Vercauteren *et al.*, 2008, 2009a; Mansi *et al.*, 2010). All hearts are registered to an initial reference image, which is updated toward the morphological average of all hearts (Guimond *et al.*, 2000). Once the transformations of all hearts toward the average cardiac shape are computed, the diffusion tensors are warped (Peyrat *et al.*, 2007) to the morphological atlas.

The processing of abnormal hearts is performed within the same framework (Lombaert *et al.*, 2011b). Firstly, the myocardia of the abnormal hearts are segmented using a minimal user interaction. Secondly, their masks are registered to the newly computed average healthy heart. The diffusion tensors are warped accordingly to the shape of the average healthy heart.

B.2.3 Comparison with Abnormal Hearts

The diffusion tensors fields from all hearts, $\{\mathbf{D}^{(i)}\}_{i=1...N}$ (with $N = 10$ healthy + 6 abnormal hearts), are warped to the morphological average of the healthy hearts (i.e., in a common reference). The Log-Euclidean metric (Arsigny *et al.*, 2006b) is used to compute efficiently the average diffusion tensor of the healthy hearts (hearts #1 to #10) with the Fréchet mean, $\overline{\mathbf{D}} = \exp\left(\frac{1}{10}\sum_{i=1}^{10}\log(\mathbf{D}^{(i)})\right)$.

The eigendecomposition of the diffusion tensor matrix $\overline{\mathbf{D}}$ gives the principal directions $\mathbf{v}_{1,2,3}$ describing the fiber structures. More precisely, the first eigenvector $\mathbf{v}_1$ gives the fiber orientation, the second eigenvector $\mathbf{v}_2$ is believed to lie within the laminar sheet and to be perpendicular to the fiber, and the third eigenvector $\mathbf{v}_3$ is assumed to give the normal of the laminar sheet.

The abnormal hearts are compared with the average healthy heart by measuring the angular deviations of the fiber structures of each heart with the average heart. The angle θ between the direction of an eigenvector $\mathbf{v}_j$ of the i^{th} heart and the direction of the corresponding average eigenvector $\overline{\mathscr{V}}_j$ is defined between 0° and 90° with:

$$\theta_j^{(i)} = \arccos\left(\frac{|\mathbf{v}_j^{(i)} \cdot \overline{\mathbf{v}}_j|}{\| \mathbf{v}_j^{(i)} \| \| \overline{\mathbf{v}}_j \|}\right) \tag{B.1}$$

The absolute value of the dot product removes the inherent ambiguity in the orientation of the eigenvectors (i.e., $|a \cdot b| = |a \cdot (-b)|$).

B.3 Results

We study the deviation of the fiber structures of each heart (healthy and abnormal) to the average structures of the healthy hearts (i.e., to the atlas). The structures in the healthy hearts are, as expected, very similar to the atlas. The histograms of the angular differences of the first, second, and third eigenvectors of the healthy hearts to the atlas (gray curves in Fig. B.2) show average modes of respectively (i.e., the curves are peaking at) $\overline{\theta}_1 = 13.03°$, $\overline{\theta}_2 = 21.76°$, and $\overline{\theta}_3 = 15.77°$. Abnormal hearts show by contrast fiber structures that have larger deviations to the atlas of healthy hearts. The histograms of the angular differences of structures show higher modes in abnormal hearts (black in Fig. B.2), with a deviation of $\overline{\theta}_1 = 20.96°$ for the fibers (i.e., first eigenvector) and of $\overline{\theta}_2 = 48.21°$ and $\overline{\theta}_3 = 34.36°$

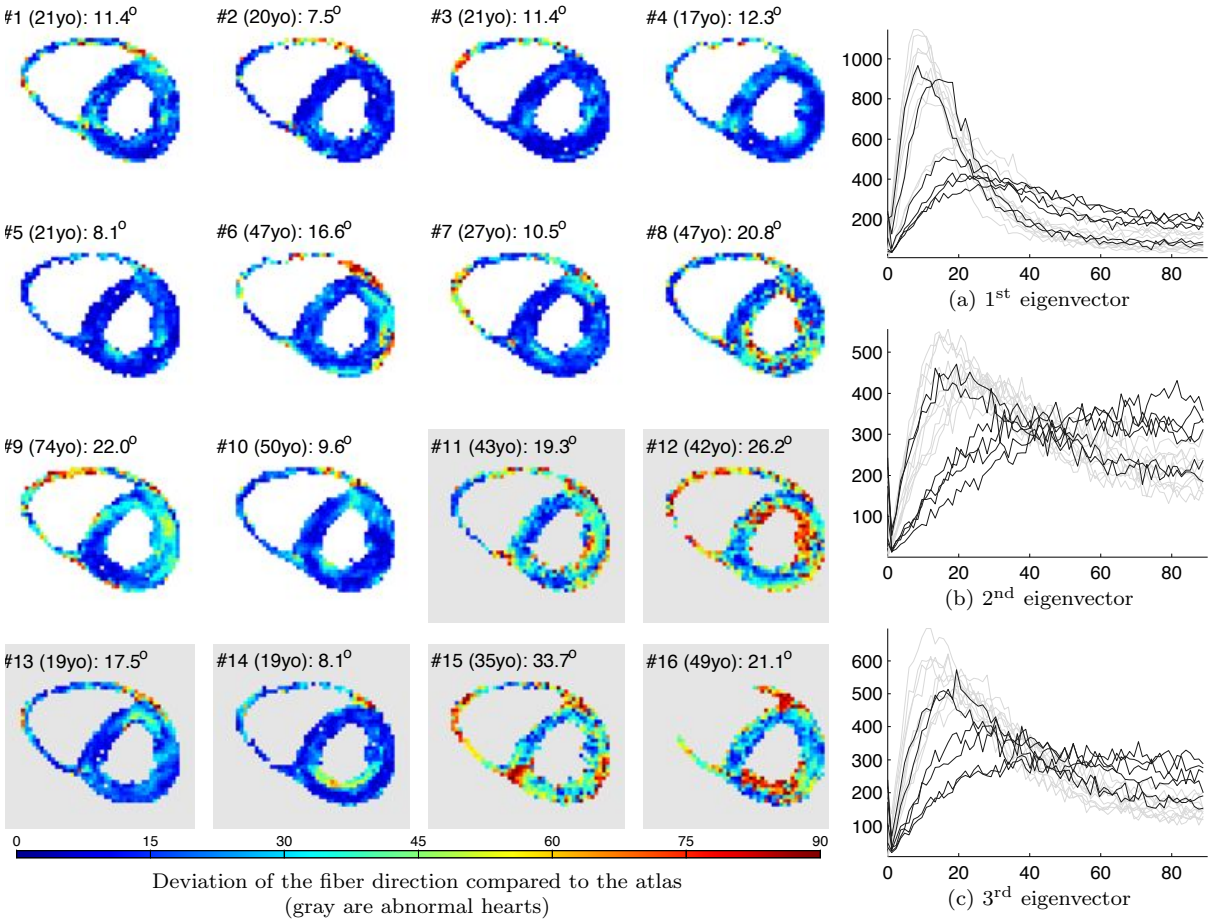

Figure B.2: *(Left)* Deviation of the fiber direction of each heart to the atlas of healthy hearts. Coloring is the angular difference in degree. Abnormal hearts are with gray background. The age of each subject is provided in each sub figure. *(Right)* Histograms of the angular variability (in degrees) of *(a)* the 1st eigenvector, *(b)* 2nd eigenvector, and *(c)* 3rd eigenvector (abnormal hearts in dark lines, healthy hearts in light lines).

for the laminar sheets (i.e., second and third eigenvector). The visualization of the angular difference in a slice of each heart shows large discrepancies in the left ventricle with localized high-variability areas for patient #12, #15, and #16 (shown in the sub-figures of Fig. B.2 with gray backgrounds). This is again confirmed when visualizing the angular difference of the second and third eigenvectors (i.e., the laminar sheets). It is to note that the registration of the right ventricle (which exhibited a very small volume) failed for the last patient (#16). The patients #13 and #14, even if classified as abnormal, presented consistently very small deviations to the average fiber structures of the healthy atlas (Fig. B.2).

B.4 Discussion and Conclusion

The question whether the variability of the cardiac fiber structures is a marker to normality or abnormality is relevant to the study of many cardiomyopathies, including left ventricular hypertrophy or myocardial disarray. In this paper, we compared the structural changes between a population of abnormal hearts and of healthy hearts. It was shown that the three eigenvectors of the diffusion tensors have measurable differences between abnormal hearts and healthy hearts. When compared to an atlas of healthy hearts, the fibers of abnormal hearts showed an angular difference of 20.96°, while the fibers of healthy hearts showed less deviation with 13.03°. The laminar sheets also showed a greater deviation and a greater variability in abnormal hearts than in healthy hearts. Even though the laminar sheet is known to be more variable than the fiber structure in humans (Lombaert *et al.*, 2011b), the difference in both populations is non negligible (deviation of the laminar sheet normal of 34.36° in abnormal hearts compared to 15.77° in healthy hearts). The abnormal hearts also experience a large fiber angle difference around trabeculae areas. A localized study might reveal the origin of such large deviance.

Nonetheless, two outliers are present (hearts #13 and #14 as shown in Fig. B.2). They were initially classified as abnormal even though their cardiac fiber structures are very similar to the average healthy heart (deviation of 17° and 8°). We hold the attention on the age of both subjects, both very young (19 years old). Furthermore, the modes in healthy hearts, i.e., the peaks of the histograms in Fig. B.2 (a), also show that the cardiac fibers are less variable in younger subjects than in older subjects. Age is thought to have an impact in the fiber structure of skeletal muscles (Lexell and Taylor, 1991). No study has yet been performed in cardiac muscles. For that matter, the mode of the angular differences of the first eigenvector (i.e., the fiber direction) was plotted against the age of each subject (Fig. B.3). The correlation factor between age and fiber variability is 0.73 when considering only the 10 healthy hearts (with a low p-value of 0.016). When considering only the abnormal hearts, the correlation factor is higher at 0.86 (with a p-value of 0.029). The estimated least-square fit lines of both population are overlaid in Fig. B.3. Before hypothesizing that the variability of the fiber

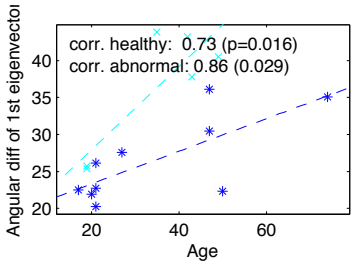

Figure B.3: Possible correlation between the age and the fiber variability (healthy hearts in blue with a correlation factor $\rho = 0.73$, abnormal in cyan with $\rho = 0.86$).

directions increases faster with age in abnormal hearts, many unknown parameters should be considered firstly (for instance, the distinction between primitive hypertrophy or secondary hypertrophy, known to occur in old subjects, is here unknown).

In conclusion, our study comparing a population of abnormal hearts and of healthy hearts showed that there are observable differences in the fiber directions in both populations. Abnormal hearts have fiber directions that are more variable and that are on average 20.96° different from the average healthy heart. Future studies will include additional hearts in order to further study these preliminary findings.

Acknowledgment

The authors wish to acknowledge the members of the Asclepios Team. The project was supported financially by the National Science and Engineering Research Council of Canada (NSERC), the Michael Smith Scholarship (CGS-MSFSS), and the EGIDE/INRIA Scholarship.