

Université Paris Cité
Ecole Doctorale 386, Sciences Mathématiques de Paris Centre
Laboratoire de Mathématiques Appliquées à Paris 5 - UMR CNRS 8145
GE Healthcare

Diffeomorphic registration with topological changes: Application to interventional radiology

PIERRE-LOUIS ANTONSANTI

Thèse de Doctorat en MATHÉMATIQUES APPLIQUÉES

13 Octobre 2022

| | | |
|--------------------------|--------------------------|--------------------|
| GARY E. CHRISTENSEN, PR | UNIVERSITY OF IOWA | Rapporteur |
| XAVIER PENNEC, DR | INRIA | Rapporteur |
| MIRZA FAISAL BEG, PR | SIMON FRASER UNIVERSITY | Examinateur |
| BARBARA GRIS, CR | CNRS | Examinatrice |
| MEGAN OWEN, MCF | LEHMAN COLLEGE | Examinatrice |
| STEFAN SOMMER, PR | UNIVERSITY OF COPENHAGEN | Examinateur |
| VINCENT JUGNON, PhD | GE HEALTHCARE | Co-encadrant |
| THOMAS BENSEGHIR, PhD | GE HEALTHCARE | Co-encadrant |
| JOAN ALEXIS GLAUNÈS, MCF | UNIVERSITÉ PARIS CITÉ | Directeur de Thèse |

Université Paris Cité
Ecole Doctorale 386, Sciences Mathématiques de Paris Centre
Laboratoire de Mathématiques Appliquées à Paris 5 - UMR CNRS 8145
GE Healthcare

Diffeomorphic registration with topological changes: Application to interventional radiology

PIERRE-LOUIS ANTONSANTI

PhD Thesis in APPLIED MATHEMATICS

October 13 2022

| | | |
|-----------------------------------|--------------------------|------------|
| GARY E. CHRISTENSEN, PROF. | UNIVERSITY OF IOWA | Reporter |
| XAVIER PENNEC, DR | INRIA | Reporter |
| MIRZA FAISAL BEG, PROF. | SIMON FRASER UNIVERSITY | Examiner |
| BARBARA GRIS, CR | CNRS | Examiner |
| MEGAN OWEN, ASSOC. PROF. | LEHMAN COLLEGE | Examiner |
| STEFAN SOMMER, PROF. | UNIVERSITY OF COPENHAGEN | Examiner |
| VINCENT JUGNON, PhD | GE HEALTHCARE | Co-Advisor |
| THOMAS BENSEGHIR, PhD | GE HEALTHCARE | Co-Advisor |
| JOAN ALEXIS GLAUNÈS, ASSOC. PROF. | UNIVERSITÉ PARIS CITÉ | Supervisor |

Remerciements

First of all, I would like to thank my referees, Gary Christensen and Xavier Pennec, for having accepted to review my thesis. Their detailed feedback allowed me to bring clarity and precision to my manuscript. I also thank the members of the jury for having accepted to be part of my examiners for the defense.

Je dois ensuite passer au français remercier mon équipe de choc, Joan, Vincent et Thomas pour m’avoir accompagné tout au long de cette thèse. Je tenais tout d’abord à remercier Joan pour ta patience, son attention et ta pédagogie, le tout dans la bonne humeur qui ont rendu ces années si riches et enthousiasmantes. Je ne te remercie pas cependant pour les jeux blancs au tennis...

Merci ensuite à Vincent, pour les discussions toujours éclairantes et le soutien, pour avoir su faire aussi bien le lien entre le laboratoire et l’entreprise et pour m’avoir lancé sur ce sujet il y a trois ans. Je souhaite à tous les doctorants de pouvoir commencer dans d’aussi bonnes conditions que celles que tu avais mises en place. Merci enfin à Thomas, et ton énergie débordante, tu avais annoncé dès le début ce qu’on ferait à la fin de la thèse. C’est un art de savoir jongler avec autant d’idées mais de réussir à rester focalisé. Merci également pour la tasse et les jeux, toujours de bon goût ! Sans vous trois, la thèse aurait semblé beaucoup plus longue.

Merci à mon équipe GE, Valérie ma manager et tout le monde pour avoir su garder le contact malgré la pandémie et les restrictions GE... Merci à Irène au laboratoire pour tes discussions et tes encouragements (« il faut être positif !! »).

Merci ensuite à mes camarades de labo, les plus anciens Pierre, Claire, Vincent, Juliana, Alex... qui ont rendu l’arrivée au labo aussi agréable ; ceux de ma promo Anton, Rémi (LB), Florian, Antoine avec qui j’ai pu avancer jusque-là. Mention particulière aux cours de Lindy-hop Anton, je n’aurais jamais pensé faire ça un jour. Merci bien sûr aux plus récents, que je ne peux pas tous citer, Zoé, Rémi (LM), Antoinettes, Thais, Sonia... Le laboratoire MAP5 est un havre, et va le rester ! Merci également à mes co-bureau de GE, Marion et Mario, pour s’être serré les coudes jusqu’au bout, et à la bande de stagiaires qui avait tout commencé en mammo...

Je remercie mes amis, d’enfance, du lycée de prépa puis d’école, et toujours du rugby, la liste est longue mais c’est un privilège de pouvoir compter sur vous. Merci tout particulièrement à la grotte ainsi qu’à Carole et Hakim.

Enfin merci à mes proches, ma deuxième famille, les Rigaud pour votre joie de vivre communicative, et votre chaleur, et à Vincent, mon autre frère avec qui dès petit il fallait bien défendre ses idées. Merci à ma première famille, les Antonsanti, pour tout. À mes grands parents Pierrot, Odile, Mamaille, pour les sourires et les tendresses, et l’affection que vous nous prodiguez, que vous compreniez ou non ce que je raconte ! J’ai grandi dans une famille entourée d’amour et de bonne humeur, et mes parents Hélène et Hervé m’ont toujours encouragé à chercher un métier qui me plaisait. C’est une grande chance et je le chéris. Merci à Binou et Ninou, je suis parti du nid le premier mais nous nous retrouvons toujours avec autant de plaisir. Lors même qu’on n’est pas le chêne ou le tilleul, ne pas monter bien haut, peut-être, mais bien entouré !

Enfin, à la patate que j’ai trouvée sans vraiment la chercher.

CONTENTS

| | |
|--|------------|
| Contents | iii |
| Résumé en Français | 1 |
| Introduction | 5 |
| 1 Interventional Radiology and Automatic Vascular Trees Annotation | 8 |
| 1.1 The Field of Interventional Radiology | 9 |
| 1.1.1 Imaging Modalities for Interventional Radiology | 10 |
| 1.1.2 Diversity of Interventional Radiology Procedures | 14 |
| 1.1.3 Prostatic Artery Embolization for Benign Prostatic Hyperplasia | 16 |
| 1.2 Modeling the Vascular Trees | 21 |
| 1.2.1 From Images to Vessels | 22 |
| 1.2.2 The Pelvic Vascular Tree | 25 |
| 1.3 Automatic Annotation of Vascular Trees: A State of the Art | 30 |
| 1.3.1 Learning-based labeling | 30 |
| 1.3.2 Atlas-based labeling | 31 |
| 1.4 Conclusion | 33 |
| 2 Large Diffeomorphic Deformations and Application to Vascular Tree Alignment | 35 |
| 2.1 Computational Anatomy | 36 |
| 2.2 Reproducing Kernel Hilbert Spaces: a Cornerstone | 38 |
| 2.3 A Space of Deformations | 44 |
| 2.3.1 Flows of Vector Fields | 45 |
| 2.3.2 Computing Large Diffeomorphic Deformation | 47 |
| 2.4 Metric Mapping: from Exact to Inexact Registration | 49 |
| 2.4.1 Spaces for the Data Attachment | 52 |
| 2.4.2 Matching Oriented Varifolds | 56 |
| 2.4.3 Optimal Transport Cost as Data Attachment Term | 59 |
| 2.5 Statistics over Deformations | 62 |
| 2.5.1 Building a Template of a Population | 62 |
| 2.5.2 PCA and LDDMM | 65 |
| 2.6 Application to the Registration of Simplified Pelvic Vascular Trees | 66 |
| 2.7 Conclusion | 71 |
| 3 Aligning a Shape Onto a Subset of a Target | 74 |
| 3.1 Missing Pieces: A Topological Challenge | 75 |
| 3.1.1 Partial Matching between Shapes: a State of The Art | 75 |
| 3.1.2 First Ideas with a Toy Example | 88 |
| 3.2 Partial Matching in the Space of Oriented Varifolds | 90 |
| 3.2.1 Adding a Local A Priori | 91 |

| | | |
|----------|--|------------|
| 3.2.2 | Normalizing the Kernels | 93 |
| 3.2.3 | Use with Rigid | 96 |
| 3.2.4 | Use with LDDMM | 96 |
| 3.3 | Regularization as Source Attachment Term | 98 |
| 3.4 | Examples of Applications | 102 |
| 3.5 | Conclusion | 105 |
| 4 | Registration of Tree-Like Shapes | 109 |
| 4.1 | Building a Space of Tree-like Shapes | 116 |
| 4.1.1 | A Tree from Leaves to Root | 116 |
| 4.1.2 | Definition of a Tree-Like Shape | 120 |
| 4.1.3 | The Space of Tree-like Shapes | 121 |
| 4.2 | Computing Geodesics in the Space of Tree-like Shapes | 126 |
| 4.2.1 | Characterization of the Geodesics | 127 |
| 4.2.2 | Geodesic Algorithm in a Space of Tree-Like Shapes | 130 |
| 4.3 | Implementation | 132 |
| 4.3.1 | Estimation of the Fréchet Mean | 132 |
| 4.3.2 | Principal Geodesic analysis in the Space of Tree-Like Shapes | 133 |
| 4.4 | A Registration Problem for Deforming Tree-like Shapes | 136 |
| 4.4.1 | A Link between Tree-like Shapes and Vascular Trees | 137 |
| 4.4.2 | Registering Vascular Trees with Hierarchical Changes | 139 |
| 4.4.3 | Dealing with the Change of Orthant | 141 |
| 4.4.4 | Experiments | 144 |
| 4.5 | Conclusion | 149 |
| 5 | Applications | 151 |
| 5.1 | LDDMM Application to Automatic Annotation of Simplified Pelvic Vascu- lature | 152 |
| 5.1.1 | Database Description | 153 |
| 5.1.2 | Transferring the Labels from the Template to the Target | 153 |
| 5.1.3 | Evaluation and Results | 155 |
| 5.2 | Registration of a Template Vascular Tree onto Real Ones | 164 |
| 5.2.1 | Database Description | 164 |
| 5.2.2 | Evaluation and Results | 165 |
| 5.2.3 | Discussion | 170 |
| 5.3 | Beyond Vascular Trees: Feature-based Multi-modality Liver Volume Regis- tration | 172 |
| 5.3.1 | CT/CBCT Volume Registration | 173 |
| 5.3.2 | Database Description | 173 |
| 5.3.3 | Liver Surface Registration with Partial Matching | 175 |
| 5.3.4 | From Surface to Volume Registration | 177 |
| 5.3.5 | Evaluation and Results | 178 |
| 5.3.6 | Discussion | 183 |
| 5.4 | Conclusion | 185 |
| | Conclusion and Perspectives | 188 |

| | |
|--|------------|
| Appendices | 192 |
| 5.5 Matrix Valued Kernels | 192 |
| 5.6 Normal Cycles | 194 |
| 5.7 Partial Matching Application to CT/CBCT volumes registration | 196 |
| 5.7.1 POIs Detailed Results per Patients | 196 |
| 5.7.2 Lesions Detailed Results per Patients | 197 |
| 5.7.3 Rotations Deformations Comparison | 198 |
| 5.7.4 Worst Case Scenario : Bad Feature Extraction | 199 |
| Bibliography | 201 |

Résumé en Français

Cette thèse s'inscrit dans le domaine des mathématiques appliquées à la radiologie interventionnelle - un domaine médical qui s'appuie sur des systèmes d'imagerie radiologique en temps réel pour réaliser des interventions peu invasives. Nous nous concentrons sur les interventions endovasculaires au cours desquelles les radiologues guident leurs outils à l'intérieur du système vasculaire du patient. Ces outils sont utilisés, par exemple, pour administrer un traitement ou à des fins de diagnostic. Pour préparer leurs procédures cliniques, les médecins réalisent une cartographie 3D des vaisseaux sanguins au niveau de l'anatomie d'intérêt. L'arbre vasculaire peut être extrait d'une image 3D, ce qui permet d'améliorer la visualisation et la planification du trajet dans l'arbre pour atteindre la zone cible.

Une procédure qui a gagné en intérêt des dernières années est l'embolisation de l'artère prostatique pour traiter l'hyperplasie bénigne de la prostate : une croissance anormale de la prostate qui touche un homme sur deux passé cinquante ans, et qui dégrade considérablement la qualité de vie au quotidien. Pour réaliser l'embolisation de l'artère prostatique, les cliniciens injectent un agent embolique dans les vaisseaux des patients alimentant l'artère prostatique. L'identification des artères dans lesquelles injecter l'agent, ainsi que des artères potentielles non ciblées, est essentielle pour le bon déroulement de la procédure.

Ainsi, le fait de pouvoir placer un nom anatomiquement sur les vaisseaux sanguins extraits des images facilite grandement les procédures endovasculaires en aidant à la compréhension de l'arbre vasculaire. En ce qui concerne la complexité de la collecte de données annotées, en particulier dans le domaine médical, les solutions d'annotation basées sur de petites bases de données qui ont été étiquetées manuellement sont d'un grand intérêt. Les méthodes dites *basées atlas* répondent exactement à ce critère, par opposition aux méthodes dites *basées apprentissage*.

Les méthodes basées atlas comparent un modèle de référence, appelé **modèle**, aux observations de la base de données. Il est ensuite utilisé pour transférer des informations (par exemple, une étiquette) du modèle aux observations. Il peut y avoir plusieurs modèles, qui forment un atlas : l'espace des modèles possible. Pour les arbres vasculaires, le modèle choisi peut être construit *ex nihilo* à partir d'une connaissance a priori de l'anatomie. Cette solution est intéressante mais nécessite une connaissance précise de l'anatomie. Une autre solution classique consiste à sélectionner des cas dans la base de données, à les étiqueter puis à les utiliser tous comme modèles. Une solution plus rapide et plus simple consiste à utiliser un seul cas que l'on étiquette manuellement. La plupart de ces approches basées sur l'atlas dans le cadre des arbres vasculaires utilisent peu de déformations pour aligner le modèle avec les observations : en effet, la complexité des arbres vasculaires en termes de nombre de branches et de géométrie rend ces alignements compliqués, surtout dans le cas de déformations non-rigides. Cependant, si nous alignions correctement le modèle sur une observation, leur comparaison serait simplifiée.

C'est exactement le point de vue de la théorie de l'anatomie computationnelle, qui consiste à étudier des populations de formes à travers les déformations qu'il faut générer pour aligner les formes entre elles. Dans ce cadre, les chercheurs ont construit des déforma-

tions lisses de l'espace dans lequel les formes existent. Ces déformations agissent ainsi sur les formes et peuvent être appliquées à d'autres objets existant dans le même espace. Lorsqu'elles sont appliquées aux modèles, les déformations fournissent un espace de modèles admissibles que nous utilisons comme **atlas**.

Dans le cadre classique du recalage, les formes ne peuvent pas être parfaitement alignées les unes avec les autres. On peut formuler un recalage dit inexact, consistant à minimiser la distance entre l'objet déformé et la cible, sous condition d'une déformation aussi régulière que possible. C'est le principe du Large Deformation Diffeomorphic Metric Mapping (LDDMM), qui génère des difféomorphismes de l'espace ambiant et que nous utiliserons tout au long de la thèse.

Cependant, les déformations étant lisses, elles n'expliquent pas les changements de topologie entre le modèle et les observations. Ces changements peuvent provenir de toutes sortes de sources, depuis la méthode d'acquisition des données et leur construction jusqu'aux individus observés eux-mêmes. L'exemple des arbres vasculaires est typique : des différences dans le nombre de branches, l'ordre des bifurcations, ou la présence ou l'absence d'artères sont autant de changements topologiques. Par la suite, nous proposons un cadre pour le recalage des arbres vasculaires sous déformations non-rigides : il s'agit d'une approche basée sur l'atlas qui s'appuie sur l'alignement des formes pour faciliter l'étiquetage automatique.

La déformation non-rigide doit être adaptée à la variabilité de la forme des arbres vasculaires, tant en termes de géométrie que de topologie. Nous adaptons donc le LDDMM à deux cas typiques de changements topologiques : le *correspondance partielle*, ou le recalage d'un arbre modèle sur une sous-partie de l'arbre cible ; et le *changements dans les bifurcations*, ordonnancement rendu possible par l'intégration de l'arbre modèle dans un espace de formes arborescentes adapté à ces changements topologiques.

Nous utilisons le recalage et la correspondance partielle du LDDMM dans des applications de radiologie interventionnelle en travaillant sur les arbres vasculaires. Tout d'abord, les LDDMM et la construction du modèle associé sont utilisés comme prétraitement pour un étiquetage automatique des arbres vasculaires basé sur un atlas. Nous profitons du fait que les déformations ne nécessitent aucune annotation des arbres vasculaires. Ainsi, la construction du modèle et de l'atlas peut être effectuée avant l'étape d'étiquetage. Nous montrons que cela rend le pipeline d'étiquetage proposé plus robuste au choix initial de la forme de référence comme modèle et améliore la performance de l'étiquetage.

Nous appliquons ensuite le cadre de correspondance partielle à deux problèmes différents : le recalage d'un arbre modèle simplifié sur des arbres complets, et le recalage de surfaces hépatiques tronquées sur des surfaces complètes. Dans cette application, nous évaluons l'intégration du modèle déformé dans la cible complète du point de vue anatomique : nous voulons évaluer si les artères déformées du modèle correspondent effectivement à un sous-ensemble des artères correspondantes dans la cible.

La deuxième application s'inscrit dans un autre cadre : le recalage de volumes multimodaux, et plus précisément le recalage de tomographies injectées préopératoires sur des volumes 3D de tomographie à faisceau conique (CBCT) en direct pour le même patient dans la région de l'abdomen. Le recalage est basé sur l'alignement d'une surface hépatique tronquée extraite de l'image CBCT sur la surface hépatique complète extraite du CT-scan correspondant, d'où

une correspondance partielle des surfaces.

L'organisation du document est la suivante : Les chapitres 1 et 2 présentent à la fois le contexte clinique et la théorie existante sur laquelle nous nous appuyons. Le chapitre 1 motive ce travail, en illustrant les multiples interactions de la radiologie interventionnelle (RI) avec les techniques d'acquisition et de traitement d'images. Nous montrons la grande variété d'applications de recalage en RI et l'importance de disposer d'outils d'annotation automatique des arbres vasculaires dans le contexte de l'embolisation de l'artère prostatique pour l'hyperplasie prostatique bénigne. Plus généralement, le fait de disposer d'outils de recalage automatique en radiologie interventionnelle est un plus, et tant les interventions peropératoires que les études cliniques pourraient bénéficier de tels outils.

Dans le chapitre 2, nous fournissons les ingrédients clés pour la construction d'outils efficaces pour l'analyse de forme dans le contexte de l'anatomie computationnelle. Dans ce chapitre, nous exploitons le cadre du Large Deformation Diffeomorphic Metric Mapping (LDDMM) et les espaces à noyaux reproduisants pour calculer des déformations réalistes et construire des statistiques sur celles-ci. Nous utilisons ces outils théoriques et informatiques bien étudiés pour l'alignement non rigide d'un arbre vasculaire modèle sur un arbre cible. Les statistiques sur les déformations sont ensuite utilisées pour construire l'atlas, dans lequel le modèle vit. Nous montrons qu'un modèle significatif peut être dérivé des statistiques, et aider au recalage sur une nouvelle observation.

Dans le chapitre 3, nous abordons le problème du recalage de formes sous la contrainte de la correspondance partielle des formes qui est un type de changement topologique. La déformation que nous recherchons est l'alignement du modèle déformé sur un sous-ensemble de la cible. Nous concevons un nouveau terme d'attachement des données dans l'espace de forme, et nous détaillons sa construction tout au long du chapitre. Nous introduisons également un terme de régularisation comparant la forme déformée à sa position initiale, afin d'éviter que le modèle déformé ne soit rétréci. Tout au long du chapitre nous illustrons les résultats sur des exemples simples de correspondance partielle, puis nous appliquons la méthode proposée à nos arbres vasculaires ainsi qu'à une surface simple.

Le chapitre 4 est consacré à un autre changement topologique entre les arbres : les changements dans l'ordre des bifurcations. En fait, les déformations difféomorphes que nous construisons avec le LDDMM ne permettent pas de tels changements topologiques, et nous cherchons à construire un cadre de recalage plus adapté. À cette fin, nous intégrons les arbres vasculaires dans un espace pertinent de formes arborescentes et décrivons une façon d'effectuer le recalage entre un modèle annoté et une cible qui est un arbre vasculaire non étiqueté. Le recalage est formulé comme un problème de minimisation et le modèle vit dans l'espace des formes arborescentes dans lequel les changements dans l'ordre des bifurcations sont disponibles.

Dans le chapitre 5, nous rassemblons les trois applications à des données réelles que nous avons menées : annotation d'arbres vasculaires pelvien simplifiés en utilisant la construction et le recalage d'atlas basés sur le LDDMM, extension au recalage d'arbres simplifiés sur des cas réels, et enregistrements CT/CBCT multi-modalité basés sur le recalage de surfaces de

foie. Dans la première application nous montrons que les recalages d'un exemple annoté sur le reste de la base de données (non annotée) permettent de construire un modèle spatialement cohérent. Cela permet d'améliorer les performances de recalage de ce modèle sur la base de données, et ainsi améliore les performances d'annotation qui en découlent. Nous comparons également la méthode d'annotation proposée à des méthodes basées apprentissage, et nous montrons que lorsque la base de données annotées est petite, la méthode proposée est largement plus performante et robuste. Dans la seconde application, nous montrons que le recalage de notre modèle de référence simplifié de l'arbre pelvien sur une base de données d'arbres complet en utilisant notre matching partiel donne des résultats prometteurs : non seulement l'inclusion du modèle dans la cible est obtenue via des déformations diffeomorphiques non rigides, mais en plus il y a un sens anatomique dans les recalages obtenus, en particulier au niveau des feuilles. Enfin l'application aux recalages de volumes de foies issus de scanners et de tomosynthèse basé surface des foies permet d'illustrer l'intérêt et la versatilité du terme de matching partiel que nous avons proposé. En effet, à travers l'évaluation sur une petite base de données et des marqueurs anatomiques à l'intérieur des foies, nous montrons que le bon recalage des volumes (associés à un même patient) sont composés d'une translation et de déformations non-rigides.

En conclusion, les méthodes basées atlas peuvent être très efficaces, en particulier dans le cas de données peu annotées, lorsqu'elles sont couplées à des modèles de déformation pertinents. Ceci est vrai même pour des formes aussi complexes que les arbres vasculaires, avec une grande variabilité géométrique ainsi que des différences topologiques importantes en termes de nombre de branches et d'ordre des bifurcations. Dans le cas de grandes bases de données annotées, d'autres solutions, comme celles basées sur l'apprentissage, devraient également être explorées, et une combinaison des deux approches permettrait de tirer le meilleur de chacune d'entre elles : statistiques et méthodes robustes, inférence rapide et bonne extraction de caractéristiques. Les perspectives sont nombreuses, à commencer par la combinaison du matching partiel et de la représentation dans un espace de formes arborescentes pour la construction d'un atlas encore plus complet. Dans ce contexte, les statistiques sur les déformations seraient plus compliquées à générer, car elles dépendraient à la fois du champ de vecteurs générant le difféomorphisme, mais aussi de la position dans l'espace de formes arborescentes. Les perspectives d'applications sont également nombreuses, en particulier en radiologie interventionnelle, où l'analyse des arbres vasculaires et la comparaison de données issues de différentes modalités d'acquisition sont essentielles.

Introduction

This thesis falls within the field of mathematics applied to interventional radiology - a medical field relying on real-time X-ray imaging systems to perform minimally invasive interventions. We focus on endovascular interventions during which radiologists navigate their tools inside the patient's vasculature. These tools are used for example to deliver a treatment or for diagnosis purposes. To prepare their clinical procedures, the physicians carry out a 3D mapping of the blood vessels at the level of the anatomy of interest. The vascular tree can be extracted from a 3D image, which allows for improved visualization and planning of the path through the tree to reach the target area. In Prostatic Artery Embolization, the clinicians inject an embolic agent into the vessels of the patients feeding the prostatic artery. The identification of the arteries in which to inject the agent, as well as the potential non-target arteries is key for the good outcome of the procedure.

Thus being able to place an anatomical name on the blood vessels thus extracted from the images greatly facilitates endovascular procedures by helping understand the tree. Regarding the complexity of collecting annotated data, especially in the medical field, annotation solutions based on small databases that have been manually labeled are of great interest. The so-called *atlas-based* methods meet exactly this criterion.

Atlas-based methods compare a reference model, called the **template**, to observations in the database. It is then used to transfer information (for example, a label) from the model to the observations. For vascular trees, the chosen model can be built *ex nihilo* from a priori knowledge of anatomy. This solution is interesting but requires precise knowledge of anatomy. Another classical solution is to select cases from the database, label them and then use them all as models. A faster and simpler solution is to use a single case. Most of these atlas-based approaches in the framework of vascular trees use few deformations to align the model with the observations: indeed, the complexity of vascular trees in terms of numbers of branches and geometry makes such alignments complicated, especially in the case of non-rigid deformations. However, if we correctly aligned the model on an observation, their comparison would be simplified.

This is exactly the point of view of the theory of computational anatomy, which consists in studying populations of shapes through the deformations that must be generated to align the shapes with each other. In this framework, the researchers have constructed smooth deformations of the space in which the shapes exist. These deformations thus act on the shapes and can be inferred to other objects existing in the same space. When applied to the templates, the deformations provide a space of admissible templates that we called the **atlas**.

In the classical framework of registration, the shapes cannot be perfectly aligned with each other. A so-called inexact registration can be formulated, consisting in minimizing the distance between the deformed object and the target, under the condition of a deformation as regular as possible. This is the principle of Large Deformation Diffeomorphic Metric Mapping (LDDMM), which generates diffeomorphisms of the ambient space and which we will use throughout the thesis.

However, the deformations being smooth, they do not explain the changes in topology between the template and the observations. These changes can come from all sorts of sources, from the data acquisition method and their construction to the observed individuals themselves. The example of vascular trees is typical: differences in the number of branches, the order of bifurcations, or the presence or absence of arteries are all topological changes. Thereafter, we propose a framework for vascular tree registration under non-rigid deformations: it is an atlas-based approach that relies on shape alignment to facilitate the automatic labeling.

The non-rigid deformation should be adapted to the shape variability of the vascular trees, both in terms of geometry and topology. Therefore we adapt the LDDMM to two typical cases of topological changes: the *partial matching*, or the registration of a template tree onto a subset of the target tree; and the *changes in the bifurcations* ordering made possible by embedding the template tree in a space of tree-like shapes adapted to such topological changes.

We use the LDDMM registration and partial matching in IR applications. First, the LDDMM and the associated template construction are used as a preprocessing to an atlas-based automatic vascular tree labeling. We take advantage of the fact that the deformations do not need any annotation of the vascular trees. Hence the template and atlas construction can be done before the labeling step. We show that it makes the proposed labeling pipeline more robust to the initial choice of the reference shape as a template and improves the labeling performance.

We then apply the partial matching framework to two different problems: the registration of a simplified template tree onto complete ones, and the registration of truncated liver surfaces onto complete ones. In this application we assess the embedding of the deformed template into the complete target from the anatomical point of view: we want to assess whether the deformed arteries of the template are actually matched to a subset of the corresponding arteries in the target.

The second application is part of another framework: the multi-modality volumes registration, and more specifically the registration of pre-operative injected CT-scans onto live Cone Beam Computed Tomography (CBCT) 3D volumes for the same patient in the abdomen area. The registration is based on the alignment of a truncated liver surface extracted from the CBCT image onto the complete liver surface extracted from the corresponding CT-scan, hence a partial correspondence of the surfaces.

The organization of the document is the following: Chapters 1 and 2 introduce both the clinical context and the existing theory on which we rely. Chapter 1 motivates this work, illustrating the multiple interactions of Interventional Radiology (IR) with image acquisitions and processing techniques. We show the importance of having automatic vascular tree annotation tools in the context of Prostatic Artery Embolization for Benign Prostatic Hyperplasia. More generally having automatic registration tools in IR is a plus, and both per operative interventions and clinical studies could benefit from such tools.

In Chapter 2 we provide the key ingredients for the construction of efficient tools for

shape analysis in the context of Computational Anatomy. In this chapter, we exploit the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework and the Reproducing Kernel Hilbert Spaces to compute realistic deformations and build statistics over them. We use these well-studied theoretical and computational tools for the non-rigid alignment of a template vascular tree onto a target one. The statistics are then used to build the atlas, in which the template lives. We show that a meaningful template can be derived from the statistics, and help the registration of an observation.

In Chapter 3, we address the problem of shape registration under the constraint of partial shape correspondence which is one type of topological change. The deformation we seek is the alignment of the deformed template onto a subset of the target. We design a new data attachment term in shape space, and we detail its construction throughout the chapter. We also introduce a regularization term comparing the deformed shape to its initial position, to prevent the deformed template from being shrunk.

Chapter 4 is dedicated to another topological change between the trees: the changes in the bifurcations ordering. In fact, the diffeomorphic deformations we build with LDDMM do not allow such topological changes, and we seek to build a more adapted registration framework. To that end, we embed the vascular trees into a relevant space of tree-like shapes and describe a way to perform registration between some annotated template and a target that is a vascular tree unlabeled. The registration is formulated as a minimization problem and the template lives in the space of tree-like shapes in which the changes in bifurcations ordering are available.

In Chapter 5 we bring together the three applications to real data that we have conducted: annotation of simplified vascular trees using LDDMM-based atlas construction and registration, extension to simplified tree registration on real cases, and multi-modality CT/CBCT registrations based on registration of liver surfaces.

INTERVENTIONAL RADIOLOGY AND AUTOMATIC VASCULAR TREES ANNOTATION

Interventional radiology is a medical sub-specialty of radiology utilizing minimally-invasive image-guided procedures to diagnose and treat diseases in nearly every organ system. The concept behind interventional radiology is to diagnose and treat patients using the least invasive techniques currently available in order to minimize risk to the patient and improve health outcomes. Source: [John Hopkins](#).

| | | |
|-------|--|----|
| 1.1 | The Field of Interventional Radiology | 9 |
| 1.1.1 | Imaging Modalities for Interventional Radiology | 10 |
| 1.1.2 | Diversity of Interventional Radiology Procedures | 14 |
| 1.1.3 | Prostatic Artery Embolization for Benign Prostatic Hyperplasia | 16 |
| 1.2 | Modeling the Vascular Trees | 21 |
| 1.2.1 | From Images to Vessels | 22 |
| 1.2.2 | The Pelvic Vascular Tree | 25 |
| 1.3 | Automatic Annotation of Vascular Trees: A State of the Art | 30 |
| 1.3.1 | Learning-based labeling | 30 |
| 1.3.2 | Atlas-based labeling | 31 |
| 1.4 | Conclusion | 33 |

1.1 The Field of Interventional Radiology

This thesis is part of the applications that image processing and applied mathematics bring to interventional radiology (IR). This discipline is a specialty of radiology and allows to obtain diagnostic images and perform minimally invasive image-guided interventions on patients with diseases affecting nearly any organ or system. Compared to classical surgery, IR provides alternative solutions causing less pain to the patients with less risk of infection and significantly reducing the length of hospital stay. There are three levels of IR procedures, classified according to their severity, the type of anesthesia, the guidance equipment, and the medical team required. The simple procedures, usually involving needle insertion, are performed by any versatile radiologist: peripheral biopsy, guided puncture, or peripheral joint infiltration. Intermediate procedures require a technical imaging platform in line with the targeted activity and cover simple angioplasties, programmed embolization, and drainage, usually involving catheter¹ and guidewire navigation in the patient. Complex procedures require specialized teams and equipment to ensure permanent care and to handle heavy procedures requiring a specific environment: emergency embolization, aortic stent-graft, carotid angioplasty.

All these procedures are done during minimally invasive interventions: the radiologist navigates tools, like needles or catheters, inside the patient's anatomy without being able to see them directly. The intermediate imaging device plays a key role in the procedure workflow, allowing the clinicians to navigate the tools as effectively as possible and assess the risks, thus mitigate the doubts regarding the patient's anatomy, and eventually assess the outcome of the procedure. During a typical intervention, interventionalists take advantage of their knowledge of the anatomy and pathology (mostly resulting from pre-operative examination) to navigate tools guided by per-operative real-time images of the patient.

Such procedures are usually classified into two categories: endovascular and extra-vascular. In the case of extra-vascular procedures, one of the main issues is the visualization of the needles, involved in numerous applications: from biopsies to tumor ablation, including cementoplasty². The endovascular procedures find many applications as well, and target almost all organs in our body: tumor chemo – or radio – embolization, prostatic arteries embolization, or brain Arterio-Venous Malformation (AVM) treatment, brain aneurysm coiling, angioplasty³ or the injection of clot-dissolving medicines to dissolve blood clots and increases blood flow to arms, legs, or organs in the body.

The type of imaging depends on the procedure. The standard interventional system, however, is the C-arm (named after its C-shape) which provides live X-ray videos used for real-time guidance. In the next section, we detail the different types of images that interventional radiologists can use before and after procedures to plan and evaluate the treatment. We also describe the intraoperative imaging systems available to them to perform the procedures.

¹Flexible thin tube that can be inserted in the body.

²Treatment involving the injection of bone cement to reinforce weakened bones, preventing pathological fractures and relieving pain in patients with osteoporosis and bone metastases.

³Placement of stents to expand the vessel at the site of a partial or entire blockage.



Figure 1.1: View of a typical interventional radiology suite (© GE Healthcare). The interventional radiologists navigate the tools with X-ray imaging guidance and must refer to the screen to assess the position of the tools inside the patient’s anatomy.

1.1.1 Imaging Modalities for Interventional Radiology

Pre and Post-operative Imaging Pre and post-operative acquisitions can be done with many systems. The computed tomography scanner (CT scan) is a medical X-ray imaging technique used in radiology to obtain detailed internal images of the patient body non-invasively for diagnostic purposes. It is usually used before an intervention to provide a 3D reconstruction of the patient’s anatomy. The imaged parts of the body range from the extremities to the head, with application to aneurysm detection, blood vessels analysis under the injection of contrast agent, diagnosis of urinary stones, or tumor detection. When the diagnostic leads to an intervention, this pre-operative CT scan can be used to assess tumors size and positions to guide ablation or to accurately understand the patient’s vasculature to identify arteries to embolize [Bil+11].

One other system of choice for diagnostic purposes is Magnetic Resonance Imaging (MRI) (see Figure 1.4.A). It uses the response of different atomic nuclei (mainly hydrogen) to the excitation generated by an external magnetic field. Compared to CT, MRI provides better contrast in images of soft-tissues, e.g. in the brain or abdomen, and does not induce any X-ray dose to the patient. MRI is widely used in hospitals and clinics for medical diagnosis, staging, and follow-up of the disease. It also allows forming of pictures of the physiological processes in the body. It is used for diagnosis, for example: to identify the cause of spinal pain and to select the appropriate treatment [Mas+05], to identify arteries of interest for procedure planning [Kim+18], to assess the ischemic areas and patient response to artery embolization [Car+21] as follow-up imaging, or for blood vessels imaging. The latter, called Magnetic Resonance Angiography (MRA), is used to generate images of arteries in order to evaluate them for stenosis (abnormal narrowing), occlusions, aneurysms, or other abnormalities. MRA is often used to evaluate the arteries of the neck and of the brain, the thoracic and abdominal aorta, the renal arteries, and the legs.

A third imaging modality is the Ultrasound (US) imaging based on the propagation of high-frequency sound waves to view inside the body (Figure 1.4.B). Because ultrasound images are captured in real-time, they can also show the movement of the body’s internal

organs as well as blood flowing through the blood vessels. Unlike X-ray imaging, there is no ionizing radiation exposure associated with ultrasound imaging. Therefore they are routinely used to assess bone fragility, listen to the fetal heartbeat, visualize blood flow through a blood vessel, organs or, visualize abdominal tissues, and so on. It is often used as pre-operative imaging for procedure planning or diagnosis, and as post-operative imaging as a follow-up exam. Ultra-sounds are also very useful during the intervention, with the major drawback of being highly operator dependent.

Per Operative Imaging The image guidance performed during IR procedures requires specific settings to allow both the acquisition of the images and the intervention of the patient. Two imaging modalities that are sometimes used during interventional procedures are interventional CT and ultrasounds, e.g. for needle guidance or tumor identification. These two imaging techniques remain however marginal in the field of IR. The standard imaging system guidance for endovascular interventions is called the C-arm, which is the central element of a classic interventional room. This system is composed of a frame in shape of a "C" called *gantry* supporting the X-Ray tube and a flat-panel detector. The gantry rotates around two- or –three orthogonal axes allowing to turn around a table on which the patient lies. These rotations provide the degree of freedom to acquire images under every possible angulation. In practice however these angulation are limited by the patient, the operators and other systems in the interventional room. To attain every part of the patient, translations are also needed and are achieved by the table which has two- or –three translation axes and sometimes two rotation axes (title and cradle). This setup is illustrated in Figure 1.2.



(a)

Figure 1.2: GE Healthcare IGS 730 interventional guiding system. Three rotational axes and three translation axes.

The tube generates a conic X-ray beam in the direction of the detector, providing the 2D real-time projection of the patient's anatomy. The X-rays are produced by a beam of electrons traveling from the cathode to the anode inside the X-ray tube. The photons are produced in random directions at the surface of the anode, called the focal spot, and modeled by a point. They are absorbed by the tube except for a small window designed to limit the

cone-beam angle. Metal filters and a square shape collimator then restrict the beam to a given region of interest, the targeted anatomy of the patient. The different tissues of this anatomy attenuate the X-rays according to the Beer-Lambert law. The remaining photons are caught by the detector and converted into an electric signal by a matrix of photo-diodes. The induced intensities are then converted by capacitors into pixel intensities forming the raw input image. The fluoroscopy denomination comes from the initial technology based on fluorescent screens that converted the X-ray photons into images. In many detectors, a scintillator is still used to convert the X-ray photons into visible ones. The images are processed according to the application and displayed on a Large Display Monitor, that can be seen in Figures 1.2 and 1.1.

The fluoroscopic mode on the interventional system allows producing low-dose live X-ray sequences, a video of the target anatomy of up to 30 frames per second. In addition to the position of the imaging system, the clinician can also control the X-ray dose levels, the post-processing applied to the raw images, and so on. A control room isolated from the X-rays and adjacent to the operating room is also available, in which the clinician and the assistants can monitor the imaging system and review the recorded and stored images. Fluoroscopy is the most common imaging technique acquired with the interventional system, and it exploits the differential absorption of X-rays by different tissues and organs. It provides real-time 2D images of the movement of a body part (like the heart) or the course that a medical instrument or dye (contrast agent) takes as it travels through the body.

Imaging during Endovascular Interventions In the scope of fluoroscopy, the traditional angiography consists in injecting via a catheter a contrast agent such as iodine-based dye that is radiopaque. The obtained images include the blood vessels as well as the other surrounding radiopaque structures. It is acquired usually with a high-dose X-ray, and the obtained images are stored for future review. In order to better capture the vasculature, Digital Subtraction Angiography (DSA) (illustrated in Figure 1.3) is a numerical technique in which a mask of the anatomy is first acquired before the injection of the contrast agent. Then images of the same area are acquired with the latter injected, and the mask – considered as background – is subtracted from the incoming images. The radiologists can play with the injection time and the amount of injected medium to be as selective as possible in the extracted structures.

When dealing with endovascular procedures, the classic approach is to perform an angiography to guide small instruments such as catheters, stents, or coils through blood vessels. The iodinated contrast agent flows through the patient's vasculature and provides real-time 2D images in which the vessels have low intensity. Other structures such as bones or other tools such as metallic screws may have similar intensities, and the DSA allows dealing with them. On the contrary, the moving structures such as the navigated tools and the vessels will appear in the images. This is the standard per-operative imaging during endovascular procedures. The contrast agent, however, quickly fades away as the clinician stops injecting the medium and the blood keeps flowing inside the vessels. Due to the toxicity of the contrast agent, one wants to avoid injecting it continuously, and roadmap solutions have been proposed: combining the previously acquired injected images with the fluoroscopy, resulting in a superimposed fixed version of the vasculature well contrasted onto a time varying image with lower vessel visibility.




Figure 1.3: Example of Digital Subtraction Angiography (DSA) of the pelvic anatomy. (a) The injected fluoroscopy before subtraction. (b) The DSA shows the difference between an initial fluoroscopic acquisition and the injected one of (a). Thus, the vessels are clearly depicted in (b).

All these solutions, however, address the problem of vessel visualization in 2D. Most of the time the projection of the patient's vasculature results in a complex structure with a lot of ambiguities, in particular when trying to distinguish real vessel bifurcations from artificial vessel intersections induced by the projection. To cope with these limitations, Cone-Beam Computed Tomography (CBCT) is now routinely used in interventional rooms for vascular imaging [Anx+98]. A CBCT acquisition consists of a rotation over 200° of the C-arm around the patient. During the rotation, a series of 2D X-ray projections are acquired (like fluoroscopy) resulting in a stack of images called a *spin*. A reconstruction algorithm then estimates a 3D image based on the linear attenuation coefficients from the spin and the system parameters, this is further described in Section 1.2. The 3D reconstruction mitigates the ambiguities and allows the acquisition of soft tissue images like classic CT-scans (the detector is a flat panel though). In addition, if the contrast agent is injected during the rotation of the acquisition system, the vessels will have a high contrast in the 3D volumes as well as other radio-opaque objects in the volume. They will also have a good spatial resolution. Like the DSA, a subtracted CBCT can also be computed. This technique is not real-time (due to the system rotation) yet it allows to mitigate the ambiguities inherent to the 2D projection of a volume.

In short, the live X-ray guidance is done with fluoroscopy, which is low-dose. All the other kinds of acquisitions are high-dose to perform a refined analysis of the patient's anatomy. Regarding the applications of IR, the discipline is constantly interacting with other modalities and specialities. In particular, the interventionalists take advantage of the multiple sources of data available and must be able to compare the patients' information both during and outside the procedures.

1.1.2 Diversity of Interventional Radiology Procedures

The interactions between IR and other modalities are numerous: the intervention can be made upon a diagnosis coming from another clinician, and the procedure planning can be based on acquisitions from a different system than the X-ray acquisitions that are the standard ones in IR... That involves two types of interactions, the ones involving different specialities and the ones between the data. Since we are focusing on interventions in rooms equipped with the C-arm system, diagnostic IR procedures are not discussed. We distinguish three types of procedures, depending on the type of performed gesture.

Needles The first type of procedure performed in IR consists in introducing a needle into the patient's anatomy. That, can be done to remove a solid body for diagnostic purposes (biopsy), or to perform a localized treatment (such as tumor ablation). The latter often aims at destroying tumors that cannot be removed by conventional surgery. These ablative procedures are the minimally invasive pendant of surgical resections. It includes cryoablation, radiofrequency and microwave ablations, and aims at navigating the tool delivering the treatment as close to the tumor as possible to ensure efficient destruction of the diseased tissue while preserving the healthy one.

As mentioned above, ultrasound is often used to guide the needles, but CBCT-based assessments can be performed in conjunction with it to verify the position of the needles. In [Mon+21] US are acquired in multimodality procedure with US/CBCT fusion for percutaneous ablation of small renal tumors. The US are used to guide the needle toward the tumor and a CBCT is then acquired to assess the position of the needle in the patient's anatomy. Differently in [Cor+19], electrochemotherapy of unresectable liver metastasis from renal cell cancer is done, and the US are merged with the CBCT to enhance the tumor inside the volume. The fusion, however, requires a learning curve, and doing it manually, either by the clinician or the assistant, is time-consuming for the inexperienced operator.

The Figure 1.4 illustrates the use case from [Cor+19] during which the procedure planning and the diagnosis were done under MRI and US imaging. The intervention was guided by CBCT/US fusion, and the post-operative follow-up by MRI. This specific case shows the different interactions between the imaging modalities. The differences in terms of tissue contrasts, the geometry of the imaging system, and the patient's anatomy that may evolve between the acquisitions (tumor growth, breathing, fluids...) make the comparison of the data even more challenging.

Unclogging Another IR procedure is the management of narrowed or clogged blood vessels. One of its application is in neurology with the treatment of acute ischemic strokes through endovascular treatment delivery. Strokes occur when blood flow to a part of the brain is interrupted as a result of a broken (hemorrhagic) or severely restricted (ischemic) vessel. First, a plain CT is acquired to exclude hemorrhage, then an angio CT is done to confirm the large vessel occlusion. The latter can be caused for instance by narrowed arteries or blood clots, and the interventionalist can perform under live X-ray a mechanical thrombectomy via a catheter navigated toward the pathological artery. Similar procedures can be done

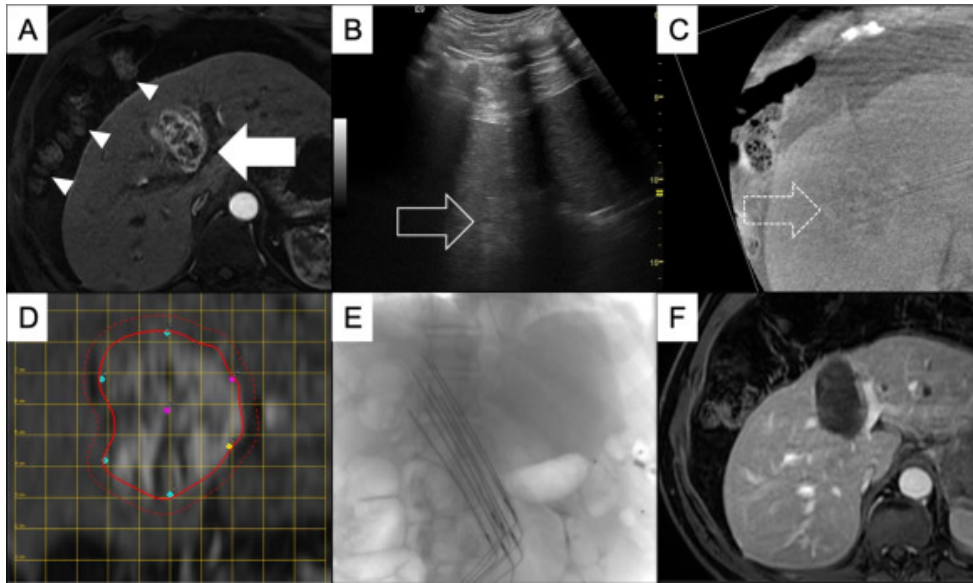


Figure 1.4: Case study from [Cor+19], the patient underwent multi-modality image guided electrochemotherapy of liver metastasis from renal clear cell adenocarcinoma. A. Preprocedural postcontrast T1-weighted MR image shows a single metastasis (arrow). B. Perprocedural ultrasound partially shows the tumor (arrow). C. Axial fused reconstruction of contrast-enhanced cone-beam computed tomography (CBCT) acquisition using electromagnetic tracking of the ultrasound probe allows three-dimensional visualization of the tumor and adequate needle placement. The dashed arrow shows enhancing metastasis. D. Individualized treatment plan, coronal reconstruction of T1-weighted MR images. E. Seven needles were inserted under multimodal imaging guidance including ultrasound and CBCT to cover the whole tumor. F. Postcontrast T1-weighted MR image in the axial plane at 2 months post-procedure demonstrates complete tumor response.

in cardiology to treat the narrowing of the coronary arteries of the heart found in coronary artery disease. The procedure called Percutaneous Coronary Intervention (PCI) combines coronary angioplasty with stenting, which is the insertion of a permanent wire-meshed tube after inflating a balloon from the angioplasty catheter to widen the blood vessel diameter.

Embolization The last type of procedure in IR consists in blocking arteries at the level of malformations (arteriovenous malformations, aneurysms) by placing a tiny metal coil to repair the ruptured aneurysms, or to place a stent-graft across an aneurysm to prevent the expansion of the defective vessel. Embolization can also be done to arteries feeding a tumor or an organ with an embolizing product, sometimes mixed with a product intended to destroy cells. For the latter, the referenced methods are the Transcatheter Arterial Chemoembolization (TACE) combining the embolic agent with a chemotherapy, and the Selective Internal Radiation Therapy (SIRT) combining it with a radiating glass or plastic microsphere using radionuclide (e.g. yttrium-90) emitting locally high energy beta radiation. The injection of embolic and destructive agents, if properly done, stops the tumor growth and kills its cells. The risks of impacting healthy tissues, however, are high, and these procedures require par-

ticular planning and assessment of the tumors as well as feeding vessels, leading to multiple image acquisitions such as SPECT and CT.

In the case of radioembolization of unresectable Hepatocellular Carcinoma (HCC), the patient is first catheterized in the IR room and injected with a low-dose radiating agent to assess where the radiating markers arrive. The patient is then sent to nuclear medicine and a Single-Photon Emission Computed Tomography (SPECT) based on gamma rays is performed allowing a direct correlation of anatomic and functional information in patients. This is also done to avoid extrahepatic deposition of the radiating microspheres as well as to assess their intrahepatic distribution. This work-up phase is an important step in the radioembolization procedure planning and the dose-finding. During the intervention, sometimes weeks later, the interventionalist can refer to the previous acquisition to verify the injection points and the position of the catheter with respect to the tumor-feeding arteries.

In such context, every computational tool for data comparison at the disposal of the interventionalists – and more generally the radiologists – is a potential gain in time for the clinicians and for the clinical outcome for the patient. In addition, these tools facilitate the communication between the interventionalists and their colleagues from other specialities, for: the description of the patient's disease, the procedure planning, the description of the delivered treatment afterward, or the evaluation of the response to treatment.

Given the exceptional diversity of procedures performed in interventional radiology, in terms of anatomy, treatment variants, or clinical gestures, it is important that tools developed are adaptable and allow different data to be combined. In particular, the development of methods that can be adapted to different anatomies – and therefore potentially to different shapes – seems indispensable. A particular IR application to the male reproductive system is the treatment of Benign Prostatic Hyperplasia through Prostatic Artery Embolization.

1.1.3 Prostatic Artery Embolization for Benign Prostatic Hyperplasia

Prostatic Artery Embolization (PAE) as a treatment of Benign Prostatic Hyperplasia (BPH) is a procedure of growing interest for IR. We have been able to work with interventional radiologists performing this procedure, and we describe it in the following section. BPH, sometimes called benign prostatic hypertrophy, is a condition in men in which the prostate gland is enlarged and not cancerous ⁴. The prostate (illustrated in Figure 1.5) is a walnut-shaped gland that is part of the male reproductive system. The main function of the prostate is to make a fluid that goes into semen. Prostate fluid is essential for man fertility. The gland surrounds the urethra at the area where the urethra joins the bladder: the bladder neck. The bladder and urethra are parts of the lower urinary tract. The prostate is in front of the rectum, just below the bladder. The urethra is the tube that carries urine from the bladder to the outside of the body. In men, the urethra also carries semen out through the penis.

The prostate goes through two main growth periods as a man age. The first occurs early in puberty, when the prostate doubles in size. The second phase of growth begins around age 25 and continues during most of a man's life. Benign prostatic hyperplasia often occurs

⁴The information comes from the National Institute of Diabetes and Digestive and Kidney Diseases.




Figure 1.5: The male reproductive and urinary systems: the prostate is located at the bladder neck, at the junction between the two systems.

with the second growth phase, most of the time in men above 40 or 50 years old. As the prostate enlarges, the gland presses against and pinches the urethra. The bladder wall becomes thicker. Eventually, the bladder may weaken and lose the ability to empty completely, leaving some urine in the bladder. The narrowing of the urethra and urinary retention—the inability to empty the bladder completely—cause many of the problems associated with benign prostatic hyperplasia.

The causes of BPH are not well understood but could be related to the decrease of testosterone production with age, or on the contrary to the production of dihydrotestosterone which would encourage the tissues to continue to grow. BPH particularly affects men after the age of 50, with a ratio of about 50% for men between 50 and 60 years old, and up to 90% for men aged 80 years and more. In particular, in addition to age, obesity, type 2 diabetes, erectile dysfunction or heart disease can lead to the development of BPH.

The symptoms of BPH are mostly due to the blocked urethra and the overworking bladder, so they are related to the difficulty to urinate normally: high urinary frequency, inability to delay urination, urinary incontinence, or on the contrary troubles starting a urine stream, weak or interrupted urine stream, urinary retention... Sometimes the symptoms also involve pain after ejaculation or during urination, urine unusual color, or smell. Consequently, an evaluation of the symptoms is made from a questionnaire called International Prostate Symptom Score (IPSS) around these problems described above. It consists of seven questions with a score from 0 – no such a symptom – to 5 – the symptom always occurs –. A specific deterioration of everyday life is the number of times one has to get up at night to urinate. From 0 to 7, the patient is mildly symptomatic, from 8 to 19 he is moderately symptomatic and from 20 to 35 he is severely symptomatic.

Complications and Treatments The complications of BPH may lead to acute urinary retention, chronic urinary retention, blood in the urine, urinary tract infection, bladder or kidney damages, or bladder stones. BPH is diagnosed on the base of personal and family medical history, a physical exam, and medical tests. During the physical exam the health care provider checks for discharge from the urethra, enlarged or tender lymph nodes in the groin, or a swollen, or tender scrotum. He or she also taps on specific areas of the patient's body and can perform a digital rectal exam. The medical tests at the health care provider's disposal are urinalysis, prostate specific antigen blood test, urodynamic test, transrectal ultrasound, and biopsy. The treatment options for BPH may include lifestyle changes, medications, minimally invasive procedures, and surgery. These solutions go from the less invasive treatment to the most invasive one. Lifestyle changes are primarily designed to avoid complications in public or during sleep periods. These include learning to hold urine for longer periods of time, or avoiding drinking when you can't easily urinate. Medications can be prescribed to stop the growth of the prostate, try to reduce its size, or reduce symptoms. Medications however may have side effects that sometimes can be serious: dizziness or fainting, rapid, pounding, or irregular heartbeat, sudden decrease, loss or blurred vision, decreased sexual drive, problems with ejaculation...

A urologist may recommend removing enlarged prostate tissue or making cuts in the prostate to widen the urethra. Urologists recommend surgery when medications and minimally invasive procedures are ineffective, symptoms are particularly severe, or complications arise. Like minimally invasive procedures, removing troublesome prostate tissue does not cure BPH. Surgery to remove enlarged prostate tissue includes Transurethral Resection of the Prostate (TURP), laser surgery, open prostatectomy, and transurethral incision of the prostate (TUIP). The surgery, however, is associated with higher risks of major adverse events such as bleeding, urinary incontinence, retrograde ejaculation, and impotence.

An alternative to conventional surgery is the minimally invasive approach. This intermediate solution can aim, in the same way as surgery, to destroy the problematic tissues of the prostate. Such approaches are image-guided and include transurethral needle ablation or transurethral microwave thermotherapy. The transurethral method involves inserting a catheter or cystoscope through the urethra to reach the prostate. Although destroying troublesome prostate tissue relieves many benign prostatic hyperplasia symptoms, tissue destruction does not cure benign prostatic hyperplasia. Prostatic stent insertion is sometimes done to widen the urethra, which can help relieve the blockage and urinary retention caused by benign prostatic hyperplasia.

Another approach called Prostatic Artery Embolization consists in clogging the feeding arteries of the prostate in the manner of tumor embolization. This intervention has received a growing interest in the last decade, and we will discuss it in the following section. As for these IR procedures, the potential complications are urinary tract infection, painful, difficult or irregular urination, urinary incontinence, blood in the urine for several days after the procedure, or sexual dysfunction. These complications subsequent to minimally invasive procedures go away within a few days or weeks, and they are less likely to have complications than surgery.

Prostatic Artery Embolization

As we have seen, for patients who do not respond favorably to drug treatments, PAE has become an important alternative to surgical approaches, its benefits to patients, and reduced cost, have been demonstrated [Bag+17; BSS19]. Prostatic Artery Embolization (PAE) is a safe procedure for the treatment symptoms due to Benign Prostatic Hyperplasia (BPH), with a major complication rate of less than 0.5% [BSS19]. In [Ray+18; Bro+22] the International Prostate Symptom Score (IPSS) median score is significantly improved after PAE at 12 months post-procedure (lower than the TURP) and causes less high-risk major adverse events.

In [Car+21] the authors propose a multicenter technical note based on several years of PAE procedures. They highlight the requirements in terms of understanding of the vascular anatomy, microcatheterization skills, and high precision in selecting the injection points to deliver the embolic agent without reflux [Ass+15b]. This is further illustrated in [Bro+22], with a PAE study using a radiopaque embolic agent, showing that in 100% of the post-procedure imaging non-target ischemia were seen, caused by unexpected embolic agent flow in the patient's vasculature. The procedures were done under angiography, and positioning was confirmed with DSA. The authors suggest therefore using a radiopaque embolic agent to assess its injection during PAE procedures. The non-target embolizations are reported to correlate with post-procedure symptoms (e.g. rectal pain). In [Car+21], however, the authors propose a standard workflow for PAE procedures, to minimize the risks of non-target embolization as well as the procedure time and the radiation dose for the patient and the clinician. This workflow is based on pre-procedural MRI or CT angiography for the understanding of the patient's anatomy, however, their lack of spatial resolution makes the complete visualization of the prostatic arteries difficult. Additionally, pre-operative US are recommended to measure the postvoid residual volume. Post-procedural MRI, 3 months and one year after the procedure, are recommended to assess the prostate reduction and the ischemic areas, pelvic US is recommended to measure the postvoid residual volume and prostate size.

The DSA offers good spatial resolution from a selective injection point. Its projective nature, however, can be misleading and does not allow the identification of small tortuous arteries. The need for multiple projections arises, increasing the radiation exposure, the procedure time, and the amount of injected contrast medium. As for that, CBCT acquisitions are encouraged [Roc+20] during the procedure. While capturing at least the same information, it also provides, in 46 to 60.8% of the cases, supplementary information such as potential non-target sites or prostatic artery origins, and anastomoses. Additionally, 3D roadmaps are recommended to perform safer, faster, and lower x-ray dose procedures.

The classic PAE consists of a bilateral successive embolization. For each side, proximal injected CBCT are performed from the internal iliac artery. These proximal CBCT are analyzed to identify arteries feeding the prostate, their pathways, and non-target vessels. This identification must be done bilaterally and with maximum attention to maximize treatment completion and reduce symptom recurrence risks. At this point, advanced planning software are suggested to facilitate the identification of the arteries and accelerate the procedure. Road-mapping and automatic vessel segmentation also facilitate the procedure planning and

visualization tools for augmented fluoroscopy.

Minimally invasive procedures allow for local treatments of pathologies with minimal side effects for the patients. They are alternatives to procedures that are sometimes burdensome for patients, or as solutions for patients who cannot undergo surgical procedures. In the case of PAE, while it does not cure BPH, the procedure has been shown to significantly reduce the symptoms of BPH for up to several years after the procedure. However, the position of the prostate in the male anatomy, and in particular the position of the prostatic artery in relation to other arteries irrigating sensitive anatomies such as the penis, the rectum or the bladder, increase the risks associated with the procedure.

Through this example of IR application, we realize the many issues related to advanced software applications for image analysis during procedures. The risks associated with artery embolization require physicians to be particularly cautious, and PAE requires a great deal of anatomical understanding as well as catheterization skills. The recommendations published in [Car+21] emphasize all precautions and the need to properly analyze the 3D volumes from the injected CBCT to understand the upcoming procedure. Despite their great expertise, the authors of the paper insist on the benefits of using advanced navigation software applications: roadmaps, virtual injections, automatic segmentation, all tools that allow a better understanding of the patient's anatomy.

The stakes for a facilitated PAE are clinical: non-targeted areas and small tortuous arteries, as well as injection points, are better identified, and interventional tools navigation is facilitated. The stakes are also economic by reducing the time of the procedures and by making the PAE easier to perform, it would decrease the number of surgical interventions and thus the time spent in the hospital for the patients. The stakes are also for communication, analysis, and follow-up of patient data: having the data analyzed during the procedure makes it easier to communicate the results or even to carry out follow-up interviews several months later.




Figure 1.6: Illustrations from [Car+17]: two DSA in the pelvic anatomy. The prostate is hardly visible, and the 2D view introduces ambiguities slowing down the procedure.

1.2 Modeling the Vascular Trees

In this thesis, we take advantage of 3D volumes acquired with interventional systems. These CBCT volumes are used in most endovascular interventions where doubts exist about the patient's anatomy in 2D projections. The benefits are a better understanding of the vascular network, and a potentially reduced dose to the patient because the acquisition of multiple injected 2D projections can generate more dose than a single CBCT. It also reduces the amount of contrast material injected into blood vessels.

Injected CBCT Workflow




Figure 1.7: CBCT angiography for Hepatocellular Carcinoma (HCC) embolization in the liver, from [Dur+18]. The vessels are clearly visible in the injected CBCT (a) under MIP, they are segmented in the volume with one click at the root (b). A manual segmentation of the tumor is done in (c) allowing a procedure planning in (d), and a navigation route. The 3D planning is superimposed on the live fluoroscopy (e).

During a typical angiography, a contrast agent is injected before the acquisition to reveal the arteries. The clinician decides which contrast agent is injected, its concentration, the delay before starting the CBCT acquisition as well as the injection rate during the acquisition. The choice of these parameters depends on the clinical application. For example, a large delay allows good opacification of tumors and the thin vessels in their vicinity that feed them.

In the pelvic anatomy, the 2D images corresponding to DSA do not allow to correctly

identify the target arteries to embolize: the prostate is hardly visible, and the dense vasculature introduces numerous ambiguities – due to the 2D projection –. The injected CBCT provides a 3D reconstruction that tackles the ambiguities and an improved visualization of the vasculature. In most of the anatomy, this solution is sufficient for the clinicians to plan their procedures.

In [Dur+18] the acquisition is done during a HCC embolization, the injected vessels are well depicted (Figure 1.7 A) in the CBCT volumes, and the vascular tree is segmented. In this example the tumor is also visible, and a Region of Interest (RoI) is manually contoured by the radiology technologist. It allows to identify the arteries feeding the tumor. Using EmboASSIST and Virtual Injection (GE Healthcare software), the clinician can assess the trajectory of the embolic agent inside the patient’s vasculature and detect potential non-target arteries.

In the pelvic anatomy however, solutions such solutions cannot be used because of the pelvic bones and the prostate that is hardly visible. The segmentation of the vessels is then key to properly visualize the tree, for instance to remove the bones from the volume. It also allows the clinicians to analyze in-depth the region of interest, and to define of the location where the treatment is to be done (as well as the potential non-target arteries close to the prostate).

The reconstruction of CBCT volumes is a whole topic of its own. For a detailed description of the volumes and their specificity, we refer to the work of [Res18]. In the following we focus on the segmentation of the vascular tree, taking the CBCT as input. We also discuss the centerline extraction and the pelvic vascular tree nomenclature.

1.2.1 From Images to Vessels

According to the [US National Cancer Institute Dictionary](#), a blood vessel is a tube through which the blood circulates in the body. Blood vessels include a network of arteries, arterioles, capillaries, venules, and veins. In our case, and to this point, the blood vessels are identified by high-intensity voxels in the 3D volumes due to the injection of a contrast agent. In order to extract the vascular tree structure from the rest of the image, a segmentation can be made resulting in binary 3D volumes in which voxels with value 1 indicate the presence of injected blood vessels, and 0 indicates any other structure or non-injected vessels. They can also be represented with the centerlines as described in the following.

Vessel Segmentation

A hessian-based filtering provides a volume in which the tubular structures such as the vessels are enhanced. The most frequently used filters, the Frangi filter, - [Fra+98] use the Hessian of the image at multiple scales to detect vessels of various radius. If I is a 3D image, it’s associated hessian matrix is given by:

$$H(I) = \begin{pmatrix} \frac{\partial^2 I}{\partial i^2} & \frac{\partial^2 I}{\partial i \partial j} & \frac{\partial^2 I}{\partial i \partial k} \\ \frac{\partial^2 I}{\partial j \partial i} & \frac{\partial^2 I}{\partial j^2} & \frac{\partial^2 I}{\partial j \partial k} \\ \frac{\partial^2 I}{\partial k \partial i} & \frac{\partial^2 I}{\partial k \partial j} & \frac{\partial^2 I}{\partial k^2} \end{pmatrix}$$

The eigenvalues of $H(I)$ at a point $q = (i, j, k)$ are denoted λ_1, λ_2 and λ_3 , with $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$.

The "vesselness" is given by the function:

$$R_F = \begin{cases} 0, & \text{if } \lambda_2 < 0. \\ \left(1 - e^{-\frac{r_1^2}{2a}}\right) \cdot e^{-\frac{r_2^2}{2b}} \cdot \left(1 - e^{-\frac{s^2}{2c}}\right), & \text{otherwise.} \end{cases} \quad (1.1)$$

with $r_1 = \frac{|\lambda_2|}{|\lambda_3|}$ that penalizes the plate-like structures, $r_2 = \frac{|\lambda_1|}{\sqrt{|\lambda_2||\lambda_3|}}$ enforcing tubular structures (opposed to blob-like ones), and $s = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$ allows a focus on areas with high contrast, a, b and c are constant parameters.

The hessian and vesselness can be computed in a multiscale fashion. Thresholding and post-processing the multiscale vesselness maps provides an accurate segmentation of the vessels. Such methods require a seed initialization, hence a manual click at the root r of the tree to select the seed point. This intermediate representation is helpful for understanding the patient's anatomy and for the procedure planning, however, it is not a sparse representation.

Centerlines Extraction

Vessels can be accurately represented by their centerlines, which may include a local radius. Such a format is much lighter to store and manipulate than their equivalent binary volume. Centerline extraction from the binary voxel segmentation S can be done through the following 4-steps:

- 1 The distance to the boundary $d_b(x) = d(x, \partial S)$ is computed for each voxel x of the segmentation.
- 2 The hessian of d_b is computed. A potential function $P_{d_b}(x)$ is built from its eigenvalues such that P takes very low values along the central axes of S and very large values elsewhere.
- 3 A front is propagated from the user indicated root point r inside S solving the eikonal equation:

$$\begin{cases} |\nabla T(x)| = P_{d_b}(x) \\ T(r) = 0 \end{cases} \quad (1.2)$$

- 4 Extremities are detected and geodesics are backtracked to the root using gradient descent in the function T .

A tree is then built from the extracted set of centerlines. The centerlines from the root to points in the vascular tree can be seen as potential navigation paths for the interventionalist's tools, or potential target path for the injected agents (in fact, the contrast agent flowed across these paths).

During the segmentation procedure, a click is made on a point of the volume which is then considered as the root of the vascular tree. In order to have an orientation, one can imagine that the indicated point corresponds to the artery from which the blood flows. The blood then flows from this point to all the other extremities of the tree, which gives a natural orientation to the structure: from the root to these extremities. The 3D segmentation will provide visualization tools, especially fusion with the 2D fluoroscopy. The centerlines representation provides geometric information that one can exploit to compute features and compare the curves.

We denote by $x_u = (x_{u,i})_{i=1,\dots,n_u}$ a set of ordered points in \mathbb{R}^d . This defines a polygonal curve interpolated by straight lines. The length of a polygonal curve is then the cumulative length of all its straight segments : $l(x_u) = \sum_{i=2}^{n_u} |x_{u,i} - x_{u,i-1}|_2$. We can now define a tree:

Definition 1. *A tree is a connected acyclic graph denoted (\mathcal{V}, Σ) with \mathcal{V} the set of nodes (or vertices) $\{0, \dots, M-1\}$ and Σ the $M \times M$ connectivity matrix with $\Sigma_{i,j} = 1$ if i is connected to j and 0 otherwise.*

This definition allows a variety of interpretation for the trees of centerlines, for instance one can see the centerlines as the nodes in a tree (\mathcal{V}, Σ) , or as the features associated with the edges connecting these nodes:

Definition 2. *Given a tree, we define an edge as any couple (i, j) such that $\Sigma_{i,j} = 1$. The set of edges in a tree is written $b = \{b_k\}_{k=0,\dots,M-2}$.*

Note that the adjacency matrix Σ can be retrieved from the edges b . We chose to represent the vascular trees as trees of centerlines, and define the 3D curves as spatial representation of the edges:

Definition 3 (Vascular Tree). *A vascular tree \mathbf{T} is a tree whose edges $b = (b_k)_k$ have a spatial representation as polygonal curve: $\mathbf{T} = (\mathcal{V}, b, x)$, with $x = (x_k)_k \in E^b$.*

From the construction of the centerlines, the nodes of a vascular tree can be seen as endpoints of the polygonal curves $x_k = (x_{k,i})_i$. The vascular trees are also binary trees: one node is connected to either one or three other nodes. As a consequence, if the edge b_j has two children b_k and b_l , then $x_{j,n_j} = x_{k,0} = x_{l,0}$.

Arteries and Labeling

In CBCT angiography, the imaged vessels are the arteries, namely the vessels feeding the organs. The blood then streams from the heart to the organs. In the pelvic anatomy, the vessels

are named accordingly to the organs they feed. In this context, one artery may correspond to one or several branches of the vascular tree, curves in the centerline representation. In the binary volumes, one artery then corresponds to a subset of the white pixels.

A label can be stored as an integer that corresponds to one artery name, and a labeling is the process of assigning a label to a subset of the vascular tree. More formally if a vascular tree is denoted $\mathbf{T} = (\mathcal{V}, b, x)$, and $L = \{1, \dots, N\}$ is a set of labels, then a labeling is an application $\mathcal{L} : L \mapsto \tilde{\mathcal{V}}$, with $\tilde{\mathcal{V}}$ a subset of \mathcal{V} . In the rest of the thesis, we will use the words **annotation** and **labeling** in an equivalent way. If there is a need to make a distinction it will be made clear in the appropriate place.

If the artery name design the vessels feeding a specific organ in the anatomy, its associated structure in the vascular tree is the image of \mathcal{L} in \mathbf{T} . Most of the time, when talking about automatic annotation, the application \mathcal{L} is surjective: one tries to assign a label to all the branches in the tree. In some cases however, it is possible to only assign label to a subset of the tree.

Manual Labeling The manual annotation of the arteries can be seen in the same way as a virtual injection: by clicking on a point in the tree, all downstream branches are identified. This can be compared to identifying all the blood vessels through which the product injected from the clicked point would flow. Note that the selected blood vessels form a subtree of the initial tree: a set of nodes, edges and centerlines that are connected and part of the complete tree. When manually labeling the trees, the technologist must identify the subtrees whose edges share the same label. Since the pelvic arteries are labeled accordingly to the organs they feed, manually annotating them consists in identifying the subtrees going to the same organ.

Two observations can be made at this point: this identification of the arteries is done on the CBCT during the interventions, either by a radiologist technician or by the radiologist himself. In addition to requiring a very good understanding of the anatomy, there are many sources of error (see Figure 1.8). It is therefore a complex and time-consuming task that is done during the procedure.

1.2.2 The Pelvic Vascular Tree

In line with the application to PAE, a specific focus on the male pelvic vasculature can be made to better understand it. During the PAE procedures, the radiologist acquires two injected CBCT : one per side of the patient. For each acquisition, the contrast agent is injected in the internal iliac artery, above the bifurcation of the anterior and posterior branches. From this point of injection, a number of arteries are clearly irrigated and visible in the injected CBCT:

- The **internal pudendal** artery branches off the internal iliac artery, the main artery of the pelvis, and supplies blood to the sex organs. The internal pudendal artery gives rise to the inferior rectal artery.




Figure 1.8: Illustration from Pr. Carnevale conference at Society of Interventional Radiology conference, 2022/06/11 - 2022/06/16, Boston USA. Annotation of arteries in the pelvic anatomy. Top: identification of the capsular artery, part of the prostatic artery. Bottom: identification of the rectal artery, that must not receive embolic agent. Left: 3D vascular tree segmented with different path from the root to end points. Right: Slice MIP, with selected branched highlighted.

- The **superior vesical** artery, supplying the upper part of the bladder. This artery often gives branches to the vas deferens (part of the male reproductive system) and can provide minor collateral circulation for the testicles.
- The **inferior vesical** artery, supplying the inferior part of the bladder. It is associated with the **prostatic** artery, feeding the prostate. Sometimes two subtrees are feeding the prostate, and the secondary tree is called the capsular artery.
- The **obturator** artery that passes forwards and downwards the lateral wall of the pelvis, and goes to the head of femur as well as muscles such as adductors.
- The **superior gluteal** artery, that supplies the muscles, as well as some skin over the sacrum.
- The **inferior gluteal** artery supplies the pelvic floor muscles. Upon exiting the pelvis

and emerging in the gluteal region, it splits into two branches to supply the sciatic nerve and the muscles and skin of the gluteal, hip and thigh regions.

- The **middle rectal** artery supplying the rectum.
- The **sacral** artery feeding the sacrum and the coccyx.
- The **iliolumbar** artery supplying the psoas major and the iliacus muscle.

Note the large number of possible anatomic variabilities. This illustrates the importance of understanding the patient's anatomy to avoid embolizing nontarget arteries. In [Ass+15b], the clinicians classify 5 types of pelvic vasculature of interest for the PAE procedures (Figure 1.9). The types 1 to 4 depends on the artery from which arises the prostatic artery. Classifying the patient anatomy in one of these four types helps the clinician understanding the vasculature and plan the embolization procedure. The fifth category is a catch-all, and includes all cases that do not fit the first four types. According to [Ass+15b], this represents about 8% of the vascular trees (among a population of 286 trees).




Figure 1.9: Illustration from [Car+17]. The four principal types of male pelvic vascular trees for PAE procedures. These illustrations focus on the arteries of interest (depending on the prostatic artery's source) and do not show all the possible arteries in the male pelvic vasculature.

The proposed classification is focused on the branching point of the prostatic artery. The total number of anatomical variants is in fact way higher and must take all the other arteries and their relative positions into account. As we will see in Chapter 4, in binary trees, the total

number of "variants", that we can call *topologies*, is of order $2N!! = 2N \times (2N - 2) \times \dots \times 1$ given N different leaves – or arteries.

The Variability of a Vascular Tree

The vascular tree of one patient can be very different from the one of another patient, and even more surprisingly, the tree on one side of the patient can vary from the one on the other side. These remarks are in the context of pelvic anatomy, however, one can keep in mind that the differences that can be listed here apply to many anatomies such as coronary, hepatic or brain vascular trees. We distinguish two kinds of differences between the vascular trees: the geometrical differences like the shape of the branches, and the regularity of the 3D curves, and the topological differences which come from the variability in terms of arteries, branching and density of branches.




Figure 1.10: Two different pelvic vascular trees. One color corresponds to one type of artery. The arteries feeding the right of the trees are the outer ones. The arteries feeding the left area are the inner ones. (a) A dense vessel tree with all the arteries (except for the middle rectal). (b) A vessel tree without obturator artery, and thin sparse interior arteries (left of the tree, from top to bottom): superior vesical, prostatic, pudendal.

Geometrical The geometrical variability of trees comes from the shape of their branches. We then talk about descriptors calculated on the curves, or functions allowing us to compare sets of curves. The variability can come from the density of points along the curves, the tortuosity of the branches, and their general shape. The pudendal artery, for example, forms a bend, almost a right angle, which makes it return to the interior of the body. This is the case in all patients (with rare exceptions). In Figure 1.10, the prostatic artery (interior of the patient, left of the figures) is much longer and more tortuous in Figure (a) than (b). The pudendal artery on the contrary has a similar shape (yet a different number of branches).

Topological Among the possible causes of topological variability across the vascular trees, we distinguish the anatomical, hierarchical and artificial reasons. The anatomical variability

is the presence or absence of certain arteries in the vascular trees. For example, the obturator is not always present on both sides of the patient. As the vascular trees are imaged on each side, one sometimes observes trees with this missing artery. It is illustrated in Figure 1.10 in which the obturator artery is present in the patient associated with the tree of (a) and not in associated with the tree of (b). This is true for the prostatic artery, the middle rectal and the sacral arteries as well. The other kind of anatomical variability is the complexity of the tree in terms of number of branches, which is also visible in Figure 1.10. Some vascular trees are dense and filled with small branches when some others contain only a few individual branches.

The other type of topological variability, in terms of tree topology, is the branching structure of the arteries characterized by [Ass+15b] for the PAE. When studying the trees as defined in 1, the branching structure can not be reduced to the information of where the prostatic artery originates. All the artery positions relative to the others must be taken into account. There are in theory a combinatorial number of possible topologies, function of the number of arteries in the tree.

The last kind of topological change comes from the vascular tree imaging pipeline, which may miss branches and arteries or introduce artifacts such as false anastomoses during the centerlines extraction. Such differences can also come from different injected CBCT acquisitions and parameters: the choice of the contrast agent, its density, the delay between the injection and the CBCT acquisition, etc.

Applied Mathematics in Interventional Radiology We realize through these multiple examples that mathematics applied to image and signal processing plays a major role in IR. Without proposing an exhaustive list of applications (for this we refer to [Gur+21]) we distinguish two major applications that can contribute a lot to the navigation of tools in the patient's anatomy: the augmented reality, using the pre-operative and per-operative supplementary information to enhance the 2D real-time fluoroscopy, and the multi-modality volumes registration providing an alignment of shapes and anatomies of interest in a common referential that is easier to read. In the case of PAE, this could be for instance the fusion of pre-operative US or MRA to better assess the position of the prostate with respect to the injected blood vessels. In order to further help the comprehension of the vascular tree and the navigation of the catheter as well as the planning of the PAE, the automatic annotation of the blood vessels in the CBCT could be of great interest for the interventionalists as well.

The multiplicity of information sources, along with need for image guidance in IR, create naturally a playground for the development of mathematical and computational applications. As many potential applications that lead today to an exponentially growing number of publications, especially since deep learning entered the scene. No matter the application or the proposed method, however, one of the essential problems remains the collection and annotation of data. This is all the more true in medical applications in which the data contains sensitive information and often requires expert annotations. These two aspects of the constitution of a consistent and representative database make it time-consuming, expensive, and sometimes unrealistic. In addition, the physicians do not always have time to dedicate to these annotations. The solutions one seeks to derive automatic tools to analyse the images may not only be learning-based and most of the time the first solutions engineered are classic

image – or data – processing tools. The results of these first algorithms are also often used as initialization for ground-truth creation, to feed future learning-based solutions. This can be done for example by using the output of the first algorithm as ground truth for training the second. However, such an approach risks introducing bias in the data, and there are methods to limit the impact of this bias in the learning procedure. It is the case of data synthesis approaches, or the semi-supervised ones. We will not discuss here the methods for boosting the training of learning based approaches, one must simply keep in mind that having a large, annotated and verified database is a luxury in most medical applications.

1.3 Automatic Annotation of Vascular Trees: A State of the Art

We have seen that understanding the pelvic vascular tree is absolutely necessary for the proper conduct of PAE interventions. More generally, the automatic annotation of tree-like structures has many clinical applications, from workflow simplification in cardiac disease diagnosis [Ezq+98; Aki+09; Cao+17; Wu+19] to intervention planning in arterio-venous malformations [Bog+13; GLS14; Rob+16; Wan+17] through lesion detection in pneumology [GBR08; Lo+11; Fer+12a].

1.3.1 Learning-based labeling

With the emergence of machine learning, many articles address the problem of automatic labeling of anatomical trees by feeding learning algorithms with a set of features describing the trees. The general framework consists in extracting features from the centerlines, learning from them to predict labels probabilities, and deciding the overall label assignment.

Early works [GBR08; Aki+09], first propose to learn from features on branches. The authors extract geometrical features (radius, length, direction, curvature...) and tree-shape features (angle between a branch and its parent) introducing a parent-child relationship. The label probabilities are learned for each branch independently using multivariate Gaussian models. Their assignments follow different rules according to the anatomy: respectively the airway tree and the coronary tree. These methods are direct but sensitive to the high variability in the data because neither the probability computation nor the final labeling guarantee anatomical consistency.

More recently, the K Nearest Neighbors (KNN) algorithm has been used in [Lo+11] to estimate the label probabilities per branch in airway trees. The approach takes into account the high variability of features and is robust to labels distributed over several clusters. Geometrical features are used too, and other tree-wise features - relative to the tree root - are introduced. The assignment is then performed bottom-up with a set of rules learned from the training set. All the previous methods are still limited in the case of high dimensional space: both multivariate Gaussian and KNN overfit in such spaces.

Moreover, other works [Hoa+11] and [Mat+14] propose the use of boosting algorithms

that search for the most discriminant features by cascading weak classifiers. That allows to introduce more features and be more accurate while preventing overfitting. In the first article [Hoa+11], the assignment follows a global optimization method which in the second [Mat+14] is performed as a maximum likelihood conditioned by the parent's labels and a set of rules specific to hepatic portal anatomy. Some articles also propose to work with a graph representation (especially for the brain vasculature). The authors of [Rob+16] compute label probabilities with an extremely randomized tree classifier. The assignment is performed through a maximum a posteriori with a set of transition rules - learned from the training set. Both previous solutions complement probability predictions with a set of strong rules implemented in the label assignment procedure.

This label assignment has been refined in [GLS14] and [Wan+17] by assuming the Markovian property: a branch label only depends on its direct neighbours. Supposing that the connectivity between labels should be more consistent across subjects than branch features only, the authors of [GLS14] use the simulated annealing scholastic relaxation as a label assignment's optimization. Authors of [Wan+17] use hidden Markov models with restricted transitions learned from the training set to provide better labels assignment.

As we saw, learning-based methods explicitly learn rules from the training set. In our context of frequent topological variations, this would require a lot of training data to cover as many anatomical variants as possible to be efficient. In practice, most of the aforementioned articles are trained on a limited size database (around 50). Because of this lack of data, very few articles propose deep learning solutions on the anatomical trees or graphs. The first applications as [AF18] were intended for binary classification of branches into arteries and veins. In [Wu+19], the authors propose to train a recurrent neural network preserving the topology of the coronary tree on a database composed of 436 annotated trees. The network is performing well in this application, automatically learning the transitions probabilities between the labels without the need for an explicit set of rules.

1.3.2 Atlas-based labeling

To be robust in the case of a limited-size database, some authors propose atlas-based methods to automatically annotate new cases. An atlas is defined as a reference model that can be built from prior knowledge [Bog+13; Cao+17] or from available annotated database [Ezq+98; Bul+06; Fer+15]. Most of the atlas-based methods follow a four steps framework: the atlas choice/construction, the registration onto the target, the estimation of the labels probabilities and the assignment.

In [Bul+06] the authors propose a point-wise solution for airway trees. There is no registration between the tree and the atlas; features are extracted at the point level. A distance between points in the feature space is computed to assign to each point a label. In order to assign labels to the branches a voting system is proposed. With these two steps, the authors try to address potential missing branches from the automatic tree extraction. On the other hand, there is no guarantee regarding the anatomical consistency along branches, and the use of a single atlas can be limiting in the case of high variability.

To take this variability into account, in [Bog+13] authors build as many graph atlases as

noticed topologies in the circle of Willis. Without any registration step, the authors calculate matching probabilities between the target and the atlas bifurcations. The assignment follows a maximum a posteriori on the likelihoods. Like learning-based methods, such a solution cannot adapt to unseen topologies.

Another approach, first described in [Ezq+98], uses a semi-automatic algorithm to manually register the coronary tree onto the target. The label assignment is performed following a branch-and-bound algorithm. The solution takes advantage of the registration: a better manual registration results in a simpler annotation solution.

Annotation in Tree Space To get rid of the need for manual interaction, the authors of [Fer+12a] have developed a non-classical topology registration method in the so-called tree-space. In such space, it is possible to compute continuous deformations between trees of different topologies. This approach is therefore robust to noise and anatomical variants and introduces a notion of distance between trees given their labels. The annotation follows a hierarchical scheme where a subtree is extracted and every labeling is tested at each iteration. The method selects the labeling minimizing the geodesic distance to the atlas' corresponding subtree. This solution is one way to cope with the limitation of using a single atlas but requires testing all possible labeling.

In [Fer+15], the authors further developed their tree-space annotation with a hierarchical assignment method. But instead of using a single case as an atlas, all annotated cases from the database are used. At each iteration, distances to all corresponding subtrees in the annotated database are computed. Once again, the labeling minimizing the sum of these distances is selected afterward. It has the advantage of taking into account all annotated trees available. On the other hand, it tends to label all the trees like the median of the training set. It also needs to compute all possible labeling in each selected subtree in addition to the distance to every atlas. These methods, testing all possible solutions, may be robust to anatomical variants, yet they are computationally expensive.

In [Gül+14] the authors use an atlas of the coronary tree defined by experts and build a two-step annotation taking advantage of some assumptions on coronary trees. As in [Fer+12a], the distances are computed in the tree space. A second cost is taken into account, based on the likelihood of the assigned labels. This second cost introduces a new a priori corresponding to the spatial distribution of the arteries relative to the heart surface. In order to speed up the optimization, the authors propose to follow the Dijkstra algorithm.

More recently in [Cao+17] the atlas is built by selecting the best example in a training set following a leave-one-in cross-validation design. To label a new tree, a rigid point-set registration is done and the labels are assigned according to a 2-step hierarchical algorithm. A direct label assignment to the main branches is performed first. Then, labeling the distal part is performed by trying all possible pairings between the atlas' and the target's branches as in [Fer+12a], although by staying in \mathbb{R}^3 the authors ensure faster results.

Recently in [Fey+19] the labeling is performed via an Optimal Transport (OT) algorithm. The algorithm is applied to brain tractograms instead of anatomical trees, yet it is close enough to this problem: the annotation is deduced from branch pairings between an atlas and a new case to annotate. The authors use the transport plan obtained for the annotation

and propose an interesting tool for automatic labeling.

In [Wu+19] the authors mention that the atlas-based methods perform well in the case of small databases. Yet, such techniques are generally sensitive to the anatomical variations and require, like the learning approaches, a set of rules or a priori information added to the labeling procedure. From this recent literature, [Cao+17; DKS18; Fey+19] computing meaningful matches between branches appears to be a promising approach to the tree-labeling problem. While a relevant registration would greatly improve the matches, the atlas-based articles working in \mathbb{R}^3 are limited to rigid deformations.

An early attempt for pulmonary trees registration has been proposed in [Gor+09], taking advantage of the current representation in a LDDMM framework to register both airway tree centerlines and lung surfaces. In another medical application of atlas-based methods, [Bai+12; Cer+13] propose to use LDDMM as registration step to segment a brain atlas. In particular, [Bai+12] compares elaborated segmentation models to registration-based approaches on small dataset. Few papers take advantage of the registration of the template onto an observation. In [Pan+16] the authors successfully register a 2D template of the airway tree onto a target, minimizing a distance in the space of varifolds. The LDDMM is selected as the best registration method, and its multi-atlas version provides them with the best segmentation results.

We see in the literature that a large variety of approaches address the automatic annotation of vascular trees. If learning based-methods – and in particular the deep learning ones – have shown increasing performance for such tasks in the past years, the lack of large, representative and annotated databases makes their use limited. On the contrary, the atlas-based approaches are suited to the limited amount of annotated data. Most of the methods however do not take advantage of the available unlabeled observations, or apply relevant automatic spatial deformations as a preprocessing to the annotation step.

1.4 Conclusion

We have seen that applied mathematics are deeply involved in the interventional radiology. From a clinical point of view, we have seen that the injected 3D CBCT reconstruction is key during endovascular IR procedures for the understanding of the vascular tree of the patient, and thus the navigation of the tools in the vessels. Today, interventional radiologists have at their disposal tools to detect and extract automatically these vessels from 3D volumes.

The automatic annotation of the branches of the extracted trees would be an asset for the smooth and shorter running of endovascular procedures, and could have many applications. It would facilitate the procedure planning, by helping the identification of target arteries and non-target ones. To that extent, it would also reduce the risks in particular when the injection of some treatment is involved. An example of application could be hiding the arteries of poor interest to the procedure and highlighting the potential arteries to treat or representing a high risk of complications. One could also provide a more focused roadmap, highlighting the vessels of interest, on top of the fluoroscopic images (guiding the operator).

The report of the procedure would also be more precise, with the potential application of indicating the different actions taken by the operator, as well as their locations within the blood vessels. This would facilitate the post-operative follow-up. In general, all the points listed above imply better technical and clinical outcomes for the patients. Regarding the benefits automatic vascular tree annotation could bring to endovascular IR, its development in a rigorous mathematical setting is key.

From the applied mathematics point of view, we emphasize the fact that a good learning-based approach always requires a large, representative, annotated and cleaned-up database. Having such a database is a big challenge, especially in medical imaging where the sensitivity of the data and the difficulty to annotate slow down the development of learning-based techniques [LY17]. In this context, solutions to the tree labeling problem should ideally work from a limited number of annotated samples. The atlas-based approaches are therefore best adapted. We identify three main stakes in the use of automatic annotation of vascular trees based on an atlas: the construction of realistic deformations to align one observation to another, the construction of a relevant atlas in which we can select a template, and the use of deformations in the framework of automatic annotation.

We believe in the paradigm: "The better the shapes are aligned, the easier it is to compare them". Thereafter, the non-rigid registration of a template annotated vascular tree onto an unlabeled one would facilitate its automatic labeling. Regarding the high complexity of the vascular trees, the template should be a synthetic and representative summary of the real observations. In order to take advantage of the available information despite the lack of annotations, we propose to derive a statistical analysis of the deformations of the template onto the rest of the unlabeled (and simplified) observations. The challenge is then to adapt the registration of the template shape to real targets, and thus to the potential topological changes that often occur when comparing different trees. To that end, we propose to search for relevant deformations that would include the simplified template in the target tree. The topological changes can also be encoded within the atlas, providing a valuable space for the template.

LARGE DIFFEOMORPHIC DEFORMATIONS AND APPLICATION TO VASCULAR TREE ALIGN- MENT

*The variability of anatomical shapes across individuals makes the study of populations challenging. One point of view is to observe the shapes as elements of the same **atlas** and consider the deformations one must generate to align two points of this atlas. **Computational Anatomy** formalizes it by studying deformation maps for the shape alignment, deriving statistics from a population of shapes, and infer on new geometric data embedded in the same space.*

| | | |
|-------|---|----|
| 2.1 | Computational Anatomy | 36 |
| 2.2 | Reproducing Kernel Hilbert Spaces: a Cornerstone | 38 |
| 2.3 | A Space of Deformations | 44 |
| 2.3.1 | Flows of Vector Fields | 45 |
| 2.3.2 | Computing Large Diffeomorphic Deformation | 47 |
| 2.4 | Metric Mapping: from Exact to Inexact Registration | 49 |
| 2.4.1 | Spaces for the Data Attachment | 52 |
| 2.4.2 | Matching Oriented Varifolds | 56 |
| 2.4.3 | Optimal Transport Cost as Data Attachment Term | 59 |
| 2.5 | Statistics over Deformations | 62 |
| 2.5.1 | Building a Template of a Population | 62 |
| 2.5.2 | PCA and LDDMM | 65 |
| 2.6 | Application to the Registration of Simplified Pelvic Vascular Trees | 66 |
| 2.7 | Conclusion | 71 |

2.1 Computational Anatomy

Computational anatomy as we see it nowadays was first formulated in the seminal work of Grenander and Miller [GM98] which made the connection between the initial theory of d’Arcy Thompson in 1917 [Tho17] and the contemporary tools both in mathematics and informatics. In [Tho17] the key idea is to see the population of shapes through the geometric transformations to align the individuals together.

Such transformations should be both computationally tractable and one should be able to apply these deformations to other shapes belonging to the same space. Moreover, they should be sufficiently regular and potentially invertible. Let us take for example the Thin Plate Splines (TPS) model proposed by [Duc77; Mei84], corresponding to the estimation of a vector field for the registration of landmarks: the model provides a deformation that only depends on the landmarks. However, there is no guarantee that the deformation is invertible, which can lead to folds in the space. In the context of computational anatomy, such deformations are not satisfactory as they do not allow for robust inference. On the contrary diffeomorphisms of the ambient space provide smooth and invertible transformations that act upon the embedded shapes. They were formalized and made computationally tractable in the 90s by the seminal works of [Tro95; CRM96; DGM98; Tro98] proposing to build the large deformations upon a group of diffeomorphisms acting on the ambient space of the shapes. It has been summarized later in [You10]. Other works introduced alternative diffeomorphic deformations that are now standard in the community: the DARTEL algorithm [Ash07] adapting the large diffeomorphic deformations and the diffeomorphic DEMONS [Ver+09].

Regarding the large variability of shapes that are being digitized today, one key step in the analysis of the shapes is to implement metrics adapted to their comparison. Getting back to the TPS example, if the landmark representation can be used for any discrete shape, the TPS requires them to be *labeled* in both shapes to align to establish correspondences. Most of the time the labeling is done manually, and such a process requires time and makes the constitution of a database even more complicated. In line with the computational anatomy framework, attempts have been made to model the shapes in a relevant space in order to derive metrics to properly compare them. It has been applied to images [Beg+05] and diffusion tensor images [Cao+06] as well as measures and spaces adapted to smooth curves and surfaces [Gla05; CT13; KCC17; RG16].

The inexact registration framework has been applied to various shapes, and most of the methods do not require any annotation of the shapes to be applied: this way one can apply them to databases without the need for manual labeling. On the other hand, the statistics over the deformations and the inference to new shapes are well adapted to statistical analysis [Pen06], and methods have been developed to build templates from sets of deformations between the shapes [Vai+04] with applications for instance to the study of brain MRI [Cur+18]. It has also been developed in the context of temporal data [Dur10].

Due to the ever increasing efficiency of deep learning methods in a large range of domains, it has also been applied to the computational anatomy framework. Several authors adapt deep-learning methods to the classical registration frameworks to provide more efficient inference. Such approaches are inspired by the Stationary Velocity Fields frame-

work using spacial transformer networks first introduced in [JSZ+15], to register the source (moving image) to the target (fixed image) by optimizing a loss function including both an image similarity term and a regularization on the transformations term. One pioneering approach is called VoxelMorph and works on (structured) 3D medical images [Dal+18]. One other slightly different application of deep-learning is to learn a spatially-varying regularizer [NKV19]. Such approaches allow estimating deformations to align shapes and derive statistics over the deformations. They compare to the classical methods such as ANTS [ATS+09] and diffeomorphic DEMONS, with similar registration performances but provide an inference time about 100 times faster on CPU ([Dal+18], Table 1).

The tools of computational anatomy have been used in many medical applications for the understanding of diseases or the modeling of organs. Due to the growing use of (medical) imaging devices, the available data are also more and more numerous. Thus, one can find for example studies of cardiovascular Magnetic Resonance Images (MRI) for the statistical shape analysis of repaired aortic coarctation arches, and allowing the extraction of 3D anatomical features and assess their variation to template features' [Gui+21]. One can also find statistical shape analysis of the population of hippocampus [Cur+18] extracted from MRI for the construction of population centroid and the disease detection, or longitudinal study of the same organ along a human lifespan [Yan+13] providing robust methods for evaluating morphological measures. Such studies find applications to the study of Alzheimer's disease. Applications can be found on retinal fundus surfaces [Lee+17] extracted from retinal fundus photographs for the purpose of glaucoma diagnosis. In [Zol+17], the ear shapes are modelled and the authors compare the models with the ear acoustic, in the perspective of ear prosthesis fabrication. Of course, the registration framework has also been applied to images [Ash07] for numerous kinds of imaging techniques, let's cite for instance brain MRI [Kle+09], cardiac MRI [Beg+04], or lung Computed Tomography (CT) [Ris+13].

Apart from the human anatomy, some authors are interested in the study of shape growth, for instance, the leaves growth with modelling of multiple global and local deformations [Lac+21] or the study of horn shapes through growth models [KT18]. In [Még+21] the authors characterize the shape deformation of gene expression to identify gene deregulation among Huntington's disease knock-in mice.

One sees that the approach of Computational Anatomy fits perfectly in the context of atlas construction and template registration for automatic annotation of vascular trees.

Chapter Organization In this chapter we introduce the minimum theoretical background needed to understand the computational tools we use in the rest of the thesis. Section 2.2 is dedicated to the construction of Reproducing Kernel Hilbert Spaces with relevant properties adapted to deformations generation and shapes comparison. In Section 2.3, we provide the main ideas of the space of deformations we use to compute diffeomorphisms. In Section 2.4, we discuss the different metrics at our disposal to derive inexact matching methods, in particular for curves and surfaces registration. In Section 2.5 we discuss the statistical tools to derive statistics over the deformations of templates on populations of shapes. They allow building a template as well as an atlas adapted to the observations in the sets.

The reader familiar with the LDDMM framework might skip these sections and go im-

mediately to Section 2.6, in which we apply the LDDMM to the registration of vascular trees. This first utilization with simplified vascular trees allows us to identify areas of improvement for the computation of realistic diffeomorphic deformations of vascular trees. Their use in an atlas-based automatic tree annotation method is developed in Chapter 5, Section 5.1.

2.2 Reproducing Kernel Hilbert Spaces: a Cornerstone

We first provide the key ingredients to the construction of Reproducing Kernel Hilbert Spaces (RKHS), which will serve as a basis for writing equations in the discrete domain for both computing the deformations and the shape's comparison. The underlying classical notions, first introduced in [Aro50] and extensively studied in the literature will be used throughout the thesis. Note that contrary to the classic RKHS framework, we are working with vector-valued function, thus matrix kernels.

Let H be a Hilbert space equipped with its inner product $\langle \cdot, \cdot \rangle_H$ and its norm $|\cdot|_H$. Its dual H' is the space containing all the continuous linear functions $\phi : H \mapsto \mathbb{R}$. The normed space structure of H' is defined by:

$$|\phi|_{H'} = \max_{\{h \in H, |h|_H=1\}} \phi(h).$$

For any $h \in H$, the function $\phi_h : h' \mapsto \langle h, h' \rangle_H$ belongs to H' , and using the Schwartz inequality in the definition of the dual norm, one can prove that $|\phi_h|_{H'} = |h|_H$.

Theorem 1 (Riesz Representation Theorem). *If $\phi \in H'$, there exists a unique $h \in H$ such that $\phi = \phi_h$.*

The proof is given, for instance, in [You10], Appendix A. The theorem provides, therefore, both a mapping between the dual space and the primal and the equivalence in the norm of the element of the dual and its *representer* in the primal that will be essential in the RKHS theory.

Definition 4 (Reproducing Kernel Hilbert Space). *Let H be a Hilbert space of vector-valued functions defined on a set Ω with values in \mathbb{R}^d . Then, H is said to be a Reproducing Kernel Hilbert Space (RKHS) if for all $x \in \Omega$ and $\alpha \in \mathbb{R}^d$, the functional $\delta_x^\alpha : u \in H \rightarrow \langle \alpha, u(x) \rangle_{\mathbb{R}^d}$ is a continuous linear form on H , in other words if $\delta_x^\alpha \in H'$.*

Note that we did not use the notion of kernel in the definition of the RKHS. For a RKHS, we can define its reproducing kernel as follows:

Proposition 1 (Reproducing Kernel). *Let H be a RKHS of functions defined on a set Ω with values in \mathbb{R}^d . There exists a unique function $K : \Omega \times \Omega \mapsto \mathcal{M}_d(\mathbb{R})$ which is called the reproducing kernel of H and satisfies:*

- 1 For all $x \in \Omega$ and $\alpha \in \mathbb{R}^d$, $K(x, \cdot)\alpha \in H$.
- 2 For all $x \in \Omega$, $\alpha \in \mathbb{R}^d$, and $u \in H$, $\langle u, K(x, \cdot)\alpha \rangle_H = \langle u(x), \alpha \rangle_{\mathbb{R}^d}$

From the definition of the RKHS, since $\delta_x^\alpha \in H'$, the Riesz representation theorem 1 gives us that there is a unique function $K_x^\alpha \in H$ such that for all $u \in H$:

$$\delta_x^\alpha(u) = \langle \alpha, u(x) \rangle_{\mathbb{R}^d} = \langle K_x^\alpha, u \rangle_H.$$

Such a function is called the representer of the evaluation functional δ_x^α . In addition, it is easy to prove that K_x^α is a linear function of α so that it can be written $K(x, \cdot)\alpha$, with $K : \Omega \times \Omega \mapsto \mathcal{M}_d(\mathbb{R})$. The functions $K(x, \cdot)\alpha$ are called the fundamental functions of the RKHS and K its associated Reproducing Kernel.

Reproducing Property Given $y \in \Omega$, $\beta \in \mathbb{R}^d$, and choosing $u = K(y, \cdot)\beta$, the second point of Proposition 1 gives:

$$\langle K(x, \cdot)\alpha, K(y, \cdot)\beta \rangle_H = \alpha^T K(x, y) \beta,$$

which is called the reproducing property.

Positivity Another fundamental property of reproducing kernels is that they are positive. It can be shown similarly to the reproducing property. We recall the definition of positive kernels in the context of matrix kernels:

Definition 5. Let Ω be a set and $E = \mathbb{R}^d$ a finite dimension euclidean space and $\mathcal{M}_d(\mathbb{R})$ the space of $d \times d$ real matrices. We say that K is a positive symmetric definite kernel on Ω with values in \mathbb{R}^d if $K : \Omega \times \Omega \mapsto \mathcal{M}_d(\mathbb{R})$ and satisfies:

- 1 For all $x, y \in \Omega$, $K(x, y) = K(y, x)^T$.
- 2 For all $N \in \mathbb{N}$, $x_1, \dots, x_N \in \Omega$ and $\alpha_1, \dots, \alpha_N \in \mathbb{R}^d$ we have:

$$\sum_{i,j=1}^N \alpha_i^T K(x_i, x_j) \alpha_j \geq 0.$$

In other words, the *block* matrix defined by $(K(x_i, x_j))_{i,j=1..N}$, $x_1, \dots, x_N \in \Omega$ is symmetric positive.

Proposition 2. Let H be a RKHS, then the reproducing kernel of H is a positive definite kernel.

Let H be a RKHS of vector-valued functions defined on a set Ω with values in \mathbb{R}^d , and K its reproducing kernel. Consider $x_1, \dots, x_N \in \Omega$ and $\alpha_1, \dots, \alpha_N \in \mathbb{R}^d$, we have :

$$\sum_{i,j=1}^N \alpha_i^T K(x_j, x_j) \alpha_j = \langle \sum_{i=1}^N K(x_i, \cdot) \alpha_i, \sum_{j=1}^N K(x_j, \cdot) \alpha_j \rangle_H \geq 0$$

From Kernels to RKHS We can now recall the following central theorem:

Theorem 2 (Moore-Aronszajn). *To any symmetric, positive definite kernel K on Ω with values in \mathbb{R}^d , there corresponds a unique RKHS H of functions from the set Ω to \mathbb{R}^d whose reproducing kernel is K .*

The proof (see for instance [Gla05] Chapter 2) relies on the construction of the vector space H_0 spanned by the fundamental functions $K(x, \cdot)\alpha$ that is in fact dense in H so H is also an RKHS of functions limit of Cauchy sequences in H_0 with the reproducing kernel K and is unique. To this point, we have seen that a RKHS theory provides a simple and elegant way to compute norms of elements in the RKHS, and most of the calculus and functions evaluations boil down to scalar products in euclidean spaces.

Example 1: Vector Splines Interpolation

As mentioned in Section 2.1 a fundamental example is the landmarks registration in the ambient space $E = \mathbb{R}^d$, most of the time $d = 2, 3$. The problem can be formulated as : Given x_1, \dots, x_N distinct points in $\Omega \subset \mathbb{R}^d$ and $y_1, \dots, y_N \in \mathbb{R}^d$ N distinct landmarks, we are looking for a vector field $v \in V$ living in a RKHS canonically embedded in $L^2(\Omega, \mathbb{R}^d)$ with minimum norm such that $v(x_i) = y_i, \forall i \in \{1, \dots, N\}$.

The solution to this problem has the form $v(x) = \sum_{i=1}^N K(x, x_i) \cdot \alpha_i$, where the $\alpha_i \in \mathbb{R}^d$ are solutions of the linear system of equations $\sum_{i=1}^N K(x_j, x_i) \cdot \alpha_i = y_j \forall j$, and moreover one has that $|v|_V^2 = \sum_{i,j=1}^N \alpha_j^T \cdot K(x_j, x_i) \cdot \alpha_i$.




Figure 2.1: Landmark pairing by solving a linear system. Left: the source points. Center: the target points. Right: the pairing and the associated deformation applied to the ambient space.

Example 2: Relaxed Vector Splines Interpolation

The result of Example 1 allows pairing labeled points with a vector field that can be applied to the whole space, however, as mentioned in the introduction such a solution raises several issues: first the landmarks must be labeled – most of the time manually –, second the deformation of the ambient space can lead to folding and non-invertible transformations as illustrated in Figure 2.1. In practice, the Vector Splines Interpolation Problem can be reformulated with soft constraints, by relaxing the exact landmark correspondence with a distance term :

$$\operatorname{argmin}_{v \in V} |v|_V^2 + \gamma \cdot \sum_{i=1}^N |v(x_i) - y_i|_{\mathbb{R}^d}^2.$$

The solution to this problem has the form $v(x) = \sum_{i=1}^N K(x, x_i) \cdot \alpha_i$, where the $\alpha_i \in \mathbb{R}^d$ are solutions of the linear system of equations $\sum_{i=1}^N \left(K(x_j, x_i) + \frac{1}{\gamma} \right) \cdot \alpha_i = y_j \forall j$. The relaxed formulation provides one way to play with the regularity of the vector field by adjusting the parameter γ , and we call this new problem the **inexact registration** framework.




Figure 2.2: Relaxed landmark pairing by solving a linear system with $\gamma = 0.5$. Left: the source points. Center: the target points. Right: the pairing and the associated deformation applied to the ambient space.

The theorem 2 will be fundamental for the construction of RKHS since it will allow the construction of spaces from the construction of their reproducing kernel. However, the choice of the RK is decisive, and the counterpart will be to build kernels with the "right

properties".

Definition 6 (Continuous embedding). *If $(X, |\cdot|_X)$, $(Y, |\cdot|_Y)$ are two normed spaces, if $X \subset Y$ and there is a constant $c > 0$ such that for all $x \in X$, $|x|_Y \leq c \cdot |x|_X$, then we say that X is continuously embedded in Y and we write $X \hookrightarrow Y$.*

Proposition 3 ([Gla05], Chapter 2, Theorem 9). *Let $K : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathcal{M}_d(\mathbb{R})$ be a positive definite kernel with its derivatives at order $\leq 2p$, $p \in \mathbb{N}$ continuous and bounded. Suppose also that for all $x \in \mathbb{R}^d$, $K(x, \cdot)$ and all its derivatives at order $\leq p$ vanish at infinity. Then the RKHS H associated with K is continuously embedded into $C_0^p(\mathbb{R}^d, \mathbb{R}^d)$.*

This last property provides guarantees about the regularity of the element of the RKHS, functions from \mathbb{R}^d to \mathbb{R}^d , depending on the regularity of the chosen kernel. This is important in the construction of the RKHS of interest since the RKHS will be used both in the construction of the deformations and in the data attachment term for the shape comparison. The regularity of the functions guarantees the regularity of the generated deformation. Moreover, the fact that the kernel vanishes at infinity ensures that the deformation gets close to identity at infinity.

We now need to provide some invariance properties for the kernels to adapt them to the deformations and shape comparison purposes.

Theorem 3 (Bochner). *Let $k : \mathbb{R}^d \mapsto \mathcal{L}(\mathbb{R}^m)$ be an integrable function and its Fourier transform \hat{k} verifying $k = \mathcal{F}^{-1}\hat{k}$. Then $K(x, y) = k(y - x)$ is positive if and only if \hat{k} is such that for any $w \in \mathbb{R}^d$, $\hat{k}(w)$ is a positive hermitian operator.*

Proposition 4. *Suppose the conditions of Theorem. 3 true. If in complement, the function $\hat{k}(\omega)$ is positive and bounded on \mathbb{R}^d , then the scalar product of the RKHS H associated with $k(x, y) = k(y - x)$ can be written:*

$$\langle f, g \rangle_H = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \bar{\hat{f}}(\omega) \hat{k}(\omega)^{-1} \hat{g}(\omega) d\omega$$

The two previous results provide a method to build kernels inducing metrics invariant to rotations and translations. This is particularly important for instance in the construction of metrics between the shapes since one would like the metric to be invariant for example to change of referential.

Scalar Radial Kernels

In practice, we are interested in comparing the shapes no matter the referential, so the rotations and translations should be isometric for the metric induced by the reproducing kernel we build. In other words, the kernels we are looking for should be invariant to rotations and translations. In theory, it is possible to build matrix kernels with the correct properties. In Appendix 5.5 we build such kernels with, in addition, divergence-free and curl-free properties. In practice, we exclusively worked with scalar kernels of the form $K(x, y) = k(x, y) \cdot Id$

with k a scalar valued function. We focus therefore on the construction of the last ones. One way to build kernels invariant to isometry [Gla05] is to set

$$K(x, y) = k(|y - x|).Id, \forall x, y \in \mathbb{R}^d$$

with $k: \mathbb{R}^+ \mapsto \mathbb{R}$ a positive scalar function. This is a convenient construction but note that the class of isometry invariant RKHS remains larger.

Proposition 5. *A function $k: \mathbb{R}^+ \mapsto \mathbb{R}$ defined as in 2.2 is a positive kernel on \mathbb{R}^d if one of the following condition is true:*

- (a) $r \mapsto k(\sqrt{r})$ is strictly monotonous.
- (b) $k(\sqrt{r}) = \int_0^\infty e^{-r^2 u^2} d\mu(u)$, with μ a positive bounded Borel measure on $[0, \infty[$.

This proposition comes from the work of Schoenberg [Sch38] and provides the characterization of a convenient class of functions to use in order to build scalar kernels with the correct invariance properties.

Scale Parameter In practice, one would like the reproducing kernel to be scale invariant. For instance, when comparing the shapes with distance in a RKHS, one would like this distance $|x - y|_H$ to be the same as $|\alpha x - \alpha y|_H$ for $\alpha \in \mathbb{R}^{+*}$. In particular, one would like $k_H(\alpha.x - \alpha.y) = k_H(x - y)$ which is impossible. A classic answer is to introduce a scale parameter $\sigma \in \mathbb{R}^{+*}$ in the kernel:

$$K_H(x, y) = k_H\left(\frac{|y - x|^2}{\sigma^2}\right).Id.$$

This scale is particularly important in practice since it defines a range of interaction between an element of the RKHS and its ambient space.

Gaussian Kernels For $\sigma > 0$, the most commonly used kernel is the Gaussian kernel:

$$K_\sigma(x, y) = e^{-\frac{|y - x|^2}{\sigma^2}}.Id.$$

Using Proposition. 4, K_σ is a positive kernel and its associated RKHS H_σ is continuously embedded in $C_0^k(\mathbb{R}^d, \mathbb{R}^m)$ for all $k \in \mathbb{N}$. We have

$$|f|_{H_\sigma}^2 = \frac{\sigma^d}{2^d \pi^{d/2}} \int_{\mathbb{R}^d} e^{-\frac{\sigma^2 |\omega|^2}{4}} |\hat{f}(\omega)|^2 d\omega$$

and two distinct σ and σ' define two distinct RKHS associated with K_σ and $K_{\sigma'}$ and $K_{\sigma'}(x, \cdot)\alpha$ belongs to H_σ if and only if $\sigma' > \frac{\sigma}{\sqrt{2}}$.

The norm being finite, it implies that the Fourier transform decreases quickly, which ensures the smoothness of the elements of H_σ . Such – well known – Gaussian kernels must

be tuned carefully in order to build a RKHS adapted to the shapes. For instance the smaller the scale, the more local the variations of the shapes are taken into account but the less far it will "see" other shapes. In the rest of the thesis, the Gaussian kernels will be extensively used for both computing the deformations of the ambient space and the distance between the shapes in the shapes spaces.

As the tuning of the scale parameter is key in the construction of relevant registration solutions, one classic approach is to use multi-scale kernels similarly to a coarse-to-fine solution:

$$K(x, y) = \frac{1}{n} \sum_{i=1}^n e^{-\frac{|x-y|^2}{(\sigma_0/s_i)^2}} . Id$$

where $(s_i)_i \in (\mathbb{R}_+^*)^n$ are the scaling coefficients and $\sigma_0 \in \mathbb{R}_+^*$ is a reference scale.

There exist other kernels that can be used such as the Cauchy kernel:

$$K(x, y) = \frac{1}{1 + \frac{|x-y|^2}{\sigma^2}} . Id.$$

In some cases (e.g. the Sobolev spaces), one can start constructing the RKHS and then deduce the associated kernel with the right properties.

2.3 A Space of Deformations

In the previous section we built the basic tools to generate reproducing kernel spaces with invariance and regularity properties. We are now interested in the construction of large deformations suitable for the statistical analysis of anatomical shapes. From [You10] we can identify the required properties for the deformations of the *ambient space* Ω , open set of \mathbb{R}^d . Getting back to the registration problem, we need that the deformation $y = \phi(x) \in \Omega$ of a point $x \in \Omega$ provides at least the following guarantees: first not to create holes in Ω so the deformation should be onto, and second it should not fold the space (see for example Figure 2.1), so it must be injective. All in all the deformations we are looking for are bijections of Ω .

One way to quantify the difference between two shapes is to quantify the amount of deformations needed to align them – thus an extrinsic quantification –. In an ideal setting, one would like to align exactly the shapes m of a set M living in Ω . The actions on the shapes belong to a certain group G of transformations Ω . In the following we will build such a group of transformations and assume that the action of G is transitive on M : for any $m_0, m_1 \in M$, there exists $g \in G$ such that $g.m_0 = m_1$.

This may look restrictive, in particular in the case of discrete shapes and regular transformations. There is however at least one typical application to the registration of one reference shape S onto a set of target shapes $M = (m_0, \dots, m_n)$. If solutions to inexact registrations $g_k \cdot S$, with $(g_k)_k \in G$ exist, the set $(g_k.S)_k$ belongs in fact to the orbit $G \cdot S$. The LD-DMM theory allows building such a group of bijective transformations, and we detail it in the following.

Definition 7. A homeomorphism of Ω is a continuous bijection $\phi : \Omega \mapsto \Omega$ such that its inverse ϕ^{-1} is continuous. A diffeomorphism – or C^1 -diffeomorphism – of Ω is a continuously differentiable homeomorphism $\phi : \Omega \mapsto \Omega$ such that its inverse ϕ^{-1} is continuously differentiable.

In order to build bijective deformations of the shapes preserving their structure (as well as their tangent structure for the manifolds), from now on the group G is chosen as a subgroup of the group of C^1 -diffeomorphisms of Ω . For a function $f \in C_0^p(\Omega, \mathbb{R}^m)$ we define the norm $|f|_{p,\infty} = \text{Sup}_{x \in \Omega} \{|f(x)| + \sum_{k=1}^p |d_x^k f|\}$ with $d_x^k f$ the k -th differential of f at x .

The intuition behind the construction of the subgroup is to start by considering "small perturbations" of elements in Ω . Consider a function $v \in C_0^1(\Omega, \mathbb{R}^d)$ such that there exists δ_0 such that for all $x \in \Omega$ and all $\delta < \delta_0$, $x + \delta v(x) \in \Omega$. One can show ([You10] chapter 8) that for a real $\epsilon > 0$ small enough, $\phi : x \mapsto x + \epsilon v(x)$ is a diffeomorphism of Ω . This basic diffeomorphism may not provide the large deformations expected, but we can retrieve them using a composition of small perturbations. Taking the limit $\epsilon \rightarrow 0$, one can see the diffeomorphism as a composition of infinitely small deformations. This corresponds to the explicit Euler scheme for the differential equation:

$$\begin{cases} \dot{x}(t) = v_t(x(t)) \\ x(0) = x_0, \end{cases} \quad (2.1)$$

The resulting diffeomorphism is therefore intrinsically linked to the flow of a time varying vector field of the ambient space.

2.3.1 Flows of Vector Fields

In the previous elementary example of the small perturbation of the identity as a first diffeomorphism, we already required some assumption for the function $v \in C_0^1(\Omega, \mathbb{R}^d)$. We will now build a space of **admissible vector fields**, and assume that Ω is \mathbb{R}^d which is more natural when looking at deformations of the ambient space.

Definition 8 (Admissible Vector Fields). A vector space V of vector fields on \mathbb{R}^d is admissible if it verifies:

1. V is a Hilbert space.
2. $(V, |\cdot|_V)$ is continuously embedded in $(C_0^1(\mathbb{R}^d, \mathbb{R}^d), |\cdot|_{1,\infty})$ the space of C^1 vector fields on \mathbb{R}^d vanishing at infinity along with all their partial derivatives: there exists a constant $c_V > 0$ such that for all $v \in V$, $|v|_{1,\infty} \leq c_V |v|_V$

From now on, we will denote by V a space of admissible vector fields. This space can be seen as the tangent space of G at identity and its norm $|\cdot|_V$ can be seen as the infinitesimal displacement cost.

We define $L_V^2 = \left\{ (v_t)_{0 \leq t \leq 1} \in V^{[0,1]} \text{ s.t. } \sqrt{\int_0^1 |v_t|_V^2 dt} < \infty \right\}$ the set of time varying vector fields with finite energy with respect to the norm on V . L_V^2 is a Hilbert space with the correct properties, and we can define the group $G_V : \{\varphi_t^v, v \in L_V^2\}$ where φ_t^v verifies the flow equation:

$$\begin{cases} \frac{\partial \varphi_t^v}{\partial t} = v \circ \varphi_t^v \\ \varphi_0^v = id, \end{cases} \quad (2.2)$$

The construction of this group is well detailed in [You10], in particular the existence, uniqueness and control of the solutions to the flow equation in G_V . We simply provide here the structure of G_V :

Theorem 4 ([Tro95]). *G_V is a group of diffeomorphisms.*

We will now explain how a right-invariant metric on G_V can be defined from the previous construction. First we can notice that if such a right-invariant metric exists on G_V , we can remark that for $\phi, \phi_1, \phi_2 \in G_V$, $d(\phi_1 \circ \phi, \phi_2 \circ \phi) = d(\phi_1, \phi_2)$, and in particular taking $\phi = \phi_1^{-1}$ we obtain $d(\phi_1, \phi_2) = d(id, \phi_2 \circ \phi_1^{-1})$. This shows that we can restrict the definition of the distance in G_V to the distance to identity:

$$d(id, \phi) := \inf_{v \in L_V^2} \left\{ \sqrt{\int_0^1 |v_t|_V^2 dt} \mid \varphi_1^v = \phi \right\} \quad (2.3)$$

which can be extended to a right invariant distance on G_V via:

$$d_{G_V}(\phi, \phi') := d(id, \phi' \circ \phi^{-1}) = \inf_{v \in L_V^2} \left\{ \sqrt{\int_0^1 |v_t|_V^2 dt} \mid \phi' = \varphi_1^v \circ \phi \right\} \quad (2.4)$$

Theorem 5 ([Tro95]). *d_{G_V} is a metric on G_V and (G_V, d_{G_V}) is complete.*

The energy of the flow equation 2.2, given by the squared distance to identity, measures the amount of deformations induced by the vector field on the whole space.

Corollary 1. *The infimum in Equation 2.3 is attained at some $v \in L_V^2$, and moreover $|v_t|_V$ is constant.*

This result provides the following theorem, and is fundamental for the construction of computable diffeomorphic solutions to the flow equation.

Theorem 6. *If V is admissible and $\phi, \phi' \in G_V$, there exists a time varying vector field $v \in L_V^2$ such that*

$$d_{G_V}(\phi, \phi') = |v|_{L_V^2}$$

and $\phi' = \varphi_1^v \circ \phi$

We can now get back to the initial problem: how to compare and align two shapes S and T in the set of shapes M living in \mathbb{R}^d ? Since we consider that all shapes belong to the orbit $G_V.S$, we can express the distance between S and T through the distance of Theorem 6 we just built in G_V :

$$d(S, T)^2 := \inf_{v \in L_V^2} \left\{ \int_0^1 |v_t|_V^2 dt \mid T = \varphi_1^v(S) \right\} = \inf_{\phi \in G_V} \left\{ d_{G_V}(id, \phi)^2 \mid T = \phi(S) \right\} \quad (2.5)$$

This term is called the energy of deformation, and from Corollary 1 we know that given v in L_V^2 minimizing this distance, $d(S, T)^2 = E(v) = E(v_0) = |v_0|_V^2$. For a given solution to the **Exact Matching** problem, we are able to easily compute the energy of deformation of the solution. We now need to explicitly build such a solution.

2.3.2 Computing Large Diffeomorphic Deformation

We have seen that the space of diffeomorphisms can be built upon a space of admissible vector fields (Definition 8), and we provided in Proposition 3 a condition on the reproducing kernel to build such admissible RKHS. In a discrete setting in the case of exact matching of landmarks, one may prove (see for instance [You10], chapter 11) that optimal vector fields for the minimization of functional E can be written by the means of the Reproducing Kernel K_V of the admissible RKHS V , as follows:

$$v_t(x) = \sum_{i=1}^N K_V(x, q_t^i) p_t^i \quad (2.6)$$

where the $q_t = (q_t^i)_i$ correspond to the trajectories of the finite number of points called *control points* of \mathbb{R}^d under the flow equation, and the $p_t = (p_t^i)_i$ are the auxiliary dual variables called *momenta*.

This result puts the reproducing kernel as the central object to model discrete shapes deformations. Classically, one may chose an appropriate kernel K_V first and then deduce the corresponding Hilbert space V . From the construction of the RKHS described in Section 2.2, we simply need to build a kernel K_V with the correct properties: K_V is continuously differentiable of order 2, all the derivatives of K_V up to the order 2 are bounded and for any $x \in \mathbb{R}^d$, the function $K_V(x, \cdot)$ and all its derivatives up to order 2 vanish at infinity. Then, for such a kernel K_V , its corresponding RKHS is an admissible space of vector fields. In practice, we will use the radial basis Gaussian kernel (Example 2.2).

One can remark that the functions $K_V(\cdot, q_t^i) p_t^i$ correspond to the fundamental functions of V similarly to the Vector Spline Interpolation Example 1, but contrary to this previous example we now generate time varying vector fields. This result can be extended to the inexact matching framework for any shapes, provided some conditions on the data attachment term that depends only on the displacement of the landmarks, as detailed in Section 2.4.

If a framework has been studied for more generic shapes, we limit the discussion to the case of landmarks trajectories, which is the typical discrete setting. In fact, the discrete shapes are considered as sets of vertices connected or not. We then apply the deformation to the vertices of these discrete shapes, and report their connectivity at the new location of the deformed vertices. The trajectories are thus parameterized by the control points and initial momenta q_t^i and p_t^i . Using the flow equation of Equation 2.2, we are looking for solutions to the **Exact Matching** problem 2.5. In the case of landmarks, the solution exists.

Shooting Along the Geodesics As shown in [MTY06], one may further derive optimality equations that must be satisfied by the trajectories q_t^i when considering deformations that minimize the cost function. These equations take the following form:

$$\begin{cases} \dot{p}_t^i = -\dot{v}_t(q_t^i)^T \cdot p_t^i \\ \dot{q}_t^i = v_t(q_t^i) \end{cases} \quad (2.7)$$

Which can be written by the mean of the reproducing kernel:

$$\begin{cases} \dot{p}_t^i = -\frac{1}{2} \nabla_{q_t^i} \left(\sum_{j=1}^N \sum_{l=1}^N \langle p_t^j, K_V(q_t^j, q_t^l) p_t^l \rangle_{\mathbb{R}^d} \right) \\ \dot{q}_t^i = \sum_{j=1}^N K_V(q_t^i, q_t^j) p_t^j \end{cases} \quad (2.8)$$




Figure 2.3: Illustration of a LDDMM registration from an atlas shape S to a target T .

They correspond to geodesic equations with respect to a specific Riemannian metric induced therefore by the Reproducing Kernel, written in Hamiltonian form inspired from the study of the dynamic of particles systems. Thus, optimal trajectories correspond to geodesics and in particular, the total energy (and the reduced Hamiltonian) is conserved over time :

$$|v_t|_V^2 = |v_0|_V^2 = \sum_{i=1}^N \sum_{j=1}^N \langle p_0^i, K_V(q_0^i, q_0^j) p_0^j \rangle = p_0^T \cdot K(q_0, q_0) p_0 = H_r(p_0, q_0) \quad (2.9)$$

The points q_0^i can be taken for instance as the discretization points of the source shape S , the deformations are then fully parametrized by the initial momenta p_0^i . More precisely, given any vectors p_0^i , one may retrieve the full corresponding time-dependent momenta p_t^i and trajectories q_t^i by integrating the geodesic equations (2.7) (this process is called **shooting**), and from this retrieve the full deformation $\phi = \varphi_1^v$ by integrating the flow equation (2.1).

Induced Metric on Landmarks When looking at the infinitesimal deformation $\eta = (\eta^i)$ of $q = (q^i)_i$, if v is solution to the interpolation problem $v(q^i) = \eta^i$ in V , with minimal $|v|_V$. Then we have that the local metric satisfies $|\eta|^2 := |v|_V^2$, which gives $|\eta|^2 = \langle K_V(q, q)^{-1} \eta, \eta \rangle$, with $K_V(q, q)$ the block matrix of the $K_V(q^i, q^j)$. We thus have the Riemannian manifold of landmarks, with metric tensor at q equal to $K_V(q, q)^{-1}$, and we can build the geodesic path and distance – the geodesic length. We retrieve the global metric on the shapes defined in Equation 2.5.

In the case of more generic shapes than the landmarks, the construction of the riemannian structure on M would require a formal construction at this stage, according to the shapes.

2.4 Metric Mapping: from Exact to Inexact Registration

The diffeomorphic alignment of the shapes presented so far is based on the assumption that differences between similar anatomical shapes may be analysed via the estimation of diffeomorphisms of the ambient space that act upon them. We have defined a space of admissible deformations and a way to compute them, however the assumption of studying populations of shapes belonging to the same orbit is too restrictive. In real world applications one can – almost – never expect a perfect mapping of one shape onto the other: there can be different samplings, noise, artifacts, missing structures across the population, etc. In practice, the registration algorithms search for the shape alignment with respect to a given metric. Such methods make the registrations robust to the real-world challenges listed, but deeply rely on the choice of the metric – or *data attachment* – and the representation of the shapes. Once again we will build metrics induced by reproducing kernels that will provide efficient formulas for computing the metrics. Thanks to their flexibility, we will also encode properties in the induced metric, for instance taking into account the orientation of the shapes or not.

In this setting, the registration illustrated in Figure 2.3 of a *source* shape S (in our case the atlas) onto a *target* shape T , is performed by minimizing a cost function:

$$J(\phi) = \gamma E(\phi) + A(\phi(S), T) \quad (2.10)$$

where A is a data attachment term that penalizes mismatch between the deformed source $\phi(S)$ and the target T . In the case of non-diffeomorphic deformations (Example 1), introducing the data attachment terms provided regularized vector splines solutions. A similar methodology can be followed for diffeomorphic framework: it will provide solutions when registering shapes that do not belong to the same orbit. Using the following theorem, one can show that the inexact registration framework can be adapted to the diffeomorphic registration framework by building data attachment terms with a certain regularity:

Theorem 7 ([Gla05] Theorem 7). *Let S, T be two shapes of Ω , and consider a matching problem of the form 2.10. If $v \mapsto A(\varphi_1^v(S), T)$ is weakly continuous from L_V^2 to \mathbb{R} ; then there always exists for $\gamma > 0$ a minimum to the matching problem.*

The result 2.9 associated with a differentiable data attachment term verifying Theorem 7 allows building an explicit forward-backward algorithm for registration called *geodesic shooting* (Algorithm 1, illustrated in Figure 2.3) introduced in [MTY06]: with $q_0 = (x_1, \dots, x_N)$

Algorithm 1 Geodesic shooting (fixed-step gradient descent).

Input: q_0, δ (step size)

Output: $\argmin_{p_0 \in (\mathbb{R}^d)^n} J(p_0)$

Set $p_0 = 0$

repeat

 Compute $(q_1, p_1) = ((q_1^i)_i, (p_1^i)_i)$ through forward integration of (2.7)

 Compute $\nabla_{q_1} A(q_1, T)$

 Compute $\nabla_{q_0} A(q_1, T)$ through backward integration

 Compute $\nabla_{p_0} J(p_0) = K_V(q_0, q_0)p_0 + \nabla_{q_0} A(q_1, T)$

$p_0 \leftarrow p_0 - \delta \nabla_{p_0} J(p_0)$

until convergence

the discretization points of the source shape to deform S and $q_1 = (q_1^i)_i$ defined as $q_1^i = \varphi_1^v(q_0^i)$. Similarly to the construction of computable spaces of diffeomorphisms, the metrics we seek for shapes comparison will also be derived from RKHS.

Example 3: Labeled Landmark Diffeomorphic Registration

In the case of the annotated landmark registration, one is given two sets of paired source $x = (x_i)_i$ and target $y = (y_i)_i$ points. One simple data attachment term is to take the sum of the squared point-wise L_2 distance :

$$A(x, y) = \sum_{i=1}^N |\varphi_1^v(x_i) - y_i|_2^2$$

Let (v_n) be a sequence in L_V^2 , weakly converging to some $v \in L_V^2$. We need to show that $A(\varphi_1^{v_n}(x), y) \rightarrow A(\varphi_1^v(x), y)$ which is true if $\varphi_1^{v_n}(x_i) \rightarrow \varphi_1^v(x_i)$. This last condition is true by definition, the Theorem 7 is verified and by setting the control points q_0^i to the initial positions x_i , the minimization function becomes

$$J(p_0) = \gamma p_0^T K(q_0, q_0) p_0 + A(\varphi_1^v(q_0), y)$$

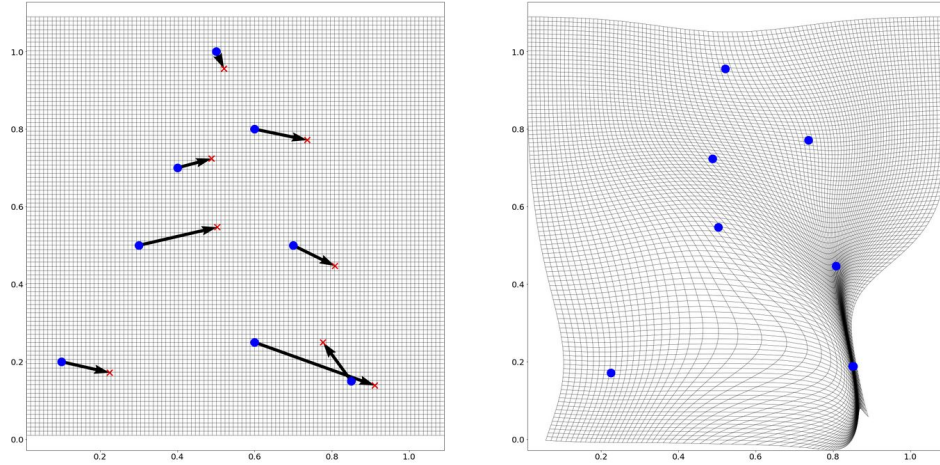


Figure 2.4: Exact landmark matching using diffeomorphic deformations with explicit correspondences. The data attachment is the point wise L_2 distance between the paired points, with $\gamma = 0$ and the scale of the Gaussian kernel $s = 0.25$. Left: the source points (blue points) and the targets (red cross) and their correspondences. Right: the output deformation applied to the ambient space.

When compared to the Vector Spline Interpolation 1, we immediately see that the deformation in the previous Example 3 does not fold the ambient space thanks to the diffeomorphic deformation *despite the parameter γ equal to 0*. The matching in this case is exact. Furthermore, the deformation is less smooth than the regularized Vector Spline Interpolation of Figure 2.2, which can be corrected by using $\gamma > 0$ or playing with the scale of the reproducing kernel. We observe in Figure 2.5 (a) that increasing γ increases the regularity of the deformations, and that the result approaches the regularized Vector Spline Interpolation, with the guarantee of a diffeomorphic deformation in addition. In Figure 2.5 (b) we have kept the same γ but we have decreased the scale of the reproducing kernel which has the




Figure 2.5: Influence of the regularization parameters on the inexact diffeomorphic registration of the landmarks. (a): $\gamma = 1$ and $s = 0.25$ and (b): $\gamma = 1$ and $s = 0.1$.

effect of decreasing the range of influence of each control point x^i on its surrounding space. We then find an almost perfect registration although γ is not zero, and the deformations are much more localized. These experiments show that the geodesic shooting method provides a particularly flexible smooth non-rigid registration method: by playing with two parameters, without even discussing the choice of the reproducing kernel one is able to generate more or less regular and local deformations. This also illustrates the importance of these parameters, which must be adjusted for each new application to find the right trade-off between the amount of deformation generated and the alignment of the shapes.

This example of annotated points gives a first way to non-rigidly align any shapes by manually discretizing them with correspondence points. This is what was done in the first applications of computational anatomy for matching landmarks [JM00]. However, as we discussed previously such a framework is not suited to automatic tools for processing the data: either it requires human time that the users, like the physicians, can not afford for instance during their procedures, or it introduces a variability inherent to the complexity of the task and the human mistakes. This is why many authors have been interested in the construction of metrics in spaces adapted to their data in order to dispense with annotations to generate deformations. The construction of such data fidelity terms deeply depends on the type of data to register: the data attachment between images may be very different from the one between curves in \mathbb{R}^d . In this manuscript we will focus on the curves of \mathbb{R}^2 and \mathbb{R}^3 and surfaces of \mathbb{R}^3 and their unions, which are the most common shape representation of the organs and vascular trees.

2.4.1 Spaces for the Data Attachment

We now introduce the data attachment terms (A in (2.10)) usually used in the diffeomorphic registration algorithms. These similarity terms should depend on the shape features one wants to consider in the registration. In the case of sub-manifolds of \mathbb{R}^d , the invariance to some reparameterization makes the distances used for the case of labeled landmarks intractable. In addition, the computing tools work with *discrete* shapes and there exist no a priori on the number of discretization points per mesh in the population of shapes, on the structure of the meshing (density of vertices for instance) or on the way it was constructed

(a population of shapes may be a concatenation of several databases obtained from different subjects). The objective is hence to build metrics robust to artifacts and noise one can get when constructing a database and adapted to the comparison of submanifolds of \mathbb{R}^d , in our case curves and surfaces. Finally, such metrics should be efficiently computable and scalable to large data, and $\phi \mapsto A(\phi(S), T)$ should be differentiable on any group of diffeomorphisms G_V .

A convenient framework inspired from the field of Geometric Measure theory [Fed69] is to compare shapes in certain spaces of generalized measures: *Currents* [Gla05], *Varifolds* [CT13] or *Normal Cycles* [RG16]. Such methods are reviewed in [Cha+20], and we provide in this section the definitions as well as some toy examples illustrating the behaviour of the different metrics and how they drive the diffeomorphisms. This corresponds to the LDDMM framework introduced in Section 2.1 based on the geodesic shooting algorithm described in Algorithm 1. The key idea is then to build some appropriate RKHS structure on spaces of measures that allow to derive computable norms. Moreover these methods have the advantage of being independent of any correspondence between the points, and they are adapted to both continuous and discrete settings.

From Labeled Points to Points Clouds

As a first example, one can adapt the problem of matching labeled points to the one of matching unlabeled points clouds by seeing them as sum of weighted Dirac in the ambient space. This setting, corresponding to distributions of points \mathbb{R}^d does not require any points correspondences or assumptions on the number of points per distribution. We are then considering elements of $\mathcal{M}_s(\mathbb{R}^d)$ the space of signed Borel measures in \mathbb{R}^d which is the dual space of $C_0(\mathbb{R}^d, \mathbb{R})$ with the total variation norm:

$$|\mu|_{TV} = \sup \left\{ \int_{\mathbb{R}^d} f d\mu, f \in C_0(\mathbb{R}^d, \mathbb{R}), |f|_\infty \leq 1 \right\}$$

We can then define the action of a transformation $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ on a test function $f \in C_0(\mathbb{R}^d, \mathbb{R})$ as:

$$\phi.f := f \circ \phi^{-1}.$$

Then the action of the transformation on a measure μ is defined by:

$$\int_{\mathbb{R}^d} f d(\phi \cdot \mu) := \int_{\mathbb{R}^d} \phi^{-1} \cdot f d\mu = \int_{\mathbb{R}^d} f \circ \phi d\mu,$$

To define a data attachment term between measures, $\mathcal{M}_s(\mathbb{R}^d)$ is continuously embedded in the dual I' of a RKHS I (See [Gla05].4.2.3). Using the reproducing property we can therefore retrieve the norm in I' with the reproducing kernel k_I . Remarking that a measure $\mu \in \mathcal{M}_s(\mathbb{R}^d)$ can be approached by a sum of Dirac masses $\tilde{\mu} = \sum_{i=1}^N a_i \delta_{x_i}$, $a_i \in \mathbb{R}$, $x_i \in \mathbb{R}^d$, the norm writes:

$$|\tilde{\mu}|_{I'}^2 = \sum_{i,j=1}^N a_i a_j k_I(x_i, x_j).$$

Proposition 6. *Supposing that I is continuously embedded in $C_0^1(\mathbb{R}^d, \mathbb{R})$ and $\mu_S, \mu_T \in \mathcal{M}_s(\mathbb{R}^d)$ with μ_S on a compact set. Then the application $v \mapsto |\varphi_1^v \cdot \mu_S - \mu_T|_{I'}^2$ is weakly continuous on L_V^2 .*

Using the Theorem 7 and the last proposition, we have a distance in the space of measures adapted to the LDDMM framework and computable thanks to the RKHS structure:

$$d(\mu_1, \mu_2) = |\mu_1 - \mu_2|_{I'}^2.$$

Provided some regularity assumption on the reproducing kernel, we have the existence and the consistency of the solutions to the minimization problem of the function

$$J(v) = \gamma \cdot E(v) + |\varphi_1^v \cdot \mu_S - \mu_T|_{I'}^2$$




Figure 2.6: Diffeomorphic registration of a point cloud seen as a sum of Diracs, with $\gamma = 1$ and $s = 0.25$ and for the reproducing kernel $k_I: s_I = 0.2$.

As illustrated in Figure 2.6 since the points are not labeled anymore, the deformation leads to a new solution and the overall solution now matches the points with a smoother diffeomorphism. The weight of each Dirac were set to 1 in this experiment considering that there is the same number of points in the source and in the target. This data representation is though adapted to more general settings: points clouds and shapes with different numbers of points and weights.

This first application shows that the space of measures seems to be adapted to the construction of metrics to compare shapes. However, the action of deformations on this metric is not geometric, so we want to use more adapted tools. In particular, one would like to take into account the tangent information inherent to the submanifold structure of the shapes. The first precursor work on the construction of a metric allowing the geometric action of diffeomorphisms is Glaunès [Gla05] who proposes to work on currents instead of measures, or in other words measures on the space of m-differential forms for submanifolds of dimension

$m \leq d$. It can be shown that the measures are in fact a particular case of the current representation. Later on, Charon [CT13] focused on the *Varifold* representation (and the framework of the functional shapes, which is out of the scope of this thesis) getting rid of the orientation of the shapes and paying more attention to thin structures while being more sensitive to the noise. The varifolds are distributions on the product $\mathbb{R}^d \times G_m(\mathbb{R}^d)$, or in other words continuous linear forms on a space of smooth test functions from $\mathbb{R}^d \times G_m(\mathbb{R}^d)$ to \mathbb{R} , where $G_m(\mathbb{R}^d)$ is the m -dimensional Grassmannian in \mathbb{R}^d . This second embedding space for the shapes introduced a representation sensitive to the thin structures which were not seen in the current representation. In a more recent work, [KCC17] proposed a slightly different version of the varifolds allowing to take into account the orientation. In this setting, the varifolds and the currents are particular cases of the oriented varifolds, and we will use this latter space to build our metrics.

Shapes Representation

We limit the notions and definitions to the essential ones and focus on the case of matching curves and surfaces, and the reader may refer to the works of Glaunès, Charon, and Roussillon [Gla05; CT13; RG16], introducing the general setting for each shape space representation of the submanifolds. The ambient space is \mathbb{R}^d with $d = 2, 3$ and since we are working with smooth curves and surfaces (or union of them) of \mathbb{R}^d , we introduce the tangent space to the shape S at x denoted $T_x S$. We also define the *orientation* of a shape by using oriented tangent spaces, represented by vectors in the unit sphere \mathbb{S}^{d-1} and written $\vec{t}(x)$ the unit tangent vector for the curves and $\vec{n}(x)$ the unit normal vector for the surfaces.

Note that apart from the dimension-1 and codimension-1 manifolds (curves and hypersurfaces of \mathbb{R}^d) one can not represent the tangent space with a unit tangent vector. One can chose the opposite orientation of a shape by using $-\vec{t}(x)$ instead of $\vec{t}(x)$. To properly define the shapes we are dealing with, we need to define the rectifiable subsets:

Definition 9. *We say that S is a k -rectifiable subset of \mathbb{R}^d if the k -dimensional Hausdorff measure of S is finite, and if there exists a countable family $(f_i)_{i \in \mathbb{N}}$ of Lipschitz maps $f_i : \mathbb{R}^k \rightarrow \mathbb{R}^d$ such that*

$$\mathcal{H}^k(S \setminus \cup_i f_i(\mathbb{R}^k)) = 0.$$

In particular, one property of the rectifiable curves and surfaces is that for almost every point in the $x \in S$, the tangent space $T_x S$ exists. In addition, we work with oriented rectifiable sets. In the particular case of vascular trees, it is indeed natural to define the orientation of the centerlines from the root to the leaves.

From Curves and Surfaces to Meshes Since we manipulate discrete shapes represented as meshes, we must define their structure.

The curves, illustrated in Figure 2.7(a), are represented as polygonal curves in \mathbb{R}^d , letting $\{x_i\}_{i=1 \dots N} \in \mathbb{R}^d$ be the discretization points, with x_i connected to x_{i+1} for $i \in \{1, \dots, N - 1\}$ we can compute the centers (c_i) of each discretization segment $[x_i, x_{i+1}]$ along with its

associated weight (the segments length) a_i and the tangent space to the center $T_{c_i}S$ spanned by the unit tangent vector $\vec{t}_S(c_i)$:

$$\begin{cases} c_i &= \frac{x_i + x_{i+1}}{2} \\ a_i &= |x_{i+1} - x_i|_{\mathbb{R}^d} \\ T_{c_i}S &= \text{span} \left(\frac{x_{i+1} - x_i}{|x_{i+1} - x_i|_{\mathbb{R}^d}} \right) = \text{span}(\vec{t}_S(c_i)) \end{cases} \quad (2.11)$$

The surfaces are represented as triangulated meshes. Let \wedge denote the cross product in \mathbb{R}^3 . Let S be a triangulated surface in $E = \mathbb{R}^3$ with N points $\{x_k\}_{k=1\dots N}$ and n triangles $(f_i^1, f_i^2, f_i^3) \in \{1, \dots, N\}^3$ for $i = 1, \dots, n$. As illustrated in Figure 2.7(b), we can compute the centers of each cell $c_i \in \mathbb{R}^3$, the (unoriented) tangent space to the shape at c_i , $T_{c_i}S \in \mathbb{S}^2$ and the surface area of each cell $a_i \in \mathbb{R}_+^*$:

$$\begin{cases} c_i &= \frac{x_{f_i^1} + x_{f_i^2} + x_{f_i^3}}{3} \\ a_i &= \frac{1}{2} \left| x_{f_i^2} - x_{f_i^1} \wedge x_{f_i^3} - x_{f_i^1} \right| \\ T_{c_i}S &= \text{span} \left(x_{f_i^1} - x_{f_i^2}, x_{f_i^1} - x_{f_i^3} \right) \end{cases} \quad (2.12)$$




Figure 2.7: Meshes structure for (a) curves and (b) surfaces.

2.4.2 Matching Oriented Varifolds

We focus now on the oriented varifolds representation introduced in [KCC17], of which classical representations such as the varifolds or the currents are particular cases. In the different examples and applications of the thesis, the oriented varifolds are the default selected shape space. If another representation is chosen, this will be specified in the description of the application. Let S be such a 1 or 2-rectifiable subset, we associate with S an oriented varifold:

$$\mu_S(\omega) = \int_{\mathbb{R}^d \times \mathbb{S}^{d-1}} \omega(x, \tau_x S) d\mu_S(x, \tau_x S) \quad (2.13)$$

$$= \int_S \omega(x, \tau_x S) d\text{vol}(x) \quad (2.14)$$

with ω a smooth test function $\mathbb{R}^d \times \mathbb{S}^{d-1} \mapsto \mathbb{R}$. In particular, the Diracs $\delta_{(x,\tau)}$ with $x \in \mathbb{R}^d$ and $\tau \in \mathbb{S}^{d-1}$ in the space of oriented varifolds act on functions of $C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$ by the relation:

$$\delta_{(x,\tau)}(\omega) = \omega(x, \tau), \quad \forall \omega \in C_0(\mathbb{R}^d \times \mathbb{S}^{d-1}).$$

If S is a smooth compact submanifold, given a parametrization $\gamma : U \mapsto \mathbb{R}^d$ with U an open subset of \mathbb{R}^n , with $n = 1$ or 2 the dimension of S , then for all $\omega \in C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$, we have:

$$\mu_S(\omega) = \int_U \omega(\gamma(u), \tau_{\gamma(u)} S) |\gamma'(u)| du$$

with $\gamma'(u) = \Lambda_{i=1}^d \frac{\partial \gamma}{\partial u_i} \in \Lambda^d(\mathbb{R}^d)$. The intuition for $|\gamma'(u)|$ is that it corresponds to the element of volume of the parallelepiped supported by the partial derivatives. This last equation shows in particular that the varifold associated with a shape S is independent to any reparametrization of S .

In a discrete setting defined in Section 2.4.1, the shapes can be associated with sums of Diracs and coded by the varifolds: $\mu_S = \sum_{i=1}^n c_i \cdot \delta_{(c_i, \tau_{c_i} S)}$. (see [KCC17] Proposition 1 for the proof of the discrete approximation).

Action of Diffeomorphism Similarly to the transport of measures in the case of point clouds, one can build a transport of varifolds through the action of diffeomorphisms in G_V .

We have thus constructed objects that can be transported by diffeomorphisms and associated with the shapes we are interested in, curves and surfaces, and their finite unions.

Kernel Metrics for Oriented Varifolds

In order to focus on a specific class of computable test functions to build our varifolds, we take advantage of the RKHS theory once more. Let W be a Reproducing Kernel Hilbert Space (RKHS) of functions defined over $\mathbb{R}^d \times \mathbb{S}^{d-1}$ continuously embedded in $C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$, the space of continuous test functions vanishing at infinity. Its dual space W' is a space of varifolds and the following proposition gives a practical way to define such a space:

Proposition 7. *Assume that we are given a positive-definite real kernel k_e on the space \mathbb{R}^d such that k_e is continuous, bounded and for all $x \in \mathbb{R}^d$, the function $k_e(x, \cdot)$ vanishes at infinity. Assume that a second positive-definite real kernel k_t is defined on the manifold \mathbb{S}^{d-1} and is also continuous. Then the RKHS W associated with the positive-definite kernel $k_e \otimes k_t$ is continuously embedded into the space $C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$.*

In the following we assume that the reproducing kernel of W is of the form $k_e \otimes k_t$, with the assumptions of proposition 7. The Diracs belong to the dual W' of the RKHS, and the dual metric satisfies:

$$\langle \delta_{x_1, \tau_1}, \delta_{x_2, \tau_2} \rangle_{W'} = k_W((x_1, \tau_1), (x_2, \tau_2)) \quad (2.15)$$

$$= k_e(x_1, x_2) k_t(\tau_1, \tau_2) \quad (2.16)$$

We associate with shape S the canonical function $\omega_S \in W$ defined for all $y \in \mathbb{R}^d$ and $\tau \in \mathbb{S}^{d-1}$ as follows:

$$\omega_S(y, \tau) = \int_S k_e(y, x) k_t(\tau, \tau_x S) d\text{vol}(x).$$

This function corresponds to the unique representer of the varifold $\mu_S \in W'$ via the Riesz representation theorem. Similarly, we define the canonical function ω_T associated with shape T . Via this representation of shapes, one may express the scalar product between the varifolds μ_S, μ_T , or equivalently between the canonical functions ω_S, ω_T as follows:

$$\langle \mu_S, \mu_T \rangle_{W'} = \langle \omega_S, \omega_T \rangle_W = \int_S \int_T k_e(x, y) k_t(\tau_x S, \tau_y T) dx dy.$$

Finally, the shape matching distance as defined in [CT13] is the following:

$$\begin{aligned} d_{W'}(S, T)^2 &= \|\omega_S - \omega_T\|_W^2 = \|\mu_S - \mu_T\|_{W'}^2 \\ &= \langle \mu_S, \mu_S \rangle_{W'} - 2\langle \mu_S, \mu_T \rangle_{W'} + \langle \mu_T, \mu_T \rangle_{W'}. \end{aligned}$$

While this separable structure does not cover all possible kernels on $\mathbb{R}^d \times \mathbb{S}^{d-1}$, it provides a large and interpretable set of kernels, independently defining the behaviour of the position (encoded by k_e) and the orientation (encoded by k_t). Similarly to the discussion of Proposition 4, a property we seek for the metric is the equivariance to rigid motion.

Proposition 8. *Within the class of separable kernels defined in Proposition 7, the metric W' is equivariant to the action of rigid transformations for kernels $k_W((x_1, \tau_1), (x_2, \tau_2)) = k_e(|x_2 - x_1|) k_t(\langle \tau_1, \tau_2 \rangle_{\mathbb{R}^d})$.*

We have seen admissible conditions for such a radial kernel k_e in Section 2.2, from [KCC17], a necessary and sufficient condition on k_t is that $k_t(u) = \sum_1^\infty a_k P_k^{(\lambda)}(u)$ with $a_k \geq 0$, $\sum_1^\infty a_k P_k^{(\lambda)}(1) < \infty$ and $P_k^{(\lambda)}$ the ultraspherical polynomials of order $\lambda = \frac{d-1}{2}$. In order to induce a metric on W' instead of a pseudometric, a sufficient condition can be added:

Proposition 9. *Let k_e and k_t be two kernels as in Proposition 7, if in addition W is dense in $C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$ and for all $\tau \in \mathbb{S}^{d-1}$ $k_t(\tau, \tau) > 0$, then $d_{W'}(S, T) = |\mu_T - \mu_S|_{W'}$ defines a distance on the set of shapes.*

The Different Shape Distances as Functions of k_t Until now we have built reproducing kernels with the right properties for the induced metric to define a distance in W' , and the generic definitions proposed to provide freedom of modelling for the oriented varifolds. In particular, depending on the choice of the reproducing kernel for the tangent space k_t , one may retrieve the particular cases of the *currents* and (unoriented) *varifolds*:

- $k_t(\tau_1, \tau_2) = \langle \tau_1, \tau_2 \rangle_{\mathbb{R}^d}$, the **current** representation proposed in [Gla05].

- $k_t(\tau_1, \tau_2) = \langle \tau_1, \tau_2 \rangle_{\mathbb{R}^d}^2$, the simplest **unoriented varifold** representation proposed in [CT13].

- $k_t(\tau_1, \tau_2) = e^{-\frac{2\langle \tau_1, \tau_2 \rangle_{\mathbb{R}^d}}{\sigma_t^2}}$ the restriction of the gaussian kernel of \mathbb{R}^d to \mathbb{S}^{d-1} , which is fast to compute and one can control the angular sensitivity through the scale parameter σ_t .

A further discussion about the choice of the different metrics as data attachment terms is provided in Chapter 3. It can be noted here that the currents allow us to have an orientation of the shapes, which is particularly interesting in the case of vascular trees, for example, which have a natural orientation from the root to the leaves. However, their lack of sensitivity to structures of opposite orientation and spatially close to each other creates disappearances or inappropriate appearances of high frequency structures. The unoriented varifolds on the contrary will be very sensitive to such structures and thus they will preserve them. However, they do not encode the orientation and therefore the natural choice that can be made to compare the shapes of centerline trees is that of oriented varifolds.

2.4.3 Optimal Transport Cost as Data Attachment Term

A second kind of data attachment term is used in Chapter 4, based on the optimal transport cost that was first used in [Fey+17] as data attachment term in the LDDMM framework.

This new type of data fidelity term designed as a transport between measures in a space of representation for our data. In the discrete framework, these measures are expressed as sum of weighted Diracs located in $\mathbb{R}^{n \times d}$ with $d = 2$ or 3 in our case:

$$\alpha = \sum_{i \in I} a_i \delta_{x_i}, \text{ and } \beta = \sum_{j \in J} b_j \delta_{y_j}.$$

with $x_i, y_j \in \mathbb{R}^{n \times d}$ the discretization points of d -dimensional curves, and $\sum_i a_i = 1$, $\sum_j b_j = 1$. The vector formulation of $(a_i)_i = \mathbf{a}$ and $(b_i)_i = \mathbf{b}$ then denotes the probability vectors associated with each measure.

First formalized by Monge in 1781 as the earth mover's problem, the original idea was to find a minimum cost transport to move a distribution to another. Following this idea the Wasserstein distance was defined by Kantorovitch [Kan60]. The key idea was to adapt the Monge problem, which could be seen as an assignment problem namely a point x_i can only be assigned to a target point $y_{\sigma(i)}$ yielding the necessity of finding correspondences between the same number of source and target points, with the same weights w . Kantorovitch introduces the notion of mass splitting, by introducing a coupling matrix $P = (P_{ij})_{i,j} \in \mathbb{R}_+^{n \times m}$ denoting the cost of a transport (the amount of mass flowing) from x_i to y_j .

Admissible couplings then read:

$$U(\alpha, \beta) = \left\{ P \in \mathbb{R}_+^{n \times m}, C \cdot \mathbf{1}_n = \left(\sum_j P_{ij} \right)_i (a_i)_i, P^T \cdot \mathbf{1}_m = \left(\sum_i P_{ij} \right)_j (b_j)_j \right\}. \quad (2.17)$$

Denoting $C = (d(x_i, y_j))_{i,j}$ the cost matrix determined by the pairwise distance between the **unit** elements of S and T , we have that a coupling is always symmetric since if $P \in U(\alpha, \beta)$, then $P^T \in U(\beta, \alpha)$. Kantorovitch OT is given by:

$$L_C(\mathbf{a}, \mathbf{b}) = \min_{P \in U(\alpha, \beta)} \sum_{i,j} C_{ij} P_{ij}.$$

This is by definition the OT between discrete measures as well as the one between their probability weights. Note that the OT can be formulated through numerous interpretations such as probabilistic viewpoint or generalized sorting, and applied to more general arbitrary measures. In this section we focus on the discrete measures that are our typical working case.

Dual Problem and Sinkhorn

The set of matrices $U(\mathbf{a}, \mathbf{b})$ in the problem 2.17 is bounded and defined by $n + m$ equality constraints which makes $U(\mathbf{a}, \mathbf{b})$ the convex hull of a finite set of matrices. The problem 2.17 is convex and admits a dual problem:

$$L_C(\mathbf{a}, \mathbf{b}) = \max_{f, g \in R(C)} \langle f, \mathbf{a} \rangle + \langle g, \mathbf{b} \rangle$$

with $R(C) = \{f \in \mathbb{R}^n, g \in \mathbb{R}^m \text{ s.t. } \forall (i, j) \in I \times J, f_i + g_j \leq C_{ij}\}$

The dual variables f, g are called Kantorovitch Potential, and provide a location of the OT solution: $P \in \mathbb{R}_+^{n \times m} \subset \{(i, j) \in I \times J, f_i + g_j = C_{ij}\}$ The solutions to the Kantorovitch problem is classically solved by regularizing the transport cost [Cut13] using the entropic regularization:

$$H(P) = - \sum_{i,j} P_{ij} \log(P_{ij} - 1).$$

This entropic regularization provides a way to minimize the cost by iteratively optimizing on the dual vectors \mathbf{f} and \mathbf{g} . In addition, their update can be formulated as matrix vector products which makes it suited to GPU implementation. Regularizing the exact optimal transport with entropic regularization speeds up the computation of OT solutions regardless the metric chosen to compare the data.

The new minimization problem then writes:

$$L_C^\epsilon(\mathbf{a}, \mathbf{b}) = \min_{P \in U(\alpha, \beta)} \sum_{i,j} C_{ij} P_{ij} - \epsilon H(P).$$

Remarkably for our framework, this problem is convex and has a unique minima. In addition, the solution of the regularized cost function is smooth with respect to the weights and positions of the Diracs, and it can be differentiated [CD14] hence used as data fidelity term [MMC16] to compare shapes represented as sums of weighted Diracs.

In particular the Sinkhorn's iterations on the dual variables are given by:

$$f_i^{k+1} = \min_\epsilon (C_{ij} - g_j^k)_j + \epsilon \log(a_i) \quad (2.18)$$

$$g_j^{k+1} = \min_\epsilon (C_{ij} - f_i^k)_i + \epsilon \log(b_j) \quad (2.19)$$

With \min_ϵ a differentiable approximation of the function *minimum*. This procedure allows to freeze one dual variable to update the other and conversely.

Unbalanced Optimal Transport

The framework presented so far suppose that the two distributions have the same total mass. In [LMS18; Chi+18a] an extension of OT to unbalanced problems between measures of different total mass is proposed. A new variational metric is introduced allowing for interpolation between the optimal transport metric and the Fisher Rao one. This distance is a spatially localized version of the OT in the sense that masses too far won't interact with each other. It provides a geodesic interpolation introducing spatial displacement as well as change of mass formulated as a convex problem.

Similarly to the case of balanced OT, the unbalanced one can be efficiently computed via entropic regularization [Chi+18b]. The generalization of the sinkhorn algorithm to point wise weighted data provides a smooth differentiable and tractable distance function in the case of unbalanced OT as well. The solution then writes:

$$L_C^{\epsilon, \rho}(\mathbf{a}, \mathbf{b}) = \min_{P \in U(\alpha, \beta)} \sum_{i,j} C_{ij} P_{ij} - \epsilon H(P) + \rho \mathbf{KL}(P \cdot \mathbf{1}_n | a) + \rho \mathbf{KL}(P^T \cdot \mathbf{1}_m | b), \quad (2.20)$$

with $\mathbf{KL}(h, p) = \sum_i h_i \log(\frac{h_i}{p_i}) - h_i + p_i$. This unbalanced formulation of the OT can be solved using a generalized version of the Sinkhorn algorithm (see [Fey+17]) and given $\epsilon > 0$ the derivatives of the function $(a, x) \rightarrow$ are given by:

$$\begin{aligned} \nabla_{\mathbf{a}} L_C^{\epsilon, \rho}(\mathbf{a}, \mathbf{b}) &= \rho \left(1 - e^{-\frac{u}{\rho}} \right) \\ \nabla_x L_C^{\epsilon, \rho}(\mathbf{a}, \mathbf{b}) &= \left(\sum_j P_{ij} \partial_1 C_{ij} \right)_i \end{aligned} \quad (2.21)$$

with P the solution of 2.20 and u the limit of the first dual variable associated with \mathbf{a} in the generalized Sinkhorn algorithm.

Normal Cycles There are many other metrics used in shape registration frameworks that are not described here. One of interest kernel metric on normal cycles on sub-manifold structures of \mathbb{R}^d , and introduced in [RG16]. The specificity of this representation is to better take singular points like bifurcations or boundaries into account. The basic idea is to consider currents on the unit normal bundle to the shape which encodes order-2 information such as the curvature. In particular, in \mathbb{R}^3 , the unit normal bundle of curves and surfaces is a surface, and the metric for these shapes is similar. More details can be found in Annex 5.6, in which we provide the basic information for curves and surfaces in \mathbb{R}^3 .

2.5 Statistics over Deformations

One of the objectives of the Computational Anatomy that has not yet been discussed is the statistical study of populations. It has been shown in many papers [Gla05; Pen06; Fer+13; Cur+18] that the statistical analysis of the deformation between the observations of a population provides anatomical information. It allows for instance to compute principal components or outlier detection regarding the dataset. In this setting one can build an admissible space of shapes – e.g. the orbit of one or several templates for the group action G_V – with respect to the deformations generated to align the observations. This space is called the **atlas** in which a representative shape called **template** can be drawn. Most of the time, the selected template is a sort of "center of mass" of the distribution.

In general, statistics can be done on data sets that do not live in a classical vector space but on a Riemannian manifold of this vector space. When we can assign to this manifold a Riemannian metric (M, g) we are able to deduce the associated local volume: $dVol(x) = \det(g(x))$ which allows us to define the integrals on the manifold (i.e. measures), and so some class of probability distributions can be written with respect to that measure.

For LDDMM, when working with sets of points in the ambient space \mathbb{R}^d , one can consider that the associated metric is given by the energy of deformation of the diffeomorphism aligning points together:

$$d(S_i, S_j) = d(Id, \phi),$$

with $\phi(S_i) = S_j$. This is the exact registration framework, one can also work with the inexact registration by using the diffeomorphism $\varphi = \arg \min_{\phi \in G_V} E(v) + A(\phi(S_i), S_j)$.

When studying deformations between shapes in a dataset, one way to reduce the computational cost and avoid computing $\frac{N(N-1)}{2}$ functions minimization, one can consider the registration of a template S onto the two shapes S_i and S_j and then compare the deformations obtained. This reduces the number of registrations to N .

2.5.1 Building a Template of a Population

The template should be an element of the atlas representative of the database. In general, the euclidean barycenter of a finite set of points lying on a submanifold of a vectorial space does not belong to this submanifold. Take for instance points $x = (x_1, \dots, x_n)$, $n \geq 2$ living on the hypersphere \mathbb{S}^{d-1} . To adapt the notion of barycenter to Riemannian manifolds, the Fréchet mean was introduced:

Definition 10 (Fréchet Mean). *Let (M, g) be a Riemannian manifold with the Riemannian distance function d and let x_1, \dots, x_n be an i.i.d data set. The sample Fréchet mean is defined as the set of minimizers of the sum of squared distances:*

$$\bar{x} = \arg \min_{y \in M} \sum_{i=1}^n d(y, x_i)^2.$$

Of course, this definition depends on the Riemannian distance function chosen. The existence and uniqueness of the solution also depend on the distance, and we have one condition of the existence for a finite set of points:

Theorem 8 (Pennec, Sommer et al. 2020, Chapter 2 [PSF20]). *If (M, g) is a complete Riemannian metric space then the Fréchet mean of any finite set of points (x_1, \dots, x_n) exists.*

It follows that given a finite set of elements on a complete Riemannian manifold for a metric one can construct a Fréchet mean but this construction depends greatly on the choice of data representation. As discussed in [Dur10], one has a metric in a chosen space to compute the data attachment term, and could construct statistics in such a space. However, these statistics would be strongly perturbed by the presence of large deformations such as changes of reference points (rigid deformations). In addition, there is no guarantee that data generated by statistics done with respect to the chosen metric provide data lying on the manifold.

On the contrary, the population of shapes can be studied by the mean of smooth deformations capturing the geometrical variations of a template in a population of shapes, and a residual distance, capturing the differences that the deformations could not explain (such as noise or topological changes). The template thus belong to an atlas corresponding to its orbit under the group action of diffeomorphisms G_V .

There are two ways to formulate this approach. First, one can see a set of shapes $(T^i)_i$ as perturbations (ϵ_i) of a template \bar{T} in the chosen representation space : $\phi(T^i) = \bar{T} + \epsilon_i$. In this case shapes in the set are co-registered into a common space, and then statistics can be performed in this space. This is the *backward* model.

Second, we can see the shapes in the set as perturbations of a template after deformation: $T^i = \phi(\bar{T}) + \epsilon_i$. In this model, the residual ϵ_i – what could not be matched by the registration – is intrinsic to the individuals in the set of shapes. It is called the *forward* model.

The forward representation introduced by Stephanie Allasonnière [AAT07] allows introducing the i.i.d perturbation in the observation space of the data T^i . This is equivalent to saying that the perturbation is intrinsically linked to the observation, whereas in the backward approach the perturbation is deformed by the diffeomorphism, so it is specific to the deformation associated with the observation. The ϵ_i are no longer identically distributed between the observations.

In particular, in the forward framework, one can notice that the observation, the deformed template and the perturbation belong to the chosen shape space. In a discrete setting the template as well as the initial momenta belong to some fixed grid/tensor of vectors in \mathbb{R}^d .

Average Shooting Let T be a fixed reference shape. We minimize the functional

$$J(p^1, \dots, p^n) = \sum_{i=1}^n A(\varphi_1^{v_i}(T), T^i) + \gamma E(v_i),$$

corresponding to the registration of T onto the observations T^1, \dots, T^n . Using this functional, the template can then be estimated with a simple mean in the vectorial space of initial

momenta that is associated with $T_{Id}G_V$. The template is then given by:

$$\bar{T} = \varphi_1^{\bar{v}}(T), \quad \bar{v} = \sum_{k=1}^m K_V(q_k, \cdot) \cdot \bar{p}_k,$$

with $\bar{p} = \frac{1}{n} \sum_{i=1}^n p^i$ and $v_i = \sum_{k=1}^m K_V(q_k, \cdot) p_k^i$.

Most of the time in this setting the reference shape is chosen as one element of the observations. It can also be a shape built a priori. The statistical analysis is then lifted to the tangent space, in which the classical linear tools are available [Dur+11]. The new template \bar{T} can be seen as the result of the average deformation of the reference shape onto the rest of the database. This can be applied iteratively following the Algorithm.2.

Algorithm 2 Average geodesic shooting (fixed-step gradient descent).

Input: q_0, δ (step size), (T_1, \dots, T_n) (the targets), k (maximum iteration)

Output: \bar{q} the template tree after average deformation

Set $p^i = 0$ for $i \in (1, \dots, n)$

Set $\bar{q} = q_0$

Set $it = 0$

repeat

 Compute (p^1, \dots, p^n) with Algorithm.1.

 Set $\bar{q}_{it+1} = \varphi^{\bar{v}}(\bar{q}_{it})$ using Equation. 2.5.1.

$it = it + 1$

until $it \geq k$

Unknown Template Coordinates An other estimation of the template through forward scheme can be done by considering the discretization points (q_k) of T as optimization variables:

$$J(q, p^1, \dots, p^n) = \sum_{i=1}^n A(\varphi_1^{v_i}(T), T^i) + \gamma \cdot E(v_i),$$

and $v_i = \sum_{k=1}^m K_V(q_k, \cdot) p_k^i$

This can be interpreted as searching both the template position and the optimal deformations of the template we are estimating onto the observations. The optimization in this setting can be done alternatively by registering first the template to all the observations, which corresponds to n classical registrations done independently, then updating the initial momenta and the template coordinates. Note that the backward model would compute the average of the observations in the space of representation of the data (in which we define the data attachment term).

Hyper Template While providing straightforward ways to estimate the template provided a set of observations, the last approach does not guarantee the existence of a solution for all the shape space representations (for example in the case of normal cycles) and does not

prevent abnormal updates of the template's coordinates such as folding. To mitigate this, a two-steps optimization can be used consisting in considering the template as part of the orbit of a *hyper template* T_{ref} under the group action G_{V_0} of diffeomorphisms which can differ from G_V . The function to minimize becomes:

$$J(p^1, \dots, p^n) = \alpha.E(v_0) + \sum_{i=1}^n A(\varphi_1^{v_i}(T), Y^i) + \gamma E(v_i),$$

with $T = \varphi_1^{v_0}(T_{ref})$. The existence of a solution in this context can be ensured by the condition $V_0 \hookrightarrow C_0^3(\mathbb{R}^d, \mathbb{R}^d)$. Using the hyper template enforces a regularity that is then given by both the initial position T_{ref} and the regularity of the diffeomorphisms in G_{V_0} .

No matter the approach for computing the template, it is then possible to compute an "average" deformation from a set of registrations from an atlas onto a set of targets, by averaging the corresponding initial momenta, and then shooting via geodesic equations to deform the atlas along this average. This process may be iterated by replacing the atlas shape with this deformed atlas and repeating the same steps. This was first introduced by [Vai+04] in order to build an average atlas of the database.

2.5.2 PCA and LDDMM

We have seen different ways to build a template seen as a statistically representative element of the observations, that can be registered to these latters up to small perturbations in the space of representation of the shapes. When fixing the template's coordinates, and computing its registration to a set of observations, the atlas could be for instance the space covered by the linear combinations of the initial momenta computed and used to generate the vector fields encoding the diffeomorphism. The approach was proposed in [Vai+04], by considering the space tangent to the space of diffeomorphisms at identity $T_{Id}G_V$ when the control points are fixed. This PCA is in fact a kernel-PCA introduced by [SSM97] where the euclidean metric is replaced by the kernel metric given by K_V , which is more natural in the space of momenta. It is therefore possible to use classical statistical concepts such as eigenvectors and eigenvalues to describe the directions of interest and their associated variance in the space of initial momenta.

As specified in equation (2.7) the optimal deformations generated by LDDMM are fully parameterized by the initial momenta when fixing the template's coordinates. Their representation in the euclidean space E allows performing classical linear statistics. In the euclidean space E , the classic PCA of a set of vectors $P = (p^1, \dots, p^n)$ (where the p^i are column vectors) is done by studying the eigenvectors of the covariance matrix $C = P P^T$. This corresponds to the case of a matrix C whose coordinates are given by the scalar product in E : $C_{ij} = \langle p_i, p_j \rangle$. The metric being induced by the reproducing kernel K_V , we can define the intermediate block squared matrix $\mathbf{K} = (K_{ij})$ with $K_{ij} = K_V(q_i, q_j)$. The matrix C can then be obtained by using the associated metric: $C_{ij} = \langle p_i, \mathbf{K} p_j \rangle$. When registering a template onto a data set, for instance following one of the procedure of mean estimation described in

2.5.1 one obtains a template with T with its associated control points $(q_k)_k$ and a set of initial moment $(p_k^i)_k$ corresponding to the registration of T on the observations (T^1, \dots, T^N) .

Given the average vector field \bar{v} , we now have the "kernel covariance matrix" Γ^V :

$$\Gamma_{ij}^V = \langle v_i - \bar{v}, v_i - \bar{v} \rangle_V = (p^i - \bar{p})^T \cdot \mathbf{K} \cdot (p^i - \bar{p}). \quad (2.22)$$

Then if V^k is the k -th eigenvector of Γ^V , the k -th principal mode is given by:

$$m^{p,k} = \bar{p} + \sum_{i=1}^n V_i^k (p^i - \bar{p}). \quad (2.23)$$

It follows that the vector field associated with the k -th mod of the deformation writes

$$m^{V^k} = \sum_i K_V(x_i, \cdot) m^{p,k}. \quad (2.24)$$

By integrating the geodesic equations from the principal modes of the vector fields, one is able to generate new deformations representative of the variations in the database.

Principal Geodesic Analysis

When the statistical analysis cannot be lifted to the tangent space at some point of the manifold, another approach is to consider the Principal Geodesic Analysis introduced in [Fle+03; Fle+04], consisting in finding a submanifold maximizing the projected variance of the observation along that submanifold. The basic idea is to build recursively an orthonormal basis in the tangent space $\tau_1, \dots, \tau_k \in T_{\bar{T}}X$ at the estimated mean \bar{T} of the observations. The i -th principal geodesic submanifold is the image of the subspace spanned by the i first tangent vectors under the exponential map.

These tangent vectors can be estimated recursively by optimizing a loss function

$$\max_{\tau \in T_{\bar{T}}X, |\tau|=1} \sum_{i=1}^N d(\bar{T}, \pi_H(T^i))$$

with $H = \exp_{\bar{T}}(\text{span}(\{\tau_1, \dots, \tau_{k-1}, \tau\}))$

2.6 Application to the Registration of Simplified Pelvic Vascular Trees

We have given the main notions for the construction and computation of smooth and relevant deformations of shapes as well as their ambient space. These deformations are furthermore adapted to a statistical study of a population of shapes, which allows the construction of a template as well as admissible deformations (and thus an atlas and potentially allow a data

augmentation). The diffeomorphic deformations are associated with a data attachment term expressed in shape space, yielding inexact registrations and computationally tractable minimization procedures. We now apply the LDDMM registration to the alignment of simplified vascular trees extracted from CBCT images of the pelvic area.

Its use in the automatic annotation of vascular trees will be detailed later in Chapter 5, Section 5.1. In this section, we will focus on applying the diffeomorphic deformations we have constructed so far in the registration of vascular trees. For this purpose, we study simplified versions of the trees so that the centerline trees have the same number of branches. These simplified vascular trees are additionally composed of the arteries of interest for the identification of the vascular tree type defined in the literature [Ass+15a] (see Figure 1.9). We focus here on the diffeomorphic deformations one can generate with LDDMM between one annotated template vascular tree and a dataset of unlabeled target vascular trees.

The shapes we consider are then vascular trees of discretized 3D curves (Definition.3). The vertices are connected, and we can use the discrete varifolds representation 2.4.2.

Registration and Optimization

In all the following experiments, we initialize the registration by aligning objects centers of mass since no prior positioning is known. To model non-rigid deformations, we define the reproducing kernel K_V of V to be a sum of Gaussian kernels

$$K_V(x, y) = \frac{1}{4} \sum_s \exp \left(\frac{-\|x - y\|^2}{(\sigma_0/s)^2} \right) . Id,$$

where $s \in [1, 4, 8, 16]$ and $\sigma_0 = 100$ is about the size of the shapes bounding boxes. In terms of data attachment scale for the space of varifolds, we use a multiscale optimization framework as well to avoid falling into local minima. This classic coarse to fine approach in the registration applications [Ris+10; TQ16; Sah20] is done by optimizing the functional of Equation.2.10 and iteratively reducing the scale σ_W of the varifolds reproducing kernel. In this application we use two scales $\sigma_W = 50$ and $\sigma_W = 25$. The optimization is done using Limited memory Broyden–Fletcher–Goldfarb–Shanno (LBFGS) algorithm [LN89], with a strong-wolfe line search strategy. Our Python implementation makes use of the libraries PyTorch [Pas+17] and KeOps [Cha+21], to benefit from automatic differentiation and GPU acceleration of kernel convolutions.

Registering all Sources onto all Targets

Two trees from the dataset can be used in the LDDMM framework as the source and target shapes. The LDDMM registration of a source tree S onto a target T is done by minimizing the functional of Equation.2.10 with the distance in the space of unoriented varifolds:

$$J(v) = \gamma |v|_V^2 + |\mu_{\varphi_1^v(S)} - \mu_T|_{W'}^2.$$

We illustrate in Figure 2.8 the LDDMM registration output using the provided parameters. For more results concerning the anatomical consistency of the registrations, we refer to

Section 5.1.3. During the LDDMM registration, the shapes do not need to be annotated. We illustrate the output of the deformation using a color-free target and the annotated deformed source shape to recall that the absence of a label in the target during automatic annotations procedures. In all the illustrations that follow the root in the labeled trees is the red leaf, usually at the middle-top of the 2D images. We see that the deformed source shape is well

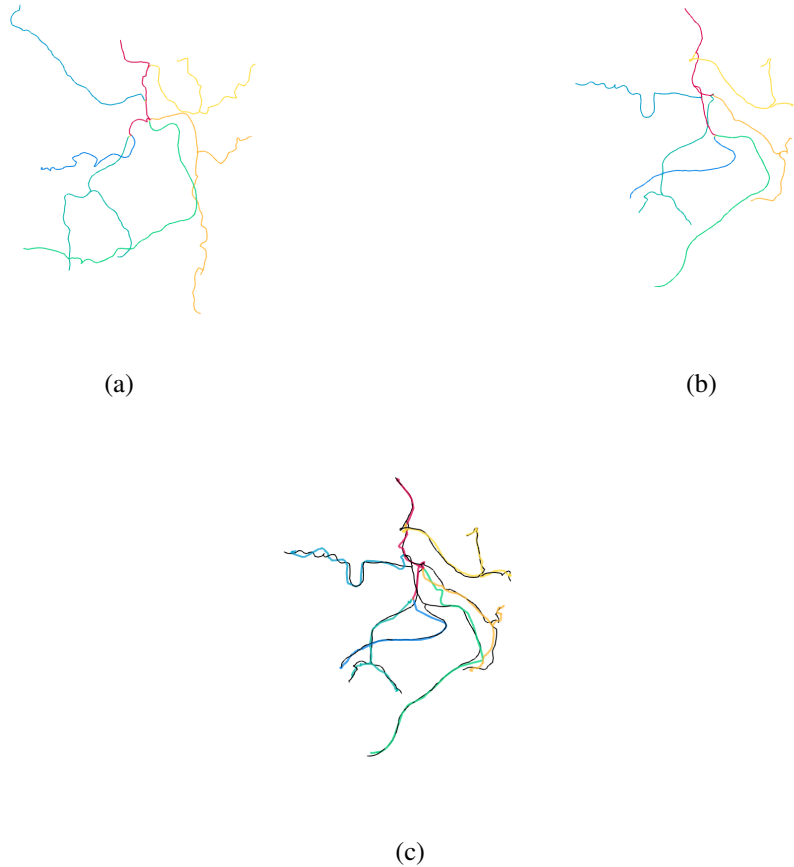


Figure 2.8: Registration of a template tree S (a) onto a target tree T (b). The labels are represented by the colormap. The output of the optimization procedure $\varphi_v^1(S)$ at the finest data attachment scale $\sigma_W = 25$ is illustrated in (c). The deformed template tree is represented as a thick, transparent structure. The target in (c) has no colormap to illustrate that it does not need to be annotated in the registration procedure.

aligned to the target. In this experiment, the leaves of the trees are correctly aligned. A lack of alignment can be observed in Figure 2.8 (c) in the proximal part of the trees, close to the root.

In a second illustration, we show a typical bad registration that illustrates the two principal sources of bad registration observed when aligning the simplified vascular trees with LDDMM. The first type of error is illustrated in Figure 2.9 (c), and is typical of the errors in the proximal parts of the trees. We observe indeed that the top left artery is abnormally distorted to match both one inner branch of the target and its real associated branch in the target. Moreover, the prostatic artery of the target (darkest blue one in Figure 2.9 (b)) is aligned with

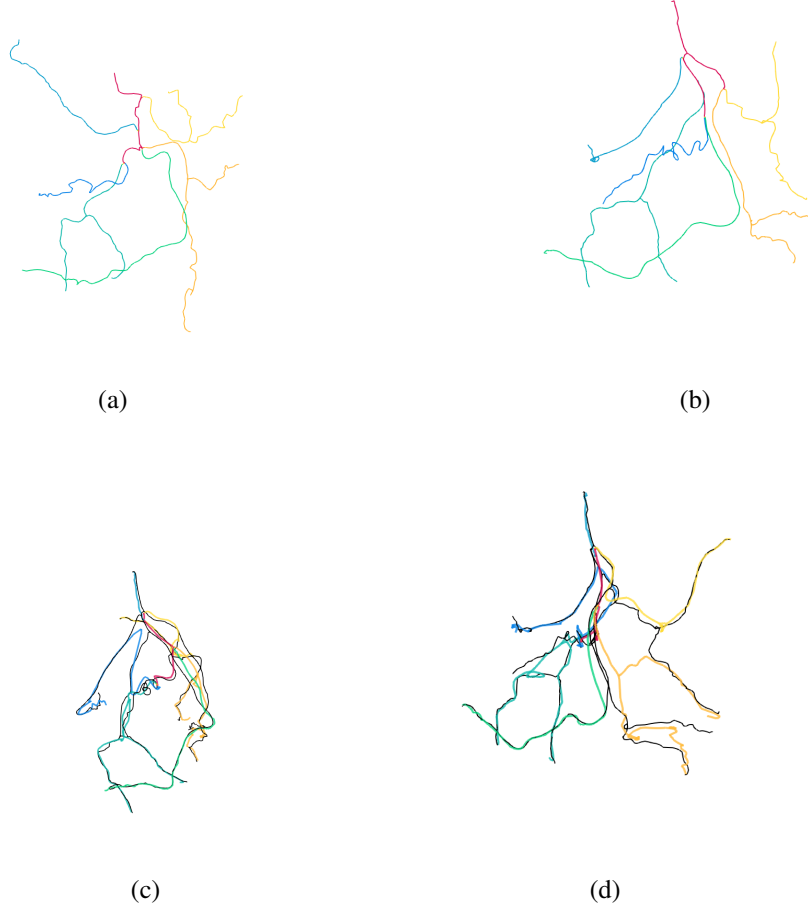


Figure 2.9: Registration of a template tree S (a) onto a target tree T (b). The labels are represented by the colormap. The inconsistent output of the optimization procedure $\varphi_v^1(S)$ at the finest data attachment scale $\sigma_W = 25$ is illustrated under a first point of view (c) and a second one (d). In (d) we observe a "jump" of one deformed branch from one artery to the other in the target.

the wrong deformed artery of the source. In Figure 2.9 (d) we see the other typical erroneous registration we observed within the registrations: a "jump" of one deformed branch from one branch of the target to another.

The two abnormal deformations illustrated in Figure 2.9 have several explanations. First of all the minimization frameworks tends to minimize the deformations cost according to the regularization parameter γ . When the initial position of the source shape with respect to the target is a bad start, the diffeomorphisms might drive the registration to a local minima. In addition, the coarse to fine approach for the data attachment term does not help avoiding such a local minima at the largest scale.

Secondly, in the inner part of the trees, we observe changes in terms of bifurcations ordering. In the target tree Figure 2.9 (b) for instance, the two gluteal arteries (most right arteries in the 2D images) branch together. On the contrary, the inferior gluteal in Figure 2.9

(a) takes source deeper in the tree. The diffeomorphisms can not explain such topological changes.

Statistics over the Deformations

We now apply the template construction described in Algorithm. 2 using the source tree of Figure 2.8(a) as initial position. We stop the template construction after 4 iterations, which results in $3 * 50 = 150$ registrations in total for the construction of the template, and 50 other LDDMM procedures to register this new template on the database. We illustrate in Figure 2.10 the template tree \bar{S} , \bar{S}_1 , \bar{S}_2 and \bar{S}_3 across the iterations of Algorithm.2. We

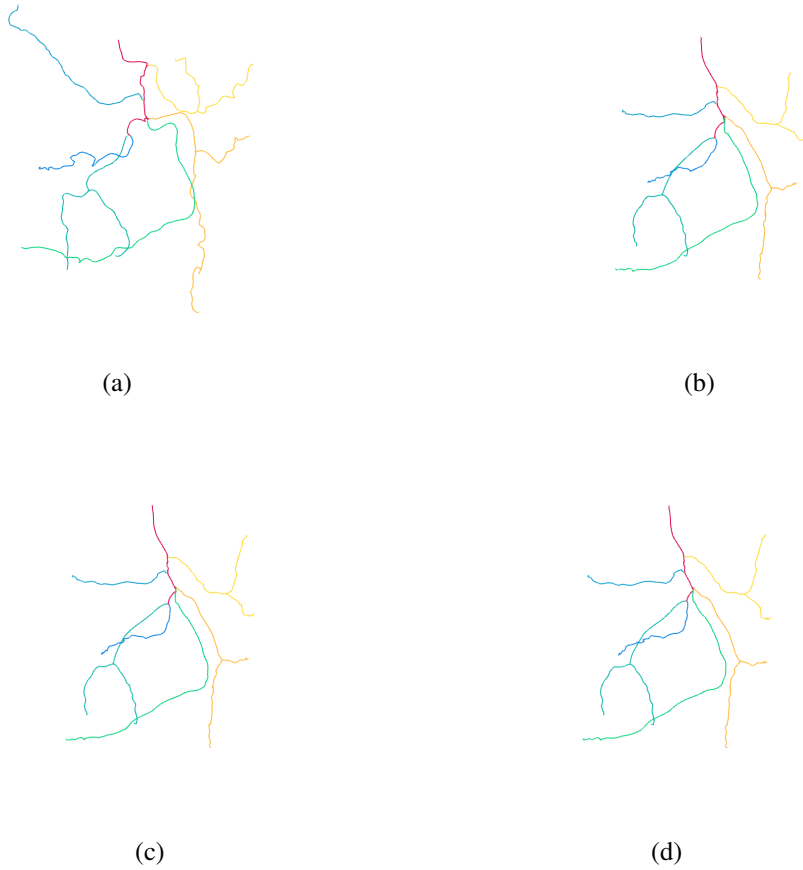


Figure 2.10: Construction of a template tree across the iterations of Algorithm. 2, with . (a): Initial shape. (b): \bar{S}_1 . (c): \bar{S}_2 . (d): \bar{S}_3 .

observe that within a few iterations the average deformation of the template vascular tree onto the database is captured. This new template captures the average geometry of the shapes in the database, and is a shape representative of the data. It is therefore a good initialization for a registration procedure. We use \bar{S}_3 in the registration to the target of Figure 2.9 as an illustration in Figure 2.11. We observe in this new figure that the registration that failed using the raw tree of the database in Figure 2.9 is now successful. No supplementary information




Figure 2.11: Registration of a template tree \bar{S}_3 (a) onto a target tree T (b). The labels are represented by the colormap. The successful output of the optimization procedure $\varphi_v^1(S)$ at the finest data attachment scale $\sigma_W = 25$ is illustrated under a first point of view (c) and a second one (d).

was used than the available shapes in the database (that did not need to be labeled). The inner branches are still slightly mismatched.

We also generate the first deformations eigen modes in Figure 2.12. These modes seem to be consistent with the variations across the database: we observe for instance strong deformations in the proximal part of the tree and the extremities, where respectively the topological and geometrical variations are the largest.

2.7 Conclusion

The LDDMM registration of one case onto the other observations of the database allows capturing spatial variability across the shapes. The construction of relevant templates helps the registration by providing a good initialization to the minimization procedure. In addition, the available statistics over the deformations have a number of applications: it could be used in data augmentation applications, or for outlier detection (such as a missing artery in the




Figure 2.12: Illustrating PCA of the initial momenta. (a) Template; (b) Shooting along the first mode; (c) Shooting along the second mode.

target).

The simplified tree template we obtained from the set of manually built examples is a realistic schematic representation of the clinicians' representation. The initial classification of the male pelvic trees was made according to the source of the prostatic artery in the tree: it depends if it starts from the superior vesicle, the obturator, the pudendal, a gluteal or "the rest" (see Figure 1.9). One limitation of this classification regarding the non-rigid registration alignment of vascular trees is the fact that each type described relatively to the prostatic artery hides a combinatorial number of other bifurcations ordering. In other terms, it suffices to switch two arteries that do not branch or take root from the prostatic artery to obtain a new tree topology. We observed that the diffeomorphic vascular tree registration could not cope with such topological changes (Figure 2.11).

In line with the idea of comparing a new vascular tree extracted from our CBCT volumes to a reference atlas, relying on a simplified vessel template seems to be a reasonable a priori: it is better to avoid trying to compare two complex trees at first, and start by comparing a simple and easily annotated template to a complete tree. The problem with this second idea is that LDDMM are then used to generate a deformation of the whole source shape on the whole target. In Figure 2.13 we apply the LDDMM registration to the alignment of a simplified tree with 17 branches on a complete tree with more than a hundred branches. The resulting deformation that has neither anatomical nor geometrical consistency, and it would require tuning the regularization parameter at each registration on a new target to get realistic deformations, which is not affordable.

Intermediate solutions such as the automatic simplification of complete trees could reduce the complexity of the problem, and must be tried before considering more complicated solutions. It would include limiting the depth of the tree or the length of the paths from the root to the extremities, or the regularization of the deformations to avoid abnormal distortions. However, it is clear that this can not be a satisfying solution as is when working on structures as complicated as vascular trees with so much topological and geometrical variability. These two observations motivate the following chapters.




Figure 2.13: Illustrating the LDDMM registration using varifold distance of (a) a simplified tree onto (b) a complete target tree. (c) The resulting deformation that has neither anatomical nor geometrical consistency.

ALIGNING A SHAPE ONTO A SUBSET OF A TARGET

We all have been confronted with a treasure map with a missing part (or at least we would have wanted it). When we try to compare this partial map to a recent one, we still manage to establish correspondences between the two. It often happens in medical imaging, where anatomical structures only partially correspond to one another in two complementary modalities (such CT and MRI). In the case of artery tree registration, this constitutes a change in topology that needs to be managed when trying to build anatomically relevant shape alignments.

| | | |
|-------|---|-----|
| 3.1 | Missing Pieces: A Topological Challenge | 75 |
| 3.1.1 | Partial Matching between Shapes: a State of The Art | 75 |
| 3.1.2 | First Ideas with a Toy Example | 88 |
| 3.2 | Partial Matching in the Space of Oriented Varifolds | 90 |
| 3.2.1 | Adding a Local A Priori | 91 |
| 3.2.2 | Normalizing the Kernels | 93 |
| 3.2.3 | Use with Rigid | 96 |
| 3.2.4 | Use with LDDMM | 96 |
| 3.3 | Regularization as Source Attachment Term | 98 |
| 3.4 | Examples of Applications | 102 |
| 3.5 | Conclusion | 105 |

3.1 Missing Pieces: A Topological Challenge

Finding shape correspondences is a standard problem in computer vision that has numerous fields of application: pattern recognition [BBK06; Bro+09; KZH13], annotation [BMV13; Fer+15] and reconstruction [Hal+20]. In medical imaging, matching an atlas and a patient’s anatomy [Fer+15], or comparing exams of the same patient acquired with different imaging techniques [Zhe+12; Bas+18], can provide critical information to physicians for both decision making and intervention planning.

As illustrated in Chapter 2 (Figure 2.8), the correct alignment of a labeled vascular tree (template) can be used as a relevant initialization for atlas based automatic annotation. Yet, most of the time, the data cannot be manually processed to the level of simplification proposed in Section 2.6 during interventional procedures. In such a case, the vascular tree obtained during the procedure is far more complex and richer than the previously built and annotated template. To perform automatic alignment of the simplified template onto the complex target, we need to deform the source to match a subset of the target. In this work, we will focus on the comparison of shapes described in Section 2.4.1, namely points clouds, curves, and surfaces of \mathbb{R}^3 and curves and points clouds of \mathbb{R}^2 .

Chapter Organization The chapter is articulated as follows: we first discuss in Section 3.1.1 the problem of partial matching for shape analysis and establish a state of the art on this topic. We then show that the previous data fidelity metrics introduced in Chapter 2 Section 2.4 and the classical LDDMM framework for the non-rigid registration of non-diffeomorphic shapes with missing parts. We also provide a first intuitive – yet unsatisfying – adaptation of the oriented varifold distance for approximating the injection of the deformed source into the target under non-rigid deformations. Afterwards, we introduce the data fidelity metric that was used in the joint work with Irene Kaltenmark, which we published in Information Processing in Medical Images¹ [Ant+21]. Section 3.3 is dedicated to the implementation of regularization terms to alleviate issues related to the combination of partial matching and non-rigid shape registration, which is an extension of the IPMI paper published in [Ant+22]. Finally, in Section 3.4, we use the proposed partial matching term for the registration of the template vascular tree constructed in Chapter 2 onto real trees. In addition, we discuss two other use-cases: for the registration of a truncated surface in \mathbb{R}^3 onto a subset of a complete target, and the converse framework, namely the deformation of a shape onto a truncated target.

3.1.1 Partial Matching between Shapes: a State of The Art

In many real-world applications, an observation often contains only a piece of the desired information. In medical imaging, it is possible to complement the missing information through other acquisitions, for example :

- Bring a third dimension from 2D images through stereoscopic acquisitions.

¹<http://ipmi2021.org/>

- Diagnose a pathology from complementary tissue visibility in CT and MRI acquisitions.
- Quantify treatment effectiveness through tumour size evolution measured on pre and post-treatment CT scans.

One condition to take advantage of two complementary acquisitions is to be able to compare them and establish a robust correspondence between the two. Because the two observations may differ in terms of content, level of details, field of view, or even acquisition quality, the correspondence needs to be robust to missing information.

The problem of shape comparison has been extensively studied during the past decades [Kai+11], with applications to shape registration, analysis, and reconstruction. It has been often associated with shape registration both in medicine and for generic shape analysis [KBD17; Sah20].

In this context, missing information can be interpreted as a problem of partial correspondence between shapes: this can be for instance related to truncation in one of the shapes coming from a limited signal (field of view, no X-ray absorption...) in one acquisition. This can also be related to the changes from one individual to the other, inducing supplementary or missing structures. In medical imaging, this can corresponds for instance to tumor growth, clinical tool insertion, or ablation of an anatomical region, which leads to topological changes in the shape of interest. Thus, establishing a robust comparison necessitates taking into account that shapes can only be partially matched.

In the specific case of partial matching, one can find two main approaches to such a problem: either by finding correspondences, sparse or dense, between structures from descriptors that are invariant to different transformations [AMC08; BMV13; Fer+15; Rod+17; Hal+19], or by looking for a deformation aligning the shapes with respect to a given metric [HC21; SBC21].

Partial Matching as Explicit Correspondences

The simplest way to compare shapes (spatially embedded in the same space) is through the search for nearest neighbors between the discretization of the source and target shapes. The fundamental distance that has been introduced with this in mind is the Hausdorff distance:

$$d(S, T) = \text{Max} \left\{ \text{Sup}_{y \in T} \left(\text{Inf}_{x \in S} |x - y| \right), \text{Sup}_{x \in S} \left(\text{Inf}_{y \in T} |x - y| \right) \right\}.$$

It has been used for instance for points cloud comparison [Gir+05], and derived for registration when coupled for instance with ICP by alternatively updating the point pairing to deform the source and compute the new distance to the target with the updated source. In the case of partial matching, this distance has been adapted by using a directed – or asymmetric – term, for instance $\delta(S, T) = \text{Sup}_{y \in T} \left(\text{Inf}_{x \in S} |x - y| \right)$. Using such a term that is not a metric anymore, one aims at including one shape into the other [Vel01]. The problem with this specific approach is that under non-rigid deformations, collapsing the source to one point

of the target minimizes the data term. A regularized version of the ICP is introduced in [Che+02] based on the least trimmed squares: this provides an outlier rejection that allows partial correspondences between the shapes.

In order to refine the shape representation, and potentially reduce the dimension of the discretized shape, most of the methods comparing the shapes in a sparse, point-wise fashion, extract representative features from the shapes, located in areas of interest that are then used to perform explicit correspondences. In [AMC08] a regularized version of the ICP selects the sets of four co-planar points in the points cloud. The early works on partial shape correspondence as reviewed in [Kai+11] rely on correspondences between points computed from geometric descriptors extracted from an isotropic local region around the selected points. These could be seen as descriptors that are characteristic and invariant to certain transformations (e.g. the SIFT descriptors in images [Low99]). Such features are usually invariant by translation and rotation, allowing to compare shapes without any prior step of rigid registration (which is a topic of its own). The method proposed in [Kai+11] is refined in [KZH13] by selecting pairs of points to better fit the local geometry using a bilateral map. This bilateral map is built starting from two points in the shape and the geodesic path connecting them in this shape. A distance map to the geodesic path is then computed, providing a region of interest seen as the points in the shapes closer than a fixed threshold to the geodesic. The features extracted can also be invariant to non-rigid transformations, as in [Rod+13] where extracted descriptors are scale invariant. Similarly to the trimmed ICP, in the case of partial matching the correspondences between the descriptors are filtered according to some metric: this allows us to restrict the number of pairings and adapt to partial correspondences [Ang+04; Hua+08]. The challenge in such approaches is to find relevant descriptors of the shapes as well as a rejection criterion for inconsistent pairings. In [Hua+08] for instance, the criterion is based on the intrinsic geodesic distance induced by the points pairings. The registration pipeline is adapted to piecewise rigid deformations.

These sparse correspondences are naturally adapted to partial matching, yet they may not be adapted to guide non-rigid registration in particular since they do not consider the shapes as a whole. Instead of comparing points or point features extracted from the shapes, some authors focus on the subset of the shapes. In the case of vascular trees, for instance, this was done in the Iterative Closest Curve procedure [BMV13]. When such a decomposition as a union of subsets is available, finding correspondences between the subsets allows for piece-wise distances and paves the way to more refined registration methods. Interestingly, the deformation cost to align the shapes can be seen as a distance between shapes itself as in [Fer+12a; Fer+15] by quantifying the amount of deformation necessary to register one shape onto the other. This is similar to the case of exact matching described in Chapter 2 with the distance in the space of tree-like shapes instead of the space of diffeomorphisms. This provides a complementary tool to the metrics used to quantify the shape's dissimilarities. This approach uses explicit correspondences between the subset of the shapes – here, the leaves – based on the annotation and the distance induced by the annotation.

In all the aforementioned works, the authors seek explicit correspondences. The increasing size of the data in the literature, in terms of shape complexity and sampling size, makes the explicit correspondences increasingly expensive as well. Different approaches mitigate the cost of computing all the possible correspondences. One classic approach is to use mul-

tiscale methods looking for correspondences at a coarse scale first and progressively refining this scale. A second and natural interpretation is to search for a minimal transportation cost between shapes. In the case of trees and graphs, [DKS18] recently proposed an Optimal Transport term between the edges of the graph seen as elastic curves. It was also used for statistics and data synthesis in [GSS21]. The optimal transport - which can be seen as generalized sorting in 2-dimensions and higher is an efficient, robust and scalable solution [Cut13; Fey+17] for pairing sparse data seen as sums of diracs measures over the discrete shapes. Yet such pairing procedures cannot be used without a good calibration of the data: they are sensitive to rigid transformations of the shapes for instance, and the deformations induced by the transport might not be realistic as well.

Dense Representation of the Shapes

In between the explicit sparse correspondences and the shapes seen as a whole, some authors compare the shapes with dense data attachment terms. Similar subsets of the shapes are then matched with respect to a given distance (deformation cost, distance between response maps...).

The early work [BBK06] and its extension [Bro+09] propose to look for the optimal trade-off between the size of the subsets of the shapes to be matched and the amount of deformations needed to align them in an embedding metric space. It is applied to surfaces and computes a regular non-rigid deformation whose cost is taken as the distance between the selected subsets. More recently, functional maps were introduced in [Ovs+12] allowing dense correspondences between shapes by transferring the problem to linear functions between spaces of functions defined over the shapes. In [Rod+17] the non-rigid partial shape correspondence is based on the Laplace-Beltrami eigenfunctions used as a prior in the spectral representation of the shapes. Similarly to the case of sparse discrete correspondences, a regularization term is introduced, and an optimization procedure alternatively optimizes the part that is mapped and the correspondences between those parts. In [Lit+17] the partial shape correspondences are computed in the spectral domain, and the partial correspondence is enforced by constraining the set of solutions of the minimization problem but requires an a priori on the ratio of partial correspondence between the truncated shape and the complete one. Recently in [Hal+19], such functional map models were adapted in a deep unsupervised framework by finding correspondences minimizing the distortion between shapes. These methods are yet limited to surface correspondences.

The correspondences, both sparse and dense, fail at considering the shapes as a whole. This can be a limitation when the goal is to generate a non-rigid registration between the shapes. One valid solution is to regularize the deformations. However, in terms of data comparison term, one would like a function that takes the geometric information of the whole shapes into account instead of focusing on sparse or limited representation.

Partial Matching without Explicit Correspondences

The methods described in the previous section highlight techniques for finding partial correspondences between the shapes by explicitly looking for matching correspondences, either sparse or dense, between the shapes. On the contrary, methods considering the shapes as a whole are usually adapted to nearly-isometric deformations but they are highly sensitive to topological changes.

In the case of image registration, metamorphosis [Fra+22] allows to make structures appear or disappear in the images, by adding a variational term encoding the intensity. This has applications for example in the study of glioblastomas (brain tumors) in MRI images of the brain. Through the registration of the images with topological changes, one can retrieve which part of the images can be aligned with diffeomorphisms, and which part cannot. A similar idea is introduced in [HC21] applied to the shape space of varifolds. The source shape is represented as a sum of weighted diracs for which the weights can vary during the optimization procedure. A penalization term is introduced to prevent inconsistent weight variations.

For graphs, another approach proposed in [SBC21] is based on the use of a mask allowing to make appear/disappear sub-parts of the source shape for partial graph matching. The advantage of this method is that it allows to make structures appear but they must be added a priori in the shape. On the other hand, it allows quite naturally to make the excess structures disappear in the deformed source instead of making them shrink in an unrealistic way. The method, however, requires a strong a priori on the data if one is to introduce structures that might appear. It could be for instance supplementary branches in the target that should be added to the source. This could be estimated automatically as a pre-processing step but would require the first step of data comparison to estimate the components of the shape that cannot be consistently matched by a realistic deformation. Thus, one has to define what a consistent matching is, as well as a realistic deformation, depending on the application.

In this work, we want to build a shape comparison function that encodes the inclusion of the whole source shape in the target. This term must be differentiable to be used in the geodesic shooting algorithm and thus be included in the LDDMM framework. Such a function would then allow working with the tools of computational anatomy described in Chapter 2.

Shape Distances are not Suited Anymore

Now that we have discussed how partial shape correspondence is addressed in the literature, we can look at the behavior of classical cost functions. In the following, we show that the terms we have used in Chapter 2 generate unrealistic deformations since they try to cover the whole target with the whole deformed source.

We propose a simple illustration of the problem with toy examples of curves in dimension $d = 2$. The shapes we build (illustrated in Figure 3.1) are: a first tree composed of 3 branches and the same tree trimmed of one branch, which in this particular case is a single curve in \mathbb{R}^2 ; a two target segments aligned and sharing the same orientation, and a single source segment

with the same orientation; and a target circle and an arc as the source shape, again, sharing the same orientation.

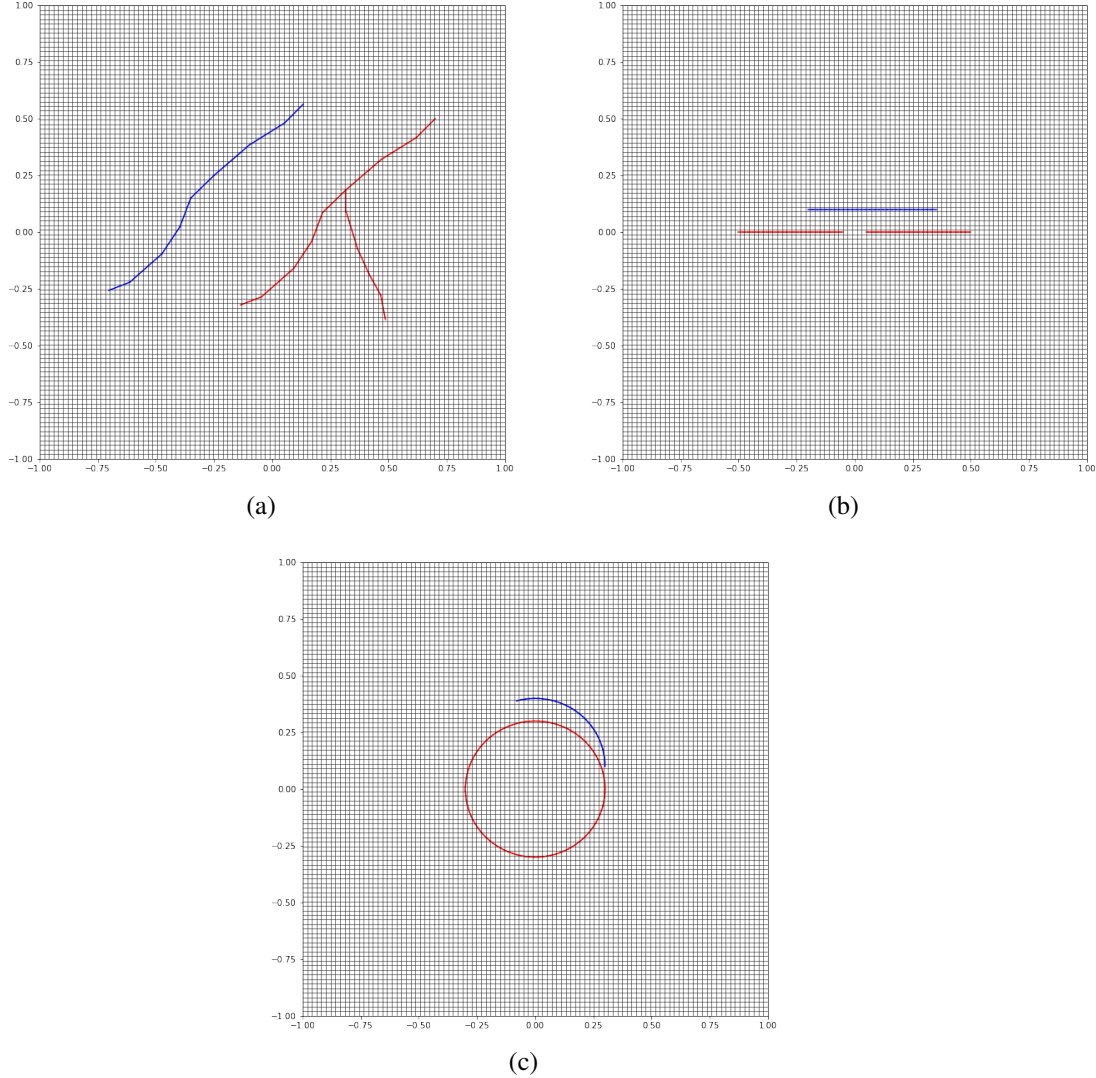


Figure 3.1: The toy examples shapes used for illustrations. The complete shape in red is the target and the blue trimmed shape resulting in a curve in \mathbb{R}^2 which is the source shape to deform.

Using such examples, we illustrate in Figure 3.3 the output of the LDDMM registration of the source in blue onto the target in red. The deformations of the ambient space are illustrated with the deformation of a 2D grid. The data attachment terms used are the distances in the space of currents, varifolds and oriented varifolds described in Section 2.4.2.

In Figure 3.2 we illustrate the registration of the full shape onto the same shape translated with the tree shape space representations described in Chapter 2 Section 2.4.2. We see that using the classic distance as data attachment term, and for the three shapes representation, the LDDMM registration is satisfying and allows to correctly align the shapes, with each branch deformed onto its corresponding branch in the target.

In Figure 3.3 (a,b), the LDDMM registration of the trimmed shape onto the complete tar-




Figure 3.2: Classic LDDMM registration between diffeomorphic shapes in (a). The source is the left-hand tree (blue) and the target is the right-hand one (red). The parameters are $\gamma = 0.01$, $\sigma_V = 0.5$ and $\sigma_W = 0.5$. (a) Using currents representation. (b) Using non-oriented varifolds representation. (c) Using oriented varifolds representation.

get guided by the current metric shows inconsistent results when looking at the deformations of the ambient space. The multi-scale approach tends to better match part of the deformed source to the target, however, the bad local minima of the largest data attachment scale can not be mitigated with the multiscale approach, and a spike is created, which is almost zero for the currents norm due to the orientation. This approach allows in fact to avoid the tuning of the scale parameter and is widely used in the literature [Sah20].

Similar observations can be made in Figure 3.3 (c,d,e,f) for the partial matching in the case of varifolds (both oriented and non-oriented) representation. The varifolds (both oriented and non-oriented) fail at providing satisfying LDDMM registration of the truncated tree onto the target. The first column of Figure 3.3 (c and e) illustrates the single scale LD-




Figure 3.3: Non-robust behavior of LDDMM registration with the distance in the space of currents (a,b), varifolds (c,d) and oriented varifolds (e,f) in the context of partial matching with the parameters $\gamma = 0.01$, $\sigma_V = 0.5$. First column (a,c,e): single scale approach with $\sigma_W = 0.5$ for the data attachment. Second column: multiscale approach with $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. Note that the output of figure (a) (resp. c, e) is the input for the scale $\sigma_W = 0.25$ of figure (b) (resp. d, f).

DMM registration: the non-oriented varifolds are sensitive to the supplementary branch, creating a spike in the deformed source to match the supplementary branch. The oriented varifolds both create a spike and the lower part of the source ends up in a local minima between the two lower branches of the target. In the case of non-oriented varifolds (Fig-

ure 3.3 (d)) with multiscale scheme, the result at initial scale $\sigma_W = 0.5$ is refined, a part of the source is correctly matched to one branch of the target while the rest of the deformed source is attracted by the two remaining curves in the target inducing strong local deformations of the ambient space. The multiscale oriented varifolds (Figure 3.3 (f)) do not provide better deformations: they seem to align the source onto a subset of the target similarly to the currents, yet the deformations of the ambient space are completely inconsistent with the expected ones (namely, a translation).

None of the distances presented here allow us to manage the missing branch in the source shape: the currents may look more robust because the supplementary branch is almost orthogonal to the rest of the structure and almost zero in the data attachment term because of the scalar product. In practice, vascular tree branches rarely are orthogonal. Yet as discussed in [Cha+20], this model of currents may create artificial thin structures (which can be noted in Figure 3.3 (a)) because they cancel in the currents space. Moreover, both varifolds terms take the additional branch into account which leads to a bad shape as well as the ambient space deformations. The orientation, however, when available, seems to be an information of interest when matching a shape onto a subset of the target.

We illustrate the LDDMM registration with the two other toy examples in Figure 3.1 (b,c). We focus on the multiscale registration to reduce the number of experiments here. On the registration of one segment onto two segments target, in Figure 3.4, the three shapes representations provide similar results: the segment is stretched to match the two targets, filling the hole between the two. When registering an arc onto a complete circle in Figure 3.5 we observe a similar result, the three models tend to guide the deformed source to match the whole target circle leading to strong deformations of the ambient space. The orientation, however, leads the oriented varifolds and the currents to similar results slightly different from the varifolds one.




Figure 3.4: Registration of a blue segment onto two red segments. Multiscale LDDMM registration with $\gamma = 0.01$, $\sigma_V = 0.5$ and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. (a) Initial positions. (b) Currents. (c) Varifolds. (d) Oriented varifolds.

Regularizing the Deformations

A first and intuitive way to deal with the inconsistent deformations is to regularize them by playing with the regularity of the vector field encoding the diffeomorphism. For instance, rigid deformations would never produce collapsed branches or the appearance of thin elongated structures. The first way to regularize the LDDMM deformations is to play with the hyper parameters σ_V the scale of the RK K_V of V (Equation 2.6) and γ the regularization weight in the cost function of Equation 2.10. In Figure 3.6 we show the influence of the regularization parameter γ on the registration of the trimmed source shape with oriented varifolds. We see that when using $\gamma = 0.1$ with the classic distance in the space of oriented varifolds and a multiscale scale approach, the obtained deformation corresponds to the




Figure 3.5: Registration of a blue arc onto a red circle. Multiscale LDDMM registration with $\gamma = 0.01$, $\sigma_V = 0.5$ and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. (a) Initial positions. (b) Currents. (c) Varifolds. (d) Oriented varifolds.

expected inclusion. The single scale approach, on the contrary, still results in inconsistent deformations. Similar behaviour can be observed using the currents representation. The solution may seem satisfying but one will see that in more complex cases (see Figure 2.13) the parameter γ would need to be tuned for every case, which is not compatible with an automatic application. For example in Figure 3.7, using the same multiscale approach with oriented varifolds distance and $\gamma = 0.1$ or $\gamma = 1$ in the two other toy examples, the registration results still tend to match the whole target shape, or to stretch the source shape. The variability in terms of registration results despite the simplicity of the examples illustrates the difficulty to adapt the regularization parameter γ to any vascular tree in a real application.




Figure 3.6: (a) Single scale LDDMM registration with oriented varifolds, and parameters $\gamma = 0.1$, $\sigma_V = 0.5$ and $\sigma_W = 0.5$. (b) Multiscale LDDMM registration with oriented varifolds, and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. Note that the output of figure (a) is the input for the scale $\sigma_W = 0.25$ of figure (b).

Divergence-Free Kernels A second solution could be using reproducing kernel encoding the admissible space for the deformations. By using matrix kernels as proposed in [MG13], one can derive properties of the vector field which will allow for example to preserve the total deformed volume.

In Chapter 2 we introduced the properties of the reproducing kernels in the generic case of vector-valued functions and then focused on the specific case of scalar kernels, so kernels that are scalar-valued functions or scalar-valued functions multiplied by the identity matrix I_d with d the dimension of the ambient space. Considering more generic class of kernels such as vector-valued ones provides the tools to construct vector fields with certain properties that are not available with scalar-valued kernels.

Such an approach gives theoretical guarantees on the nature of the generated deformations, however, two problems remain. On the one hand, the computational cost of these deformations is increased, as K_V is called at all evaluations of the geodesic shooting as many times as there are integration steps along this geodesic. In the registration of the trimmed tree example illustrated in Figure 3.8 the computing time was 32.6 seconds on CPU with the divergence free kernel and oriented varifolds distance against 6.90 seconds for the classic scalar Gaussian kernel with the same data attachment term. On the other hand, the deformations of the surrounding space are not natural. When we think about the registration of a vascular tree, the deformations generated will be very complex. Moreover, in the perspective of applying the diffeomorphisms obtained to surrounding structures, we are likely to obtain unrealistic deformations.

In summary, there are partial shape matching solutions that are adapted to realistic problems but do not necessarily allow non-rigid matching of various shapes such as unions of




Figure 3.7: Multiscale LDDMM registration with $\sigma_V = 0.5$ and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. (a) Applied to a blue arc onto a red circle, with $\gamma = 0.1$. (b) Applied to a blue arc onto a red circle, with $\gamma = 1$. (c) Applied to a blue segment onto two red segments, $\gamma = 0.1$. (d) Applied to a blue segment onto two red segments, $\gamma = 1$.

curves and surfaces (or other submanifolds in general).

On the one hand, we have classical metrics in the shapes spaces, designed for LDDMM registration. Yet since they compare the shapes as a whole, such distances do not allow non-rigid matching between shapes with topological differences. On the other hand, regularizing the deformations seems appealing at first sight but necessitates careful construction and tuning of the reproducing kernel for the diffeomorphic deformations. At the minimum of the optimization problem, the resulting distance in the shape space can be seen as the "noise" that can not be explained by the diffeomorphism. However, using the hyper parameters of the cost function may not be adapted to the complexity of the shapes in a dataset, while using




Figure 3.8: LDDMM registration of a trimmed tree onto the complete target. The reproducing kernel used is a matrix divergence free kernel conserving the total deformed volume. $\gamma = 0.01$ and $\sigma_V = 0.5$ and $\sigma_W = 0.5$.

matrix kernels increases the computational time of the deformation.

Therefore, we decided to approach the problem by adapting the shape comparison functions used in the LDDMM to encode the inclusion of a deformed source object in a target. The objective is to obtain a data attachment term that would be zero when the deformed shape is totally included in the target. In addition, to be differentiable, one would like this function to have comparable computational cost with respect to classical distances for shape matching. Ideally, such functions should also be adapted to the comparison of various shapes (basically unions of curves and surfaces) extracted from the medical images, in particular centerline trees.

3.1.2 First Ideas with a Toy Example

To build a relevant shape representation, we start from the shape space representation of oriented varifolds (Chapter 2, Section 2.4.2) which can be expressed by the mean of a reproducing kernel. Let W be a Reproducing Kernel Hilbert Space (RKHS) of functions defined over $\mathbb{R}^d \times \mathbb{S}^{d-1}$ continuously embedded in $C_0(\mathbb{R}^d \times \mathbb{S}^{d-1})$, the space of continuous test functions vanishing at infinity. Its dual space W' is a space of oriented varifolds.

In the following, we assume that the reproducing kernel of W is of the form $k_e \otimes k_t$, with the assumptions of proposition 7. We recall the dual metric:

$$\langle \delta_{x_1, \tau_1}, \delta_{x_2, \tau_2} \rangle_{W'} = k_W((x_1, \tau_1), (x_2, \tau_2)) \quad (3.1)$$

$$= k_e(x_1, x_2) k_t(\tau_1, \tau_2) \quad (3.2)$$

Here we suppose that k_e and k_t are non-negative (N.B. as functions, and not only as kernels),

excluding therefore the current representation. In order to adapt this distance for partial matching, a first and intuitive way could be to use half of the expression as follows:

$$\Delta_0(S, T) = \langle \mu_S, \mu_T - \mu_S \rangle_{W'}^2 = (|\mu_S|_{W'}^2 - \langle \mu_S, \mu_T \rangle_{W'})^2 \quad (3.3)$$

The intuition behind this definition is that if S is a subset of T , then $\mu_T - \mu_S$ is the varifold corresponding to $T \setminus S$, which is disjoint from shape S and thus roughly orthogonal to it from the varifold metric viewpoint. We illustrate in Figure 3.9 the output of the LDDMM registration with different configurations of the partial dissimilarity term $\Delta_0(\varphi(S), T)$.




Figure 3.9: First intuition of partial matching with Δ_0 as data attachment term. (a) LDDMM registration with $\sigma_W = 0.5$. (b) LDDMM registration with $\sigma_W = 0.1$. (c) Multiscale LDDMM registration with $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$.

First of all, we can see in Figure 3.9 (a) that at a large scale of data attachment, the source shape tends to stop quite far from the target. In fact, there are two phenomena that are at work here: first the fact that if the target locally has a larger weight than the source, it will contribute negatively to the proposed term and so the minimum will be reached while S is not included in T .

In particular, if we take the case of a shape S partially included in a target T : $S = S_1 \sqcup S_2$ such that $S_1 \cap T = \emptyset$ and $S_2 \cap T = S_2$. This is characterized by $\langle \mu_{S_1}, \mu_S - \mu_T \rangle_{W'} > 0$, which can be compensated by an overrated characterization of the inclusion $S_2 \subset T$ with $\langle \mu_{S_2}, \mu_S - \mu_T \rangle_{W'} < 0$, which happens if the mass of T around S_2 is larger than the mass of S .

Then if we look at the experiment in Figure 3.9 (b), we realize that at small scales σ_W , the source, and the target cannot be seen and the diffeomorphism will generate a deformation that minimizes the norm of the deformed source. The diffeomorphisms risk, if we are not careful, to shrink the source to reduce its norm (see Equation 3.3) and, if it is possible, to include it in the target. This is observed in the third Figure 3.9 (c) which uses a multi-scale registration scheme for the attachment to the data. In this experiment, the source shape is progressively included in the target while being shrunk.

This is illustrated in Figure 3.10. In this experiment, we translate the trimmed tree of Figure 3.1(b) around the complete one and compute the function Δ for each position. The point of reference of the translation is the former bifurcation: the best possible response would hence be a single 0 at the target's bifurcation. This gives a response map in which we see that the minimum value of Δ is not in the area where the trimmed tree is included in the target.




Figure 3.10: Target T in red. Left: Response map F of $\Delta(\tau(S), T)$ for $\sigma_W = 0.5$ with S the curve presented in Figure 3.1(b) and translated to each coordinates of the images. Right: the same image thresholded. The yellow area corresponds to $F \leq t$ with $t = \min(F) + 1e - 2 * (\max(F) - \min(F))$.

This result, in the case of one single supplementary bifurcation in the target, would be even worse when comparing the simplified template to a real target tree as in Figure 2.13. Thus this direct adaptation of the oriented varifolds distance is not satisfying.

3.2 Partial Matching in the Space of Oriented Varifolds

Using the shape representation in a RKHS provides convenient tools for computational shape analysis that are worth exploiting to build dissimilarity terms. In this section, we show how we deal with the mass imbalance between the deformed source and the target, and how to seek the local embedding of the source shape in the target. We also construct associated LDDMM deformations that do not cause abnormal deformations – in particular shrinkage of the source – as observed in Figure 3.9(b).

3.2.1 Adding a Local A Priori

To simplify the notation, we denote for $x, x' \in S$, $\vec{x} = (x, \tau_x S)$ and $k(\vec{x}, \vec{x}') = k_e(x, x')k_t(\tau_x S, \tau_{x'} S)$. In addition given $\vec{x} = (x, \tau)$, we denote $\omega_S(\vec{x}) = \omega_S(x, \tau)$ the unique representer of $\mu_S(\vec{x})$ with Riesz representer theorem.

Definition 11. Let $g : \mathbb{R} \Rightarrow \mathbb{R}$ defined as $g(s) = (\max(0, s))^2$. We define the partial matching dissimilarity as follows:

$$\Delta(S, T) = \int_S g(\omega_S(\vec{x}) - \omega_T(\vec{x})) dx \quad (3.4)$$

$$= \int_S g\left(\int_S k(\vec{x}, \vec{x}') dx' - \int_T k(\vec{x}, \vec{y}) dy\right) dx. \quad (3.5)$$

Note that by taking g as the identity, we would retrieve $\langle \mu_S, \mu_S - \mu_T \rangle_{W'}$. The threshold $\max(0, \cdot)$ prevents the compensation of a local mismatch by an overrated match in another area.

Proposition 10. If $S \subset T$ then $\Delta(S, T) = 0$.

Since k_e and k_t are assumed to be non negative functions, we have

$$\begin{aligned} \Delta(S, T) &= \int_S g\left(\int_S k(\vec{x}, \vec{y}) dy - \int_T k(\vec{x}, \vec{y}) dy\right) dx \\ &= \int_S g\left(-\int_{T \setminus S} k(\vec{x}, \vec{y}) dy\right) dx = 0. \quad \square \end{aligned} \quad (3.6)$$

The next proposition highlights the local nature of the dissimilarity function Δ .

Proposition 11. If $S' \subset S$, then $\Delta(S', T) \leq \Delta(S, T)$. In particular, if $\Delta(S, T) = 0$ then for any subset S' of S , $\Delta(S', T) = 0$.

Since $k \geq 0$, we have for any $S' \subset S$ and any $\vec{y} \in \mathbb{R}^d \times \mathbb{S}^{d-1}$

$$\omega_{S'}(\vec{y}) = \int_{S'} k(\vec{x}, \vec{y}) dx \leq \int_S k(\vec{x}, \vec{y}) dx = \omega_S(\vec{y}).$$

Hence, since g is an increasing function

$$g(\omega_{S'}(\vec{y}) - \omega_T(\vec{y})) \leq g(\omega_S(\vec{y}) - \omega_T(\vec{y})),$$

$$\text{and thus } \Delta(S', T) \leq \int_S g(\omega_S(\vec{x}) - \omega_T(\vec{x})) dx = \Delta(S, T). \quad \square$$

Proposition 12. $\Delta(S, T) = 0$ does not imply that $S \subseteq T$.

Proof. In this proof we exhibit a counter example. Consider the two following shapes. The source is a segment

$S_\epsilon = \{(s, \epsilon) \mid s \in [-\alpha, \alpha]\}$ slightly shifted by a step $\epsilon > 0$ above a larger target $T = \{(t, 0) \mid t \in [-\beta, \beta]\}$, with $0 < \alpha < \beta$. Since the tangent vectors are almost all equal, we can ignore k_t and consider a kernel k_e defined by a decreasing function $\rho : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as follows $k_e(x, x') = \rho(|x - x'|^2)$. Then for any such ρ , there exists $(\epsilon, \alpha, \beta)$ such that $\Delta(S_\epsilon, T) = 0$ and $S_\epsilon \cap T = \emptyset$.

We need to show that for any $x_0 \in S_\epsilon$, $\omega_{S_\epsilon}(x_0) \leq \omega_T(x_0)$ where $\omega_{S_\epsilon}(x_0) = \int_{S_\epsilon} k_e(x_0, x) dx$ and $\omega_T(x_0) = \int_T k_e(x_0, x') dx'$. Denote $x = (s, \epsilon) \in S_\epsilon$, $x_0 = (s_0, \epsilon) \in S_\epsilon$ and $x' = (t, 0) \in T$ then $\|x - x_0\|^2 = \|(s, \epsilon) - (s_0, \epsilon)\|^2 = (s - s_0)^2$ and $\|x' - x_0\|^2 = \|(t, 0) - (s_0, \epsilon)\|^2 = (t - s_0)^2 + \epsilon^2$.

We then obtain

$$\omega_{S_\epsilon}(x_0) = \int_{-\alpha}^{\alpha} \rho((s - s_0)^2) ds, \quad \omega_T(x_0) = \int_{-\beta}^{\beta} \rho((s - s_0)^2 + \epsilon^2) ds.$$

Denote these integrals $I_\alpha(s_0) = \int_{-\alpha}^{\alpha} \rho((s - s_0)^2) ds$ and $I_\beta(s_0, \epsilon) = \int_{-\beta}^{\beta} \rho((s - s_0)^2 + \epsilon^2) ds$.

The integrals are symmetric with respect to $s = s_0$ and since ρ is decreasing, we have the following inequalities:

$$\text{for any } s_0 \in [-\alpha, \alpha], \quad I_\alpha(\alpha) \leq I_\alpha(s_0) \leq I_\alpha(0), \quad (3.7)$$

$$\text{for any } s_0 \in [-\alpha, \alpha], \text{ for any } \epsilon > 0, \quad I_\beta(\alpha, \epsilon) \leq I_\beta(s_0, \epsilon) \leq I_\beta(0, \epsilon). \quad (3.8)$$

Let us now show that there exist $(\epsilon, \alpha, \beta)$ such that $I_\alpha(0) \leq I_\beta(\alpha, \epsilon)$ that is $\int_{-\alpha}^{\alpha} \rho(s^2) ds \leq \int_{-\beta}^{\beta} \rho((s - \alpha)^2 + \epsilon^2) ds$.

For α small enough and β large enough, $\int_{-\beta}^{\beta} \rho((s - \alpha)^2 + \epsilon^2) ds \geq \int_{-2\alpha}^{2\alpha} \rho(s^2 + \epsilon^2) ds$. This last integral tends to $\int_{-2\alpha}^{2\alpha} \rho(s^2) ds$ when ϵ tends to 0 and this limit is strictly larger than $I_\alpha(0)$ (with α small enough, $\rho(\alpha) > 0$). Thus, for ϵ small enough, we have $I_\alpha(0) < I_\beta(\alpha, \epsilon)$.

Thanks to eq. (3.7) and (3.8), we deduce that for any $x_0 \in S$, $\omega_S(x_0) \leq \omega_T(x_0)$. \square

\square

This example shows that if the mass of the target is larger than the mass of the source then this excess of mass can compensate the lack of alignment between the shapes. It is illustrated in Figure 3.11 with an experiment similar to Figure 3.10 in which we replaced the first proposed term with its local version of Equation 3.4.

In this illustration, we see that the local characterization of the inclusion provides a correct estimation of the local minimum area. The localization however could be better and is partly explained by the local excess of mass in the target at the bifurcation point. The other source of the extended minimum area is the square term flattening the values smaller than one.




Figure 3.11: Target T in red. Left: Response map F of $\Delta(\tau(S), T)$ for $\sigma_W = 0.5$, $k_t(u, v) = e^{\langle u, v \rangle_{\mathbb{R}^d}}$, and with S the trimmed tree presented in Figure 3.1(b) and translated to each coordinates of the images. Right: the same image thresholded. The yellow area corresponds to $F \leq t$ with $t = \min(F) + 1e - 2 * (\max(F) - \min(F))$.

3.2.2 Normalizing the Kernels

We now have a correctly localised term encoding the inclusion of a source in a target, and we are looking for a way to better deal with the mass imbalance between the shapes. To that end, we introduce a normalizing term that will penalize locally the excess of mass in the target.

Assume that $x_0 \in S$ and $y_0 \in T$ are two close points. If, around these points, the mass of T is twice the mass of S , i.e. $\omega_S(\vec{x}_0) \approx \frac{1}{2}\omega_T(\vec{y}_0)$, then the local embedding of S in T is characterized by $\omega_S(\vec{x}_0) \leq \frac{1}{2}\omega_T(\vec{x}_0)$ and more generally by $\omega_S(\vec{x}_0) \leq \frac{\omega_S(\vec{x}_0)}{\omega_T(\vec{y}_0)}\omega_T(\vec{x}_0)$. Conversely, if the mass of S is twice the mass of T , then we consider that locally $S \not\subseteq T$ (e.g. two branches of a tree should not match the same branch of a target). Hence, the criterion of Definition 11 that should be preserved : $\omega_S(\vec{x}_0) \leq \omega_T(\vec{x}_0)$ is not satisfied. These observations lead to a new dissimilarity term that encompasses these two cases.

Definition 12. Using the same threshold function g as in Definition 11, we define the **normalized partial matching dissimilarity** as follows:

$$\underline{\Delta}(S, T) = \int_S g \left(\omega_S(\vec{x}) - \int_T \min_{\epsilon} \left(1, \frac{\omega_S(\vec{x})}{\omega_T(\vec{y})} \right) k(\vec{x}, \vec{y}) dy \right) dx$$

where $\min_{\epsilon}(1, s) = \frac{s+1-\sqrt{\epsilon+(s-1)^2}}{2}$ with $\epsilon > 0$ small, is used as a smooth approximation of the $\min(1, \cdot)$ function.

Discrete Formulation The discrete version of the partial matching dissimilarity can be derived by following the same discrete setting described in Chapter 2, Section 2.4.1 for

the shape approximations. We are working with surfaces, seen as triangular meshes with vertices q_1, \dots, q_K . Each triangle f_i , $i \in [1, \dots, N_S]$ with the vertices (q_i^1, q_i^2, q_i^3) of the shape S is associated with the center $c_i = \frac{q_i^1 + q_i^2 + q_i^3}{3}$ and to the normal vector $\eta_{x_i} S = (1/2) * (q_i^2 - q_i^1) \times (q_i^3 - q_i^1)$. The unit normal vector is then $\tau_{c_i} S = \frac{\eta_{c_i} S}{|\eta_{c_i} S|}$. We define similarly the centers y_l of the target shape T and their associated normal vectors $\eta_{y_l} T$ and unit normal vectors $\tau_{y_l} T$. In the case of curve C , the unit tangent vector to C at x will also be denoted $\tau_x C$, and




Figure 3.12: Target T in red. Left: Response map F of $\underline{\Delta}(\tau(S), T)$ for $\sigma_W = 0.5$, and with S the trimmed tree presented in Figure 3.1(b) and translated to each coordinates of the images. Right: the same image thresholded. The yellow area corresponds to $F \leq t$ with $t = \min(F) + 1e - 2 * (\max(F) - \min(F))$.

(a) Top row illustrates the response map for $k_t(u, v) = \langle u, v \rangle_{\mathbb{R}^d}^2$ (varifolds). (b) Bottom row illustrates the response map for $k_t(u, v) = e^{\langle u, v \rangle_{\mathbb{R}^d}}$ (oriented varifolds.)

the center are the centers of the segments connecting the discretization points q_i .

The discrete normalized partial matching term is then written as follows:

$$\underline{\Delta}(S, T) = \sum_{i=1}^{N_S} g \left(\omega_S(\vec{x}_i) - \sum_{l=1}^{N_T} \min_{\epsilon} \left(1, \frac{\omega_S(\vec{x}_i)}{\omega_T(\vec{y}_l)} \right) k(\vec{x}_i, \vec{y}_l) \right) \quad (3.9)$$

$$\text{with } \omega_S(\vec{x}_i) = \sum_{j=1}^{N_S} k_e(x_i, x_j) k_t(\tau_{x_i} S, \tau_{x_j} S) |\eta_{x_i} S| |\eta_{x_j} S|.$$

Following the same experiment as previously, we compute the response map of $\underline{\Delta}$ in Figure 3.12 for both oriented and non-oriented varifold representation. To compare to the previous experiment in Figure 3.11, the proposed normalization provides a better localized minimum: the excess of mass at the bifurcation of the target is compensated. When comparing the two rows in Figure 3.12, we see that using the non-oriented representation the surface of the flat minimal area explodes: the orientation plays in this example an important role in the construction of a function that should be zero when the source shape is included, and not too flat around its minima.

Remark 1. *Interestingly, the minimum value of $\underline{\Delta}$ in the case of oriented varifolds never reaches 0. This comes from the introduced normalization, and whenever $\langle u, v \rangle_{\mathbb{R}^d} > 0$ we have that $k_t(u, v) \geq 1$. In such a case the normalization overcompensate the local excess of mass in the target and we have $S \subseteq T \not\Rightarrow \underline{\Delta}(S, T) = 0$.*

We show that when S is included in T , the cost function remains positive. Let $\vec{y}_0 \in T$ such that $\forall \vec{y} \in T$ we have $\omega_T(\vec{y}) \geq \omega_T(\vec{y}_0) > 0$ (\vec{y}_0 exists since ω_T is continuous and bounded, and T is compact). Let $S \subseteq T$ such that $\vec{y}_0 \in S$. We have that $\forall \vec{y} \in T$, $\omega_T(\vec{y}) \geq \omega_S(\vec{y}_0)$. Then we have:

$$\begin{aligned} \omega_S(\vec{y}_0) - \int_T \min_{\epsilon} \left(1, \frac{\omega_S(\vec{y}_0)}{\omega_T(\vec{y})} \right) k(\vec{y}_0, \vec{y}) dy &= \omega_S(\vec{y}_0) - \int_T \frac{\omega_S(\vec{y}_0)}{\omega_T(\vec{y})} k(\vec{y}_0, \vec{y}) dy \\ &= \omega_S(\vec{y}_0) \left(1 - \int_T \frac{k(\vec{y}_0, \vec{y})}{\omega_T(\vec{y})} dy \right). \end{aligned}$$

In addition, we have that $\int_T \frac{k(\vec{y}_0, \vec{y})}{\omega_T(\vec{y})} dy \leq \frac{1}{\omega_T(\vec{y}_0)} \int_T k(\vec{y}_0, \vec{y}) dy = 1$, hence:

$$\omega_S(\vec{y}_0) - \int_T \min_{\epsilon} \left(1, \frac{\omega_S(\vec{y}_0)}{\omega_T(\vec{y})} \right) k(\vec{y}_0, \vec{y}) dy \geq 0,$$

and if there is $\vec{y} \in T$ s.t. $\omega_T(\vec{y}) > \omega_T(\vec{y}_0)$, then:

$$\omega_S(\vec{y}_0) - \int_T \min_{\epsilon} \left(1, \frac{\omega_S(\vec{y}_0)}{\omega_T(\vec{y})} \right) k(\vec{y}_0, \vec{y}) dy > 0,$$

So $\underline{\Delta}(S, T) > 0$.

With that in mind, one must pay attention to the initialization and the scale σ_W of the spatial kernel k_e . We have in fact that $\underline{\Delta}(S, T)$ is minimal when $S \subseteq T$.

3.2.3 Use with Rigid

We can define a first minimization problem for a rigid registration, by minimizing the function:

$$J_{rig}(r) = \underline{\Delta}(r(S), T),$$

with r a rigid deformation composed of a translation t and a rotation R .

For any rigid deformation $\omega_{r(S)}(r(\vec{x})) = \omega_S(\vec{x})$, $\forall x \in S$, and we can thus write in the discrete setting:

$$\underline{\Delta}(r(S), T) = \sum_{i=1}^{N_S} g \left(\omega_S(\vec{x}_i) - \sum_{l=1}^{N_T} \min_{\epsilon} \left(1, \frac{\omega_S(\vec{x}_i)}{\omega_T(\vec{y}_l)} \right) k(r(\vec{x}_i), \vec{y}_l) \right)$$

We have that J_{rig} is a composition of continuous functions and that : $k(r(\vec{x}_i), \vec{y}_l) \rightarrow 0$ when the translation goes to infinity from the construction of the kernel k . We can deduce that for all rigid deformation we have $\min_{\epsilon} \left(1, \frac{\omega_S(\vec{x}_i)}{\omega_T(\vec{y}_l)} \right) k(r(\vec{x}_i), \vec{y}_l) > 0$ and $J_{rig}(r) < \sum_{i=1}^{N_S} g(\omega_S(\vec{x}_i))$. J_{rig} is continuous and bounded over the space of finite dimension of rotations and translations, the minimization problem has a solution.

3.2.4 Use with LDDMM

The framework we propose is sufficiently flexible to be embedded in a variety of inexact registration methods; and we use it in the LDDMM framework:

Proposition 13. *Let $\gamma > 0$ be a fixed parameter. The partial matching problem, which consists in minimizing over L_V^2 the function :*

$$J(v) = \gamma \int_0^1 \|v_t\|_V^2 dt + \underline{\Delta}(\phi_1^v(S), T),$$

has a solution.

Proof. From [Gla05], theorem 7, the proof boils down to showing that the mapping $v \mapsto A(v) = \underline{\Delta}(\phi_1^v(S), T)$, is weakly continuous on L_V^2 . Let (v_n) be a sequence in L_V^2 , weakly converging to some $v \in L_V^2$. We need to show that $\underline{\Delta}(\phi_1^{v_n}(S), T) \rightarrow \underline{\Delta}(\phi_1^v(S), T)$.

To simplify we denote $S_n = \phi_1^{v_n}(S)$, $S_* = \phi_1^v(S)$ and for any $\vec{x} \in \mathbb{R}^d \times \mathbb{S}^{d-1}$, $f_n(\vec{x}) = \omega_{S_n}(\vec{x}) - \int_T \min_{\epsilon} \left(1, \frac{\omega_{S_n}(\vec{x})}{\omega_T(\vec{y})} \right) k(\vec{x}, \vec{y}) dy$ and $f_*(\vec{x})$ likewise for S_* .

We then have

$$\underline{\Delta}(S_n, T) - \underline{\Delta}(S_*, T) = \mu_{S_n}(g \circ f_n) - \mu_{S_*}(g \circ f_*). \quad (3.10)$$

The area formula

$$\mu_{\phi(S)}(\omega) = \int_S \omega(\phi(x), d_x \phi(\tau_x S)) |d_x \phi|_{\tau_x S} dx,$$

leads to

$$\begin{aligned}
 & \left| \mu_{S_n}(g \circ f_n) - \mu_{S_*}(g \circ f_*) \right| \leq \int_S \left| g \circ f_n(\phi^n(x), d_x \phi^n(\tau_x S)) \cdot |d_x \phi^n|_{\tau_x S} \right. \\
 & \quad \left. - g \circ f_*(\phi(x), d_x \phi(\tau_x S)) \cdot |d_x \phi|_{\tau_x S} \right| dx \\
 & \leq \int_S \left| g \circ f_n(\phi^n(x), d_x \phi^n(\tau_x S)) \cdot |d_x \phi^n|_{\tau_x S} - g \circ f_n(\phi(x), d_x \phi(\tau_x S)) \cdot |d_x \phi|_{\tau_x S} \right| \\
 & \quad + \left| g \circ f_n(\phi(x), d_x \phi(\tau_x S)) \cdot |d_x \phi|_{\tau_x S} - g \circ f_*(\phi(x), d_x \phi(\tau_x S)) \cdot |d_x \phi|_{\tau_x S} \right| dx \\
 & \leq \int_S |g \circ f_n|_\infty \cdot \left| |d_x \phi^n|_{\tau_x S} - |d_x \phi|_{\tau_x S} \right| + |d_x \phi|_\infty |g \circ f_n - g \circ f_*|_\infty dx.
 \end{aligned}$$

Since $d_x \phi^n$ converge to $d_x \phi$, uniformly on $x \in S$ [Gla05], we only need to show that $|g \circ f_n - g \circ f_*|_\infty \rightarrow 0$. We first show that $|f_n - f_*|_\infty \rightarrow 0$. For any $\vec{x} \in \mathbb{R}^d \times \mathbb{S}^{d-1}$

$$\begin{aligned}
 f_n(\vec{x}) - f_*(\vec{x}) &= \omega_{S_n}(\vec{x}) - \omega_{S_*}(\vec{x}) \\
 &+ \int_T k(\vec{x}, \vec{y}) \left[\min_\epsilon \left(1, \frac{\omega_{S_n}(\vec{x})}{\omega_T(\vec{y})} \right) - \min_\epsilon \left(1, \frac{\omega_{S_*}(\vec{x})}{\omega_T(\vec{y})} \right) \right] dy. \tag{3.11}
 \end{aligned}$$

Since W is continuously embedded in $C_0^2(\mathbb{R}^d \times \mathbb{S}^{d-1})$, there exists c_W such that for any n , $|\omega_{S_n} - \omega_{S_*}|_\infty \leq c_W |\omega_{S_n} - \omega_{S_*}|_W$. Moreover, since v_n weakly converges to v , Corollary 1 from [CCT17] ensures that $|\omega_{S_n} - \omega_{S_*}|_W \rightarrow 0$.

Regarding the integral, since $\mathbb{R} \ni s \mapsto \min_\epsilon(1, s)$ is Lipschitz, there exists $c_\epsilon > 0$ such that

$$\begin{aligned}
 & \left| \int_T k(\vec{x}, \vec{y}) \left[\min_\epsilon \left(1, \frac{\omega_{S_n}(\vec{x})}{\omega_T(\vec{y})} \right) - \min_\epsilon \left(1, \frac{\omega_{S_*}(\vec{x})}{\omega_T(\vec{y})} \right) \right] dy \right| \\
 & \leq \int_T \frac{k(\vec{x}, \vec{y})}{|\omega_T(\vec{y})|} c_\epsilon |\omega_{S_n}(\vec{x}) - \omega_{S_*}(\vec{x})| dy \leq c_\epsilon c_W |\omega_{S_n} - \omega_{S_*}|_W \int_T \frac{k(\vec{x}, \vec{y})}{\omega_T(\vec{y})} dy \tag{3.12}
 \end{aligned}$$

Since T is compact and ω_T is continuous and strictly positive on $\vec{T} = \{(y, \tau_y T) \mid y \in T\}$, we have $c_T = \inf_{\vec{T}} \omega_T(\vec{y}) > 0$ so that $\int_T \frac{k(\vec{x}, \vec{y})}{\omega_T(\vec{y})} dy \leq \frac{\omega_T(\vec{x})}{c_T} \leq \frac{c_W |\omega_T|_W}{c_T} < +\infty$. This shows that $|f_n - f_*|_\infty \rightarrow 0$. Now, since f_* is bounded, there exists $M > 0$ such that for any n , $|f_*|_\infty + |f_n|_\infty \leq M$ and since g is locally Lipschitz, we deduce that $|g \circ f_n - g \circ f_*|_\infty \rightarrow 0$. \square

We illustrate in Figure 3.13 the LDDMM registration of the trimmed tree onto the complete target under two configurations: a single-scale approach and a multi-scale one. The single-scale approach remains unsatisfying as the deformed source still ends up far from the target. The second approach, however, shows better results. By gradually reducing the size of the scale of the spatial kernel, the deformed source is gradually included in the target as intended.

We still observe a shrinking of the source, induced by the carpool effect of the vector fields generated by Gaussian scalar reproducing kernels, and that is worsened by the minimization of the norm of the deformed source. Indeed we have the following property:

Proposition 14. *If $S_2 \subseteq S_1 \subseteq T$, then $\underline{\Delta}(S_2, T) \leq \underline{\Delta}(S_1, T)$.*

The proof is immediate, as for $\underline{\Delta}(S_2, T)$ we integrate positive terms over a subset of S_1 .




Figure 3.13: LDDMM registration under partial matching using $\underline{\Delta}$ as data attachment term with $\gamma = 0.01$ and $\sigma_V = 0.5$. (a) LDDMM registration with $\sigma_W = 0.5$. (b) Multiscale LDDMM registration with $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. Note that the output of figure (a) is the input for the scale $\sigma_W = 0.25$ of figure (b).

3.3 Regularization as Source Attachment Term

In this section we build the supplementary data attachment term that will be used in the final data attachment term for LDDMM registration in the context of partial matching. So far we only considered an objective function that would be minimal when the deformed shape is completely embedded in the target, though we immediately see that such a situation can be achieved with infinitely many different deformations and one may need additional a priori in order to derive realistic deformations. There are many ways to regularize the LDDMM for instance but we focus on encoding the regularization within the data term, which is simpler to implement and faster than encoding regularity within the diffeomorphisms (divergence free kernels for instance).

It is known from experiments that the deformations can lead to abnormal shrinkage or stretching of the deformed shapes (Figure 3.13). This phenomenon comes from two things combined: first, we introduce an asymmetric data attachment term that favors the inclusion of a deformed object in a target. In the case of non-rigid deformations, it can create a multitude of local minima. Second, in the regularization of the LDDMM model (2.9), the deformations tend to shrink the objects along the geodesics, so it is possible that the diffeomorphisms create shrinkage of the non-realistic source shape. In order to limit, this we can add a regularization term in the function J to minimize.

We illustrate a registration in Figure 3.14 the result of the LDDMM registration with the toy example used so far, but using an increased deformation scale $\sigma_V = 2$, and a regularization parameter $\gamma = 0.01$.

Using such parameters, we obtain in fact the expected inclusion of the source inside the




Figure 3.14: LDDMM registration under partial matching using $\underline{\Delta}$ as data attachment term with $\gamma = 1$ and $\sigma_V = 1$. (a) LDDMM registration with $\sigma_W = 0.5$. (b) Multiscale LDDMM registration with $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$.

target. With a single scale, the result is not perfect we have a slight shrinkage of the source, however, switching to multiscale data attachments allows us to obtain the expected inclusion. If the deformation is close to the expected one, the large σ_V does not allow to take the finest details into account during the registration. To refine the model, one can use a multiscale reproducing kernel, and this is the solution proposed in the IPMI publication [Ant+21].

We still observe that the usual regularization is not always best suited to the case of partial matching and that shrinkage may be an admissible deformation regarding the regularization parameter γ . One may want to preserve more information in the regularization, as the local or global structure of the shape. In the next section, we discuss the use of a regularization term designed to prevent the shrinking of the source shape under non-rigid deformations and formulated in the space of oriented varifolds.

Expressions The purpose of this regularization term is to prevent the deformations from shrinking or stretching the source shape. We chose to address it in the shape space of oriented varifolds as well, by controlling the norm of the deformed shape:

$$R_{global}(S, \Phi(S)) = (\|\omega_S\|_W^2 - \|\omega_{\Phi(S)}\|_W^2)^2 \quad (\text{Global}) \quad (3.13)$$

Interestingly this term can be written with the change of variable:

$$\begin{aligned} R_{global}(S, \Phi(S)) &= \left(\int_S \omega_S(\vec{x}) dx - \int_{\Phi(S)} \omega_{\Phi(S)}(\vec{y}) dy \right)^2 \\ &= \left(\int_S \omega_S(\vec{x}) - \omega_{\Phi(S)}(\Phi(\vec{x})) |d_x \phi_{|\tau_x S}| dx \right)^2 \end{aligned}$$

This enforces the conservation of the norm of the deformed shape. Yet in practice, it can lead to local deformations of one part of the shape compensated with another part (Figure 3.18). Therefore, we introduce a *local* regularization allowing us to locally preserve the mass by enforcing the terms inside the integral to be close to 0 everywhere:

$$R_{local}(S, \Phi(S)) = \int_S \left(\omega_S(\vec{x}) - \omega_{\Phi(S)}(\Phi(\vec{x})) |d_x \phi_{|\tau_x S}| \right)^2 dx \quad \textbf{(Local)} \quad (3.14)$$

Discrete formulation Similarly to the partial matching term, we can write the regularization terms (both global and local) in the discrete setting:

$$R_{global}(S, \Phi(S)) = \left(\sum_{i=1}^{N_S} \omega_S(\vec{x}_i) - \sum_{i=1}^{N_S} \omega_{\Phi(S)}(\Phi(\vec{x}_i)) \right)^2 \quad \textbf{(Discrete Global)} \quad (3.15)$$

$$R_{local}(S, \Phi(S)) = \sum_{i=1}^{N_S} \left(\omega_S(\vec{x}_i) - \omega_{\Phi(S)}(\Phi(\vec{x}_i)) \frac{|\eta_{\Phi(\vec{x}_i)} \Phi(S)|}{|\eta_{\vec{x}_i} S|} \right)^2 \quad \textbf{(Discrete Local)} \quad (3.16)$$

The overall function to minimize in this LDDMM setting is given by the formula:

$$J_{reg}(v) = \gamma \int_0^1 \|v_t\|_V^2 dt + \underline{\Delta}(\phi_1^v(S), T) + \gamma_1 \cdot R(S, \phi_1^v(S)) \quad (3.17)$$

with $R = R_{global}$ or R_{local} .

Proposition 15. *Let $\gamma > 0$ and $\gamma_1 > 0$ be two fixed parameters. The regularized partial matching problem, which consists in minimizing over L_V^2 the function J_{reg} (defined in Eq.3.3) has a solution.*

Similarly to the proof of Proposition. 13, the proof boils down to showing that the mapping $v \mapsto A(v) = \underline{\Delta}(\phi_1^v(S), T) + R_{local}(S, \phi_1^v(S))$ is weakly continuous on L_V^2 . We use the same notations, and define (v_n) a sequence in L_V^2 , weakly converging to some $v \in L_V^2$. We

only need to show that $R_{local}(S, \phi_1^{v_n}(S)) \longrightarrow R_{local}(S, \phi_1^v(S))$.

To simplify we denote $S_n = \phi_1^{v_n}(S)$, $S_* = \phi_1^v(S)$ and for any $\vec{x} \in \mathbb{R}^d \times \mathbb{S}^{d-1}$, $f_n(\vec{x}) = \omega_{S_n}(\vec{x}) - \int_T \min_{\epsilon} \left(1, \frac{\omega_{S_n}(\vec{x})}{\omega_T(\vec{y})} \right) k(\vec{x}, \vec{y}) d\vec{y}$ and $f_*(\vec{x})$ likewise for S_* .

Using the same arguments as in 13, we have that $d_x \phi^n$ converge to $d_x \phi$, uniformly on $x \in S$ [Gla05]. In addition, ω_{S_n} converges uniformly to ω_{S_*} . We can deduce that

$$R_{local}(S, \phi_1^{v_n}(S)) - R_{local}(S, \phi_1^v(S)) \longrightarrow 0. \quad \square$$

We now apply the proposed regularized partial matching with diffeomorphic deformations to the toy example proposed along this chapter. It is illustrated in Figure 3.15 in the case of a multiscale approach with the regularization parameter $\gamma = 0.1$.




Figure 3.15: LDDMM registration with $\gamma = 0.01$, $\sigma_V = 0.5$ under partial matching using the cost function of Equation 3.3 with local regularization term. (a) LDDMM registration with $\sigma_W = 0.5$. (b) Multiscale LDDMM registration with $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. Note that the output of figure (a) is the input for the scale $\sigma_W = 0.25$ of figure (b).

We observe in this experiment the correct alignment of the source shape onto its corresponding subset in the target. In particular, there is little shrinkage effect when compared to the (unregularized) partial matching with the same γ parameter of Figure 3.14.

Influence of the Regularization Term

We illustrate on the segment and the circle examples the influence of the proposed regularization term on the partial matching. In the left column of Figure 3.16 we see that without regularization, the registration ends up in a local minima and the deformed source shape is shrunken. On the contrary, the proposed regularization on the right column tends to preserve the size of the deformed source.




Figure 3.16: Multiscale LDDMM registration with $\gamma = 0.1$, $\sigma_V = 0.5$ and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. (a) Applied to a blue arc onto a red circle, with $\underline{\Delta}$. (b) Applied to a blue arc onto a red circle, with $\underline{\Delta} + R_{local}$. (c) Applied to a blue segment onto two red segments, $\underline{\Delta}$. (d) Applied to a blue segment onto two red segments, $\underline{\Delta} + R_{local}$.

In the rest of the manuscript, the term we selected for the partial matching is $\underline{\Delta} + \gamma_1 R_{local}$, with $\gamma_1 = 1$ in most of the experiments.

3.4 Examples of Applications

We have seen how we constructed the partial matching fidelity term for diffeomorphic matching, and a regularizing term both derived from the oriented varifolds representation. In order to illustrate different possible applications, we now provide some examples of diffeomorphic registration of a truncated source shape onto the subset of a complete target. We apply this to

a simplified pelvic vascular tree registered onto a real one, to the registration of a truncated sphere onto a complete one, and finally to the converse case when one wants to register a complete source onto a truncated one.

Similarly to Chapter 2, we model the non-rigid deformations using a reproducing kernel K_V of V as be a sum of Gaussian kernels. For each set of experiments we use the same hyperparameters $(\sigma_0, \sigma_W, \gamma, \gamma_1)$ to compare the influence of the regularization and for the clinical application.

Simplified Template Aligned to Real Trees As we’ve shown in Figure 2.13, the classic distance in the space of oriented varifolds is not adapted to the registration of a simplified template vascular tree onto a complete real one. To illustrate the performance of the proposed partial matching applied to the registration of the template vascular tree described in Section 5.1.3, we register it with diffeomorphic deformations onto real vascular trees that have not been processed. If we do not process the trees, we still suppose that the target trees contain all the labels of the template. The registration is initialized by a translation, aligning the root of the template with the one of the target.

We see in Figure 3.17 that the non-rigid deformations guided by partial matching with the local regularization term allow to correctly include the template tree into the target. Note that if there is no guarantee that the registration provides anatomically consistent registration, the result in this example is quite satisfying. The relevant template and atlas construction of Section 5.1.3 provides a relevant initialization of the registration. When associated with partial matching and local regularization, it results in the embedding of the deformed source in the target with smooth and relevant deformations without abnormal shrinking. A more detailed analysis of the registration of this template onto the database is developed in Chapter 5 Section 5.2.

Application to Surfaces We illustrate the influence of the local and global regularization terms with the registration of a truncated surface onto a complete one. To do so we perform a LDDMM registration using a small regularization parameter $\gamma = 10$ in the functional J_{reg} and we set $\gamma_1 = 1$. The data attachment term we use is the one proposed for the partial matching in Section 3.2.2. In Figure. 3.18, we illustrate the result of the LDDMM registration guided by the partial matching data attachment term using alternatively no regularization, the proposed global regularization and the proposed local regularization. The ideal deformation should be a translation of 33 millimeters.

We observe that without any regularization, the non-rigid deformations lead to global shrinkage of the source shape. On the contrary, the proposed regularizations both global and local prevent such shrinkage. The global one though preserves the norm of the shape at the end of the diffeomorphic deformation and does not prevent inconsistent local deformations. The most regular deformation is thus induced by the local regularization, which allows preserving the position of the relative position of the vertices and the relative orientation of the tangent vectors associated with the deformed source. It should be noted that we work at a small scale of attachment to the data and that this influences the regularization.




Figure 3.17: LDDMM registration of a template tree (a) onto a real case (b). The deformations are guided by partial matching with local regularization. In (c) is plotted the deformed template (thick tree) into the real target in transparent.

Including the target in the Deformed Source Another framework of interest is the inclusion of a truncated target into a deformed source. The application of such a framework is the exact opposite of the previous setting, and one can consider that the template is a complete shape and that the data may have missing subsets and be truncated. With the same toy example, we illustrate in Figure 3.19 the behaviour of the LDDMM registration of the complete shape onto the truncated one. In Figure (a) we see that without the proposed regularization (and a small $\gamma = 0.1$), the diffeomorphisms induce few deformations of the supplementary branch. The target is yet correctly matched. In Figure 3.19 (b) the local regularization is used in addition. In this second experiment, the supplementary branch is deformed ensuring the local conservation of $\omega_S(\vec{x})$ for $\vec{x} \in S$. The deformation is smooth: the resulting shape in fact covers the target truncated one, and the deformations are global. This experiment shows that the partial dissimilarity function we introduced in this chapter is adapted to both inclusion frameworks: the truncated source, or the truncated target. Note that compared to Figure 3.14 regularized with the σ_V scale, the deformation here is more local, indicating that the proposed regularization term allows a different degree of freedom for the deformation.




Figure 3.18: Influence of the regularization term on the non-rigid deformation of a truncated sphere onto a complete one. (a): Source (blue, opaque) and target (red, transparent) surfaces. (b-c-d): registration results. The colormap for (b-c-d) indicates the euclidean distance (in mm) of the points to their initial position before diffeomorphic deformations.

3.5 Conclusion

We propose a new way to easily encode partial matching in the space of varifolds (both oriented and non-oriented) adapted to the LDDMM framework and allowing for the registration of a shape onto a subset of a target. Dismantling the classic distance in the space of varifolds provides the tools we need to build such functions. The solutions however require careful attention in order to build consistent deformations and handle consistently the imbalance between the deformed shape, simpler, and the target one supposed to include the deformed source. Constructing a term that is no longer a distance, and reaching its minimum when the distorted source is included in the target, produces many possible minima for the partial matching function. The first and most important point for solving such problems is then to correctly initialize the minimization. A good starting point will indeed prevent numerous bad local minima while starting far from the target position will increase the risk of falling into bad minima. A second solution is to regularize the deformation, creating a bias guiding it to some realistic solution regarding the application. We chose to encode this regularization as a regularization term comparing the deformed shape to its initial position, constraining the deformed shape to locally resemble its original one. This choice has been motivated by the




Figure 3.19: LDDMM registration under partial matching using $\underline{\Delta}$ as data attachment term with $\gamma = 0.1$, $\gamma_1 = 1$ and $\sigma_V = 0.5$ and $\sigma_W \in [\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}]$. (a) LDDMM registration without regularization. (b) LDDMM registration with local regularization.

light computational cost of this term when compared to the cost of encoding more complex deformations such as the ones generated with divergence-free vector fields.

We apply the proposed partial matching with LDDMM deformations to toy examples of both curves of \mathbb{R}^2 and \mathbb{R}^3 , and to surfaces of \mathbb{R}^3 taking advantage of the adapted oriented varifold representation. Note that if no orientation is available, it is still possible to work with the non-oriented version. This new framework is then applied to the registration of the template constructed in Chapter 2 Section 2.5.1 onto real pelvic vascular trees automatically simplified to derive a fully automated registration pipeline in which the target trees are neither labeled nor manually processed.

We also show that the proposed partial matching with local regularization is adapted to the registration of a complete shape onto a truncated one. One of the advantages of registering a complete template onto truncated targets is the possibility to then find correspondences between subset of multiple truncated shapes, using the intermediate complete template.

Clinical applications, particularly in interventional radiology, could therefore benefit from the proposed realignment. Our partial matching term coupled with regularization can be used with different registration models, and it is adapted to curves and surfaces. It can therefore be used on many structures from medical images. The approach we propose is then particularly suited to the comparison of structures extracted from various images acquired with the limited field of view of the C-arm system.

In Chapter 5 we apply the proposed partial matching framework to typical use cases in interventional radiology: trees and surfaces extracted from interventional images. In the first application (Section 5.2) we study the anatomical consistency of the registration of the

simplified template pelvic vascular tree that we built in 5.1.3 onto real cases that are not pre-processed. The second application (Section 5.3) is dedicated to the partial correspondence between the livers extracted from the CBCT volumes, and the ones extracted from the CT volumes. Indeed, CBCT acquisitions have a field of view of limited size causing the portion of the imaged liver to be cropped in the reconstructed volume. To compare efficiently the volumes of both modalities, the CBCT acquired during the intervention and the CT used as pre-operative acquisition for diagnosis or procedure planning. We use the proposed LD-DMM partial matching framework to find a realistic non-rigid deformation between the liver surfaces automatically extracted from both modalities. The diffeomorphisms obtained can then be applied to the whole volume, which provides a tool for volume comparison and to take the best out of the two modalities.

Next Steps

The functions proposed in this chapter to encode the partial matching in the data attachment term gives promising results, however, there are some points that could benefit from further exploration. Indeed, we have the non-negativity of $\underline{\Delta}(S, T)$ when $S \subseteq T$ which is problematic and favors the shrinkage of the source under the action of a non-rigid deformation. For this, an interesting approach would be to manage the normalization differently by compensating for $k(\vec{x}, \vec{y}_{max})$ with $\vec{y}_{max} = \arg \max_{\vec{y} \in T} (k(\vec{x}, \vec{y}_{max}))$ instead of the integral over the whole target $\omega_T(y)$.

In the spirit of further investigating partial matching as a data attachment term, a natural extension would be to use the representation of normal cycles [RG16] which would allow to better take into account the boundaries of objects such as curves and surfaces. This shape space is therefore even more challenging since truncation creates artificial edges that one does not necessarily want to align with the target. In addition, in this study, we supposed that there was an inclusion of one shape into the other, hence injective registrations. Following the approach of [BBK06] one could search instead for an optimal trade-off between the size of the subsets of the shapes matched in both the source and the target and the amount of deformations required to register these selected sub-parts. This would be a solution to the problem of missing artery in the target, meaning that one is aiming to register a subset of the template onto a subset of the target.

Changes in the Bifurcation Ordering The partial matching allows dealing with the truncation of the template tree for non-rigid alignment onto a subset of the target. The diffeomorphisms however cannot explain the topological changes in general. As listed in Section 1.2.2, another kind of typical topological change that occurs between the trees is the variation in terms of bifurcations order which must be encoded in the deformation model if one aims at perfectly aligning the shapes. We illustrate in Figure 3.20 one registration example with partial matching in which the target arteries' connectivity differs from the template one.

This example illustrates that despite a consistent global registration, the alignment of the source shape onto the target suffers from the incapacity of the diffeomorphisms to handle the




Figure 3.20: LDDMM registration of a template tree (a) onto a real case (c). One change in the bifurcation ordering is illustrated with the obturator, highlighted in red for the source (b) and the target (d). The resulting registration is globally consistent, the proximal branches however cannot be correctly matched without further deformation of the template.

changes in terms of bifurcations ordering: the obturator branch close to the root cannot be matched to the target associated obturator branch. One way to cope with this limitation is to less regularize the deformations, which could lead to strong local and inconsistent deformations. In the next chapter, we embed the vascular tree in a new space called space of tree-like shapes introduced by [Fer+11], in which the position of the tree is encoded by the geometry of the branches as well as their connectivity. Moving in this space can provide topological changes in terms of bifurcations ordering, and we associate this tree representation with the registration framework presented so far to propose a tree registration method that allows both diffeomorphic deformations and tree topological changes.

REGISTRATION OF TREE-LIKE SHAPES

Let's picture someone driving every day from point A to point B, and who has a map to guide him. Let's also imagine that every night the road changes, junctions are switched and some roads may appear or disappear. The map on the other hand does not change. This is typically the situation of the interventional radiologist during endovascular procedures, who has a vascular tree model for a given anatomy, but who navigates his tools in a new tree for each new patient.

| | | |
|-------|--|-----|
| 4.1 | Building a Space of Tree-like Shapes | 116 |
| 4.1.1 | A Tree from Leaves to Root | 116 |
| 4.1.2 | Definition of a Tree-Like Shape | 120 |
| 4.1.3 | The Space of Tree-like Shapes | 121 |
| 4.2 | Computing Geodesics in the Space of Tree-like Shapes | 126 |
| 4.2.1 | Characterization of the Geodesics | 127 |
| 4.2.2 | Geodesic Algorithm in a Space of Tree-Like Shapes | 130 |
| 4.3 | Implementation | 132 |
| 4.3.1 | Estimation of the Fréchet Mean | 132 |
| 4.3.2 | Principal Geodesic analysis in the Space of Tree-Like Shapes | 133 |
| 4.4 | A Registration Problem for Deforming Tree-like Shapes | 136 |
| 4.4.1 | A Link between Tree-like Shapes and Vascular Trees | 137 |
| 4.4.2 | Registering Vascular Trees with Hierarchical Changes | 139 |
| 4.4.3 | Dealing with the Change of Orthant | 141 |
| 4.4.4 | Experiments | 144 |
| 4.5 | Conclusion | 149 |

In the previous chapters, we have studied spatial alignment of tree-like shapes such as the vascular trees, and the construction of atlas and templates from data sets. We have also studied the essential challenge of registering a shape onto a subset of a target and showed how to encode it in the space of Varifolds. However, as discussed in Chapter 1, the variability of the vascular trees is not just partial correspondence of the shapes but encompasses other structural and topological changes, in particular when dealing with annotated trees. In this last case, the appearance order of the bifurcations induce a *tree topology* and we call the changes in term of bifurcations ordering from one individual to another the *hierarchical changes*. The application taking these changes into account are numerous, ranging from the study of trees and plants [Wan+18] and phylogenetic trees in biology for the study of the evolution between the species [BHV01] to anatomical trees analysis [Fer+12a], brains connectivity [DDV17], and neurons [DKS18].

In the context of endovascular interventions in interventional radiology being able to take these topological changes into account is decisive. In fact, following a map and navigating tools must be done carefully, just like while driving on a road, and the physicians navigating in the patients' vessels must not make a mistake in the path they chose. It is even more important since they are injecting treatment through the vasculature of the patient. In the pelvic anatomy, there is no uniform map of the vessels, and the bifurcation order can change from one patient to another. The risks range from prolonged intervention to destruction of healthy tissue. We are therefore interested in a method allowing us to build an atlas of the patients' vasculature and encoding the changes in the bifurcations order that we call hierarchical changes. To that end, we will embed our trees in a space of tree-like shapes adapted to the hierarchical changes, and use this template representation in the registration framework.

A State of the Art

The hierarchical changes in the trees are a well-known problem in tree-structured data analysis. Their management is particularly important in three applications of interest: the annotation of the branches, the alignment or registration of one tree shape onto the other, and the statistical analysis of sets of trees with varying topologies. No matter how the authors address the hierarchical changes in their applications, it can be separated into two categories, either by treating such changes in a post-processing step or in the main step of the algorithm. The first approach takes advantage of some rules derived from a priori knowledge or observations of the database. Classical algorithms can then be used in the applications, and their outputs are used as input for the post-processing step. The management of the hierarchical changes is extrinsic to the annotated tree shape models. The second approach on the contrary takes advantage of the encoding of the hierarchical changes in the model to build algorithms adapted to the annotated tree structures. Such intrinsic models are appealing but often require a combinatorial algorithm to fully handle the hierarchical changes.

In the following, we discuss the way these hierarchical changes are managed in the three applications we identified: annotation, alignment, and statistics applied to tree structures.

Hierarchical Changes and Statistics

When studying a population of shapes, one classic approach is to compute statistics over the observations. The classic statistics are done in euclidean spaces of finite dimensions. More refined methods consider sets of points living on a manifold. In the case of trees there are different representations due to the hierarchical structure of the shapes and in particular stratified spaces and tropical ones.

Historically, these statistics in stratified spaces were intended to answer questions by biologists and statisticians regarding the space of phylogenetic trees. They model hierarchical relationships between individuals, for example, the evolutionary relationship between species or the speciation events across the evolution. The edge lengths would then represent the degree of evolutionary divergence between the species. The geometry introduced by the seminal work of [BHV01] induces a space with non-positive curvature for which there is a unique shortest path between any two points. In addition, there is a centroid for any given finite set of trees in the treespace. Their work inspired many further contributions in this space also referred to as BHV -Billera-Holmes-Vogtmann- tree space after their authors.

In 2011, an algorithm introduced by [OP11] allowed to compute in polynomial time the geodesics between any two points in the treespace. Since then, several statistical methods have been proposed to study sets of points in this space, relying most of the time on the BHV geometry. Algorithms have been developed for computing sample means [MOP15] when others have developed Principal Component Analysis (PCA) in this space [Nye14; Fer+12b; Nye+17] and confidence region for the population mean [Wil19]. Such methods allow not only to perform statistical analysis of the data despite the topological changes but also to set up data augmentation with respect to the statistics. One of the drawbacks of all these methods was the necessity to pair leaves between any two trees of the datasets.

More recently in [GO19] the authors adapt the problem to a partial set of leaves by extending the space of possible trees. This setting matches the topological differences one can encounter when dealing with prostatic vascular trees. In [FN20] the authors open new perspectives by proposing to study a tree space without annotation, however for both methods the applications are limited and still under development, with no tractable algorithm to compute statistics in such spaces.

Differently, in [Gar+21], the BHV space is embedded into a space of centered probability distributions with positive semi-definite covariance matrices. The latter space seems to be a promising lead to derive more advanced methods and transpose them to the BHV treespace. It allows the study of trees with inner edges of infinite length, denoting disconnected sets of leaves. However, it lacks theoretical tools such as the computation of the centroid of a set.

For the statistical analysis of populations of annotated trees, many tools allow taking the hierarchical changes into account by embedding the trees into spaces of representation encoding the hierarchical changes. However, it is important to keep in mind that some of the embedding spaces do not provide methods to get back to the original tree representation, in particular in the case of vascular trees.

Hierarchical Changes and Automatic Annotation

We observe three different approaches to hierarchical change management in automatic annotation applications. First, the works that do not take these changes into account for the label probability estimation, staying therefore at the point or branch level and using a set of rules to assign the labels. Second, the approaches estimating relational features that feed a learning method or enrich the atlas representation are used for the label probability estimation. Such an approach is similar to the one we used for XGBoost classifier in the first application described in Section 5.1. In most of the papers, the label assignment still follows afterward sets of rules learned from the database or a priori knowledge.

In the last kind of approaches the hierarchical changes are encoded into the distance used to compare the trees. It allows atlas-based annotation methods, and the distance is derived from the tree space representation. Note that some papers cited in the following may work on graphs instead of trees, the discussion could be extended to these structures but we only cite them as other hierarchically structured shapes, on which automatic methods for annotation are developed with topological considerations in the label probability computation and assignment.

Hierarchical Changes as Post Processing After computing probabilities of labels assignment to the branches, the post-processing approaches seek to find ways to discriminate between all possible annotations of each branch. The most direct and task-adapted approaches build sets of rules learned from observations of a database or a priori knowledge of the data on which one works [GBR08; Aki+09]. These rules then allow deciding between the different choices of annotation, considering or not the labels that have already been assigned in the tree. In [Mat+14], the assignment follows a global optimization method using a maximum likelihood conditioned by the parent's labels and a set of rules specific to hepatic portal anatomy. More recently in [Rob+16], the vessels are simultaneously segmented and annotated. The label assignment is done on edges pairs, and the probabilities are computed through an extremely randomized tree classifier. The training phase learns hard rules on the label assignment that cannot occur during the annotation: which successions of two edges pairs labels (so the labels of 3 edges) do not exist in the database, which labels do not belong to ending edges pairs and which edge pairs labels do not exist in the bifurcations. This reduces the space for possible transitions. The advantage of using explicit rules is that it's easy to implement and the output of the method is also easy to interpret. One limit, however, to learn the set of rules from the database is the lack of adaptability.

Hierarchical Changes and Features Instead of building explicit rules, relational features between the nodes or branches of the trees can be extracted. These features then provide relational information between the edges and do not require any label information. In [Tsc+05] authors build a template in the space of features using population average. Use relational features such as topological distance (in the tree) between the segments and ratio features to the parent, children, and siblings. Similarly, [GBR08] uses relational features to estimate distributions from the database and assign the labels to each branch independently. Nevertheless, the lack of topological consideration in their labeling process seems to be (according

to the authors) one of the main limitations of their results. An intermediate approach for integrating hierarchical changes to the prediction models is to learn labels probability transitions by considering features of sets of branches and labels (e.g. pairs) and predict the labels probabilities on the sets of features in the new trees to annotate. This way the topological information is encoded in the features space and provides a first label probability based on topological information.

Both previous solutions complement probability predictions with a set of hard rules implemented in the labels assignment procedure. Such an approach is also proposed in [GLS14], in which relational features as the relative diameter with the neighbors as well as the angle they form are used in the feature space to predict labels probabilities. They also compute a "connectivity matrix" for the labels, built on the training set and encoding the probability of adjacency of each pair of labels. More recently in [Wan+17] the bifurcations are labeled according to their features extracted from Circle of Willis (graph structure), seen as sequences of observations supposed to only depend on their direct neighbors. The label assignment problem is formulated with a Hidden Markov Model (HMM) with a restricted transition strategy. The transition weights are learned from the training set and combined with the probability output of an XGBoost classifier, working therefore on both feature and hierarchical knowledge extracted from the database.

Deep learning is the most standard learning method used today for image analysis, it has also been used in the context of vascular tree annotation. The hierarchical changes are implicitly learned in tree modules like Long Short Term Memory (LSTM) networks [Wu+19]. In this work, the authors propose to train a recurrent neural network preserving the topology of the coronary tree on a database, and the network is automatically learning the transitions probabilities between the labels without the need for an explicit set of rules. In [Yan+20] the trees do not need to be binary, and the hierarchical transition conditions are learned with a similar bidirectional tree LSTM, which takes as input information extracted from 3D CT images with a 3D CNN. Then the output of the bidirectional Tree LSTM is used as image domain conditional weights for a (partial residual) Graph Convolution Network (GCN) that also works on features extracted from the centerlines. This combination of information from the image domain and the centerline representation of the curves brings, according to the authors, robustness to missing branches as well as hierarchical changes. In [Zho20] the output of the network is used as a likelihood for posterior estimation. The prior comes from the rules learned in the database. Similarly authors in [Ham+21] propose an attention mechanism in a Graph Convolutional Network on features extracted from the branches of the trees and concatenated with the direct neighbors' ones. The tree representation is the dual graph of the tree (the edges become the nodes, so instead of bifurcations we have 3-nodes cycles).

In all the aforementioned methods, the management of the labels ordering is done extrinsically, either by learning labels transitions from the database or by adding hard rules during the label assignment procedure. None of these methods takes advantage of the intrinsic shape representation induced by the labeling of the tree.

Hierarchical Changes Encoded in the Distance between Trees As mentioned earlier, the intrinsic approaches embed the annotated trees into space to compare this new point to one or several other points in the embedding space. This provides a natural template-based

method, in which the template is a reference point of the space.

In [Fer+12a] authors have developed a tree-space based method to automatically annotate pulmonary trees using the distance between one annotation of the target tree and a template tree. The annotation solution selected is the one minimizing the distance of the annotated tree to the template. To reduce the computational cost, the annotation follows a hierarchical scheme where a subtree of a certain depth is extracted at each step, and every labeling is tested at each iteration. The method selects the annotation minimizing the geodesic distance to the atlas' corresponding subtree. This solution is one way to cope with the limitation of using a single atlas but requires testing all possible annotations.

The approach is extended in [Fer+15] to the use of several templates, namely all annotated cases from the database are used. At each iteration, the distances to all corresponding subtrees in the annotated database are computed. Once again the labeling minimizing the sum of these distances is selected afterward. It has the advantage of considering all available annotated trees. On the other hand, it tends to label all the trees, like the median of the set. It also needs to compute all possible annotations in each selected subtree in addition to the distance to every atlas. These methods of testing all possible solutions may be robust to anatomical variants, yet they are computationally expensive.

Similarly, in [Gül+14] the authors use an atlas of the coronary tree defined by experts and build a two-step annotation taking advantage of some assumptions on coronary trees. As in [Fer+12a], the distances are computed in the tree space. A second cost is used, based on the likelihood of the assigned labels. This second cost introduces a new a priori corresponding to the spatial distribution of the arteries relative to the heart surface.

More recently, in [Cao+17] the atlas is built by selecting the best example in a training set following a leave-one-in cross-validation design. To label a new tree, a rigid point-set registration is done, and the labels are assigned according to a 2-step hierarchical algorithm. A direct label assignment to the main branches is performed first. Then, labeling the distal part is performed by trying all possible pairings between the atlas' and the target's branches as in [Fer+12a], although by staying in \mathbb{R}^3 the authors ensure faster results.

As we have seen, many automatic annotations of tree structures address the problem of changes in the bifurcation ordering of the trees. Few of them however try to add some spatial deformation to their algorithms.

Hierarchical Changes and Registration

In terms of tree shape registration and topological changes, few papers address this problem using classic deformations. Some authors [Cao+17; CFV20] consider rigid registration of the trees or graphs: the resulting deformation is robust to topological changes. It is yet limited in terms of spatial shape matching, no precise registration of the underlying structures such as the branches is available. In [BMV13] the authors use the iterative closest curve method, and the deformations applied to the trees are both translations and rotations of \mathbb{R}^3 and 3D/2D projections.

Unknown Sets of Leaves In a tree space representation when no labeling of the leaves is available, or when the tree structures have a different number of nodes and edges, correspondences must be computed between the leaves to compare the subtrees spanned by the matched edges. In this setting statistical shape analysis of the shapes – and their deformations – can then be performed as in [FN20]. For this purpose, a tree space is considered to represent real unlabeled trees with points along 3D curves, and this has been used in [Fer+15; Gül+14; Wan+18; DKS18].

In [Fer+12a; Fer+15; Gül+14; Wan+18] the authors are looking for a mapping of the leaves by testing all the possible labeling following a coarse to the fine scheme. More recently in [DKS18], statistical shape analysis was done on neural trees following a three step method: first, the trees are aligned with translations and rotations, then a hierarchical framework is proposed by iteratively matching the trees represented with increasing depth, hence increasing complexity. The branch matching is found using the Hungarian algorithm with a transport cost formulated as a distance between square root velocity functions adapted to trees. Finally, the trees are interpolated in tree space, providing a spatial alignment of the shapes as well as topological changes.

We saw in Section 2.6 that diffeomorphic alignment of tree shapes can be computed. Such deformations provide efficient geometrical alignment of the structures but are sensitive to topological changes: risks of shrinking, inconsistent jumps of a deformed branch from one branch of the target onto another. On the contrary, to the tree space representation, this approach does not require explicit correspondences between the shapes. One limitation of such an approach is that the interpolation between the trees along the geodesic path results in interpolation between the branch coordinates in the space of representation which may not be the most realistic deformation of vascular trees.

It has been shown, however, that optimal transport plan could be efficiently computed [Cut13], adapted to diffeomorphic registration and scaled up to large point clouds (the points potentially belong to $\mathbb{R}^{n \times d}$) [Fey+17; Fey+19]. These results pave the way to a registration framework in which the atlas would provide different topologies for the template tree. The latter could then have a hierarchical structure, in the sense of bifurcations ordering, varying during the registration procedure. From the registration point of view, the target tree is the target shape to which we want to align the template, and just like the LDDMM registration framework, it does not need any labeling. Statistics on tree populations could then be derived with topological changes within the atlas.

Chapter Organization The organization of this chapter is the following: we first redefine the tree-like shapes and their underlying structure. We also provide the grounding properties of the space of tree-like shapes, allowing to compute geodesics between shapes and adapted distance between the trees. We build a tree-space from a set of possible topologies, in which we can minimize a cost function designed to compare shapes in \mathbb{R}^d . Second, we discuss the construction of the geodesics between the tree-like shapes following the algorithm proposed in [OP11]. Third, we introduce a new formulation for tree-like shapes representation providing a differentiable function to go from this first space to the spatial representation of vascular trees using Bézier curves. The shapes registration tools are then available for such vascular trees, and we can compute LDDMM registration as described in Chapter 2 and

update the position in the space of tree-like shapes through gradient back-propagation. We propose a python implementation of this registration as well as a python implementation of the tree-like shapes and the statistical tools available in such a space. Fourth and last, we discuss the potential applications and study that are yet to be explored with this new registration framework.

4.1 Building a Space of Tree-like Shapes

As observed in the literature, to perform realistic deformations of tree-like structures while allowing topological changes, we need to build a suitable representation space. Tree spaces have been built specifically to handle such cases, but they require to precisely define the choices of metric and representation of nodes and edges.

4.1.1 A Tree from Leaves to Root

We now recall the definitions of a tree in order to further build tree-like shapes that will guide the construction of the tree space throughout this section, and the registration methods in the next sections.

Definition 13. *A tree is a connected acyclic graph denoted (\mathcal{V}, Σ) with \mathcal{V} the set of nodes $\{0, \dots, M-1\}$ and Σ the $M \times M$ connectivity matrix with $\Sigma_{i,j} = 1$ if i is connected to j and 0 otherwise.*

This basic representation of a tree already provide the information regarding the **edges** of this tree, derived from the connectivity matrix:

Definition 14. *Given a tree, we define an edge as any couple (i, j) such that $\Sigma_{i,j} = 1$. The set of edges in a tree is written $b = \{b_k\}_{k=0, \dots, K}$.*

Remark 2. *A tree can be written $T = (\mathcal{V}, b)$, and one can retrieve the adjacency matrix Σ from b .*

The degree of a node is the number of nodes it is connected to, and the leaves of the tree are nodes of degree 1.

Definition 15. *Let $L = \{0, \dots, N\}$ be a set of leaves. A leaf-labeled tree on L is a tree $T = (\mathcal{V}, \Sigma)$ endowed with a bijective labeling $\mathcal{L} : L \rightarrow \mathcal{V}_1$, with \mathcal{V}_1 the degree 1 nodes of T .*

For now on, all trees we consider will be supposed to be *resolved*, meaning that they don't have degree 2 nodes. The degree 3 nodes correspond to bifurcations in the tree. As a convention, if $N+1$ is the number of leaves, we always associate the nodes $(0, \dots, N)$ to these degree-1 nodes in the tree. It does not change the generality of our tree definition, since we can always reorder the nodes.

Definition 16. *A fully resolved tree is only composed of degree 1 and 3 nodes.*

A fully resolved tree – also called binary tree – can always be constructed from trees with degree 4 nodes or higher: such nodes can be split into degree-3 ones with new intermediate edges. For instance a degree 4 node can be split into three different successions of two degree 3 nodes.

Remark 3. *The vascular trees we build from CBCT volumes (see Chapter 1 Section 1.2) are fully resolved.*

In such fully resolved trees, it is possible to count the number of edges and nodes given the number of leaves:

Proposition 16. *In a fully resolved tree with $N + 1$ leaves, there are exactly $2N$ nodes and $2N - 1$ edges.*

This can be simply deduced by fixing one leaf, and removing all the other leaves and their associated edge (a single one by definition). The remaining parents nodes become the new degree-1 nodes in the tree. One can repeat this operation until the only node remaining is the fixed one. Thus $M - 1$ edges in a fully resolved tree composed of M nodes. In addition if the number of leaves in the tree is $N + 1$, the total number of nodes M in the fully resolved tree is exactly $M = 2N$. It can be shown by recurrence by seeing that, if $\kappa(N)$ is the number of nodes in a tree with N leaves, $\kappa(1) = 0$, $\kappa(2) = 0$ and $\kappa(3) = 1$. By induction, if it is true for a given $N \geq 3$, we can see that for a tree with $N + 1$ leaves, by removing two siblings leaves we obtain a fully resolved tree with N leaves so $2N - 2$ nodes. Hence $2N = 2(N + 1) - 2$ nodes in the complete tree.

Definition 17. *A rooted tree is a tree for which we fix a reference leaf. As a convention we set the root to the node 0.*

Again this can be done without loss of generality since we can always build a rooted tree \tilde{T} with the root set to another node by permuting the nodes of T .

Proposition 17. *A rooted tree defines a natural orientation given by the paths connecting the root to the other leaves.*

In the case of vascular trees, this natural orientation could be associated with the blood stream in the vessels. We always suppose that the root is known in these trees, and that it corresponds to the point of injection of contrast medium.

Characterizing the Tree Topology

The generic Definition 13 provides the information regarding the branching structure of the tree, namely its *topology*:

Definition 18. *The topology of a tree $T = (\mathcal{V}, b)$ is the connectivity between its nodes, given by b (or Σ).*

In line with the seminal work on the subject [BHV01], we are working with a space composed of euclidean subspaces corresponding to unique topologies that are glued together. We are working with leaf-labeled trees, to characterize the topology with the connectivity between the leaves instead of all the nodes, which will be convenient for the construction of the space of tree-like shapes, we first introduce the splits:

Definition 19 (Split). *Let T be a tree and $L = \{0, \dots, N\}$ be the set of its leaves. A split $L_1|L_2 = L_1|L_1^C$ is a partition of L .*

Let $T = (\mathcal{V}, b)$ be a leaf-labeled tree, and consider an edge b_i in b . By removing the edge b_i we create two disconnected subtrees T_1 and T_2 . The split associated to an edge b_i in T is $L_1|L_2$ where L_1 (resp L_2) is the set of leaves of T_1 (resp. T_2).

Definition 20 (Compatible Splits). *Two splits $L_1|L_1^C$ and $L_2|L_2^C$ are said compatible if one of the sets : $L_1 \cap L_2$, $L_1 \cap L_2^C$, $L_1^C \cap L_2$ and $L_1^C \cap L_2^C$ is empty.*

In Figure 4.1, the two splits $\{1\}|\{0, 2, 3\}$ and $\{2\}|\{0, 1, 3\}$ are compatible. One can build incompatible splits by permuting two leaves between the complementary sets of leaves: take for instance $\{0, 1\}|\{2, 3\}$ and $\{0, 3\}|\{1, 2\}$ encoding edges that cannot exist in the same leaf-labeled tree in the example in Figure 4.1.

Remark 4. *In a given leaf-labeled tree $T = (\mathcal{V}, b)$ all the splits associated to edges in T are pairwise compatible.*

This split representation is the corner stone of the construction of tree spaces and computation of geodesics between the trees that we detail in Section 4.2. From the tree point of view, two splits $L_1|L_1^C$ and $L_2|L_2^C$ are compatible if and only if there exists a tree containing the edges associated to them.

Definition 21 (Common Edges). *Let $T = (\mathcal{V}, b)$ and $T' = (\mathcal{V}', b')$ be two trees and $L = \{0, \dots, N\}$ be the set of their leaves. Two edges $b_i \in b$ and $b'_j \in b'$ are called common edges if they are associated with the same partitions of the leaves in T and T' .*

There exists a characterization of the topologies through the description of the splits, provided in [Bun71]:

Theorem 9 (Splits Equivalence). *Let $L = \{0, 1, \dots, N\}$ be a set of leaves. Any set of pairwise compatible splits \mathcal{S} which contains the splits $\{0\}|\{0\}^C, \dots, \{N\}|\{N\}^C$ determines an unweighted (leaf-labeled) tree $T = (\mathcal{V}, b)$ on L such that \mathcal{S} is the set of splits associated with all the edges in T .*

In particular the tree determined by the minimal set of pairwise compatible splits

$$\{(0)|(0)^C, \dots, (N)|(N)^C\}$$

has a single node of degree $N+1$ and $N+1$ degree-1 nodes. Adding new pairwise compatible split to this initial set consists in adding a node of degree strictly greater than 2.




Figure 4.1: Examples of splits for the set of leaves $L = \{0, 1, 2, 3\}$.

(a)-(b): The two splits $\{1\}|\{0, 2, 3\}$ and $\{0, 1, 2\}|\{3\}$ are compatible with all the other splits (the same goes for $\{0, 1, 3\}|\{2\}$).

(c)-(d): The two splits $\{1, 2\}|\{0, 3\}$ and $\{0, 1\}|\{2, 3\}$ are not compatible.

Remark 5. *Theorem. 9 provides an alternative definition of the tree topology that does not use the inner nodes of degree 3 or more.*

We are now able to build a tree with a set of leaves and associated pairwise compatible splits containing the minimal set. On the contrary, given a leaf-labeled tree T we can deduce the set of all pairwise compatible splits associated to all the edges in T . In particular, we can remark that this new characterization of the topology allows to introduce permutations of the rows and columns of the connectivity matrix Σ in $T = (\mathcal{V}, \Sigma)$ as long as the unordered set of all pairwise compatible splits associated with T remains unchanged.

Remark 6. *In the rest of the manuscript, we will denote by $\mathcal{E} \subset b$ the set of inner edges associated with the tree $T = (\mathcal{V}, b)$.*

We illustrate in Figure 4.1 two trees represented both spatially and with the splits of the leaves.

4.1.2 Definition of a Tree-Like Shape

We seek to construct trees for which the edges have additional information beyond the nodes they connect. In the case of trees with edges that are curves in \mathbb{R}^d , this could be any representation of these curves. We define an edge representation that is fundamental in the construction of a space of tree-like shapes:

Definition 22 (Edge attributes). *Let $T = (\mathcal{V}, b)$ be a tree. Let E be a euclidean space encoding different features associated to the edges in T . We define the edge attributes as a vector $e = (e_k) \in E^b$ associated to the edges $b = \{b_k\}$ in the tree space.*

Each edge attribute e_k is then encoding the features of its corresponding edge b_k : the length, sampling points, edge parametrization... A tree-like shape is then a tree with edge attributes:

Definition 23 (Tree-Like Shape). *A tree-like shape T is a leaf-labeled tree with edge attributes written as the triplet (\mathcal{V}, b, e) .*

An example of tree-like shape with edge attributes in $\mathbb{R}^{2 \times 3}$ and $\mathbb{R}^{3 \times 50}$ is provided in Figure 4.2. We now have the definition of one tree-like shape, and we want to get to the description of the whole space containing such objects. Before getting to this construction we give one more definition on the norm of a set of edges (that we associate with the edge attributes) for a tree-like shape:

Definition 24 (Norm of a Set of Edges). *Let $T = (\mathcal{V}, b, e)$ be a tree-like shape with $e = (e_k)$ the edge attributes defined in Definition 22. Given a set of edges A in $b = \{b_k\}$, we define the norm of A :*

$$|A| = \sqrt{\sum_{b_k \in A} |e_k|^2}.$$

This norm associated with a subset of the edges will be used throughout the chapter, to compare tree-like shapes and calculate distances in the space of tree-like shapes as described in Section 4.2. The coordinates of a tree in a tree space should provide both the edge attributes and the connectivity information. When observing the trees sharing the same topology, their position in the treespace then only depends on the edge attributes:

Definition 25 (Orthant). *An orthant¹ \mathcal{O}_Σ is the euclidean space E^{M-1} in which live the edge attributes $e = (e_k)$ of a tree $T = (\mathcal{V}, b, e)$ for a fixed topology.*

In Figure 4.2 are illustrated two different tree-like shapes sharing the same topology but with different edge attributes. The use of a simple attribute such as the length allows an infinity of spatial solutions and we prefer the segment representation for illustration purpose.

¹We use the designation of orthant although in reality it corresponds to the generalization to n dimensions of the quadrant. This allows us to remain consistent with the notation of phylogenetic tree spaces which use length as an edge attribute, and thus work in actual orthants.

Remark 7. *In the case of fully resolved trees on $N + 1$ leaves, the orthant \mathcal{O}_Σ is E^{2N-1} .*

The borders of one orthant correspond to collapsed edges which, when the edge attributes live in \mathbb{R}^{md} , is a sequence of m -times the same point at the origin the associated euclidean space. In the connectivity matrix it corresponds to introducing nodes of order greater than 3 (collapsing one edge generates a trifurcation for instance). As a convention the vector of points of one edge is translated so that the position of the first point of the edge and the last one in the parent's curve coincide if the parent exists. Like authors in [FN20], we simply add a point fixed at the origin.

We are dealing with stratified spaces². In the case of space of tree-like shapes, the orthants faces are glued together with the orthants sharing edges that induce the same splits of leaves. Each edge attribute is then relative to the one of its parent edge. If the edge attributes are points, the first point of the child's attribute coincides with the last one of the parent's. We also allow an edge attribute to collapse by moving all its coordinates to 0.




Figure 4.2: Examples of tree-like shapes with edge attributes as vectors of $E = \mathbb{R}^{3 \times 50}$ (a) and one representation with edge attributes as vectors of $E = \mathbb{R}^{3 \times 2}$ (b).

4.1.3 The Space of Tree-like Shapes

Definition 26 (Space of Tree-like Shapes). *The space of tree-like shapes \mathcal{T}_N on the finite and fixed set of leaves $L = (0, \dots, N-1)$ is the space of leaf-labeled tree-like shapes $T = (\mathcal{V}, b, e)$ on L .*

Note that using the split equivalence theorem, the degree 3 nodes in $T \in \mathcal{T}_N$ do not need to be identified, since they can be constructed using the set of pairwise compatible splits defining the tree. When the orthants are $(\mathbb{R}^+)^{M-1}$, we retrieve the space of phylogenetic trees described in [BHV01] and used in a number of works on tree spaces.

²topological space decomposed into manifolds required to fit together in a certain way

Remark 8. We identify the trees in \mathcal{T}_N up to permutations of the degree 3 nodes and their associated rows and columns in the connectivity matrix Σ .

The structure of the space of Tree-like Shapes \mathcal{T}_N consists in gluing together the orthants associated to each fully resolved topology (called *maximum orthants*) and forming a connected space. We build \mathcal{T}_N upon the set of N leaves, and set the root at the leaf 0. Such a space is composed of euclidean spaces associated to each possible topology, which can be counted as the number of unique subsets (up to permutations) in the leaves $(1, \dots, N - 1)$, hence $2^{N-2} - 1$ possible splits and $(2N - 5)!! = (2N - 5) \times (2N - 7) \times \dots \times (1)$ topologies.




Figure 4.3: Moving across orthants: the tree moves toward the face of the orthant (here all the inner branches are shrunk) and the new inner branch associated to the second split grows. In order to go from T_0 to T_1 , the branch $[\nu_5, \nu_6]$ must shrink in the topology 1, associated to the leaves split $(1, 4); (2, 3)$ and grow to produce the split $(1, 2); (3, 4)$. For illustration purpose, the edge attribute associated to a branch in this figure is its length.

The orthants are arranged as follows: when the edge attributes of one edge $e_k \in E$ is 0_E the nodes this edge, then they are reduced to one node. It can be extended to any number of inner edges, the ones connecting degree 3 or more nodes. The resulting unresolved trees are then located at the intersection between three or more orthants associated with different fully resolved topologies. When shrinking a single edge of a resolved topology Σ one moves toward the codimension-1 faces of \mathcal{O}_Σ (see Figure 4.3). This results in a degree 4 node that can be resolved in 3 different ways, providing topologies Σ, Σ' and Σ'' hence two new orthants $\mathcal{O}_{\Sigma'}$ and $\mathcal{O}_{\Sigma''}$. The intersection of all the orthants of \mathcal{T}_N , called the *star tree*, is then the tree with all its internal edges collapsed: it is the leaf-labeled tree on the set of leaves with the minimal set of splits described in Theorem. 9.

Geodesics Metric Spaces

Suppose \mathbf{X} is a space in which we want to develop a statistical methodology and d a metric on \mathbf{X} . To design relevant summary statistics, we consider the paths in \mathbf{X} and their corresponding lengths.

Definition 27 (Geodesics). *A geodesic in a metric space [BH99] (\mathbf{X}, d) is defined as a path $\Gamma : [0, 1] \rightarrow \mathbf{X}$ such that for any $t, t' \in [0, 1]$ we have $d(\Gamma(t), \Gamma(t')) = |t - t'|d(\Gamma(0), \Gamma(1))$. The image of a geodesic Γ is called a **geodesic segment** in \mathbf{X} .*

In the space of tree-like shapes, this path (described in details in Section 4.2.1) is an interpolation between the edge attributes of starting and arriving points corresponding to the same splits. The rest of the edges in the starting point must collapse to let edges associated to new splits grow.

Definition 28 (Geodesic Metric Space). *(\mathbf{X}, d) is called a geodesic metric space if there exists at least one geodesic path between every two points in \mathbf{X} . If in addition this path is unique, it is called uniquely geodesic.*

If the geodesic between two points $x, y \in X$ exists, we note $\Gamma(x, y) : [0, 1] \rightarrow \mathbf{X}$ this geodesic.

Definition 29 (Path length). *If $c : [0, 1] \rightarrow \mathbf{X}$ is a path in \mathbf{X} then the length of c is*

$$l(c) = \sup_{t_0 \leq t_1 \leq \dots \leq t_n} \sum_{i=0}^{n-1} d(c(t_i), c(t_{i+1}))$$

where the supremum is taken over all possible n and partitions of the interval $[0, 1]$. The length of c is taken to be infinite when this expression is unbounded.

With the triangle inequality we have $l(c) \geq d(c(t_1), c(t_2))$, $t_1, t_2 \in [0, 1]$ for any path c . It follows from the geodesic definition Γ on \mathbf{X} that $l(\Gamma) = d(\Gamma(0), \Gamma(1))$. The geodesic segment between two points is then equivalent to the shortest path between them. The geodesic length then provides an induced intrinsic metric on the geodesic metric spaces.

Theorem 10 (Generalized Hopf-Rinow [HR31]). *If a length-metric space (M, d) of finite dimension is complete and locally compact then any two points in M can be connected by a minimizing geodesic, and any bounded closed set in M is compact.*

We thus have a characterization of geodesic metric spaces, and of the metric induced by the length of the geodesic between two points. The Hopf-Rinow theorem also gives us that geodesic completeness is equivalent to metric completeness. We now want to make sure that the space \mathcal{T}_N constructed so far is a unique geodesic metric space.

Curvature of the Space of Tree-like Shapes

To derive tools in the constructed stratified space \mathcal{T}_N , in particular the geodesic existence and uniqueness between two distinct points, we rely on its CAT(0) structure. It is characterized with the triangles in \mathcal{T}_N and their comparison triangles in the euclidean space of same dimension. Take a geodesic triangle (x, y, z) in \mathcal{T}_N , three points in \mathcal{T}_N forming the nodes of a triangle whose faces are the geodesic segments connecting x, y and z . A comparison triangle (x', y', z') is a triangle in the euclidean space preserving the edges length: $d(x, y) = d(x', y')$ and so on.

Definition 30 (Comparison Pair). *Let \mathbf{X} be a geodesic metric space, (x, y, z) be a geodesic triangle in \mathbf{X} , \mathbb{R}^2 an euclidean space of same dimension as \mathbf{X} and (x', y', z') a comparison triangle in \mathbb{R}^2 . A comparison pair (a, a') is a pair of points $a \in \mathbf{X}$, a on (x, y, z) , for example $a \in \Gamma(x, y)$ and $a' \in \Gamma(x', y')$ such that $d(x, a) = d(x', a')$.*

Definition 31. A CAT(κ) space is a metric space in whose curvature is bounded above by κ .

In order to verify the CAT(0) structure, a geodesic metric space must verify that its triangles are "at least as thin" (see Figure 4.4) than their comparison ones:

Definition 32 (CAT(0) Space). *A geodesic triangle (a, b, c) satisfies the CAT(0) inequality if, for any comparison triangle, $d(x, c) \leq d(x', c')$ for all comparison pairs (x, x') with $x \in \Gamma(a, b)$, $x' \in \Gamma(a', b')$, and so on for all comparison pairs on the other two edges. The geodesic metric space \mathbf{X} is a CAT(0) space if every geodesic triangle satisfies the CAT(0) inequality.*




Figure 4.4: euclidean triangle (left) versus CAT(0) triangle (right).

Conversely, any CAT(0) space is a unique geodesic metric space. \mathcal{T}_N is of finite dimension, locally compact since the maximal orthants are glued together according to the unresolved tree they represent at their boundaries: \mathcal{T}_N is a geodesic metric space. It is also shown in [FN20], and following the proof of [BHV01], that \mathcal{T}_N is CAT(0) hence a unique geodesic metric space.

The proof relies on the extension of the one provided in [BHV01] for the BHV_N tree space. The idea is to consider a space of tree-like shapes with an upper bound on the edge attributes norm, which makes the space compact. It can be seen as a gluing of unit hypercubes of adequate finite dimension, and remark that each cube is connected to a finite number

of other cubes: it is locally compact hence a geodesic metric space using Hopf-Rinow theorem. The space is then a cubical complex, and Gromov's theorem tells us that such spaces are $CAT(0)$ if and only if the link of their vertices (forming a simplicial complex) is flag complex. In other word, if the 1-skeleton of one of the simplex belongs to the simplicial complex, the associated simplex should also belong to the complex. In our case it is true since the boundaries of the orthants are connected by faces determining valid tree topologies – yet unresolved –, seen as the set of all pairwise compatible splits with non collapsed associated edge attributes.

We illustrate in Figure 4.5 an example of cubical complex that does not satisfy the $CAT(0)$ condition: the bottom-right simplex associated with the link of the vertices does not belong to the cubical complex.




Figure 4.5: Example of a cubical complex that is not $CAT(0)$.

All in all, we know that given two points in \mathcal{T}_N there is a unique geodesic segment between them whose length is the induced metric on \mathcal{T}_N .

Building a Space of Admissible Topologies

As discussed, provided a set of leaves $\{0, 1, \dots, N\}$ with 0 the root, there are up to $(2N - 5)!!$ possible topologies, which is intractable in practice. Considering a registration application, unlike [Fer+12b], we set a space of admissible topologies and we set the number of leaves in the space where the template lives. In the application to anatomical trees we can mitigate this by considering admissible topologies such as the extrinsic methods in the literature (4): such topologies can be defined a priori knowing the anatomy or derived from a database. In the last case it is possible to list all the unique topologies and store all the corresponding Σ . When computing the geodesic between the trees, one could cross new orthants that are not represented in the database. Such orthants being used in the construction of the template, they can be added to the space of admissible topologies. This reduces the sparsity of the constructed space while preserving its limited size. This procedure corresponds to the construction of an *Atlas*, i.e. the construction of a set of orthants in which the new trees can be built. The *Template* is then the choice of one position within this Atlas. It is important to note that the construction of the tree-space is done once, and the test of all the existing topologies would also be done offline.

Example: simplified vascular trees We show in Figure 4.6 an example of admissible topologies for a set of 49 trees with a root and 6 leaves. This example illustrates the already high variability in terms of topologies despite the few leaves in the trees. However,

some subtrees are common to all the individuals and allow us to discard numerous unlikely topologies. We can also notice the proximity in the treespace of all these trees: many of them




Figure 4.6: Example of unique topologies for vascular trees database of 50 trees and leaves $\{0, 1, \dots, 6\}$. The edge attributes in this figure are the vectors of endpoints in \mathbb{R}^3 .

find their bifurcations – nodes of degree 3 – close one to another, meaning that all these trees are close to a common boundary.

4.2 Computing Geodesics in the Space of Tree-like Shapes

So far, we described the space of tree-like shapes and provided some properties of its structure. We now discuss how to move in such a space along its geodesics segments. In a first part we characterize the geodesics in \mathcal{T}_N , then we discuss the algorithm of [OP11] for computing them, and finally we provide a python implementation of the geodesic space adapted to tree-like shapes (Definition 23).

There exists many topological distances between trees along with algorithms to compute them. The most basic topological change distance for a tree is the Nearest Neighbor Interchange (NNI) distance proposed in [Rob71] and describing one change of maximal orthant in the space of tree-like shapes. Later the Robinson-Foulds (RF), or partition distance [RF81], proposed to count the number of splits in T_1 and not in T_2 and the number of splits in T_2 that are not in T_1 . Its weighted adaptation the Weighted Robinson-Foulds (WRF) sums the

lengths of the edges associated to these splits, which corresponds to the geodesic (not unique anymore) distance in a space like BHV with an L^1 metric on the orthants. More recently two approaches (Subtree-Prune-And-Regraft distance [Hei90] and Tree Bisection and Reconnection distance [AS01]) considered the topological changes between leaf-labeled trees despite the shape (in most of the cases, the length) of the branches.

We chose to focus on the geodesic distance and the computation of the geodesic path, adapted to both topological and spatial changes in the trees. In short, the approach of [OP11] for computing the geodesic distance and path in polynomial time is to start from a simple path going through the star tree, and transforming it into shorter paths successively until the geodesic is obtained. They only consider then orthants intersected by the geodesic to limit the number of iterations creating a polynomial time algorithm identifying at each step the new orthant through which the updated path will go.

4.2.1 Characterization of the Geodesics

Before going into the details of the method, we must make several remarks. First, we rely on the algorithm proposed by [OP11], but which was initially applied on edge attributes in \mathbb{R}^+ . Similarly to [Fer+14], we adapt this algorithm to the case of a spatial representation with edge attributes as vectors of euclidean spaces. The properties for the characterization of the geodesic, however, only depend on the edges norm. In the following, the norm of the edge attributes thus depends on the space in which they live. Second, the initial algorithm works on trees having no common edges. In practice, the method can also be applied to more generic trees by applying the algorithm to each pair of subtrees sharing the same leaves but having no common edges.

Definition 33. *A path Γ is a local geodesic path if there exists some $\epsilon > 0$ so that every subpath of Γ of length inferior or equal to ϵ is the shortest path between its endpoints.*

Lemma 1. *In a CAT(0) space (Definition 32), every local geodesic is a geodesic.*

This results yields that in a CAT(0) space, the local condition is sufficient to determine the geodesic. (Proved in more generality in [BH99], Chapter II.1, Proposition 1.4.) The following proof illustrates the philosophy of the algorithm used to compute the geodesics.

The principle of the geodesic construction is then to start with a given path, locally test whether it is a geodesic and if it isn't, build a shorter one. The *path space* described in [BHV01] is a sequence of orthants containing the geodesic between two trees $T = (\mathcal{V}, b, e)$ and $T' = (\mathcal{V}', b', e')$ with the sets of inner edges \mathcal{E} and \mathcal{E}' respectively, and satisfying certain properties. These properties are given in [OP11] and we recall them here.

Assuming first that T and T' are disjoint, that is, they have no common edges. Two edges sets $A \subset \mathcal{E}$ and $B \subset \mathcal{E}'$ are compatible if every pair of splits associated with A for T and B for T' are compatible.

Proposition 18. *Let $T = (\mathcal{V}, b, e)$ and $T' = (\mathcal{V}', b', e')$ be two disjoint trees on the set of leaves L with their associated sets of inner edges \mathcal{E} and \mathcal{E}' respectively. Let $\mathcal{A} = (A_1, \dots, A_k)$*

and $\mathcal{B} = (B_1, \dots, B_k)$ be partitions of \mathcal{E} and \mathcal{E}' such that $(\mathcal{A}, \mathcal{B})$ satisfies the following property: for each $i > j$, A_i and B_j are compatible.

Then for all $1 \leq i \leq k$, $B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k$ is a compatible set and hence $\mathcal{O}_i = \mathcal{O}(B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k)$ is an orthant in tree space associated with the edge attributes of $B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k$.

The union $\bigcup_{i=1..k} \mathcal{O}_i$ is a connected path called "path space", and $(\mathcal{A}, \mathcal{B})$ its support. The shortest path between T and T' through this path space is called the "path space geodesic" for $\bigcup_{i=1..k} \mathcal{O}_i$.

Intuitively, with this proposition, one constructs a succession of connected orthants from combinations of the inner branches of T and T' that are compatible. For each i , we traverse a new orthant: some branches of T have disappeared to make room for branches of T' encoding a new connectivity. It is quite clear that the geodesic between two trees T and T' will cross orthants, and thus be the path space geodesic for a certain path space.

That is given in [BHV01], Proposition 4.1: the geodesic between disjoint trees on n leaves is a path space geodesic for some path space between T and T' . In the following, we provide the properties characterizing the path space between T and T' that contains the geodesic $\Gamma(T, T')$.

Theorem 11 ([OP11] Theorem 2.3). *Let $T = (\mathcal{V}, b, e)$ and $T' = (\mathcal{V}', b', e')$ be two tree-like shapes on the set of leaves L and \mathcal{E} and \mathcal{E}' be their associated sets of inner edges. Let Γ be the geodesic in \mathcal{T}_N between T and T' . Then Γ can be represented as a path space geodesic with support $\mathcal{A} = (A_1, \dots, A_k)$ of \mathcal{E} and $\mathcal{B} = (B_1, \dots, B_k)$ of \mathcal{E}' which satisfies Proposition 18 and the following property:*

$$\frac{|A_1|}{|B_1|} \leq \frac{|A_2|}{|B_2|} \leq \dots \leq \frac{|A_k|}{|B_k|} \quad (4.1)$$

A path space satisfying Proposition 18 and Equation 4.1 is called a proper path space, and the associated path space geodesic is called a proper path.

A proper path describes the sequence of inner edges that shrink until they completely collapsed, and by which growing inner edges they are replaced. In other words, it describes in which order the boundaries between the orthants are crossed. A proper path can be described as follow:

Theorem 12 (Adapted from [OP11], Theorem 2.4). *Let $\Gamma = (\gamma(t) : 0 \leq t \leq 1)$ be a proper path between T and T' two disjoint trees on the set of leaves L with support $(\mathcal{A}, \mathcal{B})$. Then Γ can be represented in \mathcal{T}_N with segments:*

$$\Gamma^i = \begin{cases} \left[\gamma(t) : \frac{t}{1-t} \leq \frac{|A_1|}{|B_1|} \right], i = 0, \\ \left[\gamma(t) : \frac{|A_i|}{|B_i|} \leq \frac{t}{1-t} \leq \frac{|A_{i+1}|}{|B_{i+1}|} \right], i = 1, \dots, k-1, \\ \left[\gamma(t) : \frac{t}{1-t} \geq \frac{|A_k|}{|B_k|} \right], i = k \end{cases}$$

where the points on each leg Γ^i are associated with the tree $T_i = (\mathcal{V}^i, b^i, e^i)$ on L having inner edges set $\mathcal{E}^i = B_1 \cup \dots \cup B_i \cup A_{i+1} \cup \dots \cup A_k$,

$$e_l^i = \begin{cases} \frac{(1-t)|A_j| - t|B_j|}{|A_j|} e_l, b_l \in A_j, \\ \frac{t|B_j| - (1-t)|A_j|}{|B_j|} e'_l, b'_l \in B_j \end{cases}$$

and splits

$$S_e^i = \begin{cases} L_e | L_e^C, e \in A_j \\ L'_e | L'^C_e, e \in B_j \end{cases}$$

and the length of Γ is $l(\Gamma) = (|A_1| + |B_1|, \dots, |A_k| + |B_k|)$.

Even if a geodesic must describe a proper path, the theorem 11 does not fully characterize a geodesic path space, take for instance the cone path space between T and T' , with support $(\mathcal{A}, \mathcal{B}) = (\mathcal{E}, \mathcal{E}')$ even if the trees belongs to contiguous orthants. To obtain proper geodesic path space, one needs to add intermediate orthants by making sure that the new generated path space will lead to shorter proper path between the trees. In practice this corresponds to splitting the current supports of edges for T and T' in order to combine part of them to create a new orthant with some pairs (A_i, B_i) , $A_i \in \mathcal{A}$ and $B_i \in \mathcal{B}$.

Remark 9. When moving from T to T' , adding the splits (A_i, B_i) to the support of the path space corresponds to collapsing the edges of A_i in T when moving toward this new orthant, and growing edges in B_i associated to T' once in the new orthant. This operation may not lead to a shorter path in every case, and one needs to verify that the new path is indeed shorter.

To build actual proper paths as geodesic, we need a supplementary condition:

Theorem 13 (Geodesic, [OP11] Theorem 2.5). *A proper path Γ between T and T' with support $(\mathcal{A}, \mathcal{B})$ satisfying Proposition 18 and Theorem 11 is a geodesic if and only if for every support pair (A_i, B_i) there is no nontrivial partition $C_1 \cup C_2 = A_i$ and $D_1 \cup D_2 = B_i$ such that C_2 is compatible with D_1 and $\frac{|C_1|}{|D_1|} < \frac{|C_2|}{|D_2|}$*

We do not detail the proof here, in brief it consists in supposing that the last property does not hold, and showing that there exist a shorter path so Γ cannot be a geodesic. On the contrary, assuming that Γ is not a geodesic, there must exist a new orthant which provides a shortcut and one can show that if so, the property Theorem 13 does not hold.

This last theorem gives us a way to select new orthants in order to refine the proper path space and its associated proper path, making sure that the newly selected path space will provide a "shortcut" and thus a shorter path closer to the geodesic between two trees.

We know how to build an iterative algorithm that starts from a certain path verifying the properties of a proper path that we will refine during the iterations by finding new orthants to cross. To do this, we just need to select subsets of the inner branches to be removed/pushed that will verify the Theorem 13.

4.2.2 Geodesic Algorithm in a Space of Tree-Like Shapes

When the trees T and T' have no common edges (edges encoding the same splits in both trees), the construction of the geodesic consists in finding splits of the pairs of the path support verifying Theorem 13. It can be formulated as finding a partition for some pair of set of edges (A_j^i, B_j^i) verifying the conditions 13, which can be formulated as a problem over bipartite graphs.

Definition 34 (Incompatibility Graph). *Given two sets of inner edge attributes $A \subset \mathcal{E}$ and $B \subset \mathcal{E}'$ associated to trees T and T' and their sets of leaves L , the incompatibility graph $G(A, B)$ is the bipartite graph whose vertices correspond to $A \cup B$ and whose edges correspond to the pairs $(a, b) \in A \times B$ such that $L_a|L_a^C$ and $L_b|L_b^C$ are incompatible. An independent set in $G(A, B)$ is a set of vertices having no edges of $G(A, B)$ between them.*

Lemma 2. *Two edges sets $A \subseteq \mathcal{E}$ and $B \subseteq \mathcal{E}'$ are compatible if and only if they form an independent set in $G(\mathcal{E}, \mathcal{E}')$.*

The search for shortcuts in the space of tree-like shapes can be formulated as follow:

Definition 35 (Extension Problem). *Let $A \subset \mathcal{E}$ and $B \subset \mathcal{E}'$ be two sets of inner edges of two trees T and T' . The **extension problem** consists in looking for a partition of the support verifying the Proposition 13: $C_1 \cup C_2 = A$ and $D_1 \cup D_2 = B$ such that C_2 is compatible with D_1 (i.e. $C_2 \cup D_1$ is an independent set in $G(A, B)$) and $\frac{|C_1|}{|D_1|} \leq \frac{|C_2|}{|D_2|}$.*

It follows that if this **extension problem** (Definition 35) has no solution for any (A_i, B_i) , then the current path $\Gamma(T, T')$ is the geodesic. In other words there is no more shortcut to the current proper path currently defined. It can be shown that solving this problem boils down to finding a vertex cover (C_1, D_2) for $G(A, B)$, meaning that every edge of $G(A, B)$ is incident to a vertex of C_1 or D_2 , and the minimum weight vertex cover for $G(A, B)$ has normalized weight $|C_1|^2 + |D_2|^2 < 1$. Such a problem can be formulated as Min Cardinality

Vertex Cover problem using weights as capacities on the source and sink arcs [AMO93] (Section 7.7).

The problem of finding the geodesic path between the trees can be summarized as follow:

1. Start with the cone path that is proper path Γ^0 between T and T' with support $(\mathcal{A}^0, \mathcal{B}^0)$.
2. At each step s , the proper path Γ^s with support $(\mathcal{A}^s, \mathcal{B}^s)$ satisfying Prop. 18 and Prop. 4.1 is available. We check whether it satisfies the conditions 13 by solving the extension problem 35 and if not, create a new proper path Γ^{s+1} with support $(\mathcal{A}^{s+1}, \mathcal{B}^{s+1})$ shorter than Γ^s .
3. Continue until the geodesic is found.

Theorem 14 ([OP11], Theorem 3.5). *The algorithm solves the problem of finding the geodesic between two trees can be done in $O(N^4)$ time.*

Remark 10. *Note that the geodesic construction described so far assumes that the two trees have no common edges. In practice, two trees have such common edges (in particular, the ones connecting the leaves) taken into account in the geodesic length. The method consists then in splitting the trees at the level of the common edges, solve the geodesic construction between the pairs of subtrees on the same sets of leaves. By denoting the pairs of subtrees with no compatible edges $(T_l, T'_l)_{l=1, \dots, n}$, the supports of the paths between them are $(\mathcal{A}_l, \mathcal{B}_l)_l$ with $\mathcal{A}_l = (A_1^l, \dots, A_{k_l}^l)$ and $\mathcal{B}_l = (B_1^l, \dots, B_{k_l}^l)$ and the vectors of common edges between T and T' are written b_{common}^T and $b_{common}^{T'}$. The total geodesic length then becomes:*

$$l(\Gamma) = |(|A_1^1| + |B_1^1|, \dots, |A_{k_1}^1| + |B_{k_1}^1|, \dots, |A_1^n| + |B_1^n|, \dots, |A_{k_n}^n| + |B_{k_n}^n|, |b_{common}^T - b_{common}^{T'}|)|$$

When adapting Theorem 12 to the space of tree like shapes with edge attributes in $\mathbb{R}^{p \times d}$, we thus have the following proper path description:

- At time $t \in [0, 1]$, the edge attributes e_i belonging to any A_i^l in T are linearly interpolated between their initial position at time $t = 0$ and the collapsed vector of p -times a single point at time $t_i^l = \frac{|A_i^l|}{|A_i^l| + |B_i^l|}$. It remains collapsed for $t \geq t_i^l$
- Similarly, the edges e_j associated to e_j belonging to any B_j^l in T' are linearly interpolated between the collapsed vector of p -times a single point at time $t_j^l = \frac{|A_j^l|}{|A_j^l| + |B_j^l|}$ and the edge e_j at time $t = 1$. Otherwise it is collapsed for $t \leq t_j^l$.

- The common edges along the proper path are given by $e_{\text{common}} = (1 - t).e_{\text{common}}^T + t.e_{\text{common}}^{T'}$.

Remark 11. *The complexity of the algorithm regarding the number of leaves seems unsustainable in practice, yet the arteries in vascular trees correspond to whole subtrees: the subtrees can be simplified in the space of tree-like shapes to reduce their complexity while preserving the hierarchical changes. One can either consider one subtree as one leaf for the template in the space of tree-like shapes (the subtree is a varifold with a fixed number of points and tangent vectors) or select one representative path in the subtree that will be used as the leaf. The pelvic vasculature for instance is composed of at most 17 labels, so in worst case scenario about 30 subtree-leaves, several ones sharing the same labels, when the unlabeled complete tree is composed of about 300 leaves.*

4.3 Implementation

We propose a Python implementation available on [github](#)³ of the described trees along with the geodesic construction adapted from [OP11] to apply to tree-like shapes.

To the best of our knowledge, there was no python implementation providing the basics of tree-like shapes operations. If python libraries^{4, 5} were recently published, they focus on phylogenetic trees. The other libraries are written in R, Matlab or Java and also focus on phylogenetic trees. The proposed library provides tools for tree-like shapes manipulations with different edge attributes, and allows for computing geodesic path, distance, and interpolation between trees with the same set of leaves. It is also compatible with Newick format trees encoding trees with one attribute (as the norm of the edge attribute most of the time) per edge. We also implemented statistical tools for such tree-like shapes, that we detail in Section 4.3.2. Most of the statistics rely on the geodesic distance of one tree to another point or to paths in this space. We illustrate in Figure 4.8 the output of the distance of a tree T_3 to the geodesic between trees T_1 and T_2 of Figure 4.7. As the function $t \in [0, 1] \rightarrow d(\Gamma(T_1, T_2)(t), T_3)$ is convex, we perform a golden ratio search to find its minimum.

4.3.1 Estimation of the Fréchet Mean

A first statistic we implemented is the estimation of the Fréchet mean of a set of trees in the constructed space. There are numerous works looking at the Fréchet mean of set of points in such stratified spaces as the tree spaces [Fer+12b; Bač14; Nye+17; Gar+21] and in particular illustrating the limits of using certain metrics as the Tree Edit Distance. In this work we use the Sturm Mean with the Quotient euclidean Distance [Fer+12b] in order to better capture the shapes in the case of attribute spaces in $\mathbb{R}^{d \times n}$. This Fréchet mean estimation captures the average geometry of the trees along with the average topology in the Tree Space. In practice

³<https://github.com/plantonsanti/DiffeoAndTopo.git>

⁴https://github.com/pbeerli/pathtrees_project

⁵<https://github.com/kgori>




Figure 4.7: Illustration of a geodesic between two tree-like shapes T_1 and T_2 . There are two hierarchical changes, meaning that three orthants are crossed along this geodesic. In this example, we use a fixed number of discretization points per centerline curve as edge attribute: $e = (e_i)_i$, with $e_i \in \mathbb{R}^{50 \times 3}$.

this often leads to a position close to a boundary between all the orthants meaning that the edges corresponding to splits that are not shared by the different topologies will be close to or at the 0 of their attribute space. It is illustrated in the Figure 4.6 in which the central tree is the output of the Sturm algorithm. In this example with segment edges, the average position in the space of tree-like shapes composed of 9 orthants corresponds to a shrinkage of the uncommon splits between the trees of the database, the inner edges. This mean will be used in the next section dedicated to the registration of a template tree onto a target. In Figure 4.9 we show the output of the Sturm algorithm after 200 iterations for the database of 50 vascular trees used in Chapter 2. We see that in addition to the average topology, we also capture the average position using edge attributes of 50 points in \mathbb{R}^3 .

4.3.2 Principal Geodesic analysis in the Space of Tree-Like Shapes

Similarly to the principal geodesic analysis in the case of data lying on a manifold, one can search for a geodesic in the space of tree-like shapes that will best explain a set of observations in the space of tree-like shapes. This statistic of order 1 was computed in




Figure 4.8: The space of tree-like shapes with the geodesic distance metric is CAT(0): the triangles are "skinnier" than in a classic euclidean space. First row: the distance of a reference tree T_3 to the points along the geodesic between two other distinct trees $\Gamma(T_1, T_2)(t)$, $t \in [0, 1]$. The red dot marks the position of the closest point of $\Gamma(T_1, T_2)$ to T_3 . Second row: the number of orthants crossed by $\Gamma(T_1, T_2)$ after a time t .




Figure 4.9: The Sturm mean estimate of the Fréchet mean of a set of 50 vascular trees after 200 iterations of the algorithm.

[Nye11] by firing geodesics from a selected mean in the dataset, and a greedy search in all the directions. The geodesic minimizing the sum of the geodesic distances of the trees to the geodesic was selected. As shown recently in [FN20], the principal geodesic in stratified spaces may not go through the mean of the dataset.

In [Fer+13] the method consists in searching the best geodesic minimizing the score used in [Nye11] among the geodesics connecting the pairs of points in the dataset. The output of this method on a database is illustrated in Figure 4.10: the estimated principal geodesic sampled at 5 time steps. However, in [Zha16], the authors show that in some cases it will fit poorly in comparison to unconstrained geodesics.

Algorithm 3 Sturm's Algorithm ([Stu03])

Input: (T_1, T_2, \dots, T_K) K trees.

Output: An estimate of the Fréchet mean μ of the set of input trees.

Compute the geodesic $\Gamma(T_a, T_b)$ between two trees T_a and T_b selected at random.

$\mu_0 \leftarrow \Gamma(T_a, T_b)(\frac{1}{2})$

$i \leftarrow 1$

repeat

 Select a tree T_i at random

 Compute $\Gamma(\mu_{i-1}, T_i)$

$\mu_i \leftarrow \Gamma(\mu_{i-1}, T_i)(\frac{1}{i+2})$

until Convergence or maximum iterations reached

In [Nye14] such constraints are dropped by employing a stochastic search algorithm to vary the endpoints of the geodesic in the space of phylogenetic trees. The resulting segment best fits the data with respect to a given criterion. This method however may be intractable in a space of tree-like shapes with a large number of leaves and several points per edge attributes.

Proposition 19. *Given a point $x \in \mathcal{T}_N$ and a geodesic segment $\Gamma(T_1, T_2)$ between two other points of \mathcal{T}_N , there is a unique closest point on $\Gamma(T_1, T_2)$ to x with respect to the geodesic distance metric, called the projection of x on Γ and denoted $P_\Gamma(x)$.*

Proof. The function $f_x : y \in \Gamma(T_1, T_2) \rightarrow d(x, y)$ is continuous and bounded on a compact, so the minimum exists and is achieved. Suppose now that f_x reaches its minimum at two distinct points y_1 and y_2 on $\Gamma(T_1, T_2)$. If so, take any $y \in \Gamma(y_1, y_2)$ and its comparison pair in the comparison triangle (x', y'_1, y'_2) . This comparison triangle is isosceles, so for any point y' in $\Gamma(y'_1, y'_2)$, $d(x', y') < d(x', y'_1)$. From the CAT(0) structure of the space ("skinnier triangles") we have that $d(x, y) \leq d(x', y') < d(x', y'_1) = d(x, y_1)$. Furthermore, since any geodesic is locally geodesic $x' \in \Gamma(T_1, T_2)$, and so the minimum is not reached on y_1 and y_2 , which contradicts the hypothesis: the minimum is unique. \square

Note that sometimes the projection of x on Γ is one of the endpoints. We define a function f_{par}^Γ measuring the variance of the projections of a set of points onto the geodesic segment Γ :

$$f_{par}^\Gamma((e_i)_i) = \sum_i d(\Gamma(0.5), P_\Gamma(e_i))^2$$

We also define a function f_{ort}^Γ measuring the variance of the projections of a set of points onto the geodesic segment Γ :

$$f_{ort}^\Gamma((e_i)_i) = \sum_i d(e_i, P_\Gamma(e_i))^2$$




Figure 4.10: Principal geodesic segment for a database of 50 simplified pelvic vascular trees composed of 10 leaves. The starting and end points have no leaves split in common, and the geodesic passes close to the star tree.

From [Nye11] we have the following proposition:

Proposition 20. *Given a set of points in \mathcal{T}_N , there is a geodesic segment minimizing the function f_{ort}^Γ and another one maximizing the function f_{par}^Γ .*

4.4 A Registration Problem for Deforming Tree-like Shapes

The principal problem of the described geodesic construction in Section 4.2.1 is the need to know the leaves labels, otherwise the structure (\mathcal{V}, b, e) does not describe a point in the space of tree-like shapes. Methods are derived for unlabeled trees, to find the leaves mapping of two unlabeled trees minimizing the associated geodesic distance between them. When one of the tree is annotated, this leads to automatic annotation [Fer+12a]. When dealing with a set of unlabeled trees it provides relative positions in one tree space.

We would like to be able to compare the template, seen as a point in the space of tree-like shapes, to a target vascular tree that may not be labeled. To that end we build an appropriate space of tree-like shapes, in order to be able to compare the template tree-like shape to the

vascular trees with a given data attachment term such as the distance in the space of varifolds.

4.4.1 A Link between Tree-like Shapes and Vascular Trees

In the previous section we defined the basic structures and fundamental tools to analyze tree-like shapes with edge attributes encoding the shapes representation. We must now bridge this representation with the vascular trees we are working on. We can now define our tree shape structure based on the previous definitions:

Definition 36 (Vascular Tree). *We define a vascular tree \mathbf{T} as a rooted fully resolved tree (\mathcal{V}, b) with $\mathcal{V} = \{0, \dots, M-1\}$ together with spatial coordinates x with $x = (x_k)_{k=0, \dots, M-2}$, and $x_k \in \mathbb{R}^{d \times n_k}$ associated with the edges set $b = \{b_k\}_{k=0, \dots, M-2}$. We write $\mathbf{T} = (\mathcal{V}, b, x)$.*

In Definition 36, each x_k is the spatial representation of the edge b_k as a set of sampling points. This definition of the trees is slightly different from the one introduced in Chapter 1 since we now have the topological information encoded by the connectivity matrix and the leaves.

Remark 12. *To go back and forth between the edge attributes space and the ambient one, in some cases we assume that there exists a function $\mathbf{f}_k : \mathbb{R}^{d \times n_k} \rightarrow E$ the projection of the spatial coordinates of the branches of a vascular tree onto the edge attributes space. If such an operation is possible, we denote $e = \mathcal{F}(x) = (\mathbf{f}_k(x_k))_k$. We also assume that there exists in some cases a function $\mathbf{g}_k : E \rightarrow \mathbb{R}^{d \times n_k}$ the function converting the edge attribute to spatial coordinates. Similarly, if such an operation is possible, we write $x = \mathcal{G}(e) = (\mathbf{g}_k(e_k))_k$.*

From Tree-like Shapes to Vascular Trees

One of the purpose of studying the trees through the tree space point of view is to derive statistics that take hierarchical (thus, topological) changes into account. The statistics could for instance be used for realistic data synthesis. Such an application requires the generated data to be in the initial space of representation of the database. We seek therefore a way to embed this original database into the space of tree-like shapes, but as importantly, a way to get from this space back to the original one. In the tree space literature most of the works are studying phylogenetic trees, which corresponds to edge attributes that are lengths of the edges. This edge representation does not allow to retrieve a spatial representation of the trees, since there are infinitely many ways to build a segment given one extremity and a length.

Refining the Treespace One way to easily go back and forth between a pure tree-like shape description of our trees and their spatial representation is to simplify the curves to one straight segment per edge of the tree. As we already mentioned, we are in a stratified space, and the origin of each edge attributes' space coincides with the parent's last point. It implies that the number of points encoding the trees is the number of edges plus one point for the root. This representation is used, for instance, in Figure 4.2(b). It allows to reduce the size of

the vascular tree representation while preserving, in the space of tree-like shapes, the spatial resolution of the shapes we want to align in the ambient space.

However, if it provides a convenient way to embed the initial trees in a treespace, the segment representation is poor regarding the geometric complexity of the vessels centerlines. In [Fer+12a] the authors propose a first natural way to extend the tree shape representation to tree spaces: the edge attributes are the discretization points of the centerlines, sampled with a fixed number of points. This intuitive representation allows to embed the original tree shape representation into a large tree space that is a product of large euclidean spaces containing the coordinates of all the points in the trees. Using all the points coordinates provides a bijection from the original representation to the tree space that is the identity modulo a permutation of the discretization points indices and edges connections. This representation might be adapted to the shapes complexity, and one can modify it by varying the number of discretization points per edge. However, to respect the curves' variability it requires to keep a minimum number of points that makes the total number of parameters linearly grow with the number of edges.

Edges as Bézier Curves Intermediate representations can therefore be used with a limited number of edge attributes while allowing to retrieve a spatial representation close to the ones of the original curves. One convenient curve representation is using the spline encoding and its derivatives (B-splines, Bézier curves...) summarized for instance in [DD78]. We focus in this thesis on Bézier curves, reducing the the edge attributes to a small number of coordinates encoding the splines and allowing to retrieve spatial curves representation with varying number of sampling points along the curves while reducing the size of the tree space.

Definition 37. Let $P = (P_0, \dots, P_m)$, $m \in \mathbb{N}^*$, $m \geq 2$ be $m + 1$ points in \mathbb{R}^d , the Bézier curve associated to these control points is defined with the Bernstein basis polynomials of degree m :

$$B_P(t) = \sum_{i=0}^m \binom{m}{i} (1-t)^{m-i} t^i P_i$$

with $\binom{i}{m}$ the binomial coefficients, $t \in [0, 1]$ and t^0 and $(1-t)^0$ continuously extended to 1 throughout $[0, 1]$.

Using this model, the edge attributes: $e_k = (P_0, \dots, P_m) \in \mathbb{R}^{d \times (m+1)}$ then allow to compute the x_k and retrieve the vascular tree representation. Contrary to the previous segment model, the Bézier curves proposed here introduce a representation that preserves the vascular tree curves representation while providing a edge attributes space of smaller dimension. The number of interpolation points to spatially represent the curves in the ambient space is then equal to the number of time steps used in the spline model, plus one.

We now have access to different edges representations which can be covered by the Bézier curve encoding, by varying the number of control points $m + 1$ from 2 to the number n of discretization points in the curves. In particular we retrieve the segment edge attributes by choosing $m = 1$. Contrary to the embedding function \mathcal{G} of the spatial coordinates into the

tree space, the function mapping the tree space coordinates to the spatial ones is then well defined:

$$\mathcal{F} = \mathcal{B} : E^{M-1} \rightarrow \Pi_k \mathbb{R}^{d \times n_k}, e \mapsto \mathcal{B}(e) = ((\mathbf{B}_{e_k}(t_u))_u)_k,$$

with $t_u = \frac{u}{n_k - 1}$, $u \in \{0, \dots, n_k - 1\}$.

We can therefore play with a trade-off between the model complexity and the regularity of the shapes encoded by few control points. In the rest of the document we fix the number of sampling times per curve to retrieve a spatial representation close to the one in [Fer+12a], however one could adapt the number of sampling times accordingly to the length of the curves for instance.

Remark 13. *In general, if the points of one branch of the vascular tree x_k do not belong to a Bézier curve, we cannot have $\mathcal{F} \circ \mathcal{G}(x_k) = x_k$, no matter G . In our framework however, we are only interested in the function \mathcal{F} that inputs edge attributes and outputs a spatial representation of the branches (the centerlines).*

In brief, a vascular tree is a *tree* structure $\mathbf{T} = (\mathcal{V}, b, x)$ with x the spatial coordinates in the ambient space, and a tree-like shape is a *leaf-labeled tree* structure $T = (\mathcal{V}, b, e)$ with e the edge attributes in the tree space representation. In our framework most of the time $x = \mathcal{F}(e)$, if a function allowing to go from x to e is available we will specify it.

4.4.2 Registering Vascular Trees with Hierarchical Changes

Registration by updating the edge attributes When studying the deformations between shapes of a same population of individuals, the interpolation between the discretization points mapped from one individual onto another is rarely a natural one. In this section, we take advantage of the tree-like shape representation to formulate a diffeomorphic registration framework adapted to hierarchical changes. In addition, to provide more realistic deformations of the shapes, it prevents from computing explicit leaves pairings. We write the problem as a minimization process with a gradient of the data attachment term with respect to the edge attributes (thus in the space of tree-like shapes).

Now that we have provided the main points of tree-spaces and LDDMM, we introduce the proposed method for which we seek to register an atlas built in a tree-space onto a new tree. We do not know the labels of this target tree, and we optimize both the LDDMM and the position of the template in the tree-space, as illustrated in Figure 4.11. The cost function we write is a composition of the classical LDDMM cost and a regularization in the space of tree-like shape. Given a tree $T = (\mathcal{V}, b, e)$ and its associated tree-like shape $\mathbf{T} = (\mathcal{V}, b, \mathcal{B}(e))$ we consider the following function:

$$\begin{aligned} J(p_0, e, \Sigma) = & \gamma \sum_{i=1}^M \sum_{j=1}^M \langle p_0^i, K_V(\mathcal{B}(e)_i, \mathcal{B}(e)_j) p_0^j \rangle \\ & + A(\varphi(\mathbf{T}), Target) + \gamma_2 R_{tree}(\mathbf{T}) \end{aligned} \quad (4.2)$$

With A the data attachment term of the LDDMM framework (see Section 2.4.2). This cost function is designed to enforce the diffeomorphic deformation of the template tree along

with its update within the space of tree-like shapes. We recall that Σ in the previous equation can be retrieved with the edges set b .




Figure 4.11: The overall optimization: the template tree can be represented both in a tree-space with its edges length, and with a spatial representation. The latter is registered with LDDMM onto a new target that is not annotated. From the update of the position of the vertices during LDDMM, we can deduce a new position in the corresponding orthant.

The minimization procedure is two-folded: first, a minimization within the interior of an orthant, given a fixed topology and second, the detection of the boundaries between the orthants and the change of topology.

Updating the Edge Attributes

We now express the update the edge attributes in the minimization process. Since we may apply diffeomorphic deformations to the vascular trees, we denote the discretization points $x = (x_k)_k$ associated with the branches b_k of the vascular tree $T_0 = (\mathcal{V}, b, x)$, at time $t = 0$ of the diffeomorphism. We will also denote by $x^1 = \varphi_1^v(x)$ the discretization points of the tree at time $t = 1$ of the diffeomorphic shooting. The vascular trees registration is the classic LDDMM one, and the supplementary step in the algorithm is the expression of $\nabla_e A(x^1, Target)$ with e the edge attributes of the tree-like shape $T = (\mathcal{V}, b, e)$ as a function of $\nabla_x A(x^1, Target)$.

Updating the Bézier Curves As discussed in Section 4.4.1, if we use the tree-like shape representation $e = x$, the dimension of the problem is too high. One way to reduce the dimensionality is to use the Bézier curves representation. In such a case the interpolation points of the spatial representation x in the data attachment term are also used as the control points in the LDDMM framework. We have for the edge b_k the edge attributes $e_k = (e_{k,i})_i$ for $i = 0, 1, \dots, m$ and $x_k = (x_{k,u})_u = (B_{e_k}(t_u))_u$ for $t_u = \frac{u}{n_k - 1}$ for $u = 0, \dots, n_k - 1$. We keep the general expression n_k on purpose: one could vary the number of points along the curves in the vascular tree, for instance to have a fixed sampling step. In practice we use the same $n_k = n$ for all the edges b_k in the vascular trees. The gradient with respect to the $e_{k,i}$

becomes:

$$\frac{\partial J}{\partial e_k^i} = \sum_{u=0}^{n_k-1} \binom{m}{i} (1-t_u)^{m-i} (t_u)^i \frac{\partial J}{\partial x_{k,u}^0} \quad (4.3)$$

The vertices of the tree are then updated by the formula $e = e + \delta * d_e$, where d_e is a direction of descent that depends on the optimization scheme (typically $d_e = -\nabla_e J$ for gradient descent). This additional step is used in Algorithm 4 after computing $\nabla_{q_0} A(q_1, Target)$ to update the position in the tree space. We now have a minimization problem that allows us to compare shapes in a suitable space, such as that of oriented varifolds, and an update that is done in the space of tree-like shapes.

Descent Inside an Orthant

We describe a minimization step in the interior of an orthant (synthesized in Algorithm 4), with Σ fixed. Given a tree-like shape $T = (\mathcal{V}, b, e)$, we compute its corresponding vascular tree shape $\mathbf{T} = (\mathcal{V}, b, x)$ with $x = \mathcal{B}(e)$. We can then compute the gradient of a loss function $A(x^1, Target)$ with respect to the edge attributes e .

Algorithm 4 Optimization Step Inside an Orthant (Σ fixed)

Input: Σ, e, p_0, δ (step size)

Output: e, p_0 ▷ The optimized variables

Compute $x = \mathcal{B}(e)$

Compute (x^1, p_1) ▷ Geodesic shooting of Algorithm 1

Compute $\nabla_{x^1} J(p_0, e)$

Compute $\nabla_x J(p_0, e)$ through backward integration

Compute $\nabla_e J(p_0, e)$ with Equation.4.3

Compute $\nabla_{p_0} J(p_0, \mathcal{V}, \Sigma) = K(x, x) \cdot p_0 + \nabla_{p_0} A(x^1, Target)$

Update directions of descent d_e, d_{p_0} from $\nabla_e J, \nabla_{p_0} J$

Perform a line search to select a step size δ for the updates $p_0 \leftarrow p_0 + \delta d_{p_0}$ and $e \leftarrow e + \delta d_e$.

With $K(x, x)$ the bloc matrix of the $K_V(x_{k,u}, x_{l,v})$. When moving in the interior of an orthant, the topology Σ is fixed and $\nabla_e J$ is a vector in E^{N-2} . The minimization in the interior of the orthants is a classic minimization in some euclidean space of features encoding the branches of the vascular trees. We now describe how we chose to address the change of orthant, and how to detect possible topological changes with the distance to the boundaries with the other orthants.

4.4.3 Dealing with the Change of Orthant

If the current position is at the boundary between several orthants, the gradient of J can be computed with respect to each contiguous orthant. The direction is then $\arg \max_{\Sigma_i} \|\nabla_{e, \Sigma_i} J\|_2^2$, for Σ_i the connectivity matrices associated with the orthants \mathcal{O}_{Σ_i} contiguous to the boundary. This operation is made possible by the fact that the branches that are not shared by all the

topologies contiguous to this boundary are collapsed to a single point (the associated edge attributes e_k are set to 0_E), hence not used in the shape comparison term A .

In practice, after updating the spatial coordinates, we check the distance to the other topologies $\mathcal{O}_{\Sigma'}$. The distance of the current position $T = (\mathcal{V}, \Sigma, e)$ to a new topology $\mathcal{O}_{\Sigma'}$ is given by the formula:

$$d(T, \mathcal{O}_{\Sigma'}) = \max_{i \in \text{UnCo}(\Sigma, \Sigma')} |e_i| \quad (4.4)$$

with $\text{UnCo}(\Sigma, \Sigma')$ the set of inner edges (not connected to a leaf) in T associated with leaves splits that do not exist in Σ' . In other words, it is the set of inner edges that must shrink to move to the topology Σ' . The other distance which can be used is the length of the centerlines (vascular tree representation) edges is:

$$d(T, \mathcal{O}_{\Sigma'}) = \sum_{k \in \text{UnCo}(\Sigma, \Sigma')} \sum_{i=1}^{n_k-1} |x_{k,i+1} - x_{k,i}|_2 \quad (4.5)$$

If there are topologies closer than a threshold ϵ , we test whether the current position in the tree-space should be set at the boundary between all the topologies within the range ϵ as described in Algorithm. 5. If the choice of orthant at the boundary is in fact a different from the initial one, the optimizer is reset, and the minimization starts from the new position. Otherwise, the classic optimization within the interior of the current orthant keeps going. When moving to the boundary of one orthant, we set the momenta associated to shrunk segments to zero to prevent from instabilities in the LDDMM iterations.

We emphasize the importance of the ϵ parameter that can be interpreted as seeing seeing a thick boundary between the orthants, of width 2ϵ . Too large, we stick to the boundaries, too small, we cannot cross them.

Remark 14. *Instead of pre-computing the existing orthants from our database, another solution could be to evaluate the accessible orthants on the fly and locally, by considering for instance a ball around a tree, yet this solution would suffer from the combinatorial number of solutions when getting close to the star tree (that we often get when computing the Sturm mean).*

The projection of the template onto the boundary of the current orthant is an important step of the registration procedure. In our case the edge attributes e_k associated with the edges that must shrink are set to 0. Since the edge attributes are relative to the ones of the parents, the descendant edge attributes must be translated: this way the first point of the children branches still matches the last one of the (collapsed) parent's.

Regularization Let $T = (\mathcal{V}, b, e)$ be a tree-like shape and $\mathbf{T} = (\mathcal{V}, b, x)$ be a vascular tree. In order to prevent from inconsistent updates in the space of tree-like shapes, one can add some regularization $R_{tree}(\mathbf{T})$ to the function to minimize. In practice the edge attribute representation, as segments end points or control points of a Bézier curve already provide some regularity of the 3D centerlines. One can also regularize to the spatial representation of the vascular trees branches x . There are multiple ways to regularize the shape of a 3D curve, we note three of them:

Algorithm 5 Move To Boundary: A projection method onto the boundary between several orthants.

Input: $T = (\mathcal{V}, b, e)$, $AllTopo = (b_1, b_2, \dots, b_j)$, ϵ

Output: e, b_{sel} \triangleright The new topology and coordinates in $O_{\Sigma_{sel}}$

Set $Selected = \emptyset$

Set $e^{tmp} = e$

\triangleright Search accessible orthants:

for b' in AllTopo **do**

 Compute $d = d(T, \mathcal{O}_{\Sigma'})$

if $d \leq \epsilon$ **then**

 Set $Selected \leftarrow Selected \cup b'$

end if

end for

\triangleright Project T onto the boundary between the selected topologies:

for all e_i^{tmp} s.t. $\exists b_k \in Selected, e_i^k = 0$ **do**

 Set $e_i^{tmp} = 0$

end for

Set $b_{sel} = b$

Set $norm = \|\nabla_e J(p_0, e^{tmp}, \Sigma)\|_2$

\triangleright Select a direction and possibly a new orthant:

for Σ' in $Selected$ **do**

 Compute $g = \|\nabla_{e^{tmp}} J(p_0, e^{tmp}, \Sigma')\|_2$

if $g > norm$ **then**

 Set $b_{sel} = b'$

 Set $norm = g$

end if

end for

if $b_{sel} \neq b$ **then**

 Set $e = e^{tmp}$

end if

return e, b_{sel}

- Penalizing the length of the branches. This regularization may help getting close to the orthants boundaries, stick the template to the boundary between the orthants.

- Penalizing the curvature of the curves, to enforce the smoothness of the template before diffeomorphic registration.

- Penalizing the tortuosity of the curves ($\frac{\sum_k \sum_{i=1}^{n_k-1} |x_{k,i+1} - x_{k,i}|_2}{\sum_k \sum_{i=1}^{n_k-1} |x_{k,n_k} - x_{k,0}|_2}$) to enforce the smoothness of the template vascular tree as well.

In practice we chose the regularization of the length of the curves associated to the inner

edges: it ensured numerical stability when dealing with infinitesimally small branches. We thus have a cost function penalizing the length of the inner branches in the template representation of the vascular tree. This cost is in fact added to prevent from the inconsistent updates of the edge attributes that are unconstrained otherwise.

4.4.4 Experiments

In this section we illustrate the capacity of the proposed registration method to handle topological changes between trees. The registrations are evaluated on both synthetic and real trees with various topologies. All the experiments were performed on trees in \mathbb{R}^3 , although in our first example trees are contained in a plane. We first illustrate the optimization behavior on two synthetic examples. Then, we apply the method to the registration of a template built from a database of vascular trees extracted from segmentation of Cone Beam Computed Tomography volumes.

Minimization Scheme and Parameters Similarly to the previous LDDMM experiments, we use a sum of Gaussian kernels at different scales as reproducing kernel. For each experiment, we will precise the data attachment term selected. We use a multi-scale approach for the data attachment term following [Ris+11], using $\sigma_W \in \{s, \frac{s}{2}, \frac{s}{4}\}$ with s half the size of the bounding box of our trees. The trees are first registered by aligning their root node. Note that the algorithm 4 describes for simplicity a gradient descent on the functional J with a line-search on the step size. In practice, we use more advanced optimization tools such as Limited-Memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) [LN89] for faster convergence inside the orthants.

In addition, in order to prevent from sticking to the boundary between the orthants, we add a supplementary rule. The projection onto the boundary cannot be performed until we have moved away from the last boundary on which we projected by more than ϵ , or a certain number of iterations $n_{iter} = 70$ (arbitrarily fixed) was performed.

A First Experiment

We conducted a first experiment on a toy example tree. The edge attributes of a branch b_k in this experiment are the segment extremities $e_k \in \mathbb{R}^{2 \times 3}$, that are interpolated to get the vascular tree representation $x_k \in \mathbb{R}^{n \times 3}$. The data attachment term is the distance in the space of varifolds, and the parameters of the LDDMM registration are the same in the classic approach and in the proposed one. The source and target shapes have two different topologies and are the ones illustrated in Figure 4.1. Note that in the registration setting however, only the template shape is labeled. In this toy example the template can only have one of the two topologies, no other one is available.

We illustrate the output of the registration in Figure 4.12 for both the proposed approach and the LDDMM approach. We observe that the LDDMM deforms the source shape to align the leaves correctly. The inner edge however is completely collapsed. On the contrary the




Figure 4.12: First registration experiment with the proposed approach (tree-like shape template and LDDMM registration) compared to the classic LDDMM registration.

proposed method allows to change the topology of the template, and correctly align all the branches of the source onto the target. The registration could be further refined by reducing the data attachment scale. The result however, remains satisfactory.

The collapsed edge in the case of the LDDMM registration comes from the topological changes. As the diffeomorphisms cannot explain it, and the inner edge cannot be aligned to the target, it may be collapsed to reduce the diffeomorphic cost of the deformations and the distance in the space of varifolds. Contrary to [DKS18], in our setting we do not want to find a rotation aligning the trees (this would provide a valid solution, but would fail in the sense of the labels). The diffeomorphisms we build on the contrary are adapted to such spatially localized and anatomically relevant registrations.

Second Experiment: three Available Topologies

We illustrate in Figure 4.13 the three available topologies in this experiment. Like the previous experiment, we use the segment representation for the edge attributes. The starting template is the star tree, and the target one is a fully resolved tree composed of two inner

branches and 5 leaves. The data attachment term is the oriented varifolds distance.




Figure 4.13: Second experiment: three available topologies for 5 leaves and two inner edges.

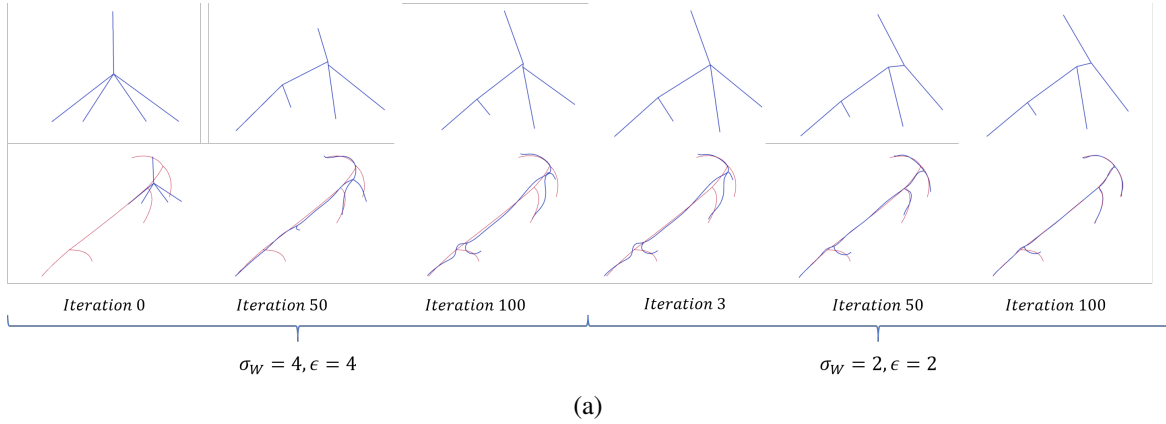


Figure 4.14: Deformations across the iterations: the registration of the template (blue) onto a target (red) starting from the star tree. A first topology is selected at the scale $\sigma_W = 4$, the correct topology is selected at the scale $\sigma_W = 2$. The first row shows the spatial representation of the template. The second row shows the target and the evolution of the global deformation both in the tree space and with the LDDMM.

The registration follows a 2-scales approach for the data attachment term. In this experiment the threshold for the projection onto the boundary between the orthants is equal to the data attachment scale. At the first scale (coarsest data attachment scale and ϵ), the correct orthant remains 1 topological change away from the current position in the tree space. When changing the data attachment scale to a finer one, we authorize a new projection onto the boundary of the orthants close enough to the current position, and this time the correct topology is selected. The evaluation of the deformation and the associated position in the tree space along the iterations is illustrated in Figure 4.14.

This second experiment highlights the importance of the data attachment scale in the minimization process: a large data attachment scale may not detect topological changes. A

too small scale however may miss structures far from each other with respect to this scale. The coarse to fine or direct multiscale approaches mitigate these scales issues.

Consequently the choice of the threshold for the projection to the boundary is also decisive. Using a too large threshold could make the borders sticky, and one would not be able to move away from them. On the contrary, too small thresholds could prevent the projection on the boundary. We decided to set the threshold to the value of the data attachment term: the branches smaller than the data scale in this case can shrink. The intuition is to say that two successive degree 3 nodes separated by a small branch (with respect to the data attachment scale) are seen as one degree 4 node.

Registration of Simplified Vascular Trees

In this experiment we register a template simplified vascular tree to targets that have the same number of leaves. The vascular trees are composed of 6 leaves, and the total number of admissible topologies is 11. We denote the tree-like shape $T = (\mathcal{V}, b, e)$. A branch b_k is represented with edge attributes as 5 control points of Bézier curves: $e_k \in \mathbb{R}^{5 \times 3}$. The associated vascular tree is then written $\mathbf{T} = (\mathcal{V}, b, x) = (\mathcal{V}, b, \mathcal{B}(e))$ and $x_k \in \mathbb{R}^{p \times 3}$, $p = 30$. The target tree $\mathbf{T}' = (\mathcal{V}', b', y)$ is sampled with the same fixed number of points p per curve.

In the following we use the optimal transport cost as data attachment defined in Chapter 2, Section 2.4.3. It is a data attachment term that can be computed between measures seen as sum of Diracs in $\mathbb{R}^{n \times d}$, which is one available representation of the branches of the vascular trees.

The optimal transport cost is given in Equation 2.20. We first perform the registration with no diffeomorphic deformation. The minimization function becomes:

$$J(e_0, \Sigma) = \gamma_1 \|\mu_{\mathbf{T}} - \mu_{\mathbf{T}'}\|_{W'}^2 + L_C^{\epsilon, p}(\alpha, \beta) + \gamma_2 R_{tree}(\mathbf{T}), \quad (4.6)$$

with $\alpha = \sum_i a_i x_i$ and $\beta = \sum_j b_j y_j$. In practice we have that $a_i b_j = \frac{1}{2N-2} \forall i, j$ with N the number of leaves in the tree (the source and target trees have the same number of leaves). The distance cost we use for the optimal transport is the L^2 norm in the euclidean space $\mathbb{R}^{p \times 3}$ with $p = 30$, normalized by $1/\sqrt{p}$.

We set $\gamma_1 = 1$ and $\gamma_2 = 0.1$

The set of targets is chosen such as the inner branches of the targets all have a length higher than $20mm$. This threshold is equal to the smallest data attachment scale for the varifolds distance. In total, 18 targets were selected, and we count 5 different topologies in this set.

Results We illustrate in Figure 4.15 the output of a successful registration. Out of the 18 targets, the minimization procedure retrieved the correct topology in 15 cases. The registration as an update of the edge attributes allow both spatial and topological alignment. The finest details of the shapes can not be matched given the number of control points for the Bézier curves. In some cases the correct topology is selected, but the spatial registration is

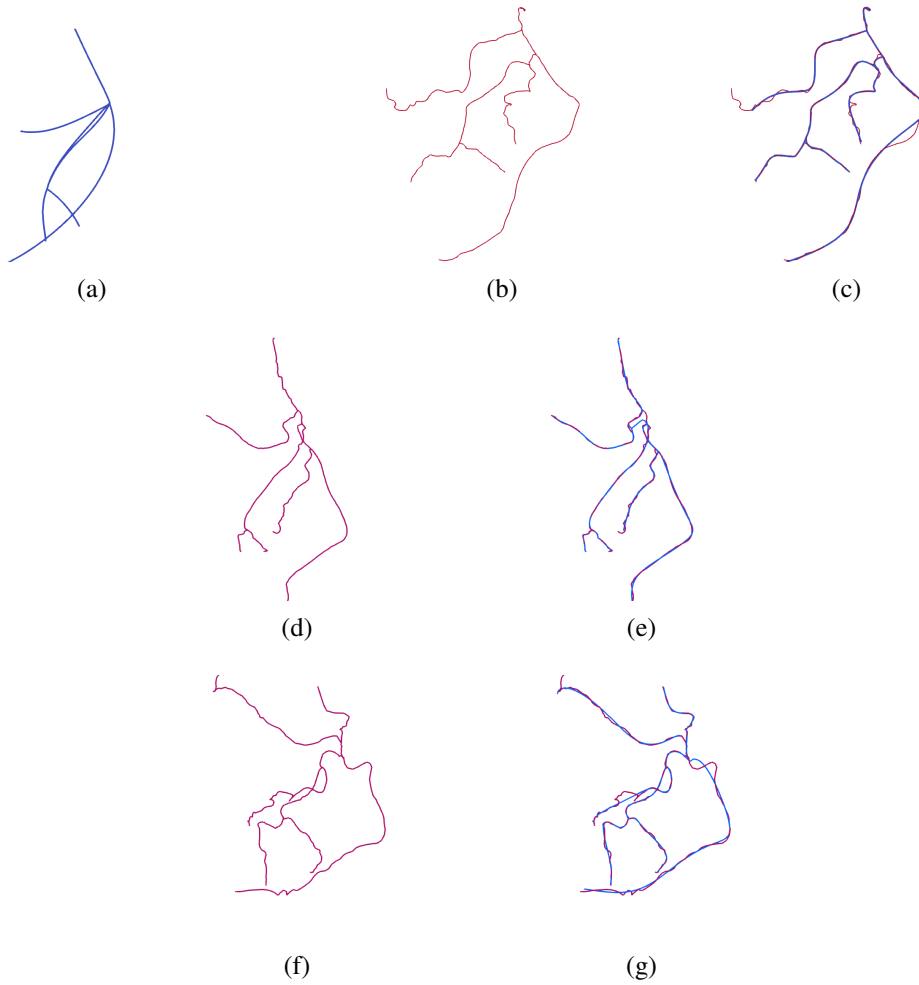


Figure 4.15: Successful registrations of the template tree (a) onto a real target vascular trees (b-d-f) by updating the coordinates of the edge attributes. The template is successfully aligned to the targets (c-e-g).

not that satisfying since we observe branches of the template matching two branches of the target. The 3 cases in which the optimization failed at retrieving the correct topology are typical cases of wrong local minimum at a coarse scale from which the finest scales does not manage to extract itself.

We applied the minimization framework with LDDMM registration as well, with less success: 12 out of 18 cases were correctly registered topologically. Spatially, the resulting diffeomorphisms are close to identity. In fact, in order to make the template cross the orthants, we use a high regularization parameter of the diffeomorphisms.

Adding the diffeomorphic shooting to the registration requires to tune the regularization parameter γ of the deformations very carefully. In fact, we observe that using a γ high, the registration provides results similar to the ones obtained with no diffeomorphic registration. On the contrary using a small γ , the template does not change its topology: the diffeomor-




Figure 4.16: (a-b) Registration procedure finding the correct topology, but introducing a "jump" in the result. (c-d) Wrong topological registration, despite a correct leaves matching.

phisms compensate the topological changes with strong deformations.

4.5 Conclusion

In this chapter we have shown that the template could be represented as a tree-like shape. By doing so, it accesses a space specific to the leaf-labeled tree structures encoding both a spatial representation (with the edge attributes) and the topology of the tree (according to the connectivity of the nodes). Using control points of a Bézier curve as edge attributes, we are able to both reduce the size of the template tree (in term of number of point per edge) and retrieve a spatial representation that corresponds to the one of the vascular trees.

When moving the template within the space of tree-like shapes, we are able to change the bifurcations ordering, thus the topology of the template. We adapt the registration framework, namely the minimization of a cost function, to adapt the position of the template in the space of tree-like shapes. This allows to manage this kind of topological changes during the registration. This update in the space of tree-like shapes is compatible with the classic LDDMM framework, and one is able to both update the template coordinates and apply a diffeomorphism of the ambient space as well.

We have applied this new framework to the registration of toy examples and simplified vascular trees, showing promising results for its use with the LDDMM registration in the case of the toy examples. In the more complex case of simplified vascular trees, the registration without the diffeomorphic shooting was successful in 15 cases out of 18.

Next Steps Using the tree-like shape representation combined with LDDMM registration in the case of real vascular trees is yet to be further explored: in particular if we are to correctly update the position within the treespace without applying too much constraints to the LDDMM.

Applications to the construction of the template, and to better vascular tree registrations are yet to be explored, which will be of great interest in future works. In the next chapter we will focus on applications of the LDDMM and the partial matching to vascular trees and liver surfaces extracted from interventional images.

APPLICATIONS

We have developed tools to generate reference maps adapted to vascular trees, and to deform them to automatically provide information and analyze new observations. It is now time to apply them to pelvic vascular trees, in order to provide interventional radiologists with anatomical information to complete their maps and help them during their interventions.

| | | |
|-------|--|-----|
| 5.1 | LDDMM Application to Automatic Annotation of Simplified Pelvic Vascu- lature | 152 |
| 5.1.1 | Database Description | 153 |
| 5.1.2 | Transferring the Labels from the Template to the Target | 153 |
| 5.1.3 | Evaluation and Results | 155 |
| 5.2 | Registration of a Template Vascular Tree onto Real Ones | 164 |
| 5.2.1 | Database Description | 164 |
| 5.2.2 | Evaluation and Results | 165 |
| 5.2.3 | Discussion | 170 |
| 5.3 | Beyond Vascular Trees: Feature-based Multi-modality Liver Volume Regis- tration | 172 |
| 5.3.1 | CT/CBCT Volume Registration | 173 |
| 5.3.2 | Database Description | 173 |
| 5.3.3 | Liver Surface Registration with Partial Matching | 175 |
| 5.3.4 | From Surface to Volume Registration | 177 |
| 5.3.5 | Evaluation and Results | 178 |
| 5.3.6 | Discussion | 183 |
| 5.4 | Conclusion | 185 |

Chapter Organization This chapter is dedicated to applications of Chapters 2 and 3 to real data. In the first section, we focus on the use of LDDMM for automatic vascular tree annotation: constructing an atlas and a realistic template used as an initial processing to perform automatic annotation.

In the second section we focus on the application of partial matching to the registration of the previously built template on real trees. We are then interested in the anatomical consistency of the obtained deformations, and we observe the influence of the different terms proposed in Chapter 3 on the template matching. If the obtained deformations do not allow to generate an annotation of the whole target tree as they are, we can use them again to compute statistics on the deformations.

Finally, in a third section, we illustrate the versatility of our work on another anatomy and different data structure. We apply partial matching to the registration of liver surfaces from different modalities. Alignments are generated between a couple of volumes of the same patient (a pre-op CT and a per-op CBCT). Due to the limited field of view of CBCT, only a part of the liver is visible, while it can be seen entirely in CT volume. We compute the diffeomorphic registration of the truncated liver surface extracted from the CBCT onto the complete one extracted from the CT. The resulting deformations are then applied to the 3D CT images from which the target surfaces are extracted.

5.1 LDDMM Application to Automatic Annotation of Simplified Pelvic Vasculature

In this first application we use the LDDMM registration method described in 2 as a pre processing for the automatic annotation of simplified vascular trees of the pelvic region based on the comparison to a template annotated tree. We also take advantage of the statistical analysis of the deformations to build an atlas in which we select a template representative of the database. This work was published in MICCAI 2020 [Ant+20] and is reported here in this section with additional illustrations and discussion. The objectives of this first application are multiple: first, to assess whether LDDMM are suitable for vascular tree registration or not. Second, to evaluate if the computational anatomy point of view and the template construction are indeed useful, both for data synthesis and for automatic annotation of vascular trees.

Clinical Context As we have seen in Chapter 1, understanding the patient's anatomy during interventional radiology procedure is essential. This is especially important for Prostatic Artery Embolization (PAE) procedures in which an embolizing agent is injected into the feeding arteries of the prostate in order to resorb it in patients with benign prostatic hyperplasia. The segmentation of blood vessels in 3D volumes of the pelvic anatomy provides a visualization of the vasculature without bones, allowing for a precise analysis of the prostate vicinity. It can facilitate for instance the definition of the location where the treatment is to be done as well as the identification of non target arteries. In this context, the automatic annotation of the vascular tree can greatly facilitate the understanding of the patient's anatomy and limit the risk of missing a vessel feeding to prostate or embolizing an extra-prostatic vessel.

5.1.1 Database Description

The vascular trees are extracted from injected CBCT 3D volumes. These images have been acquired during PAE procedures to better understand the patient's anatomy and assess the position of the target arteries (the ones feeding the prostate). We conducted experiments on a dataset of 50 pelvic vascular trees corresponding to 43 patients: some trees being the left and right vasculature of a single patient. Each tree is composed of a set of 3D curves – the centerlines – representing the center of each vessel. The centerlines are constructed as described in Chapter 1, Section 1.2.1.

While the entire vascular tree is composed of up to 500 different branches, we manually extracted a simplified tree composed of the main arteries documented in the literature [Ass+15a]. This allowed us to reduce the problem to the annotation of 17-branch trees that corresponds to the typical size of trees found in the literature [Cao+17; Wan+17]. The selected arteries of interest are the Prostatic Artery (PA), the Superior Vesicle (SV), the Obturator (O), the Pudendal (P), the Inferior Gluteal (IG) and the Superior Gluteal (SG). Following the bottom-up labeling procedure 5.1.2, we introduced an additional label called "Common Artery" (CA) for a total of 7 labels. This simplified representation still captures the anatomical variability described in [Ass+15a], as we found 28 different tree label arrangements among the 50 cases; and the variability of branches shape and position is illustrated in Figure 5.5a.

Simplified Male Pelvic Vascular Tree

In terms of tree representation, we recall that a vascular tree is defined as $T = (\mathcal{V}, \Sigma, x)$, according to the Definition 3. Σ is the adjacency $n \times n$ binary matrix, and let $\alpha \in [1, \dots, n] \times [1, \dots, n]$ denote the indices such that $\Sigma_{\alpha_0, \alpha_1} = 1$ indicates that \mathcal{V}_{α_0} is connected to \mathcal{V}_{α_1} . The template is a branch-labeled vascular tree, and the branches b_α defined in Definition 14 are labeled $l_\alpha \in \mathbb{N}$ with connections to other branches stored in Σ . Two labeled branches (b_α, l_α) and (b_β, l_β) are connected if $\alpha_0 = \beta_1$. Each branch b_α is spatially represented by a polygonal curve composed of ordered points $x_{\alpha,i} \in \mathbb{R}^3$ representing a vessel centerline.

5.1.2 Transferring the Labels from the Template to the Target

We propose an atlas-based automatic annotation pipeline illustrated in Figure 5.1 relying on the LDDMM to build the atlas and estimate realistic deformations. The data attachment term in this application is squared distance in the space of varifolds. The initial template S is chosen among the available annotated data. The control points are then taken as the discretization points of S , the $(q_i)_i$. It can be registered onto the rest of the data – labeled or not –, and used for automatic labeling which corresponds to the **LDDMM-0** step of the procedure. Following the template construction Algorithm.2, the initial template can be deformed by the estimated average deformation providing the shape \bar{S}_1 . This new shape can be once more registered to the target, and this procedure will be denoted the **LDDMM-1**. This can be repeated iteratively, thus we denote by **LDDMM-k** the k -th iteration. It is combined with

an optimal transport matching (described in Section 5.1.2) to compute relevant correspondences between atlas and target branches. The combination of these two previous steps is called **LDDMM-k+OT** in the following. In the end, to handle the topological changes between the atlas and the target, a bottom-up label assignment is performed to achieve optimal results for our pelvic vasculature labeling problem.

In our LDDMM framework, the label probabilities are computed through a voting system similar to the one described in [Bul+06]. We simplify the process by computing the distances in \mathbb{R}^3 instead of a feature space: for each point in the target, we assign the label of the nearest point in the atlas. The estimated probabilities at the branch level are the voting percentages for each label per branch. It must be emphasized that this label probabilities estimation is weak and provides no anatomical consistency. The probabilities are then strongly dependent on the registration. If one branch of the registered atlas does not match the target correctly, the points correspondences will not be relevant.

Optimal Transport for Branch Correspondences During the LDDMM atlas construction and registration, each tree is seen as one shape. Consequently, there is no assumption over branch matching and topological changes. In order to provide a relevant label assignment that takes the mutual information into account we propose to use Optimal Transport. It is convenient to compute the optimal one-to-one assignment between branches of the deformed source and the target with respect to a given distance.

Let the template tree be $S = (\mathcal{V}, \Sigma, x)$, and the target one $T = (\mathcal{V}', \Sigma', y)$. The diffeomorphic deformation is noted φ . Inspired from the works of [DKS18; Fey+19], each branch is re-sampled with 20 points and the distance matrix D between each deformed branch $\varphi(x_\alpha)$ and y_β is given by: $D_{\alpha,\beta} = \|\varphi(x_\alpha) - y_\beta\|_{\mathbb{R}^{3d}}$. We tried different numbers of points d per branch ranging from 20 to 500 with no significant impact over the matching results. Considering that our problem is of limited size (17 branches per tree), a simple Kuhn-Munkres algorithm (also called Hungarian algorithm) was used to compute the assignment solution. It consists in finding minimum weight matching in bipartite graphs by minimizing the function $\sum_{\alpha,\beta} D_{\alpha,\beta} \cdot X_{\alpha,\beta}$ with $(X_{\alpha,\beta}) \in \{0, 1\}^{17 \times 17}$ the output boolean matrix with 1 if the branches $\varphi(x_\alpha)$ is assigned to y_β , 0 otherwise. We emphasize that in this application, the optimal transport is not used as a data attachment term.

To be consistent with the alternative simple label probability estimation of our LDDMM-k pipeline, we similarly define here $\pi(b_\beta^T, l_\alpha^S) := X_{\alpha,\beta}$, although these "probabilities" are always 0 or 1 in this case. This assignment process is complementary to the LDDMM-k process since it focuses on assignment between branches while LDDMM-k focuses on the atlas construction and the registration. We will also experiment the Optimal Transport assignment without any registration (i.e. taking $\varphi = id$), which we denote **OT**.

Label Assignment

A first label assignment procedure directly takes the highest label probability for each branch. As for the probabilities estimation, it does not provide any guarantee regarding the anatomical consistency and proceeds branch by branch : one branch from the atlas can be associated




Figure 5.1: The proposed atlas-based vascular tree annotation pipeline. The annotation could be performed without atlas construction step ($k = 0$), and without registration, optimal transport or bottom-up assignment as well.

to several branches in the target. This *direct assignment* does not introduce any a priori and reflects the performance of the registration methods.

Bottom-up label assignment In practice, the vessels are labeled by the expert accordingly to the anatomy or area they irrigate [Ass+15a]. In the application to pelvic vascular tree, when two branches of different labels share a parent, this parent is called "*Common Artery*". This is the only clinical a priori we introduce in the method. To limit the effect of the topological variations between the atlas and the target, we propose a *bottom-up assignment* procedure: first for all b_α^T leaf of T , $\hat{l}(b_\alpha^T) = \arg \max_l (\pi(b_\alpha^T, l))$ then recursively every parent branch b_α^T is assigned a label with the rule:

$$\hat{l}(b_\alpha^T) = \begin{cases} l & \text{if } \hat{l}(b_\beta^T) = l \text{ for every branch } b_\beta^T \text{ child of } b_\alpha^T \\ 0 & \text{(i.e. "Common Artery"), otherwise.} \end{cases}$$

This recursive assignment procedure, although specific to this anatomy, is quite adaptable. In fact, in most of the structured tree-shaped anatomies (coronary [Aki+09; Cao+17], airway tree [Lo+11], pelvic [Ass+15a]) the branches names are also conditioned by the leaves labels. When two arteries of different labels share a common parent, this parent is either unnamed (as in our application), or named by a convention provided by the experts. The latter situation corresponds to additional conditions (as in [Lo+11]) during the assignment. One limitation of this procedure is that one risks to propagate the errors from the leaves to the root. On the other hand, there is no need to build specific rules for the anatomical variants by learning from the training data.

5.1.3 Evaluation and Results

Tree Registration

We can first observe that the deformations generated by LDDMM provide globally satisfactory results, but that we have a strong sensitivity to the scaling parameters of the kernels both for the vector fields generating the K_V deformations and for the kernels associated with the data attachment term formulated as a distance in the space of varifolds. Some registration examples are provided in Chapter 2, Section 2.6. In this section we focus on the application to automatic annotation of vascular trees.

However, we notice first errors related to the spatial position of the source branches in relation to those of the target. This may correspond to a branch of a certain label being shifted to a branch of the tree that is eligible for another label, or it may also come from a "jump" of a deformed branch that will shift to two branches of the target. The scale of the data attachment kernel in particular plays a very important role here because it defines both how far structures can be seen from each other, but conversely, if the scale of comparison is too small compared to the scale of the kernel, differences such as the gap covered by a branch jump is simply not taken into account by the metric.

A second type of error can be observed in the registration of a template tree on the rest of the database, which is what diffeomorphisms cannot explain. As discussed in Chapter 2, Section 2.4, the shapes in the database do not always belong to the same orbit for a given group of actions. In the case of diffeomorphic transformations, one can therefore never explain topological changes such as the change in the order of bifurcations or missing branches in trees. The last point has been handled through the manual construction of the database, however changes in the order of bifurcations cannot be avoided, and this leads to anatomically unrealistic alignments.

Improving Annotation Performance with Post-Processing

These observations on the behavior and limitation of the diffeomorphic deformations lead us to use post-processing to improve the annotation performances and to manage what LDDMM cannot solve alone. A first experiment consists therefore in comparing our registration-based approach (LDDMM-0) to a classical method of object matching which is the optimal transport (OT) (the Hungarian algorithm) described in 5.1.2. The third atlas-based solution is a combination of the two (LDDMM-0+OT). The results of the automatic annotation of the 49 remaining cases from a given template for all cases in the database are shown in Table 5.1.

| Method | Direct assignment | Bottom-up assignment |
|------------|----------------------------|----------------------------|
| OT | 81.3 (± 5.11) | 94.4 (± 2.79) |
| LDDMM-0 | 84.2 (± 4.43) | 93.3 (± 3.44) |
| LDDMM-0+OT | 89.3 (± 5.20) | 95.3 (± 3.20) |

Table 5.1: Average precision and standard deviation in function of the assignment procedure for the three tested atlas-based annotation approaches: LDDMM-0, OT and LDDMM-0+OT. For each method, every tree was tested as reference template, so the total number of tree annotations per tested method is $50 * 49 = 2450$.

The first observation is that the best method without the bottom-up assignment is the combination of optimal transport and LDDMM registration. This is due to the fact that the Hungarian algorithm alone fails to separate spatially close arteries while the errors of the LDDMM deformations were described in the previous section. The combination of the two therefore compensates for some of these errors and manages to improve the average accuracy of the LDDMM-0 method by 5%. However, none of the three methods managed to reach a 90% annotation accuracy without using the bottom-up post processing.




Figure 5.2: Influence of the label assignment procedure on the data annotation with LDDMM-0. (a) Direct assignment and (b) bottom-up label assignment procedure.

The improvement for the LDDMM approach of the bottom-up post-processing allows to gain 9.1% in accuracy on average, which is almost double the gain brought by OT alone. The details of the labels that have been better predicted using post processing are illustrated in Figure 5.2, and highlight that the label prediction that has been the most improved is the Common Artery class, denoted "CA", which is logical since it is the label associated with the branches close to the root for which there is the most variability. The same observations can be made for the two other approaches, with for the optimal transport more false negatives for the "CA" label compensated by the post processing. In comparison, LDDMM-0+OT reaches the best average accuracy of 95.3%. The bottom-up assignment thus homogenized the results for the three atlas-based methods, which corresponds to the fact that the methods performed well for predicting the leaf labels, and less well for the inner branches of the tree.

Improving the Template It should be noted that the deformations performed so far have not used the labels of the branches, and we can now exploit these deformations to build a template from the registration of one case to all the others. We can thus take all the registrations made in the LDDMM-0 experiment (that do not use any annotation or a posteriori verification) to build 50 templates from the 50×49 registration already made. This template construction, illustrated in Figure 5.5 (a) and (b), allows us to calculate the accuracy performances for the LDDMM-1 procedure and its derivatives. We thus obtain an average accuracy for LDDMM-1 with direct annotation of $86.6\%(\pm 4.87)$, and with the bottom-up post processing an average accuracy of $96.8\%(\pm 3.40)$. Finally the LDDMM-1+OT method with bottom-up assignment provides a score of $97.2\%(\pm 3.20)$, which is the best result so far.

We still managed to achieve an average precision level on label prediction of 95% on all branches of the database by using one case to annotate all the others. This gives a good confidence for the automatic annotation of a database with few annotated cases. Moreover, building a template from a single annotated case also improves the average accuracy over the entire database by an additional 2%. By putting all the bricks together we went from

an average accuracy of 84.2% to 97.2%, from a single case used to annotate the rest of the database. The bottom-up assignment helps atlas-based methods by bringing consistency with the pelvic vasculature. The rule is a very light a priori in the assignment process, yet specific to this anatomy: it adapts well to the high anatomical variants and is well-suited to the atlas-based methods that are limited by the atlas' topology.

Comparing to Learning-Based Methods

We have shown that LDDMMs are suitable, with the right metric for the shapes comparison, for the registration of simplified vascular trees in our database. To provide an interesting automatic annotation result, we still need a post-processing consisting of a label assignment from the leaves to the root, and we also noticed that the optimal transport between the branches instead of a simple point-by-point vote to the nearest one allowed to make our method more robust to topological changes and to registration errors. We will now compare the results obtained with learning-based methods by varying the size of the database.

Baseline Comparison To demonstrate that atlas based techniques described in Section 5.1.2 perform well compared to learning based ones in case of small size database, we implemented two classification algorithms working on branch features, KNN and XGboost, inspired from the work of [Lo+11] and [Wan+17]. These algorithms compute the labels probabilities per branch, and will be compared with the same assignment procedures described in 5.1.2. Regarding the KNN method, all numbers of neighbours between 3 and 10 were tested. We kept only the 8-Nearest-Neighbor method that was giving the best results on our database. The XGBoost parameters are provided in [Wan+17]: the maximum depth in the tree classifier is set to 8, and other parameters such as the weight regularization, the sampling rate are provided in the paper as well. To be close to the work of [Lo+11] and [Mat+14], each branch is represented by a vector composed of 41 *branch features* which can be grouped in 2 types:

- **25 geometric features:** length (1), distance between end points (1), average and standard deviation of the curvature (2), invariant moments of 3D curves (3), eigenvalues and eigenvectors of the curve covariance matrix (12), endpoint direction (3), unit orientation (3)
- **16 tree features:** number of children (1), radius ratio to parent (1), average and standard deviation of the radius ratio to children (2), average and standard deviation of vector from tree root to children endpoints (6), vector from tree root to branch's barycentre (3), vector from tree root to branch's barycentre weighted by the points' radius (3).

When testing KNN, it turned out that using all the features deteriorated significantly the performance compared to using only a subset of them. In fact, by normalizing the features to embed them in an euclidean space, the KNN gives equal importance to all of them. We decided to keep the 10 most useful features selected by XGBoost during its training.

To compare all methods, precision has been evaluated using a cross-validation over the 50 cases with a training set and test set of varying size: from 2% (1 case) to 80% (40 cases)

of the total dataset. While the notion of cross-validation is well defined for learning-based techniques, in the case of atlas-based methods, the cross-validation follows the procedure proposed by [Cao+17]: the best atlas is selected in the training set with a leave-one-in procedure and used to annotate the test set. This case is then used to annotate the test set.

Baseline comparison In Table 5.2, we provide the percentage of correctly annotated branches and its standard deviation for both direct assignment and bottom-up assignment with a 90-10 split. KNN performs poorly compared to the other methods. The good performance of XGBoost illustrates in comparison that feature selection is key for a good generalization. XGBoost confusion matrix in Figure 5.6 (c) shows that the most confused labels are either corresponding to very similar branches or to highly variable ones.

| Method | Direct assignment | Bottom-up assignment |
|---------|-------------------------|-------------------------|
| KNN | 77 (± 4.0) | 76 (± 4.0) |
| XGBoost | 94 (± 2.0) | 92 (± 2.5) |
| OT | 95 (± 3.0) | 98 (± 2.4) |
| LDDMM-0 | 91 (± 1.8) | 98 (± 1.1) |

Table 5.2: Average precision and standard deviation in function of the assignment procedure over 10 folds cross-validation, 90% of the data for training. In this case the total number of annotations tested is $10 \times 5 = 50$

When assignment is performed at the branch level, OT gives the best results. It benefits from both a branch-level representation that reduces the effects of topological variations, and a transport plan that takes into account all the branch pairings to find the optimal solution.

When using the proposed task-adapted assignment, the accuracy of atlas based-methods increases significantly (as shown in Table 5.2). For LDDMM, it increases the performance to an higher accuracy than learning based methods, and equivalent to OT. The confusion matrix of LDDMM using bottom-up assignment is summarized in Figure 5.6. With such an assignment, the precision of LDDMM for the "Common Artery" is significantly improved while the overall recall is slightly deteriorated. This is the consequence of errors at obturators, superior and inferior gluteals propagated to the parent branches. The same observations can be made for OT to a lesser degree since both cost and probabilities predictions are computed at the branch level. In addition, we can see the effect of LDDMM registration in the standard deviation: the deformations smoothing the geometrical variations, the approach is less dependent on the atlas choice.

On the contrary this dedicated assignment slightly degrades learning-based techniques. Since they consider branches independently and do not take into account the labels probabilities of the other branches, they are not impacted by a topological variation and thus the bottom-up assignation procedure has no positive impact on their accuracy. As illustrated in Figure 5.6 (c),(d) the bottom-up assignment in fact propagates XGBoost errors to the parents.

Varying the Training Set Size We have shown that atlas-based methods can be competitive with learning-based methods on "large" available database. We now want to study the influence of the training set size as presented in Figure 5.3. For each method we used the assignment technique giving the highest precision (direct assignment for training-based methods and bottom-up for atlas-based ones). As expected, under 20% of training data (10 cases) precision of learning-based methods drastically drops. KNN is outperformed by other approaches, and XGBoost seems to asymptotically reach the atlas-based performances. On the other hand the atlas-based methods with bottom-up assignment perform with very little influence of the size of the training set.

Despite topological variations, the LDDMM-1 approach generates meaningful registrations showing good results when coupled with bottom-up assignment. OT also gives relevant branch matching that provides the same level of performance with the bottom-up assignment. Consequently LDDMM-1+OT with the bottom-up assignment have the best results, particularly in the case of small training sets. We illustrate in Figure 5.3 (b) the performances of atlas-based methods for only 5 cases in the training set (Confusion matrices for this setting are provided in the complementary material). The results of LDDMM-1+OT are significantly better than each method taken independently with an average $97.6\%(\pm 0.97)$ precision.




Figure 5.3: Comparison of the annotation performance according to the training set size.

Template estimation for annotation

To illustrate the contributions of the atlas construction on the labeling of the database we computed LDDMM-0 (pure registration) and LDDMM-1 (atlas construction) by registering each of the 50 available cases onto the others. It is illustrated in Figure 5.5 (a,b). The average precision of one reference case used in LDDMM-0 to annotate the 49 other trees is $93.3\%(\pm 3.5)$ when associated to bottom-up assignment and $84.2\%(\pm 4.4)$ using direct assignment. This drop of performance illustrates the sensitivity of atlas-based methods to




Figure 5.4: Precision along iterations with one case used in atlas estimation following [Vai+04].




Figure 5.5: Building atlases via LDDMM registration. Following [Vai+04], given one atlas we compute all the registrations onto the database and the corresponding average initial momenta. The colours represent the ground truth labels, but these labels were not used to compute deformations. (a) The initial trees; (b) Shooting along average momenta, iteration 1; (c) Shooting along average momenta, iteration 2.

the choice of the atlas in the first place. The bottom-up assignment post processing allows to overcome this sensitivity: we use it in the rest of the experiments. We then select one of the worst cases in the database regarding LDDMM-0 labeling performance and iteratively build the new atlas following the LDDMM-k procedure.

We can see that performance improves with iterations, which indicates that the atlas gradually captures the database variability: it allows a better registration hence a better label assignment. This single-case solution allows to annotate the 49 cases of the database with a precision reaching $98.9\%(\pm 0.33)$ while being one of the worst with LDDMM-0. It must be emphasized that the atlas construction did not rely on any other annotated case than the one initially selected. In Figure 5.5 (c) we observe that one iteration of the atlas construction is enough to greatly improve the labeling of the entire database, then the performance slowly increases until iteration 4.




Figure 5.6: 10-fold cross validation at training set size of 5 cases, label confusion matrices (a) OT (b) LDDMM 1 (c) LDDMM 1 + OT (d) XGBoost.

Discussion

We have chosen to use the computational anatomy approach to address both the lack of annotated data, particularly sensitive in medical imaging, and to study of the database statistically. For this purpose, we based ourselves on the LDDMM algorithm which allows to generate smooth and reversible deformations of a reference object, the template, on a target. These deformations can also be applied to the surrounding space using the flow of vector fields belonging to RKHS, and thus be inferred to other objects. The data are compared using metrics in RKHS and are inspired by the theory of geometric measurement, and allow to establish distances between continuous and discrete shapes – thus particularly adapted to discrete setting. We see in this first application that the proposed pipeline allows to set up reliable automatic annotation methods using very little annotated data. We obtained a good result with a single case iteratively aligned on the rest of the database and deformed according to the average deformation. It could be further improved by using multiple templates in the annotation process. This procedure does not required any branch or point pairing, and the new template can the be aligned onto the database for a better registration and consequently, a better annotation. We insist on the fact that the deformations of a selected case on the other ones in the training set did not take advantage of the available labels. A room for improvement in the case of training sets composed of more than one individual would be to compute a label-wise data attachment term: the sum of distances between subsets of the source and

target shapes sharing the same labels.

The overall automatic annotation pipeline then requires a database of unlabeled observations, and a single labeled template. The label assignment is made more robust through a bottom-up label assignment procedure adapted to the annotation of the pelvic vasculature. We see that this proposed method outperforms by far the learning based approaches when using few annotated data. It is also competitive in the case of larger labeled training sets.

The best result we obtained yet depends on a post-processing label assignment due to the limit of the diffeomorphism model used for the deformation: we are not able, despite the most realistic diffeomorphism possible, to manage the topological changes in the registration.

One of the major drawbacks of this method is the computational time. One LDDMM registration in this application took about 240 seconds. So the registration of the template onto the rest of the database takes approximately 3 hours and 20 minutes. If there are ways to reduce the computing time by reducing the number of variables, or approximating the diffeomorphisms, most of the computing is done offline. It could be used for instance to automatically annotate the database from few examples, and then feed a learning-based approach (as a deep neural network, or a XGBoost classifier) whose inference time would be lower. In addition, the data attachment term was taken in the space of varifolds, *that are non oriented*. The orientation is a relevant feature in the space representation, and a further study should evaluate the use of the oriented varifolds representation in the data attachment term for the registration pipeline.

5.2 Registration of a Template Vascular Tree onto Real Ones

In the previous section we have applied the diffeomorphic registration to an atlas-based automatic vascular trees annotation pipeline. As a matter of fact, the classic distances are not suited to the comparison of truncated and complete shapes as illustrated in Figure 2.13, and in Figure 5.7. In particular the distance in the space of oriented varifolds is not suited to the registration of the template vascular tree built in Section 5.1.3 onto real data. We now apply the partial matching introduced in Chapter 3 to guide diffeomorphic registration of the template onto real cases.




Figure 5.7: Diffeomorphic registration of a template (a) onto a target (b) using the distance in the space of varifolds. Despite being diffeomorphic, the result (c) has neither anatomical nor geometrical consistency. The target in (c) is shown in transparent black.

5.2.1 Database Description

We focus on the particular case of registering the previously constructed template onto a subset of the target, that is the hypothesis of Chapter 3. To that end, we selected 124 vascular trees containing the 6 leaves labels of the template, namely the prostatic artery, the superior vesicle artery, the obturator artery, the superior and inferior gluteal arteries and the pudendal artery. These trees, from which the manually simplified database of the previous section was extracted, was built and annotated by an expert during 3 years. On the contrary, the template tree is composed of 17 edges. Consequently we are not automatically labeling the whole target tree, rather evaluating the anatomical consistency of the registration. The target trees are composed of 413 branches in average, which yields a fairly imbalanced registration problem. In addition, to the number of branches, the targets may have supplementary labels, such as the middle rectal, the sacral or the iliolumbar arteries. With anatomically consistent deformations, such arteries should not be matched by the deformed target.

For the sake of light notation, we will denote by S the template tree (the source), when it should be written after the template construction used in the previous section: \bar{S}_3 . The output of the registration procedure of S onto the target T is denoted $\varphi_T(S)$. We have

that $S = (\mathcal{V}, \Sigma, x)$, and given a branch b_α its associated discretization points are denoted $x_\alpha = (x_{\alpha,i})_i \in (\mathbb{R}^3)^{n_\alpha}$, with $i \in \{1, \dots, n_\alpha\}$.

Registration Parameters We initialize the diffeomorphic registration of the template onto a full target by aligning the first point of both trees. Due to the segmentation and centerlines extraction procedure, this point is supposed to be close to the real point of injection of the contrast agent. Additionally it is manually identified during the segmentation procedure and used as reference point for the construction of the centerlines. The diffeomorphisms are generated with multiscale reproducing kernel similarly to the previous experiment.

5.2.2 Evaluation and Results

In this experiment we focus on evaluating the correct embedding of the deformed template into the target from the anatomical point of view. In fact we want to assess whether the branches we align with diffeomorphisms share the same label. Additionally we want to avoid strong deformations of the template such as strong distortions similar to the ones in Figure 2.13, or shrinking observed when using the partial matching without regularization as in Figure 3.13.

To assess these two information, we focus on two metrics. The first one is the total length of the centerlines in the deformed template tree, which is to be compared to the initial total length of the template tree (833mm). It provides a first criterion to detect inconsistent deformations: the ratio of the length of the deformed template with the length of the template:

$$r_l = \frac{\sum_\alpha \sum_{i=1}^{n_\alpha-1} |\varphi(x_{\alpha,i+1}) - \varphi(x_{\alpha,i})|_2}{\sum_\alpha \sum_{i=1}^{n_\alpha-1} |x_{\alpha,i+1} - x_{\alpha,i}|_2}.$$

The second metric is based on the point-wise projection of the deformed template vertices onto the target. Let $T = (\mathcal{V}', \Sigma', y)$ be the target vascular tree and q a point in \mathbb{R}^3 , we have the projection:

$$p(q, T) = \underset{y_{\beta,j} \in y}{\operatorname{argmax}} |q - y_{\beta,j}|_2.$$

To assess the correct embedding of the deformed template into the target, we verify that the label of the deformed source points and their associated projections in the target are equal. Let $\mathcal{L}_S : x_{\alpha,i} \in S \rightarrow [0, \dots, 12]$ be the labeling of the template points (both initial and deformed) and let $\mathcal{L}_T : y_{\beta,j} \rightarrow [0, \dots, 12]$ be the labeling of the target ones. We define the evaluation function per point:

$$f(x_{\alpha,i}, T) = \begin{cases} 1 & \text{if } \mathcal{L}_S(x_{\alpha,i}) = \mathcal{L}_T(p(x_{\alpha,i}, T)) \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

We can then define the score per label $f^l(S, T) = \frac{1}{\operatorname{card}(l, S)} \sum_{x_{\alpha,i} \in S, \mathcal{L}(x_{\alpha,i})=l} f(x_{\alpha,i}, T)$ with $\operatorname{card}(l, S)$ the number of points in S sharing the label l . Similarly we define a total

score:

$$\text{score}(S, T) = \frac{1}{\text{card}(S)} \sum_{x_{\alpha,i} \in S} f(x_{\alpha,i}, T),$$

with $\text{card}(S)$ the number of points in S .

Ablation Study We compare 3 different diffeomorphic registrations of the template tree onto the targets. First we use the partial matching regularized we selected in chapter 3 (Equation 3.3 with the local regularization term of Equation 3.14). Second we use the partial matching term without any regularization, to evaluate the influence of the regularization term. Third we use the classic distance in the space of oriented varifold as data attachment term, that we combine with the local regularization of Equation 3.14. These three registrations problems allows to compare the influence of the proposed partial matching on the registration of the template onto a subset of the target as well as the influence of the local regularization we proposed on the deformations. We do not evaluated the registration of the template onto real targets in the classic LDDMM framework and simply provide visual examples to illustrate the difference with the other approaches.

We summarize in the following Table 5.3 the average projection score total of the deformed template onto the 115 target trees of the database, as well as the average length of the deformed tree after all the registrations.

| Evaluation Criterion | Average length ratio to the template (%) | Pointwise total projection score (%) |
|--|--|--------------------------------------|
| Varifolds distance with local regularization | 181 | 38.9 |
| Partial Matching without regularization | 85.0 | 71.0 |
| Partial Matching with local regularization | 98.3 | 74.1 |

Table 5.3: Average length and total score assessing the embedding of the template into the targets after deformation. In total the registration was performed onto 124 real targets.

First we observe that the classic varifold distance fails at guiding the deformations of the template onto the real targets despite the additional regularization. The average total length of the deformed template is almost multiplied by a factor of 2. The amount of points correctly registered to a point in the targets is 38.9%. The partial matching without any regularization achieves better results with an average of 70.0% of the deformed points sharing a label with their projection on the target. The total length of the deformed trees however is reduced by 125mm in average, which is about 15% of the initial length of the template. We illustrates the typical inconsistent deformations associated with these two registration frameworks in Figure 5.8.

On the contrary the partial matching term associated with the local regularization achieves an average anatomically relevant matching of 72.9%. The performance is slightly better than

the non-regularized partial matching, the total length of the deformed trees however remains close to the one of the original template (98.7%) indicating smoother deformations.

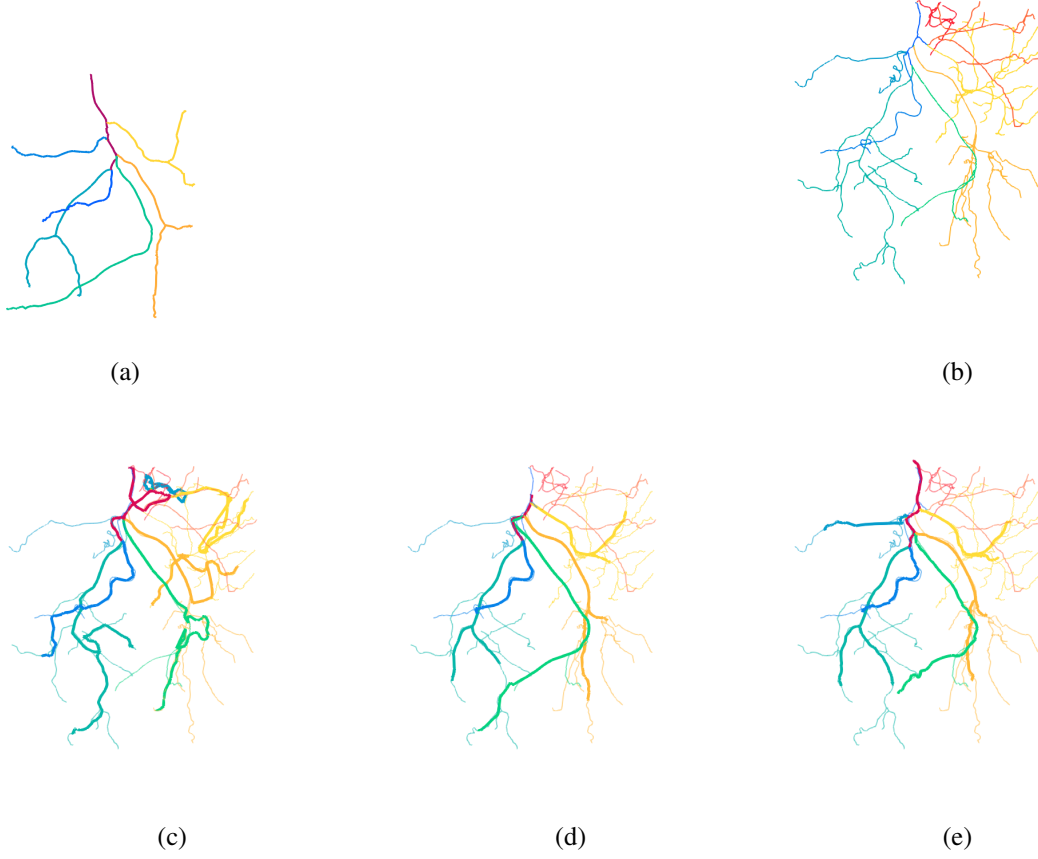


Figure 5.8: Output of the diffeomorphic registration of a template (a) onto a target (b). In (c-d-e) the deformed tree is superimposed to the transparent target. The deformation is inconsistent under varifold guidance with local regularization (c). The proposed partial matching allows the embedding of the source in the target (d), but it tends to shrink the deformed template (here, the superior vesicle artery). The partial matching with local regularization allows a consistent alignment of the source onto a subset of the target (e).

In Figure 5.8 are illustrated typical examples of diffeomorphic deformations guided by the three methods we study. The distance in the space of oriented varifolds is too high regarding the proposed local regularization we introduced, and the deformed source has been strongly distorted. The partial matching without any regularization provides more realistic deformations. The branches however are globally smaller than the template one. In this example the length ratio is $r_l = 87.8\%$, due to the shrinking of the superior vesicle artery. Finally the regularized partial matching produces regular deformations while correctly embedding the deformed template in the target.


Embedding Evaluation We now focus on evaluating the embedding of the deformed arteries of the template into subsets of the targets. We provide in Table 5.4 the projection scores

per label in the template using the three proposed registration approaches. The registration with regularized varifolds distance approach is outperformed by both partial matching approaches. The best average registration of one artery is achieved by the partial matching with the local regularization for the obturator artery. The point-wise projection score is 84.2% of the points of the source correctly matched to their pendant in the target. We observe that the obturator artery, and the superior and inferior gluteal arteries are slightly better registered with the partial matching (regularized or not). This corresponds to the arteries in the template that are composed of 3 branches. On the contrary the labels prostatic, inferior pudendal and superior vesicle arteries are showing lower point-wise projection score. The only artery for which the partial matching without regularization is performing significantly better than using the regularization is the prostatic one.


| Registration Method | Varifolds distance with local regularization | Partial Matching without regularization | Partial Matching with local regularization |
|--------------------------|--|---|--|
| Prostatic Artery | 36.0 | 65.8 | 61.1 |
| Superior Vesicle Artery | 30.0 | 56.6 | 69.3 |
| Obturator Artery | 33.7 | 79.7 | 84.7 |
| Inferior Pudendal Artery | 41.1 | 72.5 | 73.4 |
| Inferior Gluteal Artery | 45.1 | 78.1 | 77.9 |
| Superior Gluteal Artery | 58.9 | 73.4 | 78.1 |

Table 5.4: Projection score per label (in %) for the three registration procedures.


The average projection score per point in one artery strongly changes as one moves away from the root. We illustrate this for three arteries in Figure 5.9: the prostatic, the obturator and the superior gluteal arteries. The two later arteries are composed of one parent branch and two children. We observe for the prostatic artery (Figure 5.9(a)) that the average projection score for partial matching approaches increases with the depth in the tree. It is not true for the regularized varifold distance since one artery tends to be aligned to several target ones. In terms of projection scores, the two partial matching show similar trends for the three arteries. For all of them, the first branch is increasingly better matched to the target when going deeper in the tree. For the two arteries composed of 3 branches, the projection score per point does not much vary in the children branches when using the partial matching methods. On the contrary the behavior of the varifold distance varies a lot. In the case of the obturator artery, the projection score decreases with the depth for both children branches when it increases for one of the child branch of the superior gluteal arteries. In fact this particular branch is deformed to match several ones in the targets. However, the branches in this area all share the same label: superior gluteal artery. This last situation also helps the partial matching for which the projection score slightly increases as well.



(a)



(b)



(c)

Figure 5.9: PM stands for Partial Matching. The average projection scores per point for three arteries of the template: (a) the prostatic artery composed of one branch, (b) the obturator artery composed of three branches and (c) the superior gluteal artery composed of three branches. The first point is the closest to the template's root. The red dashes in (b) and (c) indicate the change of branch.

5.2.3 Discussion

The deformations we produce with the proposed method seems realistic. This is achieved thanks to the combination of the partial matching data attachment term and the local regularization. The partial matching term guides the diffeomorphic deformations to embed the template into the target. The regularization, formulated as a comparison to the initial template, preserves the relative position of the source points before and after diffeomorphism. It prevents for instance from inconsistent shrinkage of the template.

In the specific application to pelvic vascular trees, the registration of the template tree onto a subset of the target has numerous solutions. The cost function in the minimization problem has a lot of local minima. Some are avoided with the regularization of the diffeomorphic deformations and the local regularization term. The other way to prevent local minima is to refine the template: in fact the more branches in the source shape the less solutions there are to the embedding of the deformed template into the target. A single reference branch in the template as a representation of a branch of reference is for instance too restrictive. The number of solutions to the partial matching in such a case are too high.

In terms of projection scores, the only artery that is better anatomically matched without regularization is the prostatic artery. One possible explanation is the high geometrical variability of this artery. In this situation, the regularization does not allow to correctly align the deformed source onto its corresponding target artery, and the deformation seek another subset of the target.

The projection score at the points level improves as one moves away from the root of the tree. This is correlated to the tree topological changes that occur when comparing two trees with the same labels and the same number of branches. It is all the more true when comparing a simplified template to a real target composed of much more branches, leaves and labels. The leaves of the tree are however more and more spatially separated, and the labels are assigned according to the organs irrigated by the vessels. It is therefore normal that the leaves sharing the same label are better represented by their position when compared to the inner branches. In the perspective of an anatomically realistic partial matching, the leaves of the template (if the latter is well constructed) are therefore close to their targets. This is what we observe in the average scores per point that we have displayed along the curves associated with each label.

Anomaly Detection In about 60 cases of the database, the vascular tree had no obturator artery. Since it does not correspond to the actual partial matching framework we set-up, we discarded them in the application Section 5.2. In line with a real application, we would like to be able to work with all the existing most important labels despite some of them may be missing in the target from time to time. To illustrate a potential solution, we registered the template onto the 60 cases without obturator artery, and we show in Figure 5.10 two average deformations: onto the cases with the obturator and onto the ones without this artery. It appears clearly that they are different: in the case of absence of obturator in the target, the artery is shrunk by the average deformation while it does not move in the “normal” case.




Figure 5.10: Average shooting of the template (a). On a set of trees containing all the arteries of the template (b). On a set of trees that do not have the obturator artery (c). In transparent black we superimpose the initial position of the template.

This automatic detection of abnormal deformations paves the way to solutions to adapt the template tree to the new target: given a deformation belonging to the second distribution (the deformations to a target in which one artery is missing), one could remove the missing artery from the template and restart the registration procedure with the new template.

Further work A systematic study of the application of partial matching to the annotation of complete vascular trees should be carried out to assess the annotation potential of that method. Since the anatomical names of the arteries refer to entire subtrees, matching a simplified tree to a subpart of a complete one would indeed allow to extend the labels to all descending subtrees. In particular, arteries with typical shapes such as the pudendal and inferior gluteal appear to map well, and seems to be efficiently detected with the registration of the template under regularized partial matching guidance.

5.3 Beyond Vascular Trees: Feature-based Multi-modality Liver Volume Registration

So far we have only focused on the application of the proposed data attachment terms and the LDDMM registration to vascular trees and more generally curves or set of curves structures. The LDDMM framework however, as well as the space of oriented varifolds, has been developed for more diverse data namely submanifolds of \mathbb{R}^d . In line with these approaches, we proposed a partial matching data attachment term derived from the oriented varifold to build a diffeomorphic registration framework under partial matching adapted to various type of data. In particular the other usual structure extracted from 3D volumes is the surface or union of surfaces, for instance delineating a segmented volume in the 3D volumes.

This application is taken from the MELBA paper [Ant+22] in which we extended the partial matching application and applied it to the registration of CT/CBCT volumes based on the registration of a truncated liver surface extracted from one patient onto a complete one extracted from a CT volume.

As discussed in Chapter 1 Section 1.1.2, medical images are often acquired through different modalities, including ultrasound, computed tomography, and magnetic resonance imaging, each providing different and complementary information. In this context, image registration allows physicians to obtain combined inputs from different imaging modalities using for instance image comparison or fusion. The latter has been shown to be valuable in image-guided procedures, yielding less complications and decreasing radiation dose [RV16]. This section is the clinical application of the work introducing the partial matching in the space of varifolds [Ant+21] and the extensions proposed in 3.2.




Figure 5.11: Examples of slices from a CBCT volume (a) and CT volume (b) from the same patient visualized with a 3mm Maximum Intensity Projection (MIP). The liver and the tumors are only hardly visible in CBCT while the vessels are well depicted. The injected CT volume is complementary: it provides a well delineated liver.

5.3.1 CT/CBCT Volume Registration

Transcatheter directed liver therapies are part of the therapeutic arsenal of primary and secondary liver malignancies. The objective of these procedures is to locally treat the tumor and be as selective as possible (meaning placing the microcatheter used to inject the treatment as close as possible to the tumor) to preserve surrounding healthy tissues, while ensuring the destruction of the malignant cells.

These minimally invasive procedures are performed by navigating through the patient's arteries under real time 2D angiography, acquired through an imaging device called C-arm. Additionally, the latter can perform a 200 degrees rotation to allow 3D reconstruction of the patient's anatomy, called Cone-Beam Computed Tomography (CBCT) [Tac+15], to obtain a "live" 3D imaging of the patient at point of care. Performing CBCT during such procedure improves tumor detection and navigation guidance (Figure 5.11a).

Classically, preprocedural diagnostic CT scan or MRI are reviewed by the interventional radiologist to plan the procedure accordingly. The preprocedural acquisitions provide information on the entire liver anatomy, tumor burden and tumor feeding arteries that are decisive for procedural planning, such as number of tumors to be treated in one session and the dose of therapeutic agent to inject. Contrary to CT or MRI, CBCT is performed during the procedure, and the operator can compare procedural CBCT to preprocedural CT (Figure 5.11b) to ensure adequate treatment delivery.

While CBCT offers a superior spatial resolution compared to conventional CT scan, with intra-arterial injection of contrast agent providing a detailed visualization of the arteries, low contrast visibility is better in CT (Figure 5.11a). CBCT can also be subject to several artifacts such as beam hardening and motion artifacts that might decrease the CBCT performance to visualize the tumor, which is key to selective and successful treatment. A major difficulty in the fusion of a CT volume with a CBCT one comes from the fact that, unlike CT, liver is only partially visible in CBCT (due to the limited size of the field of view in the latter modality). In addition the acquisitions are taken at different times, potentially several weeks apart, and with different patient stances introducing deformations of the liver. For all these reasons, the two types of volumes are very different one from another as illustrated in Figure 5.11.

We propose a registration method based on liver surfaces¹ providing a deformation of the entire volume. To that extent we apply our partial matching dissimilarity term allowing to tackle the issue of partial correspondence between the truncated surface extracted from the CBCT and the one extracted from the CT. The registrations are evaluated on landmarks inside the liver, which were annotated by a physician.

5.3.2 Database Description

The database is composed of CBCT/CT pairs where CBCT have been acquired during hepatic arteriography and CT scans obtained at early or late arterial phase. We do not pro-

¹One of the feature that is visible in both the CT and CBCT whatever the clinical acquisition protocol, unlike vessels that necessitate a dedicated injection protocol




Figure 5.12: Examples of annotated Points of Interest on CBCT (a) and on CT (b). Examples of annotated longest axis diameter lesion visualized on CBCT (c) and on CT (d). All images come from the same patient.

vide the acquisition parameters here, yet the spatial resolution (in mm) of the volumes are $(0.45, 0.45, 0.45)$ for the CBCT and in average $(0.75, 0.75, 1.25)$ for the CT. Both were selected to show good visualization of the vessels and tumors. In total, 19 pairs of CT/CBCT liver volumes were evaluated, as the one illustrated in Figure 5.11.

Liver Segmentation

The livers were segmented in each modality using a deep neural network [MNA16] providing a binary volume in both modalities. In this application, the segmentations were evaluated by a clinical specialist and manually corrected if a major error was detected such as missing part of the liver. The underlying assumption being: staying as close as possible to clinical conditions. By doing so, we ensure that the registrations are based on features as reliable as possible. The mesh of the surfaces were then extracted and decimated, leading to meshes of approximately 10^4 points per surface. In the case of CBCT volumes, the meshes are then cut by the cylinder corresponding to the field of view using [Mus+21] resulting into truncated liver surfaces.

Points of Interest

For each patient, the branches of the proper hepatic artery visible on CT volume were annotated with Points of Interest (POIs) for evaluation purpose. Each POI was similarly annotated in the same location on the corresponding CBCT volume. Selected POIs often corresponded to arterial bifurcations that are easily identifiable on both CT and CBCT. For each pair of volumes, a physician annotated 10 POIs (Figure 5.12a, 5.12b). Because of the limited visibility of distal hepatic arteries on arterial phase CT compared to CBCT acquired during hepatic arteriography, most of POIs were located close to the bifurcation of the proper hepatic artery, thus mainly located at central parts of the livers, of importance to physicians.

Tumors Annotation

To evaluate the registrations, in addition to POIs, we annotated the longest axis diameter of a tumor according to [Gho+21] to ensure better reproducibility of the annotations across volumes (Figure 5.12c, Fig 5.12d). It was done in the axial view of the volumes for tumors that were visible in both modalities. The axis can be decomposed into 3 Tumor Points : the extremities and the center. In the database, one invasive tumor could not be annotated, reducing the number of pairs of tumors to 18. The annotated tumors were located in all the liver segments and their size varied from 9mm to 109mm. This variability in terms of position and size provides a complementary information to that of the POIs.

5.3.3 Liver Surface Registration with Partial Matching

As a first registration step, the truncated liver surfaces from CBCT were registered onto complete liver surfaces from CT scans with a LDDMM deformation model using the discrete framework described in Section 3.3. The LDDMM deformations of the truncated livers surfaces with partial dissimilarity function may lead to small shrinkage of the borders. To compensate this phenomenon in the application we added the a priori regularization of Equation 3.14 to the partial data attachment term that prevents from strong local deformations. For illustration purpose, one subject was registered twice with this model: once with the distance in the space of Varifolds, once with the normalized partial dissimilarity term (Def. 12) with the local regularization (Equation 3.14). This particular experiment is illustrated in Figure 5.13. Both results are generated with the same deformations and regularization parameters. As expected the varifolds distance leads to unrealistic deformations that tend to fill the holes in the source shape to cover the entire target. The deformations applied to the ambient space may be even more difficult to visualize. From the anatomical and medical point of view this is misleading and can not be used in a clinical application of multi-modality volumes registration. On the contrary the partial matching produces a more realistic deformation of the source onto a subset of the complete surface. We will only use and discuss this model in the following.




Figure 5.13: Registration of a truncated liver's surface (Live) from a CBCT (a) onto a complete liver's surface (Liver) from a CT (b) for Patient 2. Varifold registration (c); Partial normalized registration (Def. 12) with a local regularization (Eq. 3.14) (d). The color scale indicates the euclidean distance (in mm) of the points to their initial position before diffeomorphic deformation.

To initialize the LDDMM deformations, one classically performs a rigid registration. We find in the literature that the livers are principally deformed in translation, so we tested a set of combinations between rigid deformations and LDDMM. We selected the methods providing the best results: a translation followed by LDDMM (denoted *translation+LDDMM*) and a rigid deformation of limited angulation followed by LDDMM (denoted *rigid+LDDMM*). The rigid registration is limited to rotations between -15° and 15° around each axis that is the range of realistic rotations for the liver deformations.

In addition to these registration methods, we tested the standard rigid Iterative Closest Point (ICP) applied to the surfaces, using as data attachment term the function:

$$\Delta_{ICP}(S, T) = \frac{1}{Card(S)} \sum_{x \in S} \min_{y \in T} (|x - y|_{\mathbb{R}^d}).$$

By minimizing this term one minimizes the average distance of the source points to the target. This asymmetric function can be seen as a partial matching dissimilarity term, being equal to 0 if S is included in T .

Implementation details

LDDMM is computed with the partial matching dissimilarity term and the localized mass preservation term Eq. 3.14. In the following, J is written:

$$J(v) = \lambda_1 \int_0^1 \|v_t\|_V^2 dt + \underline{\Delta}(\phi_1^v(S), T) + \lambda_2 R_{local}(S, \phi_1^v(S)) \quad (5.2)$$

The optimization of this functional J is performed using Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm. In the LDDMM framework the cost of the deformations is controlled by the parameters λ_1 and λ_2 . To enforce smooth deformations of the surface as well as its ambient space, we set λ_1 to 10^7 , and we control the risk of shrinkage of the partial matching non-rigid registration by setting $\lambda_2 = 1$. The Reproducing Kernel of the deformations is the same as in 3.4 allowing large deformations of the shapes along with more detailed ones. To better register the shapes, we also use a multi-scale registration scheme for the data attachment terms by iterative applications with $\sigma_W = 10mm$ and $\sigma_W = 5mm$ the output of the optimization at scale 10 is used as input at the scale 5.

Computational Cost

We now provide some quantitative details about the computation time of the different steps of the registration procedure. To compute the registrations, we use a TITAN RTX Graphic Processing Unit. One rigid deformation of a $444 \times 512 \times 512$ voxel grid is computed in $9.10 \times 10^{-1}s$ when the diffeomorphic deformation takes 57.0 seconds. In terms of optimization, the ICP (on CPU) takes $4.64 \times 10^{-2}s$ for a source and a target of approximately 10^4 points. On the same data the rigid registration guided with our partial matching dissimilarity function takes $9.63 \times 10^{-1}s$ and the LDDMM optimization takes 4 minutes 34 seconds. This huge difference between LDDMM and the others methods comes from the number of parameters to optimize in its framework : the dimension of the space times the number of points in the source shape.

There are many possible ways to accelerate the LDDMM, from reducing the number of points in the source to the approximation of the diffeomorphisms with other deformations models such as successive combinations of splines. No matter the deformation, the computational cost of the partial matching dissimilarity function for the same data is $7.24 \times 10^{-3}s$.

5.3.4 From Surface to Volume Registration

The registration framework described so far focused on the spatial alignment of one shape onto another. We are interested in this application to the CT/CBCT volumes registration and both rigid deformation and LDDMM can be extended to the whole volume. To do so, the voxel grid of the source volume is deformed and interpolated on the target volume. In the case of LDDMM deformations, the voxel grid is deformed with the diffeomorphism of \mathbb{R}^3 as

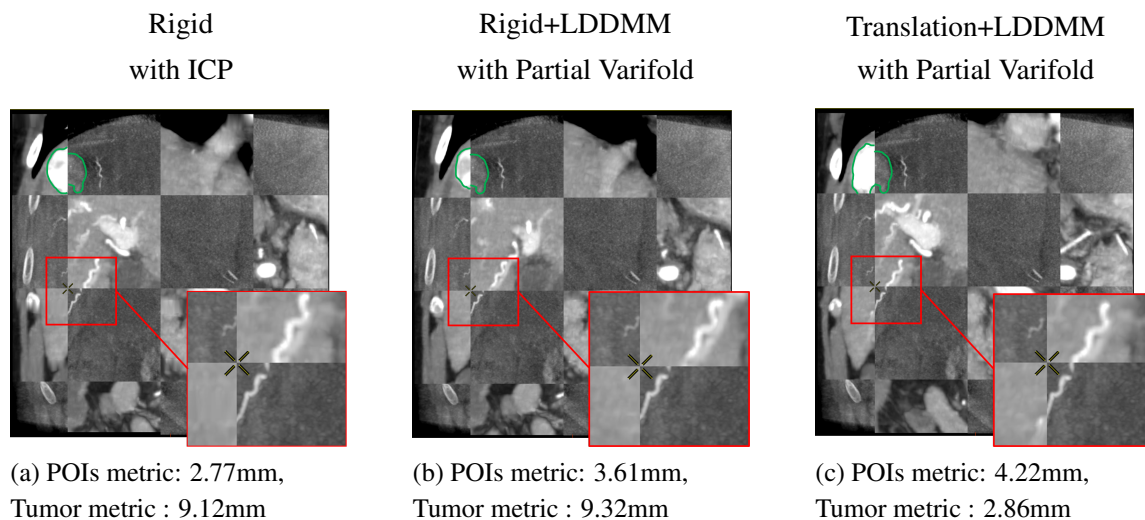


Figure 5.14: Tiled visualization of the registrations for Patient 11 for which the approaches using rigid deformations (a,b) register correctly the vessels but fail to align the tumor. The tiles of the CBCT target volume are the dark ones, those of the deformed CT volumes are the light ones.

described in 2.3.2. The values of the interpolated grid are then reported in the initial volume, providing a registration of the CT volume onto the CBCT one as illustrated in Figure 5.14.

The registrations of the volumes can be visualized to qualitatively assess the registration in the livers. To provide a 2D visualization of the results, we use in Figure 5.14 and Figure 5.15 a *tiled* representation that alternatively shows two volumes. Such visualization allows to see the continuities between the volumes. Each tile contains a 2D view of the CT volume (light tiles) or the CBCT one (dark tiles). As the dynamics of the images are very different, and the tissues that emerge differ from one modality to another, we are interested in the continuity of the emerging structures such as vessels, liver parenchyma or tumors.

5.3.5 Evaluation and Results

We recall that the key to clinical success is to register precisely the local area around the tumor despite the fact that this tumor is not segmented in CBCT clinical routine (and thus not usable in the registration procedure). Therefore the liver registration is evaluated through the euclidean distance between the deformed points of interest of the CBCT and those of the CT. It is done similarly with the tumors landmarks. Since the POIs are more centered than the tumor landmarks (see Figure 5.12), this second evaluation is significant as we will further explain in the discussion. The detailed results per case are provided in Table 5.5 and Table 5.6. In addition, we evaluate the registration at the surface level, which provides an indication of the overall registration quality, and gives the physician an additional benchmark for the comparison of CT and CBCT volumes. We first compute the barycenters registrations between the surfaces and apply the resulting translations to the volumes. We obtain an average distance of $18.3mm$ between the POIs, and of $21.5mm$ between the tumors landmarks. We show these values as red lines in the corresponding figures. They illustrate what the physicians can quickly obtain during the procedures, registering the volumes in translation by clicking one corresponding point in both modalities. As a reference method for the

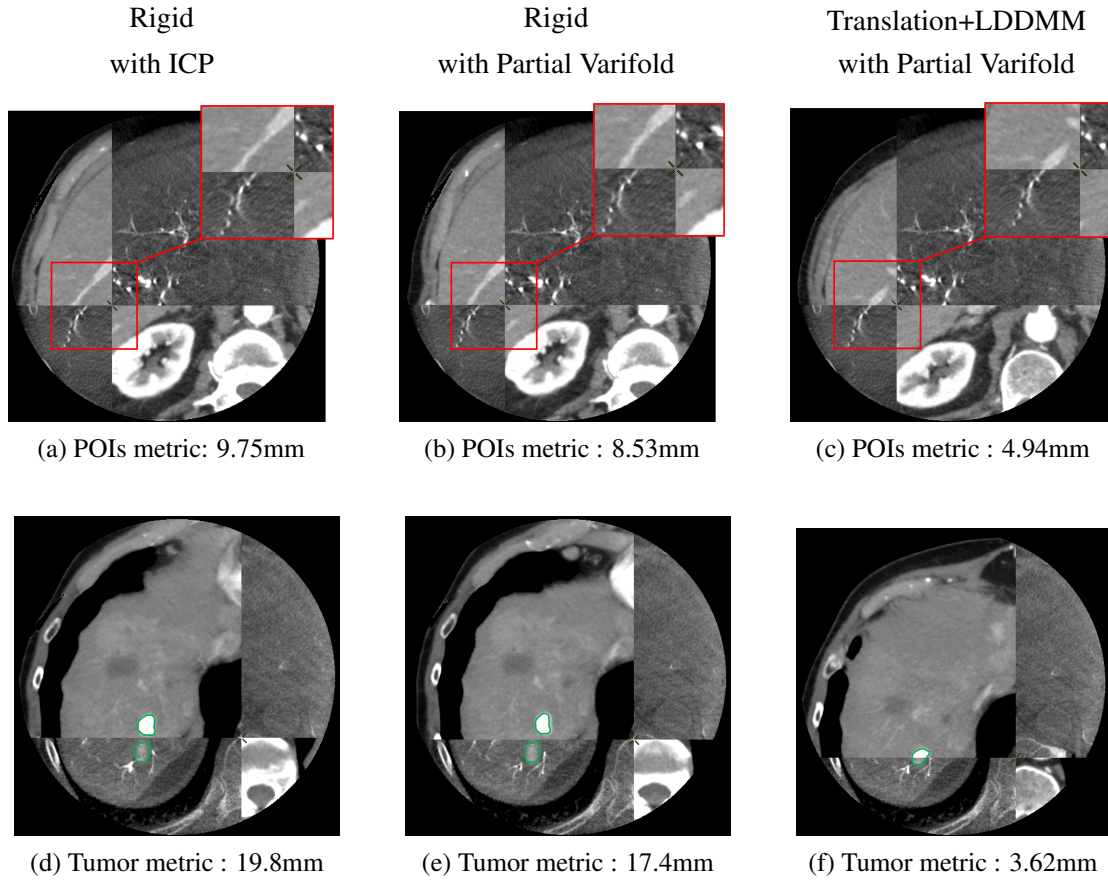


Figure 5.15: Tiled visualization of the registrations for Patient 2. The tiles of the CBCT target volume are the dark ones, those of the deformed CT volumes are the light ones. First row : sagittal view, second row : axial view.

rigid registration, we also computed the ICP registration directly based on the distance between the points of interest. This registration setting is the only one to exploit the annotated landmarks as input and it will only be used for quantitative comparison.

Therefore the liver registration is evaluated through the euclidean distance between the deformed points of interest of the CBCT and those of the CT. It is done similarly with the tumors landmarks. Since the POIs are more centered than the tumor landmarks (see Figure 5.12), this second evaluation is significant as we will further explain in the discussion. The detailed results per case are provided in Table 5.5 and Table 5.6. In addition, we evaluate the registration at the surface level, which provides an indication of the overall registration quality, and gives the physician an additional benchmark for the comparison of CT and CBCT volumes. We first compute the barycenters registrations between the surfaces and apply the resulting translations to the volumes. We obtain an average distance of $18.3mm$ between the POIs, and of $21.5mm$ between the tumors landmarks. We show these values as red lines in the corresponding figures. They illustrate what the physicians can quickly obtain during the procedures, registering the volumes in translation by clicking one corresponding point in both modalities.




Figure 5.16: Registration’s evaluation on the Points of Interest. The non-rigid LDDMM deformations based on partial matching allow robust surface registration while ensuring consistent deformation of the POIs. The red line corresponds to the average metric on POIs at initial position of registered barycenters.

Evaluation on the POIs

We computed the euclidean distance between the POIs in the target volumes and the deformed ones. The results are presented in two figures: the first one (Figure 5.16) provides a detail of the distances between the POIs as a function of the average distances from the points of the deformed surface to those of the target surface. This gives an idea of the distance between the edges of the liver after registration and the influence on the distance between the POIs. The associated box plots in Figure 5.17 provide a summary of the POIs registration results according to the method used. The right box in Figure 5.17 corresponds to the ICP rigid registration based on the POIs for the data attachment term and is used as reference. Note that this box shows that the variations cannot be explained by rigid deformations only. We observe in Figure 5.16 that the LDDMM deformations allow a consistent and robust registration of the surfaces with an average distance of the deformed source points to the target of about $2mm$ in average. This cannot be achieved by only rigid deformations guided by ICP ($4mm$ in average), but must be validated with other metric to assess the quality of the deformation applied to the whole livers volumes. In terms of POIs distances, none of the three methods illustrated in this scatter plot shows significant difference with the others, as validated in Figure 5.17.

Each of them performs differently depending on which patient they are evaluated as one can see in the detailed table in Appendix 5.7.1, but none of them stands out for the POIs metric. The best average performance $5.78mm \pm 5.32$ is achieved with the translation+LDDMM deformation guided by our partial dissimilarity term yet it is not significantly better than the




Figure 5.17: Registration’s evaluation on the Points of Interest. The rightmost box corresponds to the reference rigid registration of the POIs, hence the best possible results for rigid deformations. The red line corresponds to the average metric on POIs at initial position of registered barycenters.

rigid ICP ($6.49mm \pm 5.18$) or the rigid+LDDMM method ($6.54mm \pm 5.09$). When referring to the details in Appendix 5.7.1, we see that the LDDMM deformations significantly improve the translations (reducing by 44% the distance between the POIs). However, rigid registrations provide a poor initialization for the LDDMM deformations. By looking at the rotations angles obtained with the ICP based on the surfaces in Table 5.7 (Appendix 5.7.3), we observe a difference with those obtained with the reference rigid registration based on the POIs. The wrong rotations of the rigid ICP based on the surfaces come from the truncation of the source which can be interpreted as less constraints for the registration problem. Similar results are observed for the rigid registration guided by the partial varifold term. In such cases, the LDDMM fails to compensate for the error that causes the non-rigid deformation to start in a local minimum. It is illustrated in Figure 5.15, where the vessel and the tumor are clearly mismatched for the volumes deformed by approaches using a rigid transformation.

Inherently, we will not be able to do better than the reference method based on POIs, but we observe that the surface-based registration provides satisfactory results on the whole. When viewing the results, the points of interest are quickly found from one volume to another, even for Patient 14 which is the outlier. Even in this case (first row of Appendix 5.7.4) we can see that the structures are not so far apart visually, which illustrates a certain robustness of the registrations. In particular, the partial varifold term allows both rigid and non-rigid consistent registrations with respect to the metric on the POIs.

Figure 5.14 displays the results for Patient 11. This case corresponds to the best result in terms of POIs distance ($2.77mm$), which is achieved by the rigid method guided by the ICP. Yet this case illustrates that a good alignment of the POIs does not guarantee a good alignment of the tumors: tumors boundaries have been highlighted for better visibility.




Figure 5.18: Registration’s evaluation on the Tumors Landmarks. Only the Translation + LDDMM deformation method using Partial Matching preserves the same level of performance on the Tumors Landmarks as for the metric on POIs. The red line corresponds to the average metric on tumors using barycenters registration. The outlier (Patient 14) is excluded from this evaluation.

Evaluation on the Tumors

In the Figure 5.18 we show the metric results for the lesions landmarks that we resume per approach in Figure 5.19 and a detailed table is provided in Appendix 5.7.2. The first remarkable result is that the methods using rigid deformations fail to register the tumors closer than about $1cm$ in average while the translation+LDDMM guided by the partial varifold term maintains the same performance level as for the POIs metric with an average distance of $5.13mm$. In the scatter plot of Figure 5.18 we observe that the rigid ICP and the rigid+LDDMM are more spread than in the scatter plot with the metric on the POIs. In particular, the reference deformation optimized with the POIs does not perform well on the tumor registration. The main reason is that none of the POIs is located on the tumor. In fact the results of rigid ICP and rigid+LDDMM Partial Varifolds are similar to those of the reference rigid registration. These observations suggest that the rigid deformation, based on surfaces or POIs, is not always the global solution to the volume registration and may lead to local minimum. The non-rigid deformations driven by LDDMM do not improve rigid registration, despite the limitation of the rotation angles.

Moreover, the rigid ICP tends to minimize the average distance from the deformed source to the target. In some cases, as for patient 2 or patient 9, the truncation allows a freedom of deformation which makes the method unsuitable and leads to errors, whereas the reference method generating rigid deformations allowed to obtain good results on both POIs and tumor landmarks. In some cases (Patients 1, 2, 9 for instance), the reference rigid method is providing better results than the rigid ICP based on the livers surfaces, showing that despite




Figure 5.19: Registration’s evaluation on the tumors landmarks. The rightmost box corresponds to the application of the rigid registration computed with the POIs. The red line corresponds to the average metric on tumors using barycenters registration.

an acceptable global rigid deformation, the rigid ICP fails at retrieving the correct one using the surfaces to drive the registrations.

There is a clear difference between these methods and the registration with translation+LDDMM.

Figure 5.15 (Patient 2), second row, shows a case for which the translation+LDDMM was better overall. In such case, the rotations based on the surfaces could not explain the registration between the volumes with a poor result on the lesion metric (about $17.5mm$). On the contrary, the translation only guided by partial matching shows performance of $10.8mm$ that is further improved by $7mm$ with LDDMM for a distance between the lesion landmarks of $3.62mm$.

In Appendix 5.7.4, first row, is illustrated the worst case (Patient 14) for which the infiltrated tumor was not annotated, and the liver is hardly visible in CBCT. Comparing the results of rigid+LDDMM in Appendix 5.7.2 and Appendix 5.7.1, we see that the LDDMM deformations fail to improve significantly the rigid registration guided by the partial matching in the space of varifolds both regarding the POIs metric and the tumors one. The best initialization for the non-rigid deformations of LDDMM is the translation, providing consistent results for both metrics.

5.3.6 Discussion

We applied the proposed partial matching dissimilarity term to the registration of pre-op CT volumes on CBCT volumes acquired during the interventions based on segmented liver

surfaces. These surfaces, truncated in CBCT, present non-rigid deformations between one another due to differences in time point and patient stances in each modality. Partial matching can be used in a rigid registration process, providing results equivalent to a standard method like ICP. It is important to note that the registered surfaces in this application are relatively smooth, which favors the ICP for the rigid registration part. The deformation of the patient, though, is non-rigid: the spinal chord merely moves when the liver may be deformed. If a rigid deformation within the liver may suffice in some cases, it cannot be applied to the whole volume. Moreover, the ICP approach in the non-rigid case can lead to projecting several points on one single point of the target, and does not take into account the local orientations of the objects or their resolution. On less regular anatomies, one could expect less good performances.

The proposed Partial Matching can also be used with diffeomorphic shooting, providing the tools of computational anatomy and allowing non-rigid, yet regular, deformations. The diffeomorphic deformations lead to accurate registration of the liver surfaces (about $2mm$ on average), which gives the physician a first tool to easily compare CT and CBCT volumes during the procedure. However, the areas of interest for the procedures are also within the liver, and care must be taken to ensure that the non-rigid deformations generated from the surfaces generalize well within the liver volume.

The distance between the Points of Interest, close to the bifurcation of the proper hepatic artery, provide a first evaluation metric to the registration methods. However, since they do not cover the volume of the livers correctly, it does not allow to discriminate between the deformation models. Regarding their location, a rigid registration is sufficient to align them when it fails to extend to the tumors landmarks as illustrated in Figure 5.14. In particular, the extremities of the livers seem to be deformed by non-rigid deformations which can be explained by the better performance of the translation+LDDMM regarding the distance on the tumor landmarks. In this case, the rigid deformations with rotation lead to a local minimum that the LDDMM are unable to compensate. On the contrary, translation+LDDMM provide a consistent registration of both POIs and tumors. This last finding indicates that the correct deformations to generate to register a liver from a CBCT acquisition to the liver from the CT acquisition are translations and local non-rigid deformations without large (or even no) rotations.

Although the methods based on surface registrations are sensitive to surface extraction, we have seen through one case illustrated in Appendix 5.7.4 that the proposed registrations remain suitable, despite the outlier (Patient 14) with respect to the metric on the POIs. In particular natural the regularization of the LDDMM associated with the one proposed in this paper provide a smooth and consistent non rigid registration of the surface and its ambient space, ensuring realistic registrations of the volumes.

Regarding the computation time, we are studying a non-rigid multi-modality volume registration solution, with one of the volumes acquired during the procedure. The current computation time does not allow to use this registration as is in a real time procedure, however there are many levers to accelerate the diffeomorphic deformations, main source of computation cost. These solutions, such as limiting the number of control points (and therefore variables), could allow the use of this solution in an application used during the procedures.

This study was conducted to provide a visualization tool for the transcatheter directed liver therapies in minimally invasive procedures. The results obtained with the translation+LDDMM method thus provide robust metrics around $5mm$ on average, which is useful for physicians in the perspective of navigating their tools in the patient's anatomy to locate structures that are hardly visible in the CBCT used during the procedure. Such registrations could facilitate the physician's intervention by providing, for example through image fusion, an improved visualization of the pre-procedure CT volume tumor placed in the CBCT volume. This would allow the physician to avoid redoing a traditional CBCT to see the tumor correctly, thus limiting the X-ray dose sent to the patient and the time of the procedure.

5.4 Conclusion

We have shown that our work on atlas construction and shape registration (potentially under partial matching constraints) is suitable for many situations. We first applied LDDMM to the registration of an labeled vascular tree on a non-labeled database. These registrations allowed us to build a template that captures the variability of the data. This template construction, which is used as a new initialization of an LDDMM registration, allows to improve the deformations and thus to be robust to the choice of the reference individual at the beginning. It also allows the labels transfer from the deformed template to the targets, and an adapted post processing makes the global pipeline very efficient for the annotation of vascular trees — with the same number of leaves.

The template can also be realistically registered on more complete trees. This is made possible thanks to the partial matching formulated as a data attachment term. Combined with a dedicated regularization we designed, the partial matching guides the diffeomorphisms towards a realistic registration on a sub-part of the target. Partial matching can also be used to align a truncated surface to a complete one, which has been exploited in the context of liver surfaces from CT (complete liver) and CBCT (truncated liver) acquired volumes. The generated deformations can be applied to the volumes, which allows realistic and automatic volume alignment.

These two applications of partial matching, which are a priori very different, give confidence in the applicability of the solution for interventional radiology. This paves the way to a lot of possible applications that are yet to be explored.

Automatic Labeling At least two applications could be derived from the registration of the simplified template onto a new target. The projection of the deformed template onto a subset of the complete target provides – up to a post-processing – a corresponding *extracted* annotated subtree that can be embedded in the space of tree-like shapes in which the template lives.

Similarly to the bottom-up labeling of Section 5.1, the labels could be propagated from the annotated points in the target to the root, and to the descendants leaves. By doing so we should be able to retrieve a larger subtree. Provided the correct registration of one artery, for instance the obturator, one could extract the whole real obturator subtree.

Anomaly Detection In many works analyzing populations of anatomical shapes, the detection of outlier, or anomaly, is one key ingredient to the characterization of a disease. In the case of the registration of the template onto real observations, this could be used, for instance, to detect a missing artery or an inconsistent deformation of an artery in the template. Such automatic detection in a registration procedure could help adapting the template to the target. This could be done by detecting deformation modes corresponding to pathological cases.

Changes in Bifurcations Now that we have seen that we can register the template tree to a real target despite the large difference in term of number of branches, the next step would be to change the topology of the template during the registration as well.

The possible combinations of LDDMM and tree space representation with edge attributes (that are control points of several Bézier curves) are multiple. First, given an extracted tree in the real target from the partial matching of the template, one could deduce the topology of the extracted target tree. It is then possible to build a geodesic path between the template embedded in the tree space and the extracted subtree with the labels of the template.

Second, we have shown that both the construction of a relevant template and its diffeomorphic alignment to new targets were both key in atlas-based automatic annotation solutions. We provided theoretical tools to do it in Chapter. 4. Instead of deducing the topology a posteriori, the next step is to combine these topological changes to the registration of the template simplified trees onto the real targets.

Data augmentation Beyond pure registration purpose, one promising application of the alignment of the template onto unlabeled targets is the statistical analysis one can perform on the deformations. The relevant deformations and their principal components can exhibit shapes associated with a disease, outliers as well as provide generative tools for the construction of synthetic data. We have constructed tools for statistical analysis of both the spatial and topological changes. This is one step toward a more involved data augmentation, based on both topological changes along geodesics in the treespace and diffeomorphic deformations. Such tools could be used during the training of deep learning methods [Wu+19; Yan+20; Zho20; Ham+21] for the automatic annotation of the vascular trees or other applications as the vessels segmentation or the artifact correction in the CBCT reconstruction.

Visualization Tools One of the key steps in data comparison is the way it is displayed to the clinicians. For the navigation of the interventional tools, having a labeled 3D vascular tree could provide an augmented roadmap in the fluoroscopy with anatomical information projected on the vessels. It could also provide some filtering tools to reduce the amount of projected vessels from the CBCT volume to the fluoroscopy by removing from the segmented tree (hence in the roadmap) the additional irrelevant arteries. The application to real time images projection deserves a study on its own, which was studied in numerous works [Mar+12; BMV15].

Beyond Vascular Trees We developed these registration tools to be as versatile as possible: such tools should be adaptable to other applications and frameworks or shapes – mostly curves and surfaces of \mathbb{R}^d . The hierarchical structures that can be encoded as trees with features are numerous, and one can imagine applying the proposed framework to various data, from pulmonary trees close to the structure of the vascular ones, to proteins that are slightly different or more abstracted structures, and whose edges can be represented with features in some euclidean space.

Converse Partial Matching The registration of the CT/CBCT was done using the surfaces extracted from the segmentations. In order to be able to deform the whole volume in which the livers were embedded, we had to use a large scale for the reproducing kernel associated with the deformations. However, the control points and initial momenta encoding the diffeomorphic deformations do not need to be placed on the surface of the liver, and we could improve the volume registration by having better located deformations.

We could use for example take advantage of the multiple data living in the same ambient space: both the liver and the hepatic vascular tree can be segmented in CT and CBCT. Due to its spatial resolution the CT scan fails at imaging the small vessels. Hence the vascular tree extracted from the CT would be a truncated version of the vascular tree extracted from CBCT. In this case it would be the same vascular tree acquired with different modalities: the partial matching is perfectly adapted.

The framework, though, is converse to the one we used in Sections 5.2 and 5.3: the deformed vascular tree from the CBCT should include the target tree from extracted from the CT volume. We showed in Chapter 3, Section 3.4 that the proposed regularized partial matching was adapted to this converse situation. One could then compute LDDMM registration with two different partial terms, applied to livers surfaces and vascular trees, yielding better spatially located initial vector fields for the diffeomorphic shooting. Thus better volume deformations.

More generally, the problem of finding a trade-off between the amount of deformations needed to register shapes subset, and the size of the selected subsets in the target and source shapes is a problem that arises with a number of practical applications in computer vision and computational anatomy. It happens when the data comes from different acquisition systems, procedures or subjects (e.g two CBCT of different patients, or two CBCT of one patient at different times), or even if the data are extracted by different specialists. To provide a concrete example, take two truncated surfaces corresponding to the same organ truncated differently and deformed. Without supplementary a priori, one cannot decide which one to include in the other. On the contrary, they can be both registered to one complete target. Conversely, and for better deformations of the ambient space, one could register a complete template to truncated targets. This could be used for instance to better understand the non-rigid deformations of the livers both across the patients and under the deformations caused by the artificial breath holding during CBCT acquisitions.

Conclusion and Perspectives

In this work we focused on the problem of template vascular tree construction and registration in the context of atlas-based automatic annotation. Among the different steps of a generic atlas-based approach, we chose to study the construction of an atlas adapted to the data, and the selection of a template within this atlas. The motivation was to build a template as representative as possible of the data, labeled or not. Provided a realistic alignment of the template onto a target, the automatic atlas-based labeling is facilitated.

Automatic Vascular Trees Annotation with LDDMM registration

We first introduced a new automatic vascular trees annotation pipeline, based on the construction of an atlas and the registration of a template using the LDDMM. In this work we showed the benefits of using diffeomorphic registration, and that the atlas-based approach was adapted to the case of little annotated data.

In general the template and its atlas are built for a specific purpose, and a priori is often used for their construction. For instance to identify the prostatic artery in the pelvic vascular tree, one has to include it in the template. In the case of vascular trees, the atlas is usually built following one of the following methods:

- Keep in the template and the atlas all the arteries listed a priori or detected in the database.
- Select among the existing arteries a subset of arteries of interest that are expected in new observations.
- Select the set of arteries that are common to all the annotated observations in the database.

Such template constructions, though, do not take advantage of the available data that are not annotated. The LDDMM does not require the trees to be annotated, and the atlas construction can be done with the whole dataset of unlabeled observations by registering one case onto the rest of the observations and deriving statistics on the deformations. We showed that the template, taken as average shape obtained with the average diffeomorphic deformation of the a reference shape onto the rest of the database, and its registration, were good initialization for the automatic atlas-based annotation procedure one wants to use. In addition it allows to be robust to the initial choice of the reference shape. We still used a bottom-up label assignment post-processing that takes advantage of the a priori on the arteries labels in the pelvic vasculature.

We showed in Chapter 5 that – as expected – our proposed atlas-based approach outperforms by far the learning based approaches in the case of small training sets. To keep the comparison fair, the template in this study was constructed using only the training dataset for the learning based methods.

Most of the time the template is a simplified version of the real data. When the template has a tree structure, this often introduces topological differences with the real observations: in terms of bifurcations ordering and number of branches. Dealing with such topological changes was the second part of this work.

Diffeomorphic Registration and Topological Changes

One limitation of the diffeomorphic registration of vascular trees is the behavior in the case of topological changes between the source and the target. In fact, the diffeomorphisms tend to align the whole source shape onto the whole target. When the source shape is a truncated version of the target, if the diffeomorphisms are not properly regularized, this leads to abnormal deformations of the source branches, hence inconsistent clinical shapes correspondences. In addition, when the bifurcations ordering is different from the source to the target, the diffeomorphisms cannot explain such changes and lead to abnormal shrinking or extension of the branches. These two observations led us to adapt the data attachment term and the template representation.

Partial Matching in the case of Inclusion To address the problem of partial shape matching in term of number of branches, we formulated in Chapter 3 a new data attachment term in the space of oriented varifolds which is adapted to the inclusion of the deformed source shape into the target. The proposed partial dissimilarity term can be differentiated which makes it compatible with diffeomorphic shooting and rigid registration.

We explained the construction of the proposed function and illustrated its behavior on a toy example. To cope with the shrinkage induced by diffeomorphic deformations guided by the partial matching, we proposed a regularization term preserving varifold associated with the source shapes. The global regularization term simply aims at preserving the norm of the varifold associated with the source. The local version aims at preserving the value of this later at the vertices of the shape.

We showed that the partial matching function associated with the local regularization allows to build a non rigid registration framework to include the deformed source into the target, or conversely include the target into the deformed source. Both are one step in the direction of real applications of template registration onto real observation potentially truncated. Thanks to the varifolds representation, we applied this framework to two kind of shapes extracted from medical images: the vascular trees and livers surfaces. In both cases it shows consistent deformations of the shapes, and in the case of the livers, allows a deformation of the 3D volumes from which the livers surfaces were extracted. It thus provides a multi-modality volumes registration.

The proposed partial matching data attachment term could yet be improved by better disentangling the local interaction between the discretization points of the shapes. In particular this could prevent from an over-evaluated attraction to the center of mass of the target, or the shrinking of the source.

One of the natural extension of the partial matching we proposed would be in the space of functional shapes and normal cycles. For the latter, the definition of the inclusion would be a challenge of its own since the model work on the unit normal bundle to the shapes: the

normal cycle associated with a truncated shape is not included in the one associated with the corresponding complete shape.

Registration and Hierarchical Changes The second kind of topological changes we addressed during this work is the changes in terms of bifurcations ordering. To that end we adapted in Chapter 4 the tree-like shapes representation of [Fer+12b]. The template shape is embedded in the space of tree-like shapes in which the edge attributes are Bézier curves control points in \mathbb{R}^d allowing to go back and forth between the template and the vascular tree representation. In this space of tree-like shape, one topology is associated with one euclidean space called orthants, which are "glued" together at the boundaries.

This tree-like shape representation provides a versatile space that takes into account the topological changes and a consistent representation of the vessels centerlines. The Bézier curves provide a naturally regularized representation reducing the dimension of the orthants that facilitates the optimization in the space of tree-like shapes. To change the topology of the template, the distance of its position to the boundary of the current orthant is computed. If the template is closer to the boundary than a threshold, we move to the boundary and test each orthant contiguous to the boundary on which the template was projected. The new topology is selected as the one maximizing the gradient of the cost function with respect to the edge attributes.

This proposed minimization scheme has two main advantages: it is adapted to the registration of the template tree with potential changes in the bifurcations order, which is naturally handled during the minimization, and it does not require the target to be labeled. The second point is necessary for the tree-like shapes representation. The result of the minimization procedure can then be used to automatically annotated the target tree, or as an element of the tree-like shapes space as well.

In this space, one can derive statistics on a set of observations and geodesics between them. It is therefore perfectly adequate to the construction of an atlas for a database of vascular – or other biological – trees. The cost functions used to compare the vascular trees in this framework relies on the differentiable unbalanced optimal transport cost between the branches centerlines. The resulting transport plan is one solution to deal with the differences in terms of number of branches between the source and the target. In addition, this is one solution to the problem of comparing unlabeled tree shapes while authorizing topological changes, as it is emphasized in [FN20].

Another option would be to use different centerlines trees distance to compare the shapes. For instance [DKS18] propose a square root velocity formulation to compare neurons shapes. An elastic metric is constructed capturing the amount of bending and stretching needed to align the curves. The advantage of the tree space representation as control points of Bézier curves is the versatility of this representation, and the capacity to generate various spatial representations for the vascular trees. The distance function between the vascular trees can then be adapted, and, as long as it can be differentiated, we are able to back-propagate the gradient in the space of tree-like shapes.

Similarly, a different distance between curves could be used in the Optimal Transport mapping computation between the curves of the template tree and the ones of the target. In fact, using a simple $L2$ distance between curves sampled by a fixed number of points may not be the best metric for comparing the centerlines of the trees. In particular, the distance is

sensitive to the sampling method, and if one solution is to increase the number of sampling points, the square root velocity trees (seen as collections of square root velocity functions) and elastic metrics were designed for this purpose.

Statistics in the Space of Tree-like Shapes A challenge for statistics in tree-like shapes spaces would be to be able to compute the principal geodesics without having to fix the ends of the geodesic with points of the database. This has been done in the case of phylogenetic trees [Nye14] but is far from being transposed to the case of tree-like shapes. A particularly interesting approach would be to adapt the notion of parallel transport [You+08; LP13; Gui+21] to geodesics in such spaces. Indeed, we know how to project a point of these space on a geodesic, and we also know how to move in this space according to a vector. It would then be interesting to develop an optimization procedure on the extremities of the geodesic according to the projections of the trees of the database on this geodesic.

A second open question is the use of initial momenta of diffeomorphic shooting as edge attributes: if the edge attributes encode the whole vector field v associated with the diffeomorphism φ_t^v , the position in the space of tree-like shape the tangent to the space of diffeomorphisms at identity. In such a case, by endowing the orthants with the metric on V , one would get a stratified space that is not composed of euclidean spaces glued together anymore: it would correspond to manifolds glued together. The main issue to this approach would be the way the manifolds are glued together, and the norm of the diffeomorphism at the boundaries between the orthants.

In conclusion, atlas based method can be very effective, in particular in the case of little annotated data, when coupled with relevant deformation models. This is true even for shapes as complex as vascular trees, with important topological differences both in terms of number of branches and order of bifurcations.

In the case of large annotated databases, other solutions as learning-based ones should also be explored, and a combination of the two approaches would bring out the best in each: statistics and robust methods as well as fast inference and good features extraction.

Appendices

5.5 Matrix Valued Kernels

Vector valued kernels were introduced in the seminal work of Laurent Schwartz [Sch64] and used for instance in [CA04] to generate thin plate splines for the interpolation of dense and sparse vector fields. In [MG13] the authors focus on kernels inducing translation and rotation invariant kernels that are the ones we are interested in for geometric interpolations.

Let denote by $\tau : x \mapsto x + t$ a translation of some fixed $t \in \mathbb{R}^d$, and by $\rho : x \mapsto Rx$ a generic rotation for some fixed $R \in O(d, \mathbb{R})$ the orthogonal group over \mathbb{R} of dimension d .

Theorem 15 (Translation Invariance). *Let H be a RKHS with kernel $K : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}^{d \times d}$. The map $u \mapsto u \circ \tau$ is an isometry in H for any translation τ if and only if there exists a function $\mathbf{k} : \mathbb{R}^d \mapsto \mathbb{R}^{d \times d}$ such that $K(x, y) = \mathbf{k}(x - y)$ for all $x, y \in \mathbb{R}^d$.*

A kernel K satisfying the conditions of Theorem. 15 is called translation invariant. The authors in [MG13] also introduce the projection matrices:

$$Pr_x^\parallel := \frac{xx^T}{|x|^2} \text{ and } Pr_x^\perp := I_d - \frac{xx^T}{|x|^2}, \quad x \in \mathbb{R}^d \setminus \{0\} \quad (5.3)$$

These projections onto the vector x and its orthogonal planes will be used in the decomposition of translation and rotation invariant kernels:

Lemma 3. *Let $\mathbf{k} : \mathbb{R}^d \mapsto \mathbb{R}^{d \times d}$ be a generic matrix valued function. The following are equivalent:*

(i) *for any $x \in \mathbb{R}^d$ and $R \in O(d, \mathbb{R})$ we have that $\mathbf{k}(-x) = \mathbf{k}(x)^T$ and*

$$\mathbf{k}(Rx) = R\mathbf{k}(x)R^{-1};$$

(ii) *there exists a scalar $k_0 \in \mathbb{R}$ such that $\mathbf{k}(0) = k_0 I_d$ and two scalar functions k^\parallel and k^\perp such that for all $x \in \mathbb{R}^d$, $x \neq 0$:*

$$\mathbf{k}(x) = k^\parallel(|x|)Pr_x^\parallel + k^\perp(|x|)Pr_x^\perp.$$

Theorem 16 (Rotation Invariance). *Let H be a RKHS with a translation invariant kernel ($K(x, y) = \mathbf{k}(x - y)$ for all $x, y \in \mathbb{R}^d$). The map $v \mapsto \rho v \circ \rho^{-1}$ is an isometry in H for any rotation $\rho : x \mapsto Rx$, with $R \in O(d, \mathbb{R})$ if and only if Equation 3 holds for all $R \in O(d, \mathbb{R})$.*

The kernels \mathbf{k} inducing translation and rotation invariant inner product can thus be written under the form of Equation 3. The associated vector fields are then functions $v : x \mapsto \mathbf{k}(x)\alpha$ for $\alpha \in \mathbb{R}^d$. In the case of Gaussian scalar-valued functions $k(x) = b \exp(-c|x|^2)$ with $b, c > 0$, one can show that the function:

$$\mathbf{k}(x) = a \exp(-c|x|^2) xx^T + \left(\frac{(d-1)a}{2c} - a|x|^2 \right) \exp(-c|x|^2) I_d, \quad \text{with } a, c > 0$$

is a divergence free kernel.

Similarly, the function:

$$\mathbf{k}(x) = -a \exp(-c|x|^2) xx^T + \frac{a}{2c} \exp(-c|x|^2) I_d, \quad \text{with } a, c > 0$$

defines a curl free kernel. These kernels are illustrated in Figure 5.20.




Figure 5.20: Matrix Gaussian kernels (a) Curl free, (b) Divergence free, and (c) the sum of the two, a classic scalar kernel (2.2).

5.6 Normal Cycles

We have seen in Chapter. 2 Section. 2.4 that oriented varifolds allow us to define a framework in which currents and non-oriented varifolds are special cases. This allows us to create metrics suitable for many varieties of curves and surfaces, ranging from smooth shapes to noisy and sub-sampled ones. However, one can be interested in rectifiable sets and sub-varieties that represent the shapes of organs in anatomies, and other numerically analyzed objects. These shapes frequently have edges: this can either be due to acquisition limitations (resolution, field of view) or the objects structures themselves (ends of vessels, heart...). In such cases the matching solution using a data attachment in the space of varifolds may lead to degenerated solutions.

The specificity of this metric is to use a model of currents on the unit normal bundle of the shapes, encoding second-order information (i.e. curvature). As presented in [RG16], this representation shows better results at the areas of high curvature and at singular points such as bifurcations or curves' ends. In addition we can also derive kernel metrics on normal cycles seen as currents over the unit normal bundles.

Definition 38. Let $d \in \mathbb{N}^*$. A m -differential form is a mapping $x \mapsto \omega(x)$ from \mathbb{R}^d to the space of alternating m -linear form of \mathbb{R}^d . A m -current in \mathbb{R}^d is a continuous linear mapping from the space of continuous m -differential forms to \mathbb{R} .

The space of test functions considered for the Normal Cycles is $C_0^{d-1}(\mathbb{R}^d \times \mathbb{S}^{d-1})$, the space of continuous $(d-1)$ -differential forms vanishing at infinity. In line with the varifolds discussion, we still consider curves and surfaces in \mathbb{R}^d , $d = 2, 3$ but the notions can be extended to any m -differential forms in \mathbb{R}^d , $d \in \mathbb{N}^*$.

Definition 39. Unit Normal Bundle and Normal Cycle Consider a compact C^2 curve or surfaces \mathbf{X} with or without boundary. The unit normal bundle of \mathbf{X} is:

$$\mathbf{N}_{\mathbf{X}} = \{(x, \vec{n}) | x \in \mathbf{X}, \vec{n} \in \text{Nor}(\mathbf{X}, x)\},$$

with $\text{Nor}(\mathbf{X}, x)$ the set of all unit normal vectors of \mathbf{X} at x . The **normal cycle** of \mathbf{X} associated with $\mathbf{N}_{\mathbf{X}}$ is:

$$N(\mathbf{X})(\omega) := [\mathbf{N}_{\mathbf{X}}](\omega) = \int_{\mathbf{N}_{\mathbf{X}}} \langle \omega(x, \vec{n}) | \tau_{(x, \vec{n})} \mathbf{N}_{\mathbf{X}} \rangle d\text{vol}(x, \vec{n}),$$

with $\omega \in C_0^{d-1}(\mathbb{R}^d \times \mathbb{S}^{d-1})$.

In particular for curves and surfaces in \mathbb{R}^3 , the unit normal bundle is a 2-dimensional surface in $\mathbb{R}^3 \times \mathbb{S}^2$ and the integration is made over a surfaces: $d\text{vol}(x, \vec{n})$ can be written $dS(x, \vec{n})$. In the case of curves in \mathbb{R}^2 , the unit normal bundle corresponds to a curve in $\mathbb{R}^2 \times \mathbb{S}$.

In particular, we are interested in unions of curves and surfaces, and the additive property of the Normal Cycles requires a condition on the smoothness of the shapes, the positive reach defined in [Fed59] which encompasses the C^2 compact submanifolds with boundaries.

Proposition 21. *Consider $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$ such that \mathbf{X}_1 , \mathbf{X}_2 , and $\mathbf{X}_1 \cap \mathbf{X}_2$ are C^2 compact submanifolds with or without boundaries. Then $N(\mathbf{X}) = N(\mathbf{X}_1) + N(\mathbf{X}_2) - N(\mathbf{X}_1 \cap \mathbf{X}_2)$.*

This is specially important in the discrete setting in which the curves are unions of segments and the surfaces unions of triangles that are in fact C^2 compact submanifolds with boundaries. In general, the normal cycles allow to take into account the singular points, as the bifurcations in a tree of central curves or the extremities of curves and surfaces. Indeed for such points a sub-part of the unit normal bundle is associated, and one can imagine them as areas of high curvature.

Kernel Metrics on Normal Cycles Once again we can derive kernel metrics on normal cycles seen as currents over the unit normal bundles, and we can design reproducing kernels $k = k_e \times k_t$ on $\mathbb{R}^d \times \mathbb{S}^{d-1}$. This construction detailed in 2.4.2 allows to write the inner product between normal cycles:

$$\langle N(S), N(T) \rangle_{W'} = \int_{N(S)} \int_{N(T)} k_e(x, y) \cdot k_t(\vec{n}_2, \vec{n}_2) \langle \tau_{(x, \vec{n}_1)} \mathbf{N}_S | \tau_{(y, \vec{n}_2)} \mathbf{N}_T \rangle d\text{vol}(y, \vec{n}_1) d\text{vol}(x, \vec{n}_2) \quad (5.4)$$

This last equation is only true for C^2 compact submanifolds with or without boundaries, and is adapted to discrete curves and surfaces through the additive property. Contrary to the varifold representation, the inner product in the space of normal cycles pays attention to both normal vectors to the shapes with $k_t(\vec{n}_2, \vec{n}_2)$ and to the tangent vectors to the unit bundles with $\langle \tau_{(x, \vec{n}_1)} \mathbf{N}_S | \tau_{(y, \vec{n}_2)} \mathbf{N}_T \rangle$. These two terms introduce a supplementary sensitivity to the curvature of the shapes. When working with discrete curves and surfaces, represented (see Section. 2.4.1) with unions of points, segments (edges of triangles, segments of curves) and triangles for the faces of surface meshes. The associated unit normal bundle structures are then respectively spheres, cylinders, and triangles. The spatial integration is approximated over those structures like for varifolds using weighted Diracs. For such structures, the following property simplifies the computation of the inner product with kernel metrics:

Proposition 22. *The spherical, cylindrical and planar components are orthogonal with respect to the kernel metric on normal cycles defined in Equation. 5.4.*

So the computation can be done independently on each component.

The choice of the function k_t defines again the space of the test functions used to compare the shapes, and allows (provided the same conditions on k_e as), to define a distance $d(S, T) = |N(T) - N(S)|_{W'}$, instead of a pseudo distance.

For the linear kernel $k_t(u, v) = \langle u, v \rangle$, following the notations of Section. 2.4.1, the inner

product for discrete curves (composed of vertices and segments) is written:

$$\begin{aligned} \langle N(S), N(T) \rangle = & \frac{\pi^2}{2} \sum_{i=1}^n \sum_{j=1}^m k_e(c_i^S, c_j^T) \langle \tau_i^S, \tau_j^T \rangle \cos(\alpha_{ij})^2 \\ & + \frac{16\pi^2}{3} \sum_{i=1}^N \sum_{j=1}^M k_e(x_i^S, x_j^T) \left(1 - \frac{n_{x_i^S}}{2}\right) \left(1 - \frac{n_{x_j^T}}{2}\right) \end{aligned} \quad (5.5)$$

with $n_{x_j^T}$ the number of segments connected to the vertex x_j^T in T , and α_{ij} the angle between the tangent vectors τ_i^S and τ_j^T , n, m the number of segments in S, T and N, M the number of vertices. More details about the computation of the terms can be found in [RG16], Annex A. The first half of the term is associated to the cylindrical part of the curves, and the second term corresponds to the spherical part.

For the surfaces, the author in [RG16] Annex A.2.2 only provide a truncated version of the inner product with linear kernel:

$$\begin{aligned} \langle N(S), N(T) \rangle = & 4 \sum_{i=1}^n \sum_{j=1}^m k_e(c_i^S, c_j^T) a_i^S a_j^T \langle \tau_{c_i} S, \tau_{c_i} T \rangle + \frac{1}{3} \sum_{k=1}^N \sum_{l=1}^M k_e(x_k^S, x_l^T) \times \\ & \left[\pi(4 - 2n_{x_k^S} + 2N_{x_k^S}) - 2 \sum_{i=1}^{N_{x_k^S}} \alpha_{i, x_k^S} \right] \left[\pi(4 - 2n_{x_l^T} + 2N_{x_l^T}) - 2 \sum_{j=1}^{N_{x_l^T}} \alpha_{j, x_l^T} \right] \end{aligned} \quad (5.6)$$

with n, m the number of triangles in S, T and N, M their number of vertices. The value $n_{x_k^S}$ corresponds to the number of vertices connected to x_k^S and $N_{x_k^S}$ the number of triangles with vertex x_k^S . Finally α_{i, x_k^S} is the angle of triangle i at vertex x_k^S .

5.7 Partial Matching Application to CT/CBCT volumes registration

5.7.1 POIs Detailed Results per Patients

| Patients | ICP Rigid | Rigid | Rigid + LDDMM | Translation | Translation + LDDMM | ICP Rigid Based PoIs |
|------------|-------------|-------------|------------------|-------------|------------------------|-------------------------|
| Patient 0 | 5.18 | 4.29 | 4.15 | 18.6 | 4.71 | 2.47 |
| Patient 1 | 7.33 | 22.5 | 15.8 | 14.1 | 7.86 | 4.32 |
| Patient 2 | 9.75 | 9.88 | 8.53 | 12.8 | 4.94 | 4.92 |
| Patient 3 | 5.9 | 5.97 | 5.09 | 9.09 | 6.14 | 2.87 |
| Patient 4 | 5.04 | 5.95 | 6.05 | 8.31 | 6.89 | 5.66 |
| Patient 5 | 5.1 | 6.14 | 4.89 | 5.06 | 3.37 | 2.26 |
| Patient 6 | 4.26 | 4.42 | 6.08 | 6.91 | 10.8 | 1.28 |
| Patient 7 | 4.06 | 2.93 | 3.75 | 6.36 | 3.6 | 1.74 |
| Patient 8 | 4.72 | 3.55 | 4.76 | 5.87 | 3.67 | 2.36 |
| Patient 9 | 13.7 | 7.4 | 9.14 | 14.0 | 6.26 | 1.66 |
| Patient 10 | 3.7 | 2.93 | 2.88 | 5.25 | 2.93 | 1.65 |
| Patient 11 | 2.77 | 3.48 | 3.61 | 8.87 | 4.22 | 1.48 |
| Patient 12 | 5.97 | 5.31 | 6.02 | 9.03 | 6.3 | 3.02 |
| Patient 13 | 6.53 | 5.71 | 5.56 | 13.4 | 5.6 | 2.47 |
| Patient 14 | 18.9 | 22.3 | 22.8 | 28.1 | 13.6 | 1.69 |
| Patient 15 | 2.97 | 2.3 | 2.55 | 6.46 | 3.59 | 2.64 |
| Patient 16 | 4.12 | 3.34 | 3.75 | 4.07 | 4.4 | 2.28 |
| Patient 17 | 6.17 | 11.0 | 3.1 | 5.72 | 5.32 | 2.34 |
| Patient 18 | 7.17 | 7.05 | 5.74 | 13.9 | 5.76 | 1.89 |
| Average | 6.49 | 7.18 | 6.54 | 10.3 | 5.79 | 2.58 |
| Stdev | 3.93 | 5.83 | 4.95 | 5.9 | 2.66 | 1.18 |
| Median | 5.18 | 5.71 | 5.09 | 8.87 | 5.32 | 2.34 |

Table 5.5: Average distance between the Points of Interest per case

5.7.2 Lesions Detailed Results per Patients

| Patients | ICP Rigid | Rigid | Rigid + LDDMM | Translation | Translation + LDDMM | ICP Rigid Based PoIs |
|------------|-------------|-------|------------------|-------------|------------------------|-------------------------|
| Patient 0 | 5.19 | 3.94 | 5.11 | 17.6 | 3.31 | 13.6 |
| Patient 1 | 11.2 | 18.2 | 6.05 | 1.41 | 2.93 | 5.2 |
| Patient 2 | 19.8 | 18.8 | 17.4 | 10.8 | 3.62 | 2.21 |
| Patient 3 | 16.6 | 16.8 | 16.0 | 16.7 | 11.2 | 14.0 |
| Patient 4 | 7.14 | 9.34 | 7.92 | 9.78 | 6.91 | 13.7 |
| Patient 5 | 5.68 | 6.77 | 6.03 | 7.25 | 3.18 | 8.0 |
| Patient 6 | 4.67 | 4.32 | 4.15 | 4.37 | 5.7 | 12.0 |
| Patient 7 | 3.24 | 4.02 | 2.14 | 2.8 | 0.59 | 6.39 |
| Patient 8 | 6.76 | 6.17 | 7.48 | 6.77 | 4.93 | 5.76 |
| Patient 9 | 23.3 | 18.4 | 16.6 | 6.86 | 3.94 | 8.68 |
| Patient 10 | 13.9 | 14.4 | 15.1 | 13.0 | 8.81 | 9.15 |
| Patient 11 | 9.12 | 7.49 | 9.32 | 3.66 | 2.86 | 18.7 |
| Patient 12 | 7.56 | 8.22 | 8.0 | 6.73 | 4.48 | 4.04 |
| Patient 13 | 21.4 | 21.2 | 17.7 | 22.7 | 8.55 | 13.4 |
| Patient 14 | None | None | None | None | None | None |
| Patient 15 | 6.69 | 7.36 | 8.28 | 9.49 | 6.36 | 19.3 |
| Patient 16 | 5.31 | 4.7 | 5.03 | 6.96 | 3.3 | 6.97 |
| Patient 17 | 5.56 | 12.5 | 8.22 | 7.29 | 7.37 | 2.87 |
| Patient 18 | 7.83 | 7.75 | 7.21 | 12.0 | 4.32 | 7.44 |
| Average | 10.1 | 10.6 | 9.31 | 9.23 | 5.13 | 9.52 |
| StDev | 6.06 | 5.71 | 4.8 | 5.37 | 2.56 | 5.08 |
| Median | 7.35 | 7.98 | 7.96 | 7.27 | 4.4 | 8.34 |

Table 5.6: Average distance between the tumors landmarks per case. The invasive tumor could not be annotated for Patient 12.

5.7.3 Rotations Deformations Comparison

| Patients | Rigid ICP Based POIs | | | Rigid ICP Based Surfaces | | |
|------------|----------------------|--------|-------|--------------------------|--------|-------|
| | X | Y | Z | X | Y | Z |
| Patient 0 | -1.15 | -11.52 | 14.2 | -1.53 | -12.18 | 10.7 |
| Patient 1 | -5.58 | -11.68 | 0.8 | -9.69 | -6.51 | 3.89 |
| Patient 2 | -8.66 | -6.28 | -0.03 | -17.39 | -2.59 | 0.64 |
| Patient 3 | -13.49 | -0.11 | 2.58 | -13.03 | 0.51 | 1.33 |
| Patient 4 | -9.02 | 3.17 | -10.8 | -8.37 | 0.93 | -9.45 |
| Patient 5 | -2.37 | -4.13 | -2.43 | -4.96 | -1.93 | 1.89 |
| Patient 6 | -5.09 | -1.27 | -3.64 | -6.42 | 0.65 | -4.96 |
| Patient 7 | -6.13 | 2.37 | -5.04 | -4.85 | 0.85 | -3.43 |
| Patient 8 | -6.36 | 4.98 | 6.14 | -6.81 | 3.45 | 2.47 |
| Patient 9 | -0.77 | -5.45 | 8.57 | -3.49 | -12.56 | 6.15 |
| Patient 10 | 2.49 | 6.94 | 3.18 | -0.72 | 6.71 | 3.68 |
| Patient 11 | -10.15 | -5.21 | 1.52 | -9.34 | -3.62 | 0.27 |
| Patient 12 | 5.1 | 4.96 | -6.6 | 5.3 | 4.18 | -4.14 |
| Patient 13 | -5.25 | -2.79 | 11.2 | -2.5 | -6.93 | 10.2 |
| Patient 14 | 3.3 | -12.94 | -4.14 | -9.25 | -16.26 | 6.0 |
| Patient 15 | -2.38 | 6.39 | -5.4 | -3.88 | 6.68 | -5.35 |
| Patient 16 | 1.71 | -1.56 | 1.45 | -2.71 | -2.01 | 0.31 |
| Patient 17 | -0.93 | -1.01 | 4.24 | -4.11 | -2.33 | 7.57 |
| Patient 18 | -3.28 | 10.3 | -4.61 | 4.76 | 9.8 | 0.54 |

Table 5.7: Rigid ICP Rotation Values in degree along each axis, based on the Points of Interest and based on the Surfaces. In bold are highlighted the most important differences in terms of rotation angles.

5.7.4 Worst Case Scenario : Bad Feature Extraction

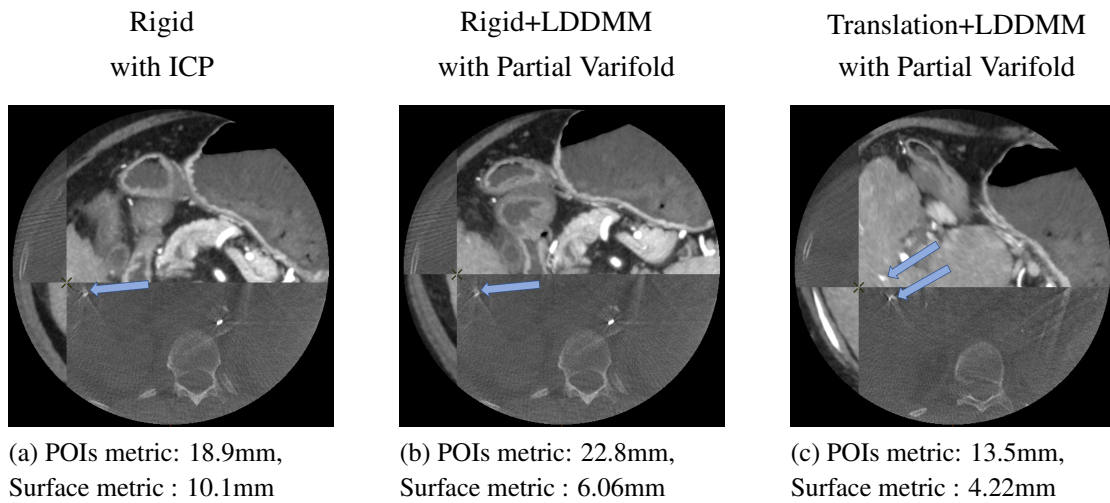


Figure 5.21: Tiled visualization of the registrations for Patient 14. The infiltrated tumor could not be annotated. The tiles of the CBCT target volume are the dark ones, those of the deformed CT volumes are the light ones.

BIBLIOGRAPHY

- [AMO93] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice hall, 1993.
- [AMC08] Dror Aiger, Niloy Mitra, and Daniel Cohen-Or. “4-Points Congruent Sets for Robust Pairwise Surface Registration”. In: *35th International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH’08)* 27 (Aug. 2008). DOI: [10.1145/1399504.1360684](https://doi.org/10.1145/1399504.1360684).
- [Aki+09] A. Akinyemi et al. “Automatic labelling of coronary arteries”. In: (Aug. 2009), pp. 1562–1566.
- [AF18] Sufian AlBadawi and Muhammad Moazam Fraz. “Arterioles and Venules Classification in Retinal Images Using Fully Convolutional Deep Neural Network”. In: (2018).
- [AAT07] Stéphanie Allasonnière, Yali Amit, and Alain Trouvé. “Towards a coherent statistical framework for dense deformable template estimation”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69.1 (2007), pp. 3–29.
- [AS01] Benjamin L Allen and Mike Steel. “Subtree transfer operations and their induced metrics on evolutionary trees”. In: *Annals of combinatorics* 5.1 (2001), pp. 1–15.
- [Ang+04] Dragomir Anguelov et al. “The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces”. In: *Advances in neural information processing systems* 17 (2004).
- [Ant+20] Pierre-Louis Antonsanti et al. “Database Annotation with Few Examples: An Atlas-Based Framework Using Diffeomorphic Registration of 3D Trees”. In: (2020), pp. 160–170.
- [Ant+22] Pierre-Louis Antonsanti et al. “How to Register a Live onto a Liver? Partial Matching in the Space of Varifolds”. In: *arXiv preprint arXiv:2204.05665* (2022).
- [Ant+21] Pierre-Louis Antonsanti et al. “Partial Matching in the Space of Varifolds”. In: (June 2021).
- [Anx+98] R Anxionnat et al. “3D angiography. Clinical interest. First applications in interventional neuroradiology.” In: *Journal of Neuroradiology= Journal de Neuroradiologie* 25.4 (1998), pp. 251–262.
- [Aro50] Nachman Aronszajn. “Theory of reproducing kernels”. In: *Transactions of the American mathematical society* 68.3 (1950), pp. 337–404.

-
- [Ash07] John Ashburner. “A fast diffeomorphic image registration algorithm”. In: *Neuroimage* 38.1 (2007), pp. 95–113.
- [Ass+15a] André Assis et al. “Pelvic Arterial Anatomy Relevant to Prostatic Artery Embolisation and Proposal for Angiographic Classification”. In: *Cardiovascular and interventional radiology* 38 (May 2015). DOI: [10.1007/s00270-015-1114-3](https://doi.org/10.1007/s00270-015-1114-3).
- [Ass+15b] André Moreira de Assis et al. “Pelvic arterial anatomy relevant to prostatic artery embolisation and proposal for angiographic classification”. In: *Cardiovascular and interventional radiology* 38.4 (2015), pp. 855–861.
- [ATS+09] Brian B Avants, Nick Tustison, Gang Song, et al. “Advanced normalization tools (ANTS)”. In: *Insight j* 2.365 (2009), pp. 1–35.
- [Bač14] Miroslav Bačák. “Computing Medians and Means in Hadamard Spaces”. In: *SIAM Journal on Optimization* 24.3 (2014), pp. 1542–1566. DOI: [10.1137/140953393](https://doi.org/10.1137/140953393).
- [Bag+17] Sandeep Bagla et al. “Cost analysis of prostate artery embolization (PAE) and transurethral resection of the prostate (TURP) in the treatment of benign prostatic hyperplasia”. In: *Cardiovascular and interventional radiology* 40.11 (2017), pp. 1694–1697.
- [Bai+12] Jordan Bai et al. “Atlas-based automatic mouse brain image segmentation revisited: model complexity vs. image registration.” In: *Magnetic resonance imaging* 30 6 (2012), pp. 789–98.
- [Bas+18] Fereshteh S. Bashiri et al. “Multi-Modal Medical Image Registration with Full or Partial Data: A Manifold Learning Approach”. In: *Journal of Imaging* 5 (Dec. 2018), p. 5. DOI: [10.3390/jimaging5010005](https://doi.org/10.3390/jimaging5010005).
- [Beg+04] Mirza Faisal Beg et al. “Computational cardiac anatomy using MRI”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 52.5 (2004), pp. 1167–1174.
- [Beg+05] Mirza Faisal Beg et al. “Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms”. In: *International Journal of Computer Vision* 61 (Feb. 2005), pp. 139–157. DOI: [10.1023/B:VISI.0000043755.93987.aa](https://doi.org/10.1023/B:VISI.0000043755.93987.aa).
- [BMV15] Thomas Benseghir, Grégoire Malandain, and Régis Vaillant. “A tree-topology preserving pairing for 3D/2D registration”. In: *International journal of computer assisted radiology and surgery* 10.6 (2015), pp. 913–923.
- [BMV13] Thomas Benseghir, Grégoire Malandain, and Régis Vaillant. “Iterative Closest Curve: A Framework for Curvilinear Structure Registration Application to 2D/3D Coronary Arteries Registration”. In: (2013). Ed. by Kensaku Mori et al., pp. 179–186.

-
- [Bil+11] Tiago Bilhim et al. “Prostatic arterial supply: demonstration by multirow detector angio CT and catheter angiography”. In: *European radiology* 21.5 (2011), pp. 1119–1126.
- [BHV01] Louis J. Billera, Susan Holmes, and Karen Vogtmann. “Advances in Applied Mathematics”. In: (2001), pp. 733–767.
- [Bog+13] H. Bogunović et al. “Anatomical Labeling of the Circle of Willis Using Maximum A Posteriori Probability Estimation”. In: *IEEE Transactions on Medical Imaging* 32.9 (Sept. 2013), pp. 1587–1599.
- [BH99] Martin R Bridson and André Haeffliger. *Metric spaces of non-positive curvature*. Vol. 319. Springer Science & Business Media, 1999.
- [Bro+09] Alexander Bronstein et al. “Partial Similarity of Objects, or How to Compare a Centaur to a Horse”. In: *International Journal of Computer Vision* 84 (Aug. 2009), pp. 163–183. DOI: [10.1007/s11263-008-0147-3](https://doi.org/10.1007/s11263-008-0147-3).
- [BBK06] Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. “Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching”. In: *Proceedings of the National Academy of Sciences* 103.5 (2006), pp. 1168–1172. ISSN: 0027-8424. DOI: [10.1073/pnas.0508601103](https://doi.org/10.1073/pnas.0508601103). eprint: <https://www.pnas.org/content/103/5/1168.full.pdf>. URL: <https://www.pnas.org/content/103/5/1168>.
- [BSS19] Andrew D Brown, Steffan F Stella, and Martin E Simons. “Minimally invasive treatment for benign prostatic hyperplasia: economic evaluation from a standardized hospital case costing system”. In: *Cardiovascular and Interventional Radiology* 42.4 (2019), pp. 520–527.
- [Bro+22] Benjamin Brown et al. “Nontarget Radiopaque Embolic Deposition during Prostatic Artery Embolization”. In: *Journal of Vascular and Interventional Radiology* 33.5 (2022), pp. 558–563.
- [Bul+06] Thomas Bulow et al. “Point based methods for automatic bronchial tree matching and labeling”. In: 6143 (2006).
- [Bun71] Peter Buneman. “The Recovery of Trees from Measures of Dissimilarity”. English. In: (1971), pp. 387–395.
- [CA04] Pascal Cachier and Nicholas Ayache. “Isotropic energies, filters and splines for vector field regularization”. In: *Journal of Mathematical Imaging and Vision* 20.3 (2004), pp. 251–265.
- [CFV20] Anna Calissano, Aasa Feragen, and Simone Vantini. “Populations of Unlabeled Networks: Graph Space Geometry and Geodesic Principal Components”. In: *MOX Report* (2020).
- [Cao+17] Q. Cao et al. “Automatic identification of coronary tree anatomy in coronary computed tomography angiography”. In: *The international journal of cardiovascular imaging* 33.11 (2017), pp. 1809–1819.

-
- [Cao+06] Yan Cao et al. “Diffeomorphic matching of diffusion tensor images”. In: *2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW’06)*. IEEE. 2006, pp. 67–67.
- [Car+21] Francisco Cesar Carnevale et al. “Advanced image guidance for prostatic artery embolization—a multicenter technical note”. In: *Cvir Endovascular* 4.1 (2021), pp. 1–7.
- [Car+17] Francisco Cesar Carnevale et al. “Anatomical variants in prostate artery embolization: a pictorial essay”. In: *Cardiovascular and interventional radiology* 40.9 (2017), pp. 1321–1337.
- [Cer+13] Can Ceritoglu et al. “Computational Analysis of LDDMM for Brain Mapping”. In: *Frontiers in neuroscience* 7 (Aug. 2013), p. 151. DOI: [10.3389/fnins.2013.00151](https://doi.org/10.3389/fnins.2013.00151).
- [CCT17] Benjamin Charlier, Nicolas Charon, and Alain Trouvé. “The Fshape Framework for the Variability Analysis of Functional Shapes”. In: *Found Comput Math* 17, 287–357 (2017).
- [Cha+21] Benjamin Charlier et al. “Kernel Operations on the GPU, with Autodiff, without Memory Overflows”. In: *Journal of Machine Learning Research* 22.74 (2021), pp. 1–6. URL: <http://jmlr.org/papers/v22/20-275.html>.
- [CT13] Nicolas Charon and Alain Trouvé. “The Varifold Representation of Nonoriented Shapes for Diffeomorphic Registration”. In: *SIAM Journal on Imaging Sciences* 6.4 (2013), pp. 2547–2580. DOI: [10.1137/130918885](https://doi.org/10.1137/130918885). eprint: <https://doi.org/10.1137/130918885>.
- [Cha+20] Nicolas Charon et al. “12 - Fidelity metrics between curves and surfaces: currents, varifolds, and normal cycles”. In: (2020). Ed. by Xavier Pennec, Stefan Sommer, and Tom Fletcher, pp. 441–477. DOI: <https://doi.org/10.1016/B978-0-12-814725-2.00021-2>. URL: <http://www.sciencedirect.com/science/article/pii/B9780128147252000212>.
- [Che+02] D. Chetverikov et al. “The Trimmed Iterative Closest Point algorithm”. In: 3 (2002), 545–548 vol.3. DOI: [10.1109/ICPR.2002.1047997](https://doi.org/10.1109/ICPR.2002.1047997).
- [Chi+18a] Lenaïc Chizat et al. “An interpolating distance between optimal transport and Fisher–Rao metrics”. In: *Foundations of Computational Mathematics* 18.1 (2018), pp. 1–44.
- [Chi+18b] Lenaïc Chizat et al. “Scaling algorithms for unbalanced optimal transport problems”. In: *Mathematics of Computation* 87.314 (2018), pp. 2563–2609.
- [CRM96] Gary E Christensen, Richard D Rabbitt, and Michael I Miller. “Deformable templates using large deformation kinematics”. In: *IEEE transactions on image processing* 5.10 (1996), pp. 1435–1447.

-
- [Cor+19] FH Cornelis et al. “Multimodal image-guided electrochemotherapy of unresectable liver metastasis from renal cell cancer”. In: *Diagnostic and Interventional Imaging* 100.5 (2019), pp. 309–311.
- [Cur+18] Claire Cury et al. “Statistical Shape Analysis of Large Datasets Based on Diffeomorphic Iterative Centroids”. In: *Frontiers in Neuroscience* 12 (Nov. 2018), p. 803. ISSN: 1662-453X. DOI: [10.3389/fnins.2018.00803](https://doi.org/10.3389/fnins.2018.00803). URL: <https://www.frontiersin.org/article/10.3389/fnins.2018.00803/full> (visited on 04/06/2022).
- [Cut13] Marco Cuturi. “Sinkhorn distances: Lightspeed computation of optimal transport”. In: *Advances in neural information processing systems* 26 (2013).
- [CD14] Marco Cuturi and Arnaud Doucet. “Fast computation of Wasserstein barycenters”. In: (2014), pp. 685–693.
- [Dal+18] Adrian V Dalca et al. “Unsupervised learning for fast probabilistic diffeomorphic registration”. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer. 2018, pp. 729–738.
- [DD78] Carl De Boor and Carl De Boor. *A practical guide to splines*. Vol. 27. springer-verlag New York, 1978.
- [Duc77] Jean Duchon. “Splines minimizing rotation-invariant semi-norms in Sobolev spaces”. In: *Constructive theory of functions of several variables*. Springer, 1977, pp. 85–100.
- [DKS18] Adam Duncan, Eric Klassen, and Anuj Srivastava. “Statistical shape analysis of simplified neuronal trees”. In: *The Annals of Applied Statistics* 12.3 (2018), pp. 1385–1421.
- [DGM98] Paul Dupuis, Ulf Grenander, and Michael I Miller. “Variational problems on flows of diffeomorphisms for image matching”. In: *Quarterly of applied mathematics* (1998), pp. 587–600.
- [Dur+18] Jeremy C Durack et al. “Assessment of automated cone-beam CT vessel identification software during transarterial hepatic embolisation: radiation dose, contrast medium volume, processing time, and operator perspectives compared to digital subtraction angiography”. In: *Clinical radiology* 73.12 (2018), 1057–e1.
- [DDV17] Daniele Durante, David B Dunson, and Joshua T Vogelstein. “Nonparametric Bayes modeling of populations of networks”. In: *Journal of the American Statistical Association* 112.520 (2017), pp. 1516–1530.
- [Dur+11] Sandy Durrleman et al. “Optimal data-driven sparse parameterization of diffeomorphisms for population analysis”. In: *Biennial International Conference on Information Processing in Medical Imaging*. Springer. 2011, pp. 123–134.
- [Dur10] Stanley Durrleman. “Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution”. PhD thesis. Université Nice Sophia Antipolis, 2010.

-
- [Ezq+98] N. Ezquerro et al. “Model-guided labeling of coronary structure”. In: *IEEE Transactions on Medical Imaging* 17.3 (June 1998), pp. 429–441.
- [Fed59] Herbert Federer. “Curvature measures”. In: *Transactions of the American Mathematical Society* 93 (1959), pp. 418–491.
- [Fed69] Herbert Federer. *Geometric measure theory*. Springer, 1969.
- [FN20] Aasa Feragen and Tom Nye. “Statistics on stratified spaces”. In: (2020), pp. 299–342.
- [Fer+12a] Aasa Feragen et al. “A hierarchical scheme for geodesic anatomical labeling of airway trees”. In: (2012). Ed. by Nicholas Ayache et al., pp. 147–155.
- [Fer+15] Aasa Feragen et al. “Geodesic Atlas-Based Labeling of Anatomical Trees: Application and Evaluation on Airways Extracted From CT”. In: *IEEE Transactions on Medical Imaging* 34 (2015), pp. 1212–1226.
- [Fer+14] Aasa Feragen et al. “Geodesic atlas-based labeling of anatomical trees: Application and evaluation on airways extracted from CT”. In: *IEEE transactions on medical imaging* 34.6 (2014), pp. 1212–1226.
- [Fer+11] Aasa Feragen et al. “Means in spaces of tree-like shapes”. In: (2011), pp. 736–746. DOI: [10.1109/ICCV.2011.6126311](https://doi.org/10.1109/ICCV.2011.6126311).
- [Fer+12b] Aasa Feragen et al. “Toward a theory of statistical tree-shape analysis”. In: *IEEE transactions on pattern analysis and machine intelligence* 35.8 (2012), pp. 2008–2021.
- [Fer+13] Aasa Feragen et al. “Tree-space statistics and approximations for large-scale analysis of anatomical trees”. In: (2013), pp. 74–85.
- [Fey+19] Jean Feydy et al. “Fast and Scalable Optimal Transport for Brain Tractograms”. In: *MICCAI 2019*. Shenzhen, China, Oct. 2019. URL: <https://hal.telecom-paristech.fr/hal-02264177>.
- [Fey+17] Jean Feydy et al. “Optimal Transport for Diffeomorphic Registration”. In: *Lecture Notes in Computer Science* (2017), pp. 291–299. ISSN: 1611-3349. DOI: [10.1007/978-3-319-66182-7_34](https://doi.org/10.1007/978-3-319-66182-7_34). URL: http://dx.doi.org/10.1007/978-3-319-66182-7_34.
- [Fle+03] P Thomas Fletcher et al. “Gaussian distributions on Lie groups and their application to statistical shape analysis”. In: (2003), pp. 450–462.
- [Fle+04] P Thomas Fletcher et al. “Principal geodesic analysis for the study of nonlinear statistics of shape”. In: *IEEE transactions on medical imaging* 23.8 (2004), pp. 995–1005.
- [Fra+22] Anton François et al. “Weighted Metamorphosis for registration of images with different topology.” In: *10th International Workshop on Biomedical Image Registration*. 2022.

-
- [Fra+98] Alejandro F Frangi et al. “Multiscale vessel enhancement filtering”. In: (1998), pp. 130–137.
- [Gar+21] MK Garba et al. “Information geometry for phylogenetic trees”. In: *Journal of mathematical biology* 82.3 (Feb. 2021), p. 19. ISSN: 0303-6812. DOI: [10.1007/s00285-021-01553-x](https://doi.org/10.1007/s00285-021-01553-x). URL: <https://europepmc.org/articles/PMC7884381>.
- [GLS14] Sahar Ghanavati, Jason P. Lerch, and John G. Sled. “Automatic anatomical labeling of the complete cerebral vasculature in mouse models”. In: *NeuroImage* 95 (2014), pp. 117–128.
- [Gho+21] Mario Ghosn et al. “Prediction of overall survival in patients with hepatocellular carcinoma treated with Y-90 radioembolization by imaging response criteria”. In: *Diagnostic and Interventional Imaging* 102 (Aug. 2021), pp. 35–44.
- [GBR08] Bram van Ginneken, Wouter Baggerman, and Eva van Rikxoort. “Robust Segmentation and Anatomical Labeling of the Airway Tree from Thoracic CT Scans”. In: *Medical image computing and computer-assisted intervention : MIC-CAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention* 11 (Feb. 2008), pp. 219–26.
- [Gir+05] Daniel Girardeau-Montaut et al. “Change detection on points cloud data acquired with a ground laser scanner”. In: *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* 36.3 (2005), W19.
- [Gla05] Joan Glaunès. “Transport par difféomorphismes de points, mesures et de courants pour la comparaison de formes et l’anatomie numérique.” 2005.
- [Gor+09] Vladlena Gorbunova et al. “Curve-and surface-based registration of lung ct images via currents”. In: *Second International Workshop on Pulmonary Image Analysis*. 2009, pp. 15–25.
- [GM98] Ulf Grenander and Michael I. Miller. “Computational Anatomy: An Emerging Discipline”. In: *Q. Appl. Math.* LVI.4 (Dec. 1998), pp. 617–694. ISSN: 0033-569X. URL: <http://dl.acm.org/citation.cfm?id=309082.309089>.
- [GO19] Gillian Grindstaff and Megan Owen. “Representations of Partial Leaf Sets in Phylogenetic Tree Space”. In: *SIAM Journal on Applied Algebra and Geometry* 3.4 (2019), pp. 691–720. DOI: [10.1137/18M1235855](https://doi.org/10.1137/18M1235855). eprint: <https://doi.org/10.1137/18M1235855>. URL: <https://doi.org/10.1137/18M1235855>.
- [Gui+21] Nicolas Guigui et al. “Cardiac motion modeling with parallel transport and shape splines”. In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. IEEE. 2021, pp. 1394–1397.

-
- [Gül+14] Mehmet A. Gülsün et al. “CTA Coronary Labeling through Efficient Geodesics between Trees Using Anatomy Priors”. In: *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention* 17 Pt 2 (2014), pp. 521–8.
- [GSS21] Xiaoyang Guo, Anuj Srivastava, and Sudeep Sarkar. “A Quotient Space Formulation for Generative Statistical Analysis of Graphical Data”. In: *Journal of Mathematical Imaging and Vision* 63.6 (2021), pp. 735–752.
- [Gur+21] Martina Gurgitano et al. “Interventional radiology ex-machina: impact of artificial intelligence on practice”. In: *La radiologia medica* 126.7 (2021), pp. 998–1006.
- [Hal+19] O. Halimi et al. “Unsupervised Learning of Dense Shape Correspondence”. In: (2019), pp. 4365–4374. DOI: [10.1109/CVPR.2019.00450](https://doi.org/10.1109/CVPR.2019.00450).
- [Hal+20] Oshri Halimi et al. “The Whole Is Greater Than the Sum of Its Nonrigid Parts”. In: *CoRR* abs/2001.09650 (2020). arXiv: [2001.09650](https://arxiv.org/abs/2001.09650). URL: <https://arxiv.org/abs/2001.09650>.
- [Ham+21] Nils Hampe et al. “Graph attention networks for segment labeling in coronary artery trees”. In: *Medical Imaging 2021: Image Processing*. Vol. 11596. International Society for Optics and Photonics. 2021, p. 115961I.
- [Hei90] Jotun Hein. “Reconstructing evolution of sequences subject to recombination using parsimony”. In: *Mathematical biosciences* 98.2 (1990), pp. 185–200.
- [Hoa+11] Bui Huy Hoang et al. “A study on automated anatomical labeling to arteries concerning with colon from 3D abdominal CT images”. In: 7962 (2011).
- [HR31] Heinz Hopf and Willi Rinow. “Ueber den Begriff der vollständigen differentialgeometrischen Fläche”. In: *Commentarii Mathematici Helvetici* 3.1 (1931), pp. 209–225.
- [HC21] Hsi-Wei Hsieh and Nicolas Charon. “Weight metamorphosis of varifolds and the LDDMM-Fisher-Rao metric”. In: *arXiv preprint arXiv:2112.04644* (2021).
- [Hua+08] Qi-Xing Huang et al. “Non-rigid registration under isometric deformations”. In: *Computer Graphics Forum*. Vol. 27. 5. Wiley Online Library. 2008, pp. 1449–1457.
- [JSZ+15] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. “Spatial transformer networks”. In: *Advances in neural information processing systems* 28 (2015).
- [JM00] Sarang C Joshi and Michael I Miller. “Landmark matching via large deformation diffeomorphisms”. In: *IEEE transactions on image processing* 9.8 (2000), pp. 1357–1370.
- [KZH13] Oliver van Kaick, Hao Zhang, and Ghassan Hamarneh. “Bilateral Maps for Partial Matching”. In: *Computer Graphics Forum (CGF)* (Sept. 2013). DOI: [10.1111/cgf.12084](https://doi.org/10.1111/cgf.12084).

-
- [Kai+11] Oliver van Kaick et al. “A Survey on Shape Correspondence”. In: *Computer Graphics Forum* 30.6 (2011), pp. 1681–1707. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2011.01884.x>.
- [KCC17] Irene Kaltenmark, Benjamin Charlier, and Nicolas Charon. “A General Framework for Curve and Surface Comparison and Registration With Oriented Vari-folds”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (July 2017).
- [KT18] Irène Kaltenmark and Alain Trouvé. “Estimation of a Growth Development with Partial Diffeomorphic Mappings”. In: *Quarterly of Applied Mathematics* 77 (Nov. 2018), pp. 227–267.
- [Kan60] Leonid V Kantorovich. “Mathematical methods of organizing and planning production”. In: *Management science* 6.4 (1960), pp. 366–422.
- [KBD17] András P Keszei, Benjamin Berkels, and Thomas M Deserno. “Survey of non-rigid registration tools in medicine”. In: *Journal of digital imaging* 30.1 (2017), pp. 102–116.
- [Kim+18] Alexander Y Kim et al. “Utility of MR angiography in the identification of prostatic artery origin prior to prostatic artery embolization”. In: *Journal of Vascular and Interventional Radiology* 29.3 (2018), pp. 307–310.
- [Kle+09] Arno Klein et al. “Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration”. In: *Neuroimage* 46.3 (2009), pp. 786–802.
- [Lac+21] Leander Lacroix et al. “IMODAL: creating learnable user-defined deformation models”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 12905–12913.
- [LY17] Choong Ho Lee and Hyung-Jin Yoon. “Medical big data: promise and challenges”. In: *Kidney research and clinical practice* 36.1 (2017), p. 3.
- [Lee+17] Sieun Lee et al. “Age and glaucoma-related characteristics in retinal nerve fiber layer and choroid: Localized morphometrics and visualization using functional shapes registration”. In: *Frontiers in neuroscience* 11 (2017), p. 381.
- [LMS18] Matthias Liero, Alexander Mielke, and Giuseppe Savaré. “Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures”. In: *Inventiones mathematicae* 211.3 (2018), pp. 969–1117.
- [Lit+17] Or Litany et al. “Fully spectral partial shape matching”. In: *Computer Graphics Forum*. Vol. 36. 2. Wiley Online Library. 2017, pp. 247–258.
- [LN89] Dong C. Liu and Jorge Nocedal. “On the limited memory BFGS method for large scale optimization”. In: *MATHEMATICAL PROGRAMMING* 45 (1989), pp. 503–528.
- [Lo+11] Pechin Lo et al. “A Bottom-up approach for labeling of human airway trees”. In: (2011).

-
- [LP13] Marco Lorenzi and Xavier Pennec. “Geodesics, parallel transport & one-parameter subgroups for diffeomorphic image registration”. In: *International journal of computer vision* 105.2 (2013), pp. 111–127.
- [Low99] David G Lowe. “Object recognition from local scale-invariant features”. In: 2 (1999), pp. 1150–1157.
- [Mar+12] Primož Markelj et al. “A review of 3D/2D registration methods for image-guided interventions”. In: *Medical image analysis* 16.3 (2012), pp. 642–661.
- [Mas+05] Salvatore Masala et al. “MRI and bone scan imaging in the preoperative evaluation of painful vertebral fractures treated with vertebroplasty and kyphoplasty”. In: *in vivo* 19.6 (2005), pp. 1055–1060.
- [Mat+14] Tetsuro Matsuzaki et al. “Automated anatomical labeling of abdominal arteries and hepatic portal system extracted from abdominal CT volumes”. In: *Medical Image Analysis* 20 (Nov. 2014).
- [Még+21] Lucile Mégret et al. “Shape deformation analysis reveals the temporal dynamics of cell-type-specific homeostatic and pathogenic responses to mutant huntingtin”. In: *Elife* 10 (2021), e64984.
- [Mei84] Jean Meinguet. “Surface spline interpolation: Basic theory and computational aspects”. In: *Approximation Theory and Spline Functions*. Springer, 1984, pp. 127–142.
- [MG13] Mario Micheli and Joan Alexis Glaunes. “Matrix-valued kernels for shape deformation analysis”. In: *arXiv preprint arXiv:1308.5739* (2013).
- [MOP15] Ezra Miller, Megan Owen, and J. Scott Provan. “Polyhedral Computational Geometry for Averaging Metric Phylogenetic Trees”. In: *Adv. Appl. Math.* 68.C (July 2015), pp. 51–91.
- [MTY06] Michael I. Miller, Alain Trouvé, and Laurent Younes. “Geodesic Shooting for Computational Anatomy”. In: *Journal of Mathematical Imaging and Vision* 24.2 (Mar. 2006), pp. 209–228. ISSN: 1573-7683. DOI: [10.1007/s10851-005-3624-0](https://doi.org/10.1007/s10851-005-3624-0). URL: <https://doi.org/10.1007/s10851-005-3624-0>.
- [MNA16] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. “V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation”. In: *2016 fourth international conference on 3D vision (3DV)* (Oct. 2016), pp. 565–571. DOI: [10.1109/3DV.2016.79](https://doi.org/10.1109/3DV.2016.79).
- [Mon+21] Lorenzo Monfardini et al. “Real-time US/cone-beam CT fusion imaging for percutaneous ablation of small renal tumours: a technical note”. In: *European Radiology* 31.10 (2021), pp. 7523–7528.
- [MMC16] Grégoire Montavon, Klaus-Robert Müller, and Marco Cuturi. “Wasserstein training of restricted Boltzmann machines”. In: *Advances in Neural Information Processing Systems* 29 (2016).

-
- [Mus+21] Marco Musy et al. *marcomusy/vedo: 2021.0.7*. Version v2021.0.7. Nov. 2021. DOI: [10.5281/zenodo.5655358](https://doi.org/10.5281/zenodo.5655358). URL: <https://doi.org/10.5281/zenodo.5655358>.
- [NKV19] Marc Niethammer, Roland Kwitt, and Francois-Xavier Vialard. “Metric learning for image registration”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 8463–8472.
- [Nye+17] Tom M W Nye et al. “Principal component analysis and the locus of the Fréchet mean in the space of phylogenetic trees”. In: *Biometrika* 104.4 (Sept. 2017), pp. 901–922.
- [Nye14] Tom M.W. Nye. “An Algorithm for Constructing Principal Geodesics in Phylogenetic Treespace”. In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 11.2 (2014), pp. 304–315. DOI: [10.1109/TCBB.2014.2309599](https://doi.org/10.1109/TCBB.2014.2309599).
- [Nye11] Tom MW Nye. “Principal components analysis in the space of phylogenetic trees”. In: *The Annals of Statistics* (2011), pp. 2716–2739.
- [Ovs+12] Maks Ovsjanikov et al. “Functional Maps: A Flexible Representation of Maps between Shapes”. In: *ACM Trans. Graph.* 31.4 (July 2012). ISSN: 0730-0301. DOI: [10.1145/2185520.2185526](https://doi.org/10.1145/2185520.2185526).
- [OP11] Megan Owen and J. Scott Provan. “A Fast Algorithm for Computing Geodesic Distances in Tree Space”. In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 8.1 (2011), pp. 2–13. DOI: [10.1109/TCBB.2010.3](https://doi.org/10.1109/TCBB.2010.3).
- [Pan+16] Yue Pan et al. “Current-and varifold-based registration of lung vessel and airway trees”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 2016, pp. 126–133.
- [Pas+17] Adam Paszke et al. “Automatic Differentiation in PyTorch”. In: *NIPS Autodiff Workshop*. 2017.
- [Pen06] Xavier Pennec. “Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements”. In: *Journal of Mathematical Imaging and Vision* 25.1 (2006), pp. 127–154.
- [PSF20] Xavier Pennec, Stefan Horst Sommer, and Tom Fletcher, eds. *Riemannian Geometric Statistics in Medical Image Analysis*. English. 1. United States: Academic Press, 2020. ISBN: 9780128147252. DOI: [10.1016/C2017-0-01561-6](https://doi.org/10.1016/C2017-0-01561-6).
- [RV16] M. Rajagopal and AM. Venkatesan. “Image fusion and navigation platforms for percutaneous image-guided interventions”. In: *Abdom Radiol (NY)* 41 (Apr. 2016), pp. 620–628. DOI: [10.1007/s00261-016-0645-7](https://doi.org/10.1007/s00261-016-0645-7).

-
- [Ray+18] Alistair F. Ray et al. “Efficacy and safety of prostate artery embolization for benign prostatic hyperplasia: an observational study and propensity-matched comparison with transurethral resection of the prostate (the UK-ROPE study)”. In: *BJU International* 122.2 (2018), pp. 270–282.
- [Res18] Aymeric Reshef. “Dual-rotation C-arm cone-beam tomographic acquisition and reconstruction frameworks for low-contrast detection in brain soft-tissue imaging”. PhD thesis. Paris, ENST, 2018.
- [Ris+11] L. Risser et al. “Simultaneous Multi-scale Registration Using Large Deformation Diffeomorphic Metric Mapping”. In: *IEEE Transactions on Medical Imaging* 30.10 (Oct. 2011), pp. 1746–1759. ISSN: 1558-254X. DOI: [10.1109/TMI.2011.2146787](https://doi.org/10.1109/TMI.2011.2146787).
- [Ris+13] Laurent Risser et al. “Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions”. In: *Medical image analysis* 17.2 (2013), pp. 182–193.
- [Ris+10] Laurent Risser et al. “Simultaneous fine and coarse diffeomorphic registration: application to atrophy measurement in Alzheimer’s disease”. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer. 2010, pp. 610–617.
- [Rob+16] David Robben et al. “Simultaneous segmentation and anatomical labeling of the cerebral vasculature”. In: *Medical Image Analysis* 32 (2016), pp. 201–215.
- [Rob71] David F Robinson. “Comparison of labeled trees with valency three”. In: *Journal of Combinatorial Theory, Series B* 11.2 (1971), pp. 105–119.
- [RF81] David F Robinson and Leslie R Foulds. “Comparison of phylogenetic trees”. In: *Mathematical biosciences* 53.1-2 (1981), pp. 131–147.
- [Roc+20] A Rocha et al. “Advantages of using cone-beam computed tomography over digital subtraction angiography to identify prostatic arteries in prostatic artery embolization”. In: *J Vasc Interv Radiol* 31.3 Suppl (2020), S31.
- [Rod+17] E. Rodolà et al. “Partial Functional Correspondence”. In: *Computer Graphics Forum* 36.1 (2017), pp. 222–236. DOI: <https://doi.org/10.1111/cgf.12797>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12797>.
- [Rod+13] Emanuele Rodolà et al. “A Scale Independent Selection Process for 3D Object Recognition in Cluttered Scenes”. In: *International Journal of Computer Vision* 102 (Mar. 2013). DOI: [10.1007/s11263-012-0568-x](https://doi.org/10.1007/s11263-012-0568-x).
- [RG16] Pierre Roussillon and Joan Alexis Glaunès. “Kernel Metrics on Normal Cycles and Application to Curve Matching”. In: *SIAM Journal on Imaging Sciences* 9.4 (2016), pp. 1991–2038. DOI: [10.1137/16M1070529](https://doi.org/10.1137/16M1070529). eprint: <http://dx.doi.org/10.1137/16M1070529>. URL: <http://dx.doi.org/10.1137/16M1070529>.

-
- [Sah20] Yusuf Sahillioğlu. “Recent advances in shape correspondence”. In: *The Visual Computer* 36.8 (2020), pp. 1705–1721.
- [Sch38] Isaac J Schoenberg. “Metric spaces and completely monotone functions”. In: *Annals of Mathematics* (1938), pp. 811–841.
- [SSM97] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. “Kernel principal component analysis”. In: (1997). Ed. by Wulfram Gerstner et al., pp. 583–588.
- [Sch64] Laurent Schwartz. “Sous-espaces hilbertiens d’espaces vectoriels topologiques et noyaux associés (noyaux reproduisants)”. In: *Journal d’analyse mathématique* 13.1 (1964), pp. 115–256.
- [Stu03] Karl-Theodor Sturm. “Probability Measures on Metric Spaces of Nonpositive Curvature, Heat Kernels and Analysis on Manifolds”. In: *Graphs, and Metric Spaces* (Jan. 2003), pp. 357–390.
- [SBC21] Yashil Sukurdeep, Martin Bauer, and Nicolas Charon. *A new variational model for shape graph registration with partial matching constraints*. 2021. arXiv: [2105.00678](https://arxiv.org/abs/2105.00678) [math.OC].
- [Tac+15] V. Tacher et al. “How I do it: Cone-beam CT during transarterial chemoembolization for liver cancer”. In: *Radiology* 274 (Feb. 2015), pp. 320–334. DOI: [10.1148/radiol.14131925](https://doi.org/10.1148/radiol.14131925).
- [TQ16] Mingzhen Tan and Anqi Qiu. “Large deformation multiresolution diffeomorphic metric mapping for multiresolution cortical surfaces: a coarse-to-fine approach”. In: *IEEE Transactions on Image Processing* 25.9 (2016), pp. 4061–4074.
- [Tho17] D’Arcy Wentworth Thompson. *On Growth and Form*. Cambridge University Press, 1917. DOI: [10.1017/CBO9781107325852](https://doi.org/10.1017/CBO9781107325852).
- [Tro95] Alain Trouvé. “An Infinite Dimensional Group Approach for Physics based Models in Pattern Recognition”. In: *International Journal of Computer Vision - IJCV* (Jan. 1995).
- [Tro98] Alain Trouvé. “Diffeomorphisms groups and pattern matching in image analysis”. In: *International journal of computer vision* 28.3 (1998), pp. 213–221.
- [Tsc+05] J. Tschirren et al. “Matching and anatomical labeling of human airway tree”. In: *IEEE Transactions on Medical Imaging* 24.12 (Dec. 2005), pp. 1540–1547.
- [Vai+04] M. Vaillant et al. “Statistics on diffeomorphisms via tangent space representations”. In: *NeuroImage* 23 (2004). Mathematics in Brain Imaging, S161–S169. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2004.07.023>. URL: <http://www.sciencedirect.com/science/article/pii/S1053811904003957>.
- [Vel01] Remco C Veltkamp. “Shape matching: Similarity measures and algorithms”. In: (2001), pp. 188–197.

-
- [Ver+09] Tom Vercauteren et al. “Diffeomorphic demons: Efficient non-parametric image registration”. In: *NeuroImage* 45.1 (2009), S61–S72.
- [Wan+18] Guan Wang et al. “The shape space of 3D botanical tree models”. In: *ACM Transactions on Graphics (TOG)* 37.1 (2018), pp. 1–18.
- [Wan+17] Xingce Wang et al. “Automatic Labeling of Vascular Structures with Topological Constraints via HMM”. In: (2017).
- [Wil19] Amy Willis. “Confidence Sets for Phylogenetic Trees”. In: *Journal of the American Statistical Association* 114.525 (2019), pp. 235–244. DOI: [10.1080/01621459.2017.1395342](https://doi.org/10.1080/01621459.2017.1395342).
- [Wu+19] Dan Wu et al. “Automated anatomical labeling of coronary arteries via bidirectional tree LSTMs”. In: *International journal of computer assisted radiology and surgery* 14.2 (2019), pp. 271–280.
- [Yan+20] Han Yang et al. “Cpr-gcn: Conditional partial-residual graph convolutional network in automated anatomical labeling of coronary arteries”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 3803–3811.
- [Yan+13] Xianfeng Yang et al. “Evolution of hippocampal shapes across the human lifespan”. In: *Human brain mapping* 34.11 (2013), pp. 3075–3085.
- [You10] L. Younes. *Shapes and Diffeomorphisms*. Applied Mathematical Sciences. Springer Berlin Heidelberg, 2010. ISBN: 9783642120558. URL: <https://books.google.fr/books?id=SdTbtMGgeAUC>.
- [You+08] Laurent Younes et al. “Transport of relational structures in groups of diffeomorphisms”. In: *Journal of mathematical imaging and vision* 32.1 (2008), pp. 41–56.
- [Zha16] Haojin Zhai. “Principal component analysis in phylogenetic tree space”. PhD thesis. The University of North Carolina at Chapel Hill, 2016.
- [Zhe+12] Xin Zhen et al. In: 57.21 (Oct. 2012), pp. 6807–6826. DOI: [10.1088/0031-9155/57/21/6807](https://doi.org/10.1088/0031-9155/57/21/6807).
- [Zho20] Chen Zhou. “A Hybrid Approach for Coronary Artery Anatomical Labeling in Cardiac CT Angiography”. In: *Journal of Physics: Conference Series*. Vol. 1642. 1. IOP Publishing. 2020, p. 012020.
- [Zol+17] Reza Zolfaghari et al. “Kernel principal component analysis of the ear morphology”. In: (Mar. 2017), pp. 481–485. DOI: [10.1109/ICASSP.2017.7952202](https://doi.org/10.1109/ICASSP.2017.7952202). URL: <http://ieeexplore.ieee.org/document/7952202/> (visited on 04/06/2022).

Recalages diffeomorphiques avec changements de topologie: application à la radiologie interventionnelle

Résumé Cette thèse de mathématiques appliquées s'inscrit dans le domaine de la radiologie interventionnelle, en particulier les interventions endovasculaires minimalement invasives. Afin de guider leurs outils à l'intérieur des patients, les praticiens peuvent s'appuyer sur des logiciels avancés, par exemple pour analyser un arbre vasculaire 3D et identifier les vaisseaux à traiter. Dans le cas du traitement de l'hyperplasie bénigne de la prostate par embolisation de l'artère prostatique, il est important d'identifier le type de chaque artère afin de limiter les risques d'embolisation incomplète ou non-désirée.

Comme il est difficile d'annoter les données médicales, les solutions techniques ne doivent s'appuyer que sur de petites bases de données pour être utilisables. Les méthodes dites "basées atlas" répondent exactement à ce critère. Cependant, peu d'entre elles exploitent l'information disponible non étiquetée et, à cause de la forme complexe des arbres vasculaires, les déformations non rigides pour aligner les arbres sont rarement envisagées. Pourtant de tels recalages favorisent le transfert automatique de l'étiquetage de l'arbre déformé vers une base de données non annotée. Nous nous appuyons sur la théorie de l'anatomie computationnelle et des LDDMM pour l'analyse de l'arbre vasculaire pelvien. Nous montrons sur un premier exemple d'arbres simplifiés qu'en utilisant un seul cas annoté, dit "template", et recalé sur l'ensemble de la base de données, on peut construire un atlas réaliste capturant la variabilité géométrique des observations. L'atlas une fois aligné sur un autre arbre est utilisé pour l'annoter et atteindre une précision de $98.9\%(\pm 0.33)$ sur une base de 49 arbres.

Cependant, en passant à des données vasculaires complètes deux problèmes se posent : 1. les deux arbres à recaler n'ont pas le même nombre de branches et ne peuvent donc pas exactement être mis en correspondance ; 2. deux arbres à recaler présentent dans la majorité des cas des changements topologiques qui ne peuvent pas être gérés par des LDDMM. Ces deux points nous conduisent d'abord à formuler le problème de l'inclusion d'une forme dans une autre comme un terme d'attache aux données. Nous proposons également un terme de régularisation comparant l'objet déformé et sa position initiale, et permettant de contrôler les déformations induites par les diffeomorphismes. Nous appliquons cette méthode au recalage du template de l'arbre pelvien sur des arbres réels. Nous l'appliquons aussi à celui de surfaces de foies tronquées sur des surfaces complètes pour un recalage de volumes issus de deux modalités d'imagerie différentes.

Pour gérer les changements topologiques, notre template est plongé dans un espace adapté. Nous pouvons alors faire varier sa topologie en créant des changements dans l'ordre des bifurcations au cours du recalage sur un arbre cible. Jusqu'à présent ces recalages dans cet espace ne pouvaient s'effectuer que si tous arbres étaient annotés. Grâce à une procédure d'optimisation sur la position du template nous pouvons effectuer son recalage sur des arbres non annotés. Les LDDMM peuvent y être associés pour combiner des recalages diffeomorphiques et topologiques qui sont appliqués à des exemples jouets. La combinaison de ces méthodes offre de nombreux outils pour les méthodes basées atlas même dans le cas de forts changements topologiques.

Mots-clés : Radiologie Interventionnelle, Anatomie Computationnelle, Arbres Vasculaires, LDDMM, Changements de Topologie.

Diffeomorphic Registration with Topological Changes: Application to Interventional Radiology

Abstract This thesis falls within the field of mathematics applied to Interventional Radiology, and more precisely minimally invasive endovascular interventions. To guide their tools inside the patients, practitioners can rely on advanced software to analyze, for example, a 3D vascular tree and identify the vessels to be treated. When treating benign prostatic hyperplasia with prostatic artery embolization, it is particularly important to identify the type of each artery in the tree in order to limit the risk of incomplete or unwanted embolization.

Because medical data is difficult to annotate, solutions to automatize this annotation process may only rely on small databases to be developed. The so-called "atlas-based" methods meet exactly this criterion, by transferring labels from an annotated tree to a new one through registration. However, few techniques exploit the available unlabeled information and, because of the complex shape of vascular trees, non-rigid deformations to align the trees are rarely considered. We rely on the theory of computational anatomy and Large Deformation Diffeomorphic Mapping (LDDMM) for the analysis of the pelvic vascular tree. In a first example on simplified trees, we register a single annotated case, called a "template", on the whole database to then build a realistic atlas capturing the geometric variability of the observations by using deformation statistics. The atlas, once aligned on another tree, is used to annotate it with an accuracy of $98.9\%(\pm 0.33)$ on a database of 49 trees.

When considering complete vascular data instead of simplified ones, two problems arise: **1.** the two trees to be aligned do not have the same number of branches and therefore cannot be exactly matched; **2.** in most cases, trees to be aligned present topological changes in terms of bifurcation ordering that cannot be modeled by LDDMM. The first problem pushes us to adapt the registration to obtain the inclusion of a deformed shape in another. We also propose a regularization term comparing the deformed object with its initial position and allowing to control the deformations induced by the diffeomorphisms. We apply this method to the registration of a simplified pelvic artery tree on complete and complex real trees. Thanks to the flexibility of the shape representation and the LDDMM framework, we also apply it to the registration of truncated liver surfaces on complete surfaces for a registration of volumes from two different imaging modalities.

To manage topological changes in terms of bifurcation ordering, our template is immersed into a suitable space called the "tree space". In this space, one can then vary its topology by creating changes in the order of bifurcations during the registration on a target tree. Until now, these changes in such space could only be performed if all the trees were annotated. Thanks to an optimization procedure on the position of the template, we can perform its registration on non-annotated trees. The LDDMM can be associated to combine diffeomorphic and topological registration which are applied to toy examples. The combination of these methods offers many tools for atlas-based methods even in the case of strong topological changes: for unlabeled database analysis, automatic annotation, data augmentation.

Key words: Interventional Radiology, Computational Anatomy, Vascular Trees, LDDMM, Topological Changes
