



HAL
open science

Quantitative analysis of healthy and abnormal fetal brain development using deformation models

Fleur Gaudfernau

► **To cite this version:**

Fleur Gaudfernau. Quantitative analysis of healthy and abnormal fetal brain development using deformation models. Mathematics [math]. Université Paris Cité, 2023. English. NNT : 2023UNIP7272 . tel-04453836v2

HAL Id: tel-04453836

<https://theses.hal.science/tel-04453836v2>

Submitted on 11 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Université
Paris Cité

Université Paris Cité

École Doctorale de Sciences Mathématiques de Paris-Centre (386)
Centre de Recherche des Cordeliers, Equipe HeKa (Health data- and
model- driven Knowledge Acquisition)

Quantitative analysis of healthy and abnormal fetal brain development using deformation models

Par FLEUR GAUDFERNAU

Thèse de Doctorat de MATHÉMATIQUES APPLIQUÉES

Dirigée par STÉPHANIE ALLASSONNIÈRE

Et par ERWAN LE PENNEC

Présentée et soutenue publiquement le 19 décembre 2023

Devant un jury composé de:

FRANÇOIS ROUSSEAU	Professeur	IMT Atlantique	Rapporteur
COLIN STUDHOLME	Professeur	University of Washington	Rapporteur
ALAIN TROUVÉ	Professeur	Université Paris Saclay	Examinateur
JULIE DELON	Professeur	Université Paris Cité	Examinatrice
GUILLAUME AUZIAS	Chercheur	Aix-Marseille Université	Examinateur
STÉPHANIE ALLASSONNIÈRE	Professeure	Université Paris Cité	Directrice de thèse
ERWAN LE PENNEC	Professeur	Ecole Polytechnique	Directeur de thèse

Titre. Analyse quantitative du développement cérébral foetal sain et anormal à l'aide de modèles de déformation.

Résumé. L'analyse quantitative des images du cerveau foetal est une question difficile liée à des problématiques de manque de données, qualité d'image hétérogène et variabilité inter-sujets élevée. L'Anatomie Computationnelle offre des méthodes non-supervisées pour caractériser des structures anatomiques normales ou pathologiques en calculant la transformation qui déforme un objet de référence en un autre objet. Le but de cette thèse est d'adapter des outils géométriques pour l'analyse de la croissance du cerveau foetal dans le cadre du Large Diffeomorphic Deformation Metric Mapping. Tout d'abord, nous présentons des contributions méthodologiques pour corriger le problème des minima locaux irréalistes qui se pose quand on déforme des objets complexes en grande dimension. Pour cela, nous introduisons une stratégie d'optimisation "coarse-to-fine" qui décompose les champs de déformation dans une base d'ondelettes de Haar et filtre les objets à déformer. Cette stratégie multiéchelle est appliquée à des modèles de recalage et d'estimation d'atlas cross-sectionnel ainsi qu'à l'estimation de trajectoires temporelles. Grâce à des images de cerveau foetal, nous démontrons que cette stratégie multiéchelle produit des images de référence et des déformatins plus naturelles. Ensuite, nous construisons une pipeline pour caractériser à une échelle globale la variabilité anatomique des cerveaux foetaux avec une agénésie du corps calleux. Les déviations par rapport au développement normal sont quantifiées en recalant chaque sujet à un cerveau de référence. Les différences anatomiques liées à l'âge sont corrigées en transportant ces déformations dans un espace commun, ce qui permet d'identifier des modes de déformation caractéristiques des cerveaux anormaux. Finalement, nous proposons de construire un atlas du développement cortical sain au cours de la grossesse. La croissance moyenne du cortex est modélisée par une trajectoire géodésique par morceaux, et la variabilité anatomique autour de cette moyenne est estimée grâce à des outils de recalage et de transports parallèle. Cette méthode de construction d'atlas est comparée à la méthode classique, dite "kernel regression", en évaluant le degré de réalisme du processus de gyrification modélisé. Les outils développés peuvent être généralisés à différents types de données, structures anatomiques et questions de recherche. Ils ont la capacité de modéliser des changements anatomiques complexes et d'identifier des sujets déviant fortement de la normalité.

Mots clés. Imagerie par Résonance Magnétique, Surface Corticale, Imagerie Cérébrale Foetale, Atlas Cérébral, Modèle de Template Déformable, Géométrie Riemannienne, Régression Géodésique, Optimisation Multiéchelle, Ondelettes.

Title. Quantitative analysis of healthy and abnormal fetal brain development using deformation models.

Abstract. The quantitative analysis of fetal brain images is a challenging task due to data scarcity, image quality issues and high intersubject variability. Computational Anatomy provides convenient and unsupervised methods to characterize normal and pathological anatomical structures by computing the deformation that warps a reference object onto another. The goal of this thesis is to adapt geometric tools to the analysis of fetal brain growth within the framework of the Large Diffeomorphic Deformation Metric Mapping. First, we present methodological contributions to address the unrealistic local minima issue that arises when warping high-dimensional or complex objects. We introduce a coarse-to-fine optimization strategy whereby the deformation fields are decomposed into a Haar-like wavelet basis and the objects are filtered. This multiscale strategy is applied to registration and cross-sectional atlas estimation models as well as to temporal atlas estimation. Using fetal brain images, we demonstrate that our multiscale strategy produces more natural template images and deformations. Then, we design a pipeline to characterize the anatomical variability of fetal brains with agenesis of the corpus callosum at a global scale. Deviations from normal development are quantified by registering each subject to an age-matched healthy template brain and age-related anatomical differences are smoothed out by transporting the deformations to a common space, allowing the identification of deformation modes characterizing abnormal brains. Finally, we propose to build an atlas of cortical surface development during pregnancy. The average cortical growth is modelled by a piecewise geodesic trajectory and the anatomical variability around the mean is estimated using registration and parallel transport. This atlas construction method is compared to the state-of-the-art method, i.e. kernel regression, by assessing the accuracy of the modelled gyrification process. The developed tools are generalizable to different data types and anatomical structures and have the ability to model complex time-dependent anatomical changes and to identify subjects with high deviation from normality.

Keywords. Magnetic Resonance Imaging, Cortical Surface, Fetal Brain Imaging, Brain Atlas, Deformable Template Model, Riemannian Geometry, Geodesic Regression, Multiscale Optimization, Wavelets.

Résumé long en français

Cette thèse s'inscrit dans le cadre de l'analyse géométrique de données d'Imagerie par Résonance Magnétique (IRM) foetales cérébrales. L'analyse quantitative des IRM du cerveau foetal pendant la grossesse est une question difficile liée à des problématiques de manque de données, qualités d'image hétérogènes (mouvements du foetus pendant l'acquisition, tissus maternels autour du cerveau foetal) et variabilité inter-sujets élevée. L'Anatomie Computationnelle offre des méthodes non-supervisées pour caractériser des structures anatomiques normales ou pathologiques en calculant la transformation qui déforme un objet de référence en un autre objet. Le but de cette thèse est d'adapter des outils géométriques pour l'analyse de la croissance du cerveau foetal dans le cadre du Large Diffeomorphic Deformation Metric Mapping (LDDMM) qui modélise les transformations entre objets par des champs de vecteurs dépendants du temps. Dans une première partie, nous présentons des contributions méthodologiques pour pallier au problème des minima locaux irréalistes qui se pose lorsque l'on déforme des structures anatomiques en grande dimension. Dans le cadre du LDDMM, les champs de vecteurs sont régularisés par un noyau Gaussien qui définit un espace de Hilbert à noyau reproduisant et paramétrisés par un nombre fini de moments vectoriels attachés à des points contrôles. La largeur du noyau, définie par l'utilisateur, détermine la distance entre les points contrôles et donc l'échelle des déformations de l'espace : un petit noyau implique un nombre élevé de points contrôles et donc des déformations fines, promptes à favoriser des minima locaux irréalistes ; à l'inverse, un noyau large implique un nombre réduit de paramètres, des déformations plus lisses mais aussi moins précises. Il revient donc à l'utilisateur de trouver le meilleur compromis -quand il existe- entre la précision et le réalisme des déformations. Pour pallier à ce problème, nous introduisons une stratégie d'optimisation multiéchelle, ou "coarse-to-fine", qui décompose les champs de déformation dans une base d'ondelettes de Haar. Plus précisément, la transformée en ondelettes est appliquée au gradient des champs de vecteurs : au début de l'optimisation, les coefficients de détail sont silencés, ce qui force l'algorithme à estimer des transformations lisses ; dès que l'algorithme est proche de la convergence, les coefficients de détails sont progressivement réincorporés dans le signal afin d'estimer des transformations de plus en plus fines. Ce transfert d'information des échelles lisses aux échelles plus fines assure une initialisation plus pertinente des déformations à chaque échelle. Cette reparamétrisation des champs de vecteurs peut être considérée comme une régularisation spatiale supplémentaire, dans le sens où elle préserve la formulation LDDMM et en particulier la structure de Hilbert des champs de vecteurs et donc l'algorithme numérique efficace de calcul des gradients. En particulier, cette régularisation additionnelle n'augmente pas la complexité de l'algorithme. La stratégie multi-échelle est appliquée à deux tâches d'Anatomie Computationnelle : l'estimation d'atlas cross-sectionnel et le recalage (cas particulier de l'estimation d'atlas). Nous effectuons des expériences de validation croisées sur trois jeux de données de difficulté croissante : des images manuscrites du chiffre deux, des personnages dessinés à la main et des IRM foetales avec agénésie du corps calleux. Sur les données d'apprentissage, une image représentative de la population est estimée ; cette image est ensuite recalée vers chacune des images de l'ensemble de test. Les expériences montrent que sur tous les jeux de données, l'algorithme multiéchelle est capable d'estimer des images moyennes op-

timales d'un point de vue qualitatif et quantitatif par rapport à l'algorithme d'origine, tout en évitant les minima locaux irréalistes lors qu'une paramétrisation fine est utilisée. Dans cette thèse, on s'intéresse également à la modélisation de la croissance du cerveau foetal, et donc à des modèles d'estimation de trajectoires temporelles. La stratégie multiéchelle est étendue de deux manières : 1- elle est appliqués à des modèles de régression géodésique par morceau, qui autorisent des changements de dynamique temporelle 2- la stratégie multiéchelle sur les champs de vecteurs est combinée à une représentation multiéchelle des objets : filtrage gaussien pour les images et filtrage laplacien pour les maillages. La capacité de ce nouvel algorithme à générer des déformations réalistes est évaluées sur trois jeux de données temporelles : des personnages dessinés à la main, des images du cerveau foetal à différents âges gestationnels et des maillages de la surface corticale du cerveau foetal à différents âges issus de la base de données publique FeTa. Pendant la phase d'entraînement, une régression géodésique représentative de la croissance moyenne de la population est estimée ; pendant la phase de test, un recalage estime la déformation entre chaque image test et une image extraite de la régression géodésique au même âge. L'algorithme multiéchelle double est comparé à l'algorithme d'origine, associés soit à une régression géodésique simple (une seule dynamique), soit à une régression par morceaux (plusieurs dynamiques temporelles). Ces expériences permettent de démontrer que la stratégie multiéchelle double associée à une régression par morceaux permet d'estimer des déformations réalistes dans des contextes divers. Elle est en particulier utile pour modéliser des modifications anatomiques complexes. Dans une seconde partie, nous appliquons les méthodes d'Anatomie Computationnelle développées à l'analyse quantitative du cerveau foetal. La première application s'attache à caractériser les cerveaux de fœtus avec agénésie du corps calleux, une anomalie congénitale à l'issue clinique incertaine. Contrairement aux études actuelles qui se concentrent sur une structure cérébrale spécifique, nous développons ici une pipeline d'analyse d'images ciblant le volume cérébral entier. 38 IRM issues de fœtus sains et 73 IRM issues de fœtus avec agénésie du corps calleux sont rétrospectivement sélectionnés à l'hôpital Trousseau. Pour chaque sujet, une image volumétrique du cerveau est obtenue de manière semi-automatisée en combinant des outils de pré-traitement open-source pour la segmentation, la reconstruction volumétrique et la réorientation. Une trajectoire continue représentative de la croissance du cerveau sain est obtenue par régression géodésique. Chaque cerveau foetal est ensuite associé et recalé à un cerveau template sain issu de cette trajectoire pour quantifier ses déviations par rapport au développement normal. Les déformations sont transportées parallèlement jusqu'à un espace commun, ce qui revient à réduire les différences d'âges entre les fœtus. Les modes de déformations spécifiques à la pathologie étudiée sont identifiés à l'aide d'une analyse en composantes principales et d'un Support Vecteur Machine entraîné à discriminer les sujets sains des sujets pathologiques. Le mode de déformation qui discrimine le mieux les sujets sains des sujets pathologiques combine des altérations cérébrales courantes chez les fœtus avec anomalie du corps calleux. Ce travail est prometteur pour l'exploration quantitative des cerveaux de fœtus anormaux à l'aide de modèles de déformation, en particulier pour l'identification d'anomalies anatomiques corrélées à une issue clinique défavorable. Finalement, nous proposons de construire un atlas du développement cortical sain au cours de la grossesse. Pour cela, nous combinons les outils développés dans les parties précédentes pour proposer une alternative à la méthode classique de construction d'atlas, dite

« kernel regression », qui estime chaque cerveau de référence à un âge donné indépendamment des autres. Notre modèle, basé sur une régression géodésique par morceaux, estime un unique objet de référence à un âge donné et la trajectoire qui déforme ce cerveau au cours du temps et représente la dynamique de croissance moyenne de la population au cours de la grossesse. La variabilité anatomique autour de cette moyenne est estimée grâce à des outils de recalage et de transport parallèle. A la différence de l'analyse précédente, ici l'on estime la trajectoire moyenne et la variabilité autour de cette moyenne de manière simultanée et non pas séquentielle. Ce modèle est évalué sur 65 maillages de surfaces corticales entre 20 et 37 semaines de gestation : 36 sont issues de la base de données publique FeTa et 29 sont issues de la plateforme Lumière. Le pré-traitement des données repose sur NiftyMIC pour la reconstruction volumétrique des images et les outils du projet dHCP pour la segmentation et l'extraction des surfaces corticales. Notre modèle est comparé à la kernel regression par validation croisée. La performance est quantifiée en calculant des mesures de distance (distance de Hausdorff, distance varifold) entre chaque sujet et un cerveau extrait de la trajectoire moyenne au même âge. Afin d'évaluer la justesse du processus de gyrification représenté, on compare également les mesures de courbures des atlas estimés à celles des sujets de test et d'entraînement (indice de gyrification, aire de la surface corticale, courbure moyenne, indice de forme). Les résultats montrent que la kernel regression modèle le plissement cortical de manière un peu plus juste, ce qui était attendu puisque ce modèle est conçu pour avoir de bonnes performances à un âge gestationnel donné. En revanche, notre méthode continue permet de garantir la cohérence temporelle du processus modélisé. Notre modèle peut non seulement construire des objets anatomiques de référence et leur segmentation, mais aussi analyser de manière fine des dynamiques de croissance, positionner n'importe quel sujet par rapport à la croissance moyenne, déplacer ledit sujet dans l'espace et le temps et finalement caractériser la variabilité anatomique des populations saines et pathologiques. Une perspective intéressante serait d'intégrer des données postnatales au sein de notre modèle afin d'étudier des périodes développementales plus longues. En conclusion, cette thèse propose un changement de perspective par rapport à l'analyse du cerveau foetal en se concentrant sur l'aspect dynamique de la croissance prénatale. De manière plus générale, cette thèse s'intéresse à la problématique de l'analyse de structures anatomiques dans un contexte difficile, présentant une variabilité importante de formes et d'apparences, ce qui soulève des problématiques méthodologiques et applicatives variées dépassant le seul cadre de l'analyse d'images pré-natales. Les outils développés peuvent être généralisés à différents types de données, structures anatomiques et questions de recherche. Ils ont la capacité de modéliser des changements anatomiques complexes et d'identifier des sujets déviant fortement de la normalité.

Remerciements

Mes remerciements vont naturellement à mes directeurs de thèse, Stéphanie et Erwan. Stéphanie pour m'avoir fait confiance avec un sujet de thèse riche et original, pour son sens de l'écoute, et sa capacité à dénicher de nouveaux collaborateurs ou des connexions à d'obscurs serveurs. Erwan pour sa disponibilité, sa bonne humeur et sa polyvalence mathématique et programmatique qui a contribué à rendre ce manuscrit présentable. A tous les deux, merci pour le temps passé à m'expliquer la géométrie Riemannienne en première année et les ondelettes en deuxième année (à vrai dire, je ne me suis pas rendue compte que je faisais des maths); merci surtout pour vos encouragements, votre enthousiasme communicatif et pour votre éthique de l'encadrement. C'est une vraie chance en tant que doctorant.e que de se sentir autorisé.e à "échouer" autant qu'à réussir. J'espère que vous continuerez à encadrer ensemble ou séparément des doctorants dans le futur !

Je suis reconnaissante à Olivier Colliot et Alain Trouvé pour avoir accepté de faire partie de mon Comité de Suivi ainsi que pour leurs conseils et remarques toujours pertinents sur mon travail.

Je tiens également à remercier les équipes médicales sans qui j'aurais déformé des chiffres et des bonhommes pendant trois ans. A l'hôpital Trousseau, merci à Eléonore Blondiaux grâce à qui cette thèse a démarré sur une question scientifique passionnante. Merci à Aurélie, stagiaire motivée et motivante avec qui j'ai eu grand plaisir à travailler. A l'hôpital Necker, merci à toute l'équipe de la plateforme Lumière pour leur enthousiasme et pour m'avoir donné l'opportunité de collaborer sur un sujet tout aussi captivant. Merci enfin aux patientes qui ont rendu ce travail et de nombreux autres possibles.

A l'Institut Pasteur, merci à l'équipe de mon stage de M2 que j'ai eu le plaisir de continuer à côtoyer par Zoom interposé tout au long de cette thèse, et notamment Guillaume et Aline.

Je suis reconnaissante aux membres de l'équipe Heka dont la bonne humeur ont contribué à rendre mon travail plaisant et en particulier aux doctorants petits et grands de l'équipe Stéphanie et du bureau 305 de Parisanté(Campus): Vianney, Grand Clément, Solange, Petit Clément, Pierre, Louis, Linus et Agathe. Merci également à Alexandre et Vianney pour m'avoir aidée à mettre cette thèse sur les rails.

Je remercie également toutes les personnes m'ayant accompagnée pendant ces trois années de thèse. Merci à Cécile qui a eu la tâche ingrate de relire deux chapitres de cette thèse. En souvenir de très beaux voyages aux quatre coins de la France et de discussions interminables, je remercie Marie, Cécile, Julien et Milo, même si à mon grand regret nous avons échoué à perdre le quatrième élément dans la campagne de l'Indre. Enfin, je remercie mes parents et ma soeur pour avoir hébergé une doctorante et deux félins fauteurs de trouble, stoïquement supporté la vue de cerveaux déformés, pour m'avoir transmis leur goût du travail, mais aussi et surtout pour leur soutien sans faille.

CONTENTS

List of abbreviations	1
List of mathematical notations	2
1 General introduction	3
1.1 <i>In vivo</i> analysis of the fetal brain	3
1.2 Deformation models	4
1.3 Bridging the gap	5
I Multifaceted Multiscale Strategies for Realistic Object Mapping	9
2 Introduction to the Large Deformation Diffeomorphic Metric Mapping framework	10
2.1 Shape spaces	11
2.2 Notions of Riemannian Geometry	12
2.3 The Large Deformation Diffeomorphic Metric Mapping framework	14
2.3.1 Building large deformations	14
2.3.2 Sparse parameterization of diffeomorphisms	15
2.3.3 Optimization	16
2.3.4 Shape attachment	17
2.4 Computational Anatomy models	18
2.4.1 Cross-sectional atlas	19
2.4.2 Registration	20
2.4.3 Geodesic regression	20
2.4.4 Piecewise geodesic regression	20
3 Multiscale deformations for atlas estimation	22
3.1 Motivation	22
3.2 Related work	24
3.2.1 Multiscale representations of deformations	25
3.2.2 Multiscale strategies in the LDDMM framework	25
3.2.3 Multiscale strategies and sparse parameterization of diffeomorphisms	27
3.3 Wavelet-based multiscale atlas estimation	28
3.3.1 Overview	28
3.3.2 Model of diffeomorphic deformations	29
3.3.3 Reparameterization of the initial velocity fields	30
3.3.4 Multiscale optimization algorithm	36
3.3.5 Multiscale optimization algorithm with local adaptation	39
3.3.6 Alternative multiscale algorithm	40
3.4 Experiments: cross-sectional atlas estimation	40
3.4.1 Handwritten digits	41
3.4.2 Toy data	42
3.4.3 Fetal brain images	44

3.5	Additional experiments	46
3.5.1	Local adaptation extension	46
3.5.2	Alternative multiscale scheme	46
3.6	Discussion	48
4	Versatile multiscale strategies	50
4.1	Motivation	50
4.2	Methods	52
4.2.1	Overview	52
4.2.2	Model of diffeomorphic deformations	52
4.2.3	Dual multiscale optimization	53
4.3	Experiments - cross-sectional atlasing	57
4.3.1	Comparison of the different multiscale strategies	57
4.3.2	Computational requirements	58
4.4	Experiments - temporal modelling	58
4.4.1	Regression on toy data	61
4.4.2	Regression of fetal brain images	63
4.4.3	Regression of cortical surface meshes	65
4.5	Discussion	67
II	Exploring healthy and pathological fetal brain growth	71
5	The quantitative analysis of fetal brain images - an overview	72
5.1	The growing field of fetal MRI	73
5.1.1	Basic principles of MRI	73
5.1.2	A brief history of fetal MRI	73
5.1.3	Use in a clinical setting	75
5.1.4	Challenges	76
5.2	Fetal MRI Preprocessing	77
5.2.1	Brain extraction	77
5.2.2	Volume reconstruction	78
5.2.3	Tissue segmentation	79
5.3	Quantitative analysis	80
5.3.1	Volumetric analyses	80
5.3.2	Cortical folding	81
5.3.3	Spatiotemporal atlases	83
6	Analysis of the anatomical variability of pathological fetal brains	86
6.1	Motivation	86
6.2	Related work	88
6.3	Dataset	89
6.3.1	Data description	89
6.3.2	Image preprocessing pipeline	91
6.3.3	Final dataset	93
6.4	Shape analysis pipeline	94
6.4.1	Geometrical tools	94

6.4.2	Shape analysis pipeline	96
6.5	Results	100
6.5.1	Healthy brains versus brains with CCA	100
6.5.2	Complete CCA versus Partial CCA	101
6.6	Discussion	103
6.6.1	Fetal MRI preprocessing	105
6.6.2	Applying deformation models to fetal brains	105
6.6.3	The anatomical variability of fetal brains with CCA	106
6.6.4	Partial and complete CCA differentially affect brain anatomy	107
6.7	Conclusion	107
7	A spatio-temporal atlas of the developing cortex	109
7.1	Motivation	110
7.2	Data	111
7.2.1	FeTa dataset	112
7.2.2	Fondation Lumière	113
7.2.3	Image preprocessing	113
7.3	Methods	115
7.3.1	Deformation model	115
7.3.2	Model of cortical surface development	117
7.3.3	Statistical model	119
7.3.4	Optimization procedure	120
7.3.5	Kernel regression model	122
7.3.6	Performance assessment	123
7.4	Preliminary results	123
7.4.1	Atlases visual comparison	124
7.4.2	Atlas evaluation	125
7.4.3	Influence of the hyper-parameters	127
7.5	Discussion	128
8	General conclusion	131
8.1	Multiscale optimization(s)	131
8.2	Analysis of fetal MRI	132
8.3	Fetal brain analysis pipeline	132
8.4	Prenatal brain development atlas	133
8.5	Perspective	134
	Supplementary Material	135
A	Supplementary material of Chapter 3	135
A.1	Characters dataset	135
A.2	Algorithmic implementations of the transforms FWT and IWT	136
A.3	Multiscale algorithm with local adaptation	142
A.3.1	Local adaptation extension	142
A.3.2	Experiments: atlas estimation on handwritten digits with local adaptation	145
A.4	Alternative multiscale scheme	147
A.4.1	Modifying the renormalization step	147

A.4.2	Additional experiments - alternative multiscale algorithm	150
A.5	Additional results - fetal brain MRI	153
B	Supplementary material of Chapter 4	154
B.1	Extrinsic mesh smoothing	154
B.2	Dual multiscale Algorithm	156
B.3	Tuning the number of components in geodesic regression	157
B.4	Additional results - regression on toy data	158
B.5	Additional results - regression on fetal brain images	160
C	Supplementary material of Chapter 6	162
C.1	Fetal MRI processing pipeline	162
C.1.1	Semi-automated brain extraction and volume reconstruction pipeline	162
C.1.2	Volumetric image postprocessing pipeline	164
C.2	Additional results - deformation modes	166
C.2.1	First mode of deformation (healthy brains versus CCA)	166
C.2.2	Second mode of deformation in MR intensities (healthy brains versus CCA)	167
C.2.3	First mode of deformation in MR intensities (Complete versus partial CCA)	168
D	Supplementary material of Chapter 7	169
D.1	Volume reconstruction pipeline (Lumière dataset)	169
D.2	Manual correction of extracted cortical surfaces	170
D.3	Modification of parallel transport	171
D.3.1	Parallel Transport algorithm	171
D.3.2	Parallel Transport example	172
D.4	Optimization procedure	173
D.4.1	Hyper-parameters setting	173
D.4.2	Bayesian prior log likelihoods	173
D.4.3	Initialization of the parameters	174
D.5	Temporal discrepancy in kernel regression	175
D.6	Hyper-parameters and curvature measures	175
D.6.1	Deformation and varifold kernels	175
D.6.2	Number of components	176
	Reference	177

LIST OF FIGURES

1.1	Early <i>in vivo</i> representations of the fetus	4
2.1	The Riemannian Exponential (panel (a)) and Parallel Transport and Exp Parallelization (panel (b)) on a sphere.	13
2.2	Parameterization of a vector field with control points and momentum vectors	14
2.3	Computing shape deformations from an initial object and a set of initial momentum vectors.	16
2.4	Varifold representation of a cortical surface mesh	17
2.5	Representation of the cross-sectional atlas, registration, geodesic regression and 2-component piecewise geodesic regression on toy data.	19
3.1	Dilemma linked to the choice of the deformation kernel on a simple registration example between two cortical surface meshes	24
3.2	Overview of the coarse-to-fine optimization strategy	29
3.3	The one-dimensional Haar wavelet	31
3.4	Decomposition of a 4-by-4 grid in two bases	33
3.5	Original and multiscale strategies applied to a registration example	39
3.6	Estimation of the template image by the original and multiscale algorithms on the dataset of handwritten digits with different values of σ_g	41
3.7	Estimation of the template image by the original and multiscale algorithms on the toy dataset with different parameters σ_g	43
3.8	Atlas estimation by the original and multiscale algorithms on the dataset of fetal brain images	47
4.1	Multiscale and dual multiscale strategies applied to a registration example	55
4.2	Residual error rate of the training and test phases	58
4.3	Number of iterations until convergence for each atlas estimation (training) and registration (test) experiments	59
4.4	Estimation of the template image by the original, multiscale momenta, multiscale objects and dual multiscale algorithms on the toy dataset	60
4.5	Three-step evaluation procedure illustrated on the dataset of cortical surface meshes.	61
4.6	Geodesic regression results on the toy dataset for different values of σ_g	63
4.7	Results of the evaluation procedure on the dataset of fetal brains images during the first fold of cross-validation	66
4.8	Comparison of one-component and piecewise trajectories associated to multiscale optimization on cortical surface meshes	68
4.9	Comparison of the original and multiscale strategies associated to piecewise regression for $\sigma_g = 3.5$	69
5.1	Evolution of fetal MRI over the years.	74
6.1	Brain T2-weighted MR images of fetuses from Hôpital Trousseau in sagittal view.	89
6.2	Histograms of the subjects gestational ages at the date of MRI acquisition.	90
6.3	Steps of the semi-automated brain extraction and volume reconstruction pipeline	92

6.4	Histogram of the year of MRI acquisition for images with different reconstruction outcomes	93
6.5	Shape analysis pipeline.	97
6.6	Results from PCA and SVM-based classification (healthy fetuses versus fetuses with CCA)	101
6.7	Second mode of deformation extracted from PCA, distribution of the subjects scores on the related component and volume changes of several structures.	102
6.8	Results from PCA and SVM-based classification (fetuses versus with complete CCA versus fetuses with partial CCA)	103
6.9	First mode of deformation extracted from PCA, distribution of the subjects scores on the related component and volume changes of several structures	104
7.1	Comparison of the main atlas building method with our proposed method	111
7.2	Histograms of the subjects gestational ages at the date of MRI acquisition	112
7.3	Steps of the preprocessing pipeline.	114
7.4	Example of a piecewise trajectory $\gamma(t)$ with $P = 3$ components.	117
7.5	Construction of an individual trajectory from the representative trajectory $\gamma(t)$	119
7.6	Visual comparison of the spatiotemporal atlases build with our method and kernel regression	125
7.7	Comparison of five curvature measures between the subjects and two spatiotemporal atlases	127
7.8	Comparison of the atlases templates at age 36 GW with two subjects of age 36 GW	128
A.1	Dataset of manually designed characters.	135
A.2	The FWT Algorithm applied to a 3-by-4 array	137
A.3	Estimation of the template image by the original, multiscale and multiscale with local adaptation algorithms on the dataset of handwritten digits with 3 different parameters σ_g	146
A.4	Atlas estimation by the original and multiscale algorithms on the dataset of fetal brain images with $\sigma_g = 7$	153
B.1	Smooth varifold representation of a cortical surface mesh	155
B.2	Residual error rate versus the number of components on a piecewise geodesic regression task with the original algorithm	157
B.3	Geodesic regression results on the toy dataset for $\sigma_g = 10, \sigma_g = 8, \sigma_g = 6, \sigma_g = 5$, and $\sigma_g = 4$	159
C.1	Volume-reconstructed fetal brain images excluded from the analysis.	162
C.2	Volumetric image postprocessing pipeline	163
C.3	Illustration of the fetal MRI postprocessing pipeline on three examples	165
C.4	First mode of deformation extracted from PCA and distribution of the subjects scores on the related component.	166
C.5	Second mode of deformation extracted from PCA and distribution of the subjects scores on the related component.	167
C.6	First mode of deformation extracted from PCA and distribution of the subjects scores on the related component	168
D.1	Comparison of brain volumes reconstructed with NiftyMIC with an outlier rejection threshold of 0.8 and 0.5	169
D.2	Cortical surface meshes extracted with the dHCP pipeline before and after manual correction and smoothing.	170

D.3	Parallel transport of a cortical surface mesh	172
D.4	Temporal discrepancy in kernel regression.	175
D.5	Comparison of five curvature measures between the subjects and three spatiotemporal atlases build with different values of σ_g and σ_w at each gestational week. . . .	175
D.6	Comparison of five curvature measures between the subjects and four spatiotemporal atlases build with different number of components at each gestational week. .	176

LIST OF TABLES

3.1	Performance of the original and multiscale algorithms on the dataset of handwritten digits during the training and test phases.	42
3.2	Performance of the original and multiscale algorithms on the dataset of artificial characters during the training and test phases	44
3.3	Performance of the original and multiscale algorithms on the dataset of fetal brains images during the training and test phases.	45
4.1	SSIM values over five folds of cross-validation for atlas estimation, geodesic regression and registration on the toy dataset	62
4.2	SSIM values over five folds of cross-validation on the dataset of fetal brain images .	64
4.3	Residual error rate over five folds of cross-validation on the dataset of cortical surface meshes	67
5.1	Fetal cortical surface spatiotemporal atlases in the literature	84
5.2	Fetal brain spatiotemporal atlases in the literature	85
6.1	Quantitative analyses of fetal brains with CCA.	89
6.2	MR acquisition parameters of the dataset.	91
7.1	MR acquisition parameters of the dataset.	113
7.2	Curvature measures.	124
7.3	Residual error for our method and kernel regression, computed as the varifold distance between each subject and the closest template	126
7.4	Average Hausdorff distance between each subject and the closest template for our method and kernel regression	126
A.1	Performance of the original, multiscale and multiscale with local adaptation algorithms on the dataset of handwritten digits during the training and test phases . . .	145
A.2	Performance of the alternative multiscale algorithm with different values of ρ and σ_g during atlas estimation on the toy dataset	151
A.3	Performance of alternative multiscale algorithm with different values of ρ and σ_g during registration on the toy dataset	152
B.1	Mean \pm standard deviation of relative residual error over five folds of cross-validation for atlas estimation, geodesic regression and registration on the toy dataset	158
B.2	Mean \pm standard deviation of relative residual error over five folds of cross-validation for atlas estimation, geodesic regression and registration on the dataset of fetal brain images.	160
B.3	SSIM values between the templates estimated during geodesic regression and Gholipour et al. atlas [74] at each gestational week.	161

List of abbreviations

CCA	Corpus Callosum Agenesis
CHD	Chronic Heart Disease
CNN	Convolutional Neural Network
EPI	Echo Planar Imaging (a type of MR sequence)
FWT	Forward Wavelet Transform
GW	Gestational Weeks
HASTE	half-Fourier single-shot turbo spin-echo (T2-weighted MR sequence)
ICA	Independent Component Analysis
IWT	Inverse Wavelet Transform
LDDMM	Large Diffeomorphic Deformation Metric Mapping
MRI	Magnetic Resonance Imaging
PCA	Principal Component Analysis
RKHS	Reproducing Kernel Hilbert Space
SNR	Signal-to-Noise ratio
SRR	Super Resolution Reconstruction
SVM	Support Vector Machine
SVR	Slice-to-Volume Registration

List of mathematical notations

\odot	Element-wise matrix multiplication
$\ \cdot\ _V$	Norm of the space V
$a_{s,k}$	Wavelet average coefficient with with scale s and position k
c_k	k^{th} control point parametrizing a diffeomorphic transformation
$d_{s,k}^o$	Wavelet detail coefficient with scale s , position k and orientation o
$\mathcal{E}xp_{p,t_0,t}(v)$	Riemannian exponential map at point p of the tangent vector v
I_d	Identity transformation
k_g	Number of control points parametrizing a velocity field
K_g	Deformation Gaussian kernel defining a RKHS
M	Riemannian manifold
$\mathcal{M}_{m,n}$	Space of real matrices with m lines and n columns.
$\{O\}$	Varifold representation of a surface mesh O
$P_{\gamma,t_0,t}(\omega)$	Parallel transport of the vector ω along the curve γ between $\gamma(t_0)$ and $\gamma(t)$
r_j	Total residuals value at iteration j
$T_p M$	Tangent space at p of the manifold M
v_t	Instantaneous velocity field at time t , with $t \in [0, 1]$
α_k	k^{th} momentum vector parametrizing a diffeomorphic transformation
$\eta_\gamma^\omega(t)$	Exp-parallelization of the trajectory γ by ω
$\gamma(t)$	Trajectory on a Riemannian Manifold taken at time t
ρ	Normalisation factor
$\phi(t)/\phi_1$	Diffeomorphic transformation at time t / at time $t = 1$
$\phi(t) \star O$	Action of a diffeomorphic deformation on the shape O
$\Phi_{s,k}$	Scaling function rescaled by 2^s and translated by $k2^s$
$\psi_{s,k}$	Wavelet function rescaled by 2^s and translated by $k2^s$
σ_t	Noise parameter (trade-off between data attachment and regularity)
σ_g	Width of the Gaussian kernel K_g
ω	Space shift, i.e. momentum parametrizing a diffeomorphism

GENERAL INTRODUCTION

Contents

1.1 <i>In vivo</i> analysis of the fetal brain	3
1.2 Deformation models	4
1.3 Bridging the gap	5

1.1 *In vivo* analysis of the fetal brain

The development of the central nervous system begins in the third week of gestation and continues through childhood and adulthood [107]. The intricate and dynamic process of fetal brain growth is controlled by the coordinated interaction of environmental and genetic factors. It involves a variety of different mechanisms at different scales and speed: neural tube formation and division, migration of cortical neurons, synapse formation, emergence of transient structures and cortical folding into sulci and gyri [107].

Constraints inherent in the fetal organism, restrictive legislation and social attitudes [50] have delayed the study of the fetal anatomy. In 1511, Leonardo da Vinci was the first scientist to accurately draw a human fetus in utero (Figure 1.1a), and around 1600, Girolamo Fabrici one of the first to dissect the placenta and fetuses [77] (Figure 1.1b). Centuries later, the emergence of medical imaging techniques provided a unique insight into *in vivo* fetal development. With the advent of ultrasonography and Magnetic Resonance Imaging (MRI) [129], it became possible to visualize the fetal anatomy with a remarkable level of detail. In a clinical setting, fetal MRI is used as a complement to ultrasonography to confirm a diagnosis of brain anomaly. For researchers, it is an excellent modality to understand prenatal brain growth and better characterize diseases.

Research in the field has been challenged by a number of factors, including but not limited to the encasement of the fetus in the mother's body, ethical and safety concerns, elevated cost of scanners, and unpredictable fetal movements during scanning [197]. The main consequence is the lack of available, high quality, homogeneous data, which have limited fetal brain research to two-dimensional measurements [218]. The recent development of ultrafast scanning and motion correction techniques considerably galvanized research around fetal brain development [192, 181]. Provided that the necessary preprocessing steps are performed, it is now possible to obtain high-resolution, 3-dimensional images of the fetal brain [56]. In the last decade, quantitative studies have blossomed,



(a) The fetus in the womb, sketches and notes on reproduction (1511), Leonardo da Vinci



(b) Drawings of human fetuses, *De formato foetu* (1600), Girolamo Fabrici [77]

Figure 1.1: Early *in vivo* representations of the fetus

performing volumetric analyses of the fetal brain, investigation of the cortical folding process [34], and spatiotemporal modelling of the fetal brain growth [86].

Analysis of the prenatal brain is a growing but newly born field of research with many obstacles to overcome, namely the lack of open-source data and complex and heterogeneous growth patterns. As a result, there is a need for new methods specifically adapted to the challenges of fetal brain MRI. In this thesis, we propose to perform quantitative analysis of the fetal brain using geometrical deformation models, which offer very interesting features:

- the ability to analyze small datasets in an unsupervised setting;
- versatility regarding the type of objects analyzed, e.g. images or meshes;
- the ability to analyze different types of datasets, e.g. cross-sectional, time-series or longitudinal data;
- explicability and visual interpretability.

1.2 Deformation models

To analyze medical data, two main approaches are available. Supervised learning focuses on building models that classify data or predict outcomes using a priori expert knowledge. With the increasing availability of high-quality medical images, supervised learning methods, specifically deep learning ones, have become very popular for they are able to tackle a variety of tasks such as disease diagnosis, pathology localization and structure segmentation [41]. When medical images are scarce or unlabelled, when the research question is less specific, or when model explicability is imperative, one must

turn to unsupervised methods which focus on discovering relationships or clusters of patterns without a-priori knowledge involved.

Computational Anatomy is an interdisciplinary field rooted in the idea from d’Arcy Thompson to describe differences between shapes through deformations of the ambient space [227]. Compared to classical volumetric analyses, deformable brain mapping can provide precise and localized information about structural deviations from a reference anatomy, whether they are age-related, disease-related or even normal anatomical variations. In this thesis, we are interested in two main applications of Computational Anatomy:

- The deformable template model [6, 54]: given a cross-sectional dataset (e.g. fetal brains imaged at the same gestational stage), the *average* features of the population are summarized by a reference shape called template, and the anatomical *variability* is summarized by the template-to-subject deformations.
- Temporal modelling [65, 30]: given a time-series dataset (e.g. fetal brains imaged at different gestational stages), we seek to model the temporal evolution of a phenomenon by computing the deformation of an estimated reference shape. Here, the focus is not on a single average shape, but on its time-dependent changes.

The choice of the mathematical framework that models and regularizes deformations occurring between anatomical structures is of prime importance. In this thesis, we rely on the Large Deformation Diffeomorphic Metric Mapping (LDDMM) setting [233, 31], a powerful method for computing *diffeomorphic* transformations, which are considered as geodesics on a Riemannian manifold. It has convenient properties such as the tangent-space representation of diffeomorphisms input, which fosters the use of linear statistical tools, but also useful geometrical tools such as parallel transport [135], which enables to transport deformations along geodesic paths.

The LDDMM framework has been exploited for a variety of clinical applications. To name but a few: estimation of reference shapes [176] and normal anatomical variability [174] of subcortical brain structures in healthy individuals; hippocampal shape classification [238] and longitudinal modelling [177, 44] in Alzheimer’s disease; characterization of early thalamus shape changes in frontotemporal dementia [40]; characterization of subregional atrophy of subcortical structures in prodromal Huntington disease [256]. With few exceptions [58, 163], the LDDMM framework has been mostly applied to the modelling of neurodegenerative diseases. In this thesis, we propose to extend its scope of application to the other end of the age spectrum, i.e. to the fetal development period.

1.3 Bridging the gap

In this thesis, we aim to bridge the gap between deformations models and the analysis of healthy and pathological fetal brains. Unlike common practices in Applied Mathematics, our goal is not to develop new mathematical models, validate them on toy data and subsequently apply them to clinical questions. Rather, we adopt a more practical approach where we seek to to *adapt* and *improve* existing methods to the analysis of the fetal brain.

The clinical application is not necessarily the end point of our work; it may actually be its driving force as in Chapters 6 and 7. That is not to say that the methods developed in this thesis are limited to specific applications. In fact, we are willing to introduce open-source tools that can be generalized to different questions. To put it simply, our goals are twofold: to provide practical, straightforward methods that meet clinical needs, but also to develop innovative mathematical approaches that may be of use outside a specific clinical question. Depending on the task at hand, the weights granted to each concern may vary.

In this thesis, we work with different types of shapes: 2-dimensional images (hand-written digits and toy data), 3-dimensional gray-level MR images, and cortical surface meshes. In particular, we rely on three rich datasets consisting of T2-weighted fetal brain MR images and comprising a single observation per subject:

- Hôpital Trousseau, France: MR images from 61 healthy fetuses and 118 fetuses with agenesis of the corpus callosum, acquired during clinical routine (Chapters 3 and 6);
- Fondation Lumière¹, France: 87 MR images from healthy fetuses, acquired in a clinical research platform (Chapter 7);
- Fetal Tissue Annotation and Segmentation Challenge² dataset [168]: 120 publicly available brain volumes from healthy and pathological fetuses collected in two different institutions (Chapters 4 and 7).

This thesis is divided into two main parts. In the first part, we take a methodological approach that aims at improving the flexibility and efficiency of Computational Anatomy models:

- **Chapter 2: Introduction to the Large Deformation Diffeomorphic Metric Mapping framework.** In this chapter, we provide the reader with an overview of the mathematical tools employed in this thesis, and the mathematical foundations and intuitions behind them.
- **Chapter 3: Multiscale deformations for atlas estimation** tackles a well-known problem in object mapping and Computational Anatomy, namely the unrealistic local minima issue. We develop a wavelet-based multiscale reparameterization of the velocity fields defining the transformations and implement a coarse-to-fine optimization strategy to enhance the performance of cross-sectional atlas estimation. We evaluate the performance of the multiscale versus original strategies by applying the algorithm to open-source digit images, toy data and pathological fetal brain images. This work has been submitted for publication.
- In **Chapter 4: Versatile multiscale strategies**, we extend our previous work by developing a dual coarse-to-fine strategy based on a multiscale representation of deformations and objects. Further, we apply the multiscale strategy to different Computational Anatomy models, namely atlas estimation, geodesic regression and piecewise geodesic regression. Our algorithm is applied to two complex regression

¹<http://fondation-lumiere.org/>

²<https://feta.grand-challenge.org/>

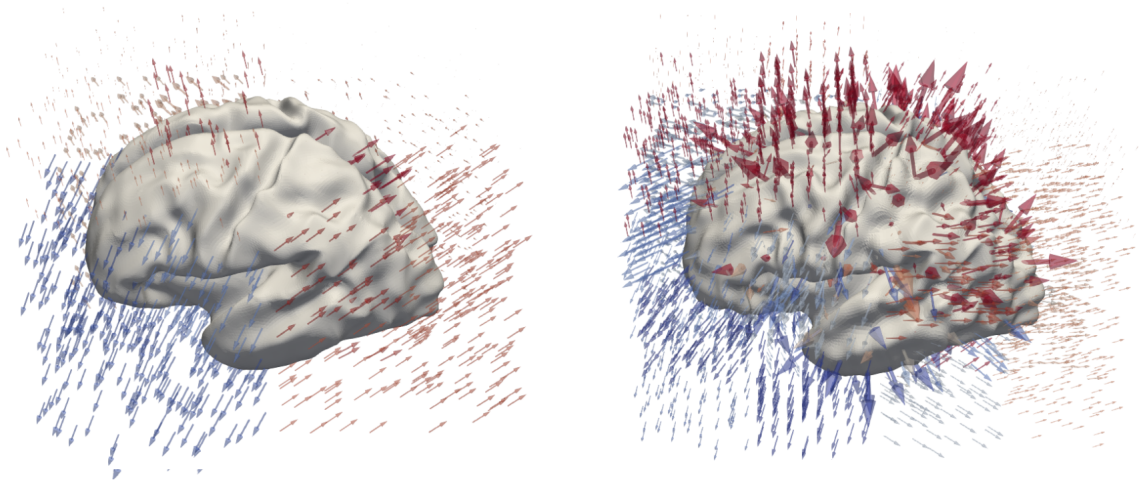
problems: modelling the time-dependent evolution of toy data and the growth pattern of healthy fetal brains. This work was presented at SPIE Medical Imaging 2023 [66] and at the French Colloquium of Signal and Image Processing 2023 [67]. An experiment on fetal brains was subject to a poster presentation at the Fetal, Infant and Toddler Neuroimaging Group 2022.

In a second part, we focus on specific clinical research questions that aim at characterizing the growth patterns of healthy and pathological fetal brains. Of note, we combine the previously developed methodologies to construct practical tools specifically adapted to the specificities of fetal brain MRI:

- **Chapter 5: The quantitative analysis of fetal brain images.** In this chapter, we review the existing work on the quantitative analysis of the fetal brain. We first take a historical perspective to provide the reader with a clear view of the inherent challenges of fetal MRI and explain why this is a recent field of research. Then, we describe the different approaches for the preprocessing of fetal MRI, which are essential for performing automated analyses. Finally, we detail the different types of research about the fetal brain, namely analyses of brain volumes, cortical folding and brain growth.
- **Chapter 6: Analysis of the anatomical variability of fetal brains with Corpus Callosum Agenesis.** Relying on a large dataset of clinically acquired pathological fetal brain images, we develop a shape analysis pipeline based on several Computational Anatomy tools and models, namely geodesic regression, registration and parallel transport. This pipeline allows to compare subjects to an age-matched reference brain, transport subjects to a common space, extract the anatomical variability in the dataset and identify modes of deformation related to abnormal corpus callosum. We also characterize, for the first time, complete versus partial agenesis of the corpus callosum. This work has been presented at the Perinatal, Preterm and Paediatric Image Analysis Workshop at MICCAI 2021 [68], for which it won the Best Paper Award.
- **Chapter 7: A spatio-temporal atlas of the developing cortex.** In this chapter, we combine Computational Anatomy tools used in Chapter 6 and optimization strategies introduced in Chapters 3 and 4 to propose a new integrative model for modelling cortical folding in the fetal period. Set in a Bayesian framework, our method simultaneously estimates an average trajectory composed of several dynamics and the subjects space-shifts, i.e. the anatomical variability of healthy fetal brains. We apply our model to cortical meshes extracted from two fetal MRI datasets. Our method is compared with the state-of-the-art kernel regression by computing curvature measures to assess the accuracy of the modelled gyrification process. This chapter will be converted into an article for submission. Preliminary results have been the subject of poster presentations at the French Colloquium of Artificial Intelligence in Biomedical Imaging 2023 and AI4Health Summer School 2023.

Part I

Multifaceted Multiscale Strategies for Realistic Object Mapping



INTRODUCTION TO THE LARGE DEFORMATION DFFEOMORPHIC METRIC MAPPING FRAME- WORK

Contents

2.1	Shape spaces	11
2.2	Notions of Riemannian Geometry	12
2.3	The Large Deformation Diffeomorphic Metric Mapping framework	14
2.3.1	Building large deformations	14
2.3.2	Sparse parameterization of diffeomorphisms	15
2.3.3	Optimization	16
2.3.4	Shape attachment	17
2.4	Computational Anatomy models	18
2.4.1	Cross-sectional atlas	19
2.4.2	Registration	20
2.4.3	Geodesic regression	20
2.4.4	Piecewise geodesic regression	20

In this chapter, we will examine the statistical analysis of shapes through the lens of the geometrical analysis of the brain. This chapter does not intend to provide an exhaustive description of the subject. Mathematical notions are introduced in a simple and accessible manner whenever possible.

A **shape** is the geometry of a structured object (image, mesh, landmark, tensor, etc.) modulo its position, orientation, and size. Usual shapes do not have Euclidean geometry, hence algebraic operations and usual statistical tools cannot be applied to them. Thus, in this chapter, we will try to answer the question: **how can we model varying objects and quantify their variability?**

2.1 Shape spaces

To perform the geometrical analysis of shapes, several interrelated questions have to be answered.

How do we represent shapes, i.e. how do we define a *shape space*?

The most intuitive, albeit not the most practical, way of considering shapes is through their *intrinsic* characteristics. In that sense, a shape space can be seen as the mathematical space of the objects we study, i.e. the space \mathbb{R}^d of *real* objects. In 1917, d'Arcy Thompson [228] introduced a novel definition of shape spaces, later formalized by Grenander [83]: instead of studying the intrinsic characteristic of shapes, he proposed to study the relationships *between* shapes and to model these links through deformations of the ambient space. According to this definition, a shape space \mathcal{S} is an *abstract* representation in which every possible shape is represented by a specific point [122]. The structure of \mathcal{S} -multidimensional, non linear- is given by the range of possible deformations that the shapes of interest can undergo; it is a quotient space from which rigid deformations are excluded, at least in theory. By construction, it naturally inherits the structure of a Riemannian manifold (see Section 2.2). Mathematically, this means that \mathcal{S} derives from the action of a group of transformations \mathcal{G} acting transitively on \mathcal{S} . For instance, in the well-known deformable template model [6], all the possible deformations of a reference object $x_0 \in \mathcal{S}$ are given by the unique orbit $\mathcal{G}x_0$, providing us with a deformation-based coordinate system whose origin is x_0 .

How do we measure shape differences?

Applying algebraic operations to shapes would be meaningless -subtracting the intensities of two brain MR images makes little sense- but measuring their *geometric similarity* is still relevant. Following the work of d'Arcy Thompson [227], we consider the distance between shapes as the difficulty to deform one shape onto another. More precisely, we suppose that the group of transformations \mathcal{G} is endowed with a right-invariant metric d_G , allowing us to compute distances between transformations. By projecting this metric onto the orbit via the group action, we also endow \mathcal{S} with a *metric* d [255]:

$$\forall (x, y) \in \mathcal{S}, d(x, y) = \inf_{g \in \mathcal{G}} \{d_G(I_d, g) \mid g.x = y\}$$

In essence, this defines the metric on \mathcal{S} as the length of the shortest path (according to the metric d_G) between shapes. Hence, to compute distances in \mathcal{S} , one must compute geodesics in the group \mathcal{G} .

How do we model shape deformations in an anatomically relevant way?

In other words, how do we define the group \mathcal{G} ? Rigid and affine deformations are insufficient to express the deformation of complex anatomical structures such as the brain, hence the need to define deformations belonging to higher dimension groups. As for most organs, deformations characterizing brains -healthy ones at least- should not add or remove portions of the shape. Ideally, we want to restrict ourselves to smooth and topology-preserving transformations.

A widely-used framework constructs small deformations ϕ through the application of displacement vector fields v : $\phi = Id_\Omega + v$ [11]. This setting guarantees the regularity of transformations but, unfortunately, not their invertibility: as soon as the variations of v become too large, folds and holes are created, which is incompatible with brain deformations. The framework that we have used in this thesis, the Large Deformation Diffeomorphic Metric Mapping (LDDMM) setting [233, 153, 53, 13] is built on the idea of generating *large* deformations all the while controlling their amplitude (i.e. their regularity). Its mathematical foundations are introduced in Section 2.3.

2.2 Notions of Riemannian Geometry

We introduced the notion of shape spaces, which have the structure of a Riemannian manifold. This allows us to rely on some tools from differentiable geometry. Simply put, the idea of differential geometry is to generalize Euclidean geometry to shapes that are not flat (such as the sphere embedded in \mathbb{R}^3). In this thesis, while we will not explicitly manipulate complex geometrical notions, it is important to introduce the tools underlying the models that we will use to deform shapes.

Riemannian Manifold. Just like a sphere is locally equivalent to its tangent planes, a Riemannian manifold is locally equivalent to a Euclidean vector space. A Riemannian manifold (M, g) is a differentiable manifold with a **metric**, which defines the cost of displacements on the manifold and determines its structure. More formally, it is a smooth manifold whose metric is a smoothly varying inner product g_p on the tangent spaces T_pM at each point $p \in M$.

Geodesic path. Thanks to the metric, one can compute the *length* of any curve $\gamma : t \in [0, 1] \rightarrow \gamma(t) \in M$ on the manifold. A **geodesic** is the generalization of a straight line, the path realizing the minimum length (according to the metric) between two points x and y . The **distance** $d(x, y)$ between x and y is defined as the length of this unique geodesic.

Exponential Map. Let $p \in M$ be a point on the manifold M and v a vector belonging to the **tangent space** T_pM at that point. The Riemannian Exponential, illustrated in Figure 2.1a, is the operator mapping each vector $v \in T_pM$ to the *unique* geodesic $\gamma_{p,v}$ which goes through p at time t_0 with the velocity vector v .

$$\mathcal{Exp}_{p,t_0,t} : v \in T_pM \rightarrow \gamma_{p,v}(t) \in M$$

where T_pM the set of all tangent vectors of M at p . For the sake of simplicity, we will write the usual Riemannian Exponential map between $t_0 = 0$ and $t = 1$ as $\mathcal{Exp}_p(v)$.

The Exponential map is easily computed by solving a second order differential equation. Computing the Exponential is similar to shooting an arrow, which depends only on the position of the bow (the point p) and the orientation and amplitude of the arrow (the velocity vector v). In the following, this operation will be referred to as **geodesic shooting**. It will be used to compute the deformation of a shape under the action of a

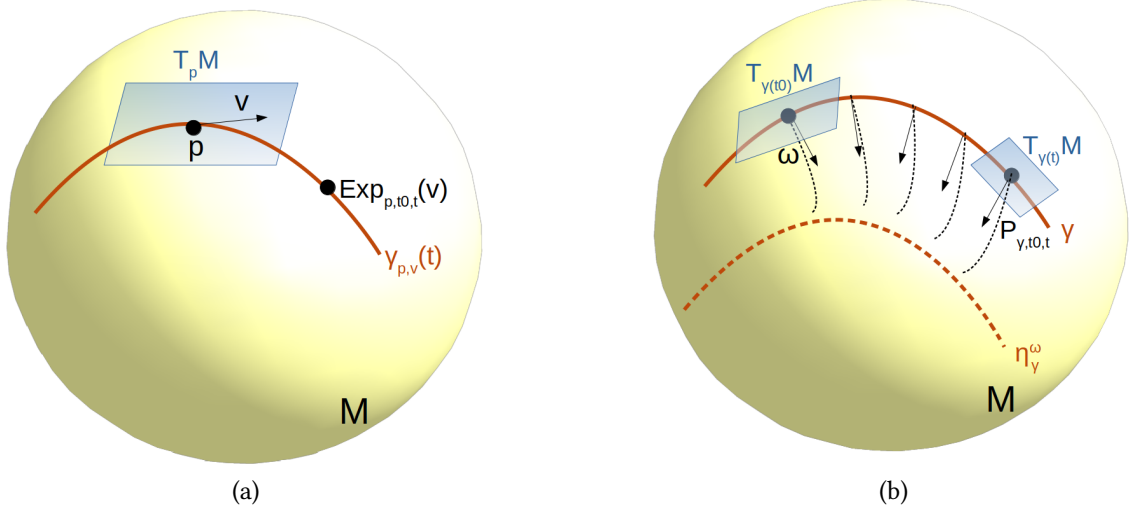


Figure 2.1: The Riemannian Exponential (panel (a)) and Parallel Transport and Exp Parallelization (panel (b)) on a sphere.

velocity field.

Parallel Transport. Parallel transport, illustrated in Figure 2.1b is a differential geometry notion which generalizes the notion of parallels to Riemannian manifolds. It translates a deformation defined by the direction vector $\omega \in T_{\gamma(t_0)} M$ along a curve γ over the manifold M . Given $(t_0, t) \in \mathbb{R}^2$, we define the parallel transport of ω between the points $\gamma(t_0)$ and $\gamma(t)$ along γ as

$$P_{\gamma,t_0,t}(\omega) \in T_{\gamma(t)} M.$$

The computation of parallel transport may be very expensive; in this thesis, we will use an efficient numerical scheme proposed by Louis et al. [135]. Once ω is transported along γ , we need to perform geodesic shooting to obtain the curve parallel to γ . Thus, we define the *exp-parallel variation* of γ along ω as the curve

$$\eta_\gamma^\omega : t \rightarrow \text{Exp}_{\gamma(t)}(P_{\gamma,t_0,t}(\omega)).$$

Let us take a practical example: a geodesic curve γ describes the growth of a brain image B_1 between t_0 and t_1 and we want to transpose the evolution of B_1 onto another brain image B_2 . B_1 and B_2 are both defined at time t_0 (i.e. they have the same age, or are at the same stage of a disease). Through geodesic shooting, we compute the diffeomorphism that best transform B_1 onto B_2 . The set of momenta ω describing this transformation is parallel transported along γ , and the Exponential of the transported vector gives us the exp-parallel trajectory of γ along ω , i.e. the growth trajectory of B_2 between t_0 and t_1 .




Figure 2.2: Parameterization of a vector field (blue arrows) with control points (red points) and momentum vectors (red arrows) using a kernel of width $\sigma_g = 4$.

2.3 The Large Deformation Diffeomorphic Metric Mapping framework

2.3.1 Building large deformations

The LDDMM framework [233, 153, 53, 13] is a generalization of the small deformations setting in which transformations are built by concatenating infinitesimal small deformations. The resulting transformations are *diffeomorphic*, i.e. they are smooth transformations with smooth inverse preserving connected sets and smoothness of anatomical features. Diffeomorphisms are particularly well adapted to the study of anatomical shapes as they are topology-preserving and can be inverted.

More precisely, deformations are constructed by integrating the differential flow equation:

$$\begin{cases} \frac{dx(t)}{dt} = v_t(x(t)) , \\ x(0) = x_0 . \end{cases} \quad (2.1)$$

where $v_t \in V$ is an instantaneous velocity field belonging to a Hilbert space V and x can be seen as a particle moving along the curve $x(t)$ in the domain of interest D .

This model builds a flow of diffeomorphisms: $\phi_t : x_0 \rightarrow x(t)$ is solution to the flow equation at any time $t \in [0, 1]$, and $\phi_t(x_0)$ is the position at time t of a particle that was at position x_0 at time $t = 0$. We have $\phi_0 = I_d$, and the diffeomorphism of interest ϕ_1 is the end point of the path $x(t)$ at $t = 1$:

$$\forall x_0 \in D, \quad \phi_1(x_0) = x(1) .$$

Note that for any time t , ϕ_t is indeed a diffeomorphism provided that the related velocity field is regular enough, i.e. that it is continuous squared integrable.

2.3.2 Sparse parameterization of diffeomorphisms

Finally, we need to define an appropriate norm $\|\cdot\|_V$ for the Hilbert space V . To this end, we restrict ourselves to vector fields that belong to a Reproducing Kernel Hilbert Space (RKHS) [18] V defined by a kernel K_g . We also rely on the work of Durrleman et al. [54] to introduce a discrete parameterization of the velocity fields: we assume that the initial velocity field v_0 can be decomposed as a *finite* linear combination of the RKHS basis vector fields. The weights of the decomposition of a given deformation onto this basis are given by a set of momentum vectors $(\alpha_k(0))_k$ attached to k_g control points $(c_k(0))_k$.

$$v_0(x) = \sum_{k=1}^{k_g} K_g(x, c_k(0)) \alpha_k(0).$$

In this work, K_g is the Gaussian kernel: $K_g(x, y) = \exp(-\frac{\|x-y\|^2}{\sigma_g}) I_d$, where σ_g stands for the kernel width and I_d for the identity matrix. The decomposition of v_0 is represented in Figure 2.2.

Vectors fields that define geodesic deformations are those that minimize the distance $\int_0^1 \|v_t\|_V^2 dt$, i.e. the total kinetic energy used to go from the identity map ϕ_0 to the end point of the path ϕ_1 . In Miller al. [152], it is proved that the vector fields that define geodesic deformations with respect to the norm $\int_0^1 \|v_t\|_V^2 dt$ keep the same structure along time and write according to:

$$v_t(x) = \sum_{k=1}^{k_g} K_g(x, c_k(t)) \alpha_k(t), \quad (2.2)$$

where for any time t , $\alpha_k(t)$ is the k^{th} momentum vector attached to the point $c_k(t)$.

Furthermore, the trajectory of the control points $(c_k(t))_k$ and momentum vectors $(\alpha_k(t))_k$ is described by the Hamiltonian system equations [152]:

$$\begin{cases} \frac{dc_k(t)}{dt} = \sum_{l=1}^{k_g} K_g(c_k(t), c_l(t)) \alpha_l(t) \\ \frac{d\alpha_k(t)}{dt} = - \sum_{l=1}^{k_g} d_{c_k(t)} (K_g(c_k(t), c_l(t)) \alpha_l(t))^t \alpha_k(t) \end{cases} \quad (2.3)$$

with initial conditions $c_k(0) = c_{0,k}$ and $\alpha_k(0) = \alpha_{0,k}$ for all $1 \leq k \leq k_g$.

Finally, one verifies that the kinetic energy along geodesic paths is preserved over time, i.e. $\forall t \in [0, 1], \|v_t\|_V = \|v_0\|_V$. This implies that a geodesic transformation is fully parameterized by the initial velocity field v_0 . Hence, estimation of the diffeomorphism ϕ_1 boils down to a *geodesic shooting* problem. The system is deterministic and we only need to optimize the initial conditions. In this work, the position of the control points is fixed on a regular grid of spacing σ_g and so we only need to optimize $\alpha_0 = (\alpha_k(0))_k$.

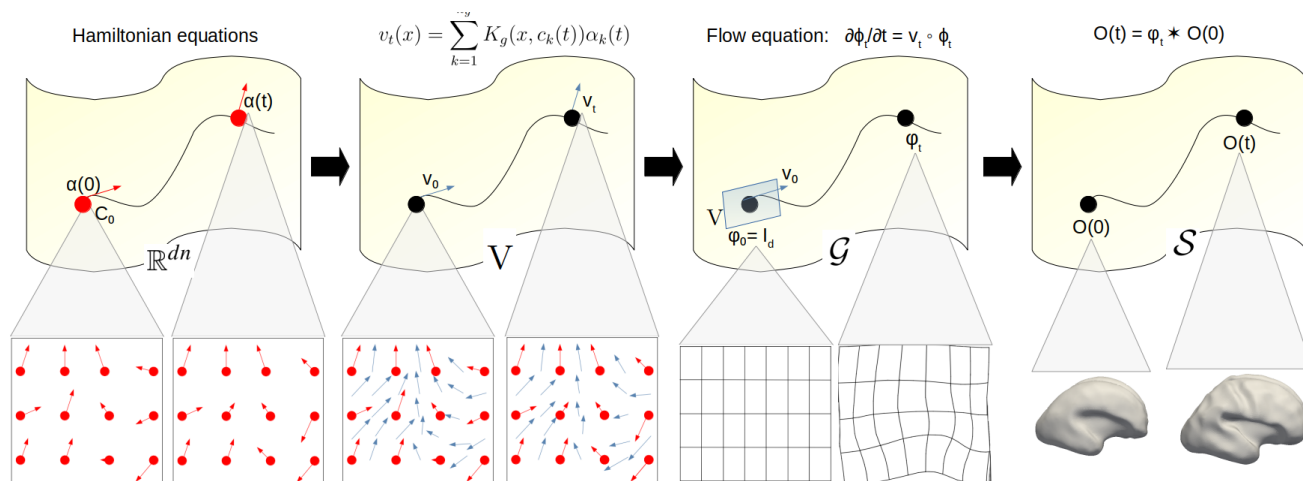


Figure 2.3: Computing shape deformations from an initial object $O(0)$ and a set of initial momentum vectors $\alpha(0)$. 1- Given $\alpha(0)$ and the initial control points c_0 , integration of the Hamiltonians (Equation (2.3)) gives the trajectory of the momentum vectors $\alpha(t)$. 2- The velocity field v_t is computed by interpolating the momentum vectors with Equation (2.2). 3- Integration of the flow equation (Equation (2.1)) gives a flow of diffeomorphisms $(\phi_t)_{t \in [0,1]}$. 4- Finally, ϕ_t is applied to the object $O(0)$, giving the deformed shape $O(t)$ at any time t .

2.3.3 Optimization

In Computational Anatomy, one seeks to match a reference object O_{ref} to one (or a series of) object(s) $(O_i)_{1 \leq i \leq N}$ observed at times $(t_{i,j})_{1 \leq i \leq N, 1 \leq j \leq M_i}$. In our particular setting, we perform *inexact* matching and seek to minimize a cost function expressing a trade-off between matching accuracy and regularity of the transformation(s). In very generic terms, this function can be written as:

$$E(O_{ref}, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N \left[\sum_{j=1}^{M_i} \frac{d(O_i(t_{i,j}), O_{ref} \circ \phi_i^{-1}(t_{i,j}))}{\sigma_\epsilon} + \text{Reg}(\phi_i) \right] \quad (2.4)$$

where:

- N stands for the number of subjects
- M_i denotes the number of observations from subject i
- ϕ_i is the i^{th} diffeomorphism matching O_{ref} to O_i and parameterized by the momentum vectors $\alpha_{0,i}$ and the fixed control points c_0 : $\phi_i(t_{i,j}) = \text{Exp}_{c_0,0,t_{i,j}}(\alpha_{0,i})$
- σ_ϵ is a parameter controlling the trade-off between the two terms
- d measures the distance between the objects to be matched and Reg is a regularity term that enforces spatial constraints on the deformations
- Reg is a regularity term that enforces spatial constraints on the deformations.

We choose $\text{Reg}(\phi_i)$ as the total kinetic energy along the geodesic path related to ϕ_i : $\text{Reg}(\phi_i) = \int_0^1 \|v_t\|_V^2 dt = \|v_{0,i}\|_V^2$. With the discrete parameterization chosen for the

velocity fields (i.e. Equation (2.2)), the norm $\|\cdot\|_V$ can easily be computed and thus

$$\text{Reg}(\phi_i) = \sum_{j=1}^{k_g} \sum_{k=1}^{k_g} \alpha_{j,i}(0) K_g(c_{j,i}, c_{k,i}) \alpha_{k,i}(0). \quad (2.5)$$

In this thesis, optimization is performed through gradient descent. Computing the gradient with respect to all the parameters is not a trivial task, but an efficient numerical scheme has been proposed in Durrleman et al. [54]. This algorithm relies heavily on the fact that the norm of the vector fields is the one of the RKHS and that the vector fields that are solution to the problem remain a *finite* sum of kernels at all times (Equation (2.2)). This enables us to solve a finite-dimensional problem even though the functions we are looking for have infinite dimension. Given a set of momentum vectors α_0 , this strong structural property allows to efficiently deform shapes and compute the cost function by performing the following steps, which are illustrated in Figure 2.3: integrating Equation (2.3) gives the evolution of the momenta over time, the velocity field v_t at any time t is computed with Equation (2.2), the flow of diffeomorphisms $(\phi_t)_{t \in [0,1]}$ is obtained by solving Equation (2.1), and the template object O_{ref} is deformed with the flow: $\phi_1 \star O_{ref}$. It is then straightforward to infer the distance between the deformed template and the target objects and to compute the regularity term with Equation (2.5). Finally, one can compute the gradient of the cost function with respect to the template, and use a backward integration along time to compute the gradient with respect to the momentum vectors.

As we shall see in Chapter 3, we are interested in the practical applications of this algorithm and any improvements to the deformation model will be made in the spirit of preserving this efficient numerical scheme.

In this thesis, shape deformations will be estimated using the open-source software Deformetrica [23].

2.3.4 Shape attachment




Figure 2.4: Varifold representation of a cortical surface mesh (zoomed). Normal vectors $(n_k)_k$ are indicated by red arrows located at the cell centers $(x_k)_k$.

To evaluate how close two shapes are and compute the data attachment term in the cost function E , a metric is needed. In this thesis, we will work with 2D and 3D images and surface meshes.

For images, we will use the the L^2 **distance**:

$$d(I_1, I_2)^2 = \sum_{x \in \Omega} (I_1(x) - I_2(x))^2,$$

where we consider images as functions defined on the domain Ω .

Similarly, for landmarks or meshes with point correspondence, the sum of squared differences is available. In many cases though, meshes have a different number of points. Different ways of computing distances have been proposed for this particular case. In this thesis, we will use the **varifold distance** [28].

Let W and its dual W' be Hilbert spaces. To a surface O , we associate the varifold $\{O\} \in W'$, which represent O as the distribution of its points with unit normal vectors attached. An example is available in Figure 2.4. The varifold distance between two meshes O_1 and O_2 is defined as

$$\begin{aligned} d_W(\{O_1\}, \{O_2\})^2 &= \|\{O_1\} - \{O_2\}\|_{W'}^2 \\ &= \|\{O_1\}\|_{W'}^2 + \|\{O_2\}\|_{W'}^2 - 2 \times \langle \{O_1\}, \{O_2\} \rangle_{W'} \\ &= \sum_{k=1}^{m_1} \sum_{l=1}^{m_1} K_w(x_k^1, x_l^1) \frac{(n_k^1 \cdot n_l^1)^2}{\|n_k^1\| \|n_l^1\|} + \sum_{k=1}^{m_2} \sum_{l=1}^{m_2} K_w(x_k^2, x_l^2) \frac{(n_k^2 \cdot n_l^2)^2}{\|n_k^2\| \|n_l^2\|} \\ &\quad - 2 \times \sum_{k=1}^{m_1} \sum_{l=1}^{m_2} K_w(x_k^1, x_l^2) \frac{(n_k^1 \cdot n_l^2)^2}{\|n_k^1\| \|n_l^2\|}. \end{aligned}$$

where K_w is a Gaussian kernel of width σ_w , m_1 and m_2 are the number of cells (triangles) in the meshes O_1 and O_2 (respectively), $(x_k^1)_k$ and $(x_k^2)_k$ the cell *centers* of O_1 and O_2 and $(n_k^1)_k$ and $(n_k^2)_k$ the cell *normals* of O_1 and O_2 .

The kernel width σ_w plays an important role, since differences between meshes at scales smaller than σ_w are smoothed. Note that the varifold distance is an alternative to the current distance [236], which takes into account the orientation of the normal vectors.

Regarding optimization, the gradient of the varifold distance with respect to the vertex positions is simple to compute, unlike other correspondence-free distances like the Hausdorff distance [101].

2.4 Computational Anatomy models

In this section, we introduce some common and less common Computational Anatomy models, all of which will be used in this thesis. A summary of these models is presented in Figure 2.5. They are available in the software Deformetrica [24]. These models, formulated here within a frequentist framework, optimize a cost function of the form of Equation (2.4), composed of a trade-off between a fidelity term quantifying how well the deformation(s) of a template object fit the observed object(s), plus a regularity term.




Figure 2.5: Representation of the cross-sectional atlas, registration, geodesic regression and 2-component piecewise geodesic regression on toy data. The parameters to optimize are the template-to-subject(s) deformations (orange arrows) and the template O_{ref} (images outlined in blue). For clarity, in atlas estimation and registration we dropped the subscript "1" of the diffeomorphism " ϕ ".

2.4.1 Cross-sectional atlas

The goal of atlas estimation is to perform population analysis on shapes, more specifically to extend the notions of *mean* and *variance* to shapes. We consider a set of N objects $(O_i)_{1 \leq i \leq N}$ of dimension d , all observed at the same time (there is one observation per subject, thus $\forall i M_i = 1$). We assume that each object O_i is a smooth deformation of a template O_{ref} representative of the average anatomy, plus an additive random white noise ϵ_i [6]:

$$O_i = O_{ref} \circ \phi_{1,i}^{-1} + \epsilon_i, \quad \forall i \in [1, n] \quad (2.6)$$

where $\phi_{1,i}$ is the i^{th} template-to-subject deformation, and $O_{ref} \circ \phi_{1,i}^{-1}$ denotes the action of the diffeomorphic deformation on the template.

The goal of atlas estimation is to estimate the template object O_{ref} and the N template-to-subject deformations $(\phi_i)_{1 \leq i \leq N}$. To do so, we seek to minimize a cost function E :

$$E(O_{ref}, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N \left(\frac{d(O_i, O_{ref} \circ \phi_{1,i}^{-1})^2}{\sigma_\epsilon^2} + \|v_{0,i}\|_V^2 \right), \quad (2.7)$$

where $\phi_{1,i} = \text{Exp}_{c_0}(\alpha_{0,i})$.

The cross-sectional atlas builds a *coordinate system* where the template-to-subject deformations $(\alpha_{0,i})_{1 \leq i \leq N}$ position each subject with respect to the reference shape O_{ref} . Computing statistics on high-dimensional objects, or on the diffeomorphisms that characterize them, is not straightforward, but the coordinates $(\alpha_{0,i})_{1 \leq i \leq N}$ belong to the *tangent space* at the template shape: this allows us to apply linear statistical tools such as principal component analysis to the dataset [175, 174, 236]

Methodological improvements will be made to the estimation of the cross-sectional atlas in Chapter 3.

2.4.2 Registration

Registration estimates the optimal transformation between two objects O_{ref} and O_1 . It can be seen as a particular case where the template object is fixed and $N = 1$. The cost function writes as follows:

$$E(\alpha_0) = \frac{d(O_1, O_{ref} \circ \phi_1^{-1})^2}{\sigma_\epsilon^2} + \|v_0\|_V^2 \quad (2.8)$$

where $\phi_1 = \mathcal{Exp}_{c_0}(\alpha_0)$.

2.4.3 Geodesic regression

Geodesic regression can be seen as the generalization of linear regression to shapes. Given M observations $(O_1(t_1), \dots, O_1(t_M))$ from a single subject ($N = 1$), one seeks to estimate a template object O_{ref} at age t_0 (analogous to an intercept parameter), along with the mean geodesic trajectory $\phi(t)$ that deforms O_{ref} [65].

The flow of diffeomorphisms applied to the template shape $t \rightarrow O_{ref} \circ \phi(t)^{-1}$ is the one that best depicts how the observations evolve with time (i.e. that is as close as possible to the input shapes). Geodesic regression optimizes the following cost function:

$$E(O_{ref}, \alpha_0) = \sum_{j=1}^M \left(\frac{d(O_1(t_j), O_{ref} \circ \phi(t_j)^{-1})^2}{\sigma_\epsilon^2} \right) + \|v_0\|_V^2. \quad (2.9)$$

where $\phi(t_j) = \mathcal{Exp}_{c_0, t_0, t_j}(\alpha_0)$.

c_0 and α_0 are analogous to an intercept and slope parameter, respectively.

In practice, in this thesis we will also apply geodesic regression to observations coming from different subjects, with a single observation per subject. Compared to the cross-sectional atlas model, this can be seen as a replacement of a single reference shape by a curve, i.e. the trajectory of the reference shape.

2.4.4 Piecewise geodesic regression

Piecewise geodesic regression [30, 44] is an extension of geodesic regression which allows changes in trajectory dynamics. The mean trajectory $\phi(t)$ is decomposed into P trajectories $(\phi^l(t))_{1 \leq l \leq P}$ called *geodesic components* and associated to a set of rupture times $t_R = [t_1, t_2, \dots, t_{P-1}]$, with $t_1 < \dots < t_{P-1}$. The template shape O_{ref} , either optimized or fixed, is defined at time $t_0 \in t_R$. O_{ref} , t_R and t_0 can be optimized, but the number of components P has to be fixed for the problem not to be ill-posed.

In each interval $[t_l, t_{l+1}]$ of t_R , the average trajectory $\gamma(t)$ is the l^{th} geodesic compo-

nent:

$$\gamma(t)(O_{ref}^1) = \phi^1(t) \star O_{ref}^1 \mathbb{1}_{[-\infty, t_1]}(t) + \sum_{l=2}^{P-1} \phi^l(t) \star O_{ref}^l \mathbb{1}_{[t_{l-1}, t_l]}(t) + \phi^P(t) \star O_{ref}^P \mathbb{1}_{[t_P, +\infty]}(t) \quad (2.10)$$

where :

- the l^{th} diffeomorphism ϕ^l is parameterized by the set of momentum vectors $\alpha_{0,l}$ and the fixed control points c_0 :

$$\forall l \in [1, P] \quad \phi^l(t) = \mathcal{Exp}_{c_0, t_l, t}(\alpha_{0,l})$$

- $\forall l \in [1, P]$, O_{ref}^l is the representative shape at the rupture time t_l :

$$O_{ref}^l = \phi^{l-1}(t_l) \star O_{ref}^{l-1}$$

To ensure the continuity of the trajectory, we impose that $\forall l \in [1, P - 1]$, $\phi^l(t_l) = \phi^{l+1}(t_l)$.

In this thesis, geodesic regression and its variants may sometimes be referred to as *temporal atlasing*. This term is not to be confused with *longitudinal atlasing* [201, 200], a more complex model where one estimates individual geodesic trajectories that are deformations of an average trajectory. This latter model requires longitudinal data, i.e. repeated observations of the same subjects, which translates as $N > 1$ and $M_i > 1$ for at least one subject i . As specified in Chapter 1, in this thesis we will work only with datasets comprising one observation per subject.

Methodological improvements will be made to geodesic regression and its variants in Chapter 4.

MULTISCALE DEFORMATIONS FOR ATLAS ESTIMATION

Contents

3.1	Motivation	22
3.2	Related work	24
3.2.1	Multiscale representations of deformations	25
3.2.2	Multiscale strategies in the LDDMM framework	25
3.2.3	Multiscale strategies and sparse parameterization of diffeomorphisms	27
3.3	Wavelet-based multiscale atlas estimation	28
3.3.1	Overview	28
3.3.2	Model of diffeomorphic deformations	29
3.3.3	Reparameterization of the initial velocity fields	30
3.3.4	Multiscale optimization algorithm	36
3.3.5	Multiscale optimization algorithm with local adaptation	39
3.3.6	Alternative multiscale algorithm	40
3.4	Experiments: cross-sectional atlas estimation	40
3.4.1	Handwritten digits	41
3.4.2	Toy data	42
3.4.3	Fetal brain images	44
3.5	Additional experiments	46
3.5.1	Local adaptation extension	46
3.5.2	Alternative multiscale scheme	46
3.6	Discussion	48

This chapter is an extended version of a manuscript submitted for publication [69]: *Wavelet-Based Multiscale Initial Flow For Improved Atlas Estimation in the Large Diffeomorphic Deformation Model Framework* by Fleur Gaudfernau and Eleonore Blondiaux, Stéphanie Allassonnière and Erwan Le Pennec.

3.1 Motivation

When working with collections of images that are instances of the same anatomical object, one faces the question of how to faithfully model the mean and variability over

these structures. Atlas estimation is a method to achieve such modeling: an estimate of the average shape is given in the form of a template image, which represents the invariants across the population, i.e. shared anatomical features, and the variability is given by deformations from the template space to each subject’s space, which express how these common features vary within the population [84].

Atlas estimation has many applications in the field of medical image analysis. The template image can be used as reference to describe average anatomical structures or serve as a tool to automatically segment new subjects. Variations around the template may be used to characterize pathological deviations from normality [68] or to isolate subgroups in the population [44]. Atlases can also be defined in a spatio-temporal fashion to characterize normal or pathological changes, such as brain growth across gestation [73] (this topic will be addressed in Chapter 7).

As in registration problems, the choice of the deformation function describing the template-to-subject transformations is of primary importance [164]. To account for the intra- and inter-subject anatomical variability in clinical images, non-linear deformations are mandatory. Here, we work with diffeomorphic transformations, which are high-dimensional, smooth and invertible functions with smooth inverse that preserve the topology of anatomical images.

In this thesis, we work within the framework of the Large Deformation Diffeomorphic Metric Mapping (LDDMM) [233, 153, 31], in which objects are deformed through diffeomorphic transformations of the whole ambient space. The group of possible transformations forms a Riemannian manifold of infinite dimension and parameterizes a flexible representation of deformations. Since such models have a very high number of degrees of freedom and parameters, the search space is narrowed to a small subset of deformations. As described in Section 2.3.3, this is typically done by optimizing a cost function that comprises a distance between the deformed template and each subject, plus an energy term acting as spatial regularizer [54]. The latter constrains diffeomorphisms to be geodesics on the Riemannian manifold, i.e. the shortest paths between the identity map and the diffeomorphism of interest according to a regularizing metric.

The choice of this regularizer is critical as it restricts the range of transformations defined by the model [215]. Specifically, it constrains the estimated deformations to a single scale. A large kernel is likely to produce smooth deformations but less accurate matches, while a fine kernel will generate more accurate yet unnatural deformations. This dilemma is illustrated in Figure 3.1 on a registration example. As clinical images often present high variability at several scales, one might be tempted to increase the number of parameters in the model, i.e. to use many control points and a small kernel. However, fine kernels make large displacements more expensive than small ones, and such over-parameterization will likely trap the optimization procedure in a local minimum, achieving a reasonable numerical solution that is qualitatively bad, especially in terms of template realism. Finding an ideal trade-off is time-consuming and sometimes impossible, especially when working with complex images whose variability cannot be summarized by single-scale deformations. To overcome such problem, hierarchical algorithms have been widely used in the field of image registration [164, 155]: they first solve the registration problem at coarse scales and transfer the solution to increasingly

fine scales to refine the transformation. These strategies avoid more efficiently trapping the algorithm in local minima related to unrealistic transformations. Except for Tan et al. [223], they have yet to be applied in the LDDMM framework.

In this chapter, we will present a new multiscale strategy to avoid unrealistic local minima in the LDDMM framework.

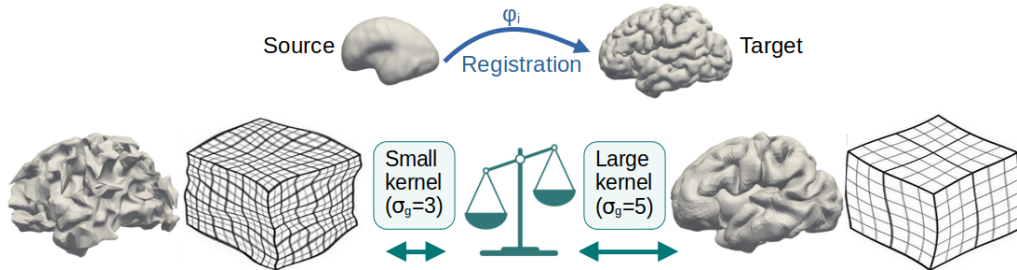


Figure 3.1: Illustration of the dilemma linked to the choice of the deformation kernel on a simple registration example between two cortical surface meshes from fetuses. The source and target structures are 24 and 35 gestational weeks, respectively. The bottom row shows the deformed source object and the corresponding deformation grid obtained with two different kernel widths.

3.2 Related work

Even though multiscale image registration has been studied repeatedly in the literature, it has rarely been extended to population analysis. As registration is a special case of atlas estimation with a fixed template image, in the following we will review both registration and atlas estimation methods that have a multiscale property.

Multiscale strategies differ in the way they handle the various scales of the optimization: some methods favor a coarse-to-fine fashion while others handle coarse and fine scales simultaneously. We will first review coarse-to-fine strategies with a focus on those relying on a wavelet-based decomposition. Then, we will survey the multiscale strategies that have been introduced in the LDDMM framework.

Coarse-to-fine or hierarchical optimization strategies seek to solve the registration problem at progressively increasing resolutions, with the objective of decreasing the computational cost and finding a more accurate solution [155]. The search area is first restricted to coarse functions and the results are progressively refined, with the parameters estimated at the previous coarser level propagated to the next finer level. Such strategies can be coarse-to-fine with regard to the deformation field, and/or with regard to the registered images. We will not dwell on the latter case, which is out of the scope of this chapter.

3.2.1 Multiscale representations of deformations

Signal representation provides numerous mathematical tools to represent signals, and many registration algorithms modelled deformations using basis functions from the Fourier and Wavelet transforms. Ashburner et al. [10] used the Fourier transformation to compress the representation of the deformation fields and accelerate calculations. In the same vein, Zhang et al. [263] introduced a Fourier representation of the velocity fields in the LDDMM framework. Retaining only low frequency coefficients yielded a low dimensional representation of the deformations and a significant reduction of the computational cost. While they did not implement a multiscale optimization scheme, the idea of reparameterizing the velocity fields with different basis functions is akin to ours. Christensen et al. [32] parameterized displacement fields by a Fourier basis and estimated frequency coefficients in a coarse-to-fine fashion.

However, experiments showed that modelling deformations using a wavelet basis provides better spatial regularization compared to using a Fourier basis [7]. Indeed, as the wavelet transform gives the location and orientation of the signal frequencies, it provides a convenient basis for hierarchical optimization strategies. Wavelet-based deformation models cover a variety of wavelet types (e.g. Haar, Cai Wang and (BV, L^2)) and deformations (e.g. displacement vectors, B-splines and elastic deformations). Displacement vectors [251, 27] and later free form deformations [219] were described and optimized in a multi-resolution fashion through the Cai Wang wavelet. Gefen et al. [70] modeled elastic deformations using finite-supported, semi-orthogonal wavelet functions. Topology-preserving displacement fields were modelled by polynomial spline basis functions and controlling the Jacobian of the transformation [158, 161]. Recently, wavelets (BV, L^2) were employed to generate a hierarchical representation of diffeomorphisms in the hyperelasticity framework and perform coarse-to-fine optimization [46]. Of note, the wavelet transform has also been used in deformation analysis with various aims such as data visualization and compression, again encompassing various types of wavelet bases such as the Cauchy-Navier wavelet [138], the non-linear Morphlet wavelet [113] and the spline wavelet [128].

3.2.2 Multiscale strategies in the LDDMM framework

In the LDDMM framework, the choice of the spatial regularizer restricts the range of possible deformations to those occurring at a single scale, which often proves unrealistic [190]. Thus, a variety of papers have focused on increasing the flexibility of the deformation model. Two main strategies can be identified:

- Coexisting flows of different scales are estimated simultaneously [190, 21, 215, 85, 223]
- Flows of different scales are composed and estimated sequentially in a coarse-to-fine manner [156, 154].

In the category of multiscale coexisting flows, Risser et al. [190, 21] first introduced a multi-kernel extension of the LDDMM framework by writing the deformation flow as

a weighted sum of Gaussian kernels whose widths are specified by the user. Weights are tuned in a semi-automatic manner during a pre-registration step. In this framework, the RKHS structure of the velocity fields is lost, and a new definition of the norm is used to ensure efficient computation of the flow. A spatially-varying version of this framework, the kernel bundle [215], used sparsity priors to allow the weights of the kernel mixture to vary across spatial locations. Even though this algorithm proved efficient on the registration of landmark points, the increase in computational cost restricts its application to registration problems involving few parameters. Using an algorithm for the multiresolution decomposition of surfaces, the kernel bundle framework was also combined with a coarse-to-fine strategy wherein the resolution of cortical surface meshes is progressively increased along with that of the deformation field [223]. Unfortunately, this strategy can only handle meshes. Multi-kernel approaches were further combined with deep learning optimization in order to learn a local regularizer from the data [160, 211]. Set in the framework of time-independent velocity fields, Niethammer et al. [160] learned the pre-weights of the mixture of Gaussian kernels through a Convolutional Neural Network and jointly optimized the deformation and network parameters by stochastic gradient descent. Shen et al. [211] conducted a similar work with both spatially- and time-varying velocity fields. These methods increase significantly the complexity of the mathematical model, and several optimization procedures are required to tune the networks parameters, the kernel pre-weights and the deformation parameters. As in the kernel bundle framework, the tuning of additional parameters is heavy on the computational cost, which dampens the application of such algorithms to high-dimensional images such as volumetric MRIs.

Of note, a related approach, based on modular deformations, enables the user to impose spatially-varying constraints on the deformation field [85]. Large diffeomorphic deformations are built by superimposing deformations modules, which encode local geometrical transformations, making it possible to construct diffeomorphisms from multiple scale flows. As in the kernel bundle framework, the space of vector fields is equipped with an adapted norm. The need for prior knowledge about the deformation modules limits the practical application of the algorithm.

A second and less explored axis of research constructed a hierarchical representation of deformations, based on non-coexisting vector flows of increasing resolution, which are estimated independently and then composed. In a theoretical paper, Modin et al. [156] extended (BV, L^2) wavelets to express diffeomorphisms as a composition of deformations of increasingly fine scales, which can be seen as a series of LDDMM steps. Despite the potential of this approach, the authors did not perform numerical experiments. A similar approach [154] constructed diffeomorphisms by composing a series of multiscale vector fields, which enables to progressively refine the deformation. Contrary to multi-kernel approaches, such strategies perform optimization in successive RKHS of increasingly finer resolution, in the spirit of coarse-to-fine strategies. Finally, a sequential multiscale approach was developed to perform registration and template estimation of ear shapes [265]. Based on the current representation of shapes [236], the method reduces iteratively the size of the kernels controlling the scales of the shape matching term and of the deformation field. Running the algorithm for n scales amounts to running n single-scale optimizations, which in practice is prohibitive - at least when performing

template estimation on clinical images. However, it should be noted that, unlike the previous methods, this approach relies on an "optimization trick" rather than a complexification of the deformation model.

As we shall see in the following, our coarse-to-fine approach is more closely related to the one that composes multiple scale flows, in the sense that we perform optimization sequentially in sub-spaces of increasing resolution. However, our algorithm differs from the previous ones by the fact that our work brings changes to the optimization procedure rather than the deformation model: the multiscale structure is only used for the initial velocity field and so that the velocity fields are still defined, at core, by a single-scale RKHS. This simplifies the implementation of our algorithm while preserving the efficient optimization scheme of [54]. It should also be noted that the above-mentioned algorithms perform registration tasks, while we are interested in improving optimization for several models: registration, atlas estimation and geodesic regression.

3.2.3 Multiscale strategies and sparse parameterization of diffeomorphisms

In this thesis, we compute object deformations using the discrete parameterization of diffeomorphisms introduced by Durrleman et al. [54] (see Section 2.3.2). We have introduced in Section 3.1 the general difficulty of choosing a single appropriate kernel size when matching objects. Here, we will dwell more on how this issue presents itself in our chosen framework.

We recall that velocity fields defining geodesic transformations write as the convolution of momentum vectors located on a grid of control points, where the convolution operator is a kernel K_g . There is a dependency between the scale of the kernel σ_g and the number of parameters, as a constant vector field has to be well approximated by the finite sum: thus, the spacing between control points has to remain close to the kernel size. This is why we set the controls points on a regular grid of spacing σ_g . Here, the kernel K_g not only controls the amount regularity of the deformations, but also the number of parameters to optimize: large kernels imply few control points and therefore smoother vector fields, whereas small σ_g foster finer transformations and more accurate matches.

Further, as the optimization problem is not convex, the gradient descent algorithm converges towards a solution that depends on the initialization. In the case of atlas estimation, the model is initialized with N template-to-subject null momenta and a template object provided by the user. The more the numbers of subjects and parameters are high, the more complex the energy landscape of the problem becomes, and the more difficult it will be to find a solution distant from the initialization.

As observed in Durrleman et al. [54], the algorithm may converge towards unrealistic local minima when the number of control points and thus of parameters is too large. In this original approach, a first step towards the estimation of multiscale deformations was taken: both the *number* and *position* of the control points were numerically optimized. A L1 sparsity prior was applied to the momenta in order to inactivate points in areas with

low variability. Unfortunately, the width of the kernel remains the same and in practice, the vector fields can only be set to zero in image areas of null intensity. Another issue is that the position of the control points cannot change significantly if the distance between control points is low. Yet, it is exactly when the number of parameters is high that the risk of converging towards an unrealistic local minima is higher.

One idea would be to change locally the scale of the kernel so that we can estimate non-evenly smooth vector fields, i.e. vector fields with spatially varying scales. Even if such parameterization can be written, one loses the RKHS structure and thereby the ease of computation.

In this chapter, we wish to address these two related issues: the dependency of the algorithm on the initialization, which restrains the number of parameters that can be properly optimized, and the difficulty in estimating vector fields with spatially-varying regularity. In the next section, we will describe a reparameterization of the vector fields which enables us to impose smoothness constraints on the deformations and progressively relax them in a coarse-to-fine fashion. In this way, the algorithm can cope with non-evenly smooth transformations while using a small kernel and remaining in the original RKHS setting.

3.3 Wavelet-based multiscale atlas estimation

In this chapter, we propose a multiscale atlas estimation procedure based on a Haar-like wavelet representation of the initial velocity fields. This strategy has the advantage of making the algorithm less dependent on the initialization while favoring more multi-scale, hence natural deformations. We rely on the finite parameterization of the velocity fields as a linear combination of RKHS basis elements [54] within the LDDMM setting. Importantly, our strategy enables us to preserve this structural assumption and the efficient numerical scheme that follows. We will show that our algorithm generates more natural template images as well as higher stability regarding the initialization.

3.3.1 Overview

For the sake of clarity, we provide the reader with an overview of our multiscale strategy. In the *original* algorithm for atlas estimation, we optimize a cost function E (see Section 3.3.2). Two types of parameters are optimized through gradient descent: the template image O_{ref} and N sets of momentum vectors $\alpha_{0,i}$ that parameterize the template-to-subjects velocity fields $v_{0,i}$. The *original* optimization iterates between two classical steps (note that the subscripts i have been dropped for clarity):

1. Computation of the gradients ΔE_{α_0} and $E_{O_{ref}}$
2. Parameters update: $\alpha_0(j+1) \leftarrow \alpha_0(j) - h \times \Delta E_{\alpha_0}$; ditto for the template O_{ref} .

In the *multiscale* strategy, summarized in Figure 3.2, we modify the optimization of the momentum vectors by replacing step 2. with the following steps:




Figure 3.2: Overview of the coarse-to-fine optimization strategy. For the sake of clarity, a **single** gradient descent iteration is represented during the coarser steps of the algorithm. $\alpha_0(j)$: momentum vectors defined in the RKHS basis at iteration j . ΔE_{α_0} : gradient of the cost function with regard to α_0 . $\beta_0(j)$: multiscale representation of $\alpha_0(j)$ in the wavelet basis. ΔE_{β_0} : gradient of the cost function with regard to β_0 .

- (i) We use the wavelet transform (Section 3.3.3) to obtain a multiscale representation of the gradient of E with regard to the momenta: $\Delta E_{\beta_0} \leftarrow FWT(\Delta E_{\alpha_0})$
- (ii) The wavelet coefficients in ΔE_{β_0} whose scale is smaller than a current scale S_j are set to zero (Section 3.3.4)
- (iii) The coordinates of the initial velocity fields v_0 in the wavelet basis are updated: $\beta_0(j+1) \leftarrow \beta_0(j) - h \times \Delta E_{\beta_0}$
- (iv) The coordinates of v_0 in the RKHS basis are recovered with the Inverse Wavelet Transform: $\alpha_0(j+1) \leftarrow IWT(\beta_0(j+1))$.

3.3.2 Model of diffeomorphic deformations

In this section, we briefly recall how we model and estimate diffeomorphic deformations to perform atlas estimation. We consider a set of N objects $(O_i)_{1 \leq i \leq N}$ of dimension d , and we assume that each object O_i is a smooth deformation of a template O_{ref} , plus an additive random white noise ϵ_i : $O_i = O_{ref} \circ \phi_i^{-1} + \epsilon_i$, $\forall i \in [1, n]$, where ϕ_i is the i^{th} template-to-subject deformation, and $O_{ref} \circ \phi_i^{-1}$ denotes the action of the diffeomorphic deformation on the template.

We work in the LDDMM framework, where diffeomorphisms are constructed by following the streamlines of time-dependent velocity fields v_t (see Equation (2.1)). We further assume that the velocity fields write as the Gaussian convolution of momentum

vectors $(\alpha_k(0))_k$ located at control points $(c_k)_k$.

$$v_0(x) = \sum_{k=1}^{k_g} K_g(x, c_k) \alpha_k(0),$$

where K_g is a Gaussian kernel of width σ_g and the position of the control points is fixed on a regular grid of spacing σ_g .

Vectors fields that define geodesic deformations with respect to the norm $\int_0^1 \|v_t\|_V^2 dt$ keep the same structure along time and so:

$$v_t(x) = \sum_{k=1}^{k_g} K_g(x, c_k) \alpha_k(t), \quad (3.1)$$

where $\alpha_k(t)$ is the k^{th} momentum vector attached to the control point $c_k(t)$ at time t .

Furthermore, the trajectory of the control points and momentum vectors is described by the Hamiltonian system equations (Equation (2.3)).

The kinetic energy along geodesic paths is preserved over time, i.e. $\forall t \in [0, 1], \|v_t\|_V = \|v_0\|_V$. This implies that a geodesic transformation is fully parameterized by the initial velocity field v_0 and thus by its initial momentum vectors $(\alpha_k(0))_k$.

In the specific case of atlas estimation, we seek to minimize a cost function E :

$$E(O_{ref}, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N \left(\frac{d(O_i, O_{ref} \circ \phi_{1,i}^{-1})^2}{\sigma_\epsilon^2} + \|v_{0,i}\|_V^2 \right), \quad (3.2)$$

where O_{ref} is the template object to be estimated, $(\alpha_{0,i})_{1 \leq i \leq N}$ are the N sets of momentum vectors parameterizing the template-to-subject deformations $(\phi_i)_{1 \leq i \leq N}$, $\|v_{0,i}\|_V^2$ is the norm of the i^{th} velocity field related to ϕ_i , and σ_ϵ is a parameter controlling the trade-off between data attachment and regularity.

We recall that optimization is performed through gradient descent using a numerical scheme which relies heavily on the fact that the vectors fields v_t are expressed at any time as a finite sum of RKHS basis elements (Equation (3.1)). More details about the optimization procedure are available in Section 2.3.3.

3.3.3 Reparameterization of the initial velocity fields

Since our algorithm relies on the multiscale structure of the wavelet transform, we introduce the definition and properties of the continuous Haar Wavelet representation. We recall how this construction can be extended to a representation of discrete signals defined on a grid, and we demonstrate how this can be used to obtain a Haar-like representation of the initial velocity fields.

3.3.3.1 The continuous Haar wavelet transform

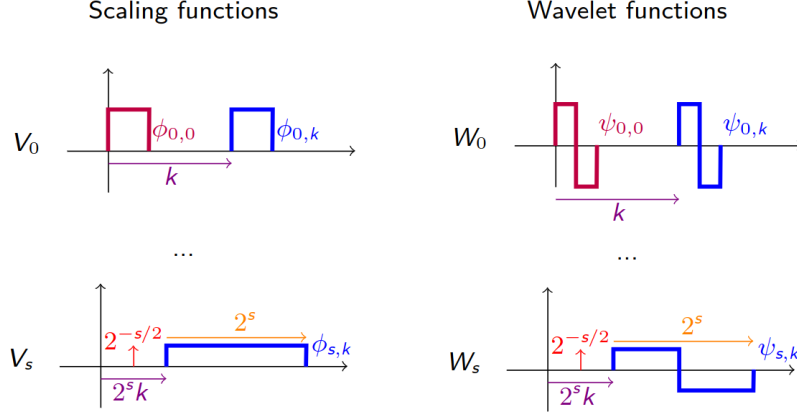


Figure 3.3: The one-dimensional Haar wavelet. At scale 0, V_0 is the space of piecewise constant functions of size 1. A basis for V_0 , i.e. $\{\Phi_{0,k}\}_k$, is obtained by translating the scaling function $\Phi_{0,0} = \psi^L$ by factors $k \in \mathbb{Z}$. A basis for W_0 , i.e. $\{\psi_{0,k}\}_k$, is obtained by translating the mother wavelet function $\psi_{0,0} = \text{psi}^H$ by factors $k \in \mathbb{Z}$. At scale s , a basis for the space V_s , i.e. $\{\Phi_{s,k}\}_k$, is obtained by dilating $\Phi_{0,0}$ by 2^s , translating it by $2^s k$ and normalizing it by $2^{-s/2}$. A basis for W_s , i.e. $\{\psi_{s,k}\}_k$, is obtained by performing the same operations on $\psi_{0,0}$.

Here, we describe the decomposition of a real signal f defined on the space \mathbb{R}^d into a Haar Wavelet basis [142, 141]. The wavelet representation decomposes f into a linear combination of basis functions which have different resolutions, locations and orientations. This representation relies on a collection of embedded spaces V_s that contain functions said of scale s . In the case of the Haar Wavelet, V_s is the space of piecewise constant functions on a regular grid of size 2^s (see an illustration for $d = 1$ in Figure 3.3). Any function f can be approximated in this space by computing a local average: the mean value in each sub-square of the grid. If we define the one-dimensional piecewise constant function

$$\psi^L(z) = \begin{cases} 1 & \text{for } 0 \leq z < 1 \\ 0 & \text{otherwise} \end{cases},$$

and the d dimensional scaling function Φ by

$$\Phi(x) = \prod_{i=1}^d \psi^L(x_i)$$

with $x \in \mathbb{R}^d$, approximating f at scale s amounts to projecting f onto the space spanned by the orthonormal family

$$\{\Phi_{s,k}(x) = 2^{-sd/2} \Phi(2^{-s}x - k)\}_{k \in \mathbb{Z}}$$

where $\Phi_{s,k}$ is the scaling function rescaled by 2^s , translated by $k2^s$ and normalized by $2^{-sd/2}$.

When transitioning from the approximation at scale s to the approximation at the coarser scale $s + 1$, some details of f are lost. These details belong to the orthogonal complement W_{s+1} of the space V_{s+1} in V_s . A basis of this space can be obtained by defining the piecewise constant function

$$\psi^H(z) = \begin{cases} 1 & \text{for } 0 \leq z < 0.5 \\ -1 & \text{for } 0.5 \leq z < 1 \\ 0 & \text{otherwise} \end{cases}$$

and the d dimensional oriented wavelet functions

$$\psi^o(x) = \prod_{i=1}^d \psi^{o_i}(x_i)$$

where $x \in \mathbb{R}^d$, $o \in \{H, L\}^d$ and $\exists i, o_i = H$. Indeed, one can verify that an orthonormal basis of W_s is given by

$$\{\psi_{s,k}^o(x) = 2^{-sd/2} \psi^o(2^{-s}x - k)\}_{k \in \mathbb{Z}, \exists i, o_i = H}.$$

where $\psi_{s,k}^o$ is the wavelet function of orientation o rescaled by 2^s and then translated by $k2^s$. Note that functions ψ^L and ψ^H act respectively as low and high pass filters. Their combination yields oriented high pass filters, e.g. for $d = 2$, there are three wavelet functions ψ^{HL} , ψ^{LH} and ψ^{HH} that express details of the signal along vertical, horizontal and diagonal orientations (respectively).

We can thus decompose any function f in V_s in the two following ways:

$$\begin{aligned} f &= \sum_k a_{s,k} \Phi_{s,k} \\ &= \sum_k a_{s+1,k} \Phi_{s+1,k} + \sum_{o,k} d_{s+1,k}^o \psi_{s+1,k}^o. \end{aligned}$$

As $\sum_k a_{s+1,k} \Phi_{s+1,k}$ belongs by construction to the space V_{s+1} , we can further decompose it into a projection onto V_{s+2} and a projection onto W_{s+2} . Repeating this scheme up to scale S leads to the following multiscale decomposition of f :

$$f = \sum_k a_{S,k} \Phi_{S,k} + \sum_{s'=s+1}^S \sum_{o,k} d_{s',k}^o \psi_{s',k}^o.$$

The classical wavelet construction is concluded by letting s go to $-\infty$, enabling one to decompose any measurable bounded function in such bases. Finally, we observe in particular that any function $\Phi_{s',k}$ or $\psi_{s',k}^o$ with $s' \geq s$ is in V_s and therefore can be

(a) Classical decomposition into unit functions

(b) Multiscale Haar wavelet basis

Figure 3.4: Decomposition of a 4-by-4 grid in two bases. The letters a and d refer to approximation and detail coefficients, respectively. Subscripts indicate the *scale* of the coefficient and the x - y *position* of the related wavelet function. Superscripts indicate the *orientation* of the wavelet function. In the grids, empty cells denote null values.

decomposed as a linear combination of $\Phi_{s,k'}$:

$$\Phi_{s',k} = \sum_{k'} \gamma_{s',k,s,k'} \Phi_{s,k'} \quad (3.3)$$

$$\psi_{s',k}^o = \sum_{k'} \gamma_{s',k,s,k'}^o \Phi_{s,k'} \quad (3.4)$$

where $\gamma_{s',k,s,k'}$ and $\gamma_{s',k,s,k'}^o$ are some fixed real numbers.

More importantly, going from the decomposition in V_s to the decomposition in the spaces V_S and $(W_{s'})_{s < s' \leq S}$ corresponds to a change of basis and thus to a discrete operation going from the coefficients $(a_{s,k})_k$ to the coefficients $(a_{S,k})_k$ and $(d_{s',k}^o)_{s < s' \leq S, k, o}$. This transformation is called the Forward wavelet transform (FWT) and its inverse the Inverse wavelet transform (IWT). Both can be computed directly on the coefficients without relying on the continuous basis functions.

3.3.3.2 Haar wavelet applied to grids

In our algorithm, rather than using the continuous Haar wavelet decomposition, we will employ the related *discrete* Haar wavelet decomposition on a d dimensional grid $\llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket$. This transform corresponds to a discrete change of basis related to the continuous Haar transform.

More precisely, to any discrete function $(a_{0,k})_k$ on the grid, one can associate the continuous function f of V_0 defined by

$$f = \sum_{k \in \llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket} a_{0,k} \Phi_{0,k},$$

where the $(a_{0,k})_k$ are interpreted as the approximation coefficients of f at scale 0.

By construction, f is a piecewise constant function on the related continuous grid. Using the FWT algorithm up to scale S , this function can be decomposed as follows:

$$f = \sum_k a_{S,k} \Phi_{S,k} + \sum_{s'=s+1}^S \sum_{o,k} d_{s',k}^o \psi_{s',k}^o.$$

One verifies that $a_{S,k} = 0$ when $k_i < 0$ or $k_i 2^S > K_i$ and $d_{s',k}^o = 0$ when $k_i < 0$ or $k_i 2^{s'} > K_i$. Thus, these sums have a finite number of coefficients. Further, when the K_i are powers of 2, i.e. $K_i = 2^{\zeta_i}$, we impose that the decomposition cannot exceed a maximum scale $S_{max} = \min_i(\zeta_i)$, so that the scaling function support remains within $\prod_i \llbracket 0, K_i \rrbracket$. Therefore, when $S \leq S_{max}$, the previous equation reduces to:

$$f = \sum_{0 \leq k_i < 2^{\zeta_i - S}} a_{S,k} \Phi_{S,k} + \sum_{s'=s+1}^S \sum_{o, 0 \leq k_i < 2^{\zeta_i - s'}} d_{s',k}^o \psi_{s',k}^o,$$

which corresponds exactly to an orthonormal change of basis.

To summarize, from the coefficients $(a_{0,k})_k$, one can compute the coefficients $(a_{S,k})_{0 \leq k < 2^{S_{max} - S}}$ and $(d_{s',k}^o)_{s+1 \leq s' \leq S, 0 \leq k < 2^{S_{max} - s'}, o}$ with the FWT algorithm and perform the inverse operation with the IWT algorithm.

3.3.3.3 Algorithmic implementation of the Wavelet Transforms

We denote by M_{FWT}^1 and M_{IWT}^1 the $\prod_i K_i \times \prod_i K_i$ matrices associated to the transformations FWT and IWT so that:

$$\begin{aligned} \text{FWT} : \mathcal{M}_{K_1, \dots, K_d} &\rightarrow \mathcal{M}_{K_1, \dots, K_d} \\ X &\rightarrow M_{FWT}^1 X_{flat} \quad \text{with } X_{flat} \in \mathcal{M}_{\prod_i K_i} \end{aligned}$$

$$\begin{aligned} \text{IWT} : \mathcal{M}_{K_1, \dots, K_d} &\rightarrow \mathcal{M}_{K_1, \dots, K_d} \\ Y &\rightarrow M_{IWT}^1 Y_{flat} \quad \text{with } Y_{flat} \in \mathcal{M}_{\prod_i K_i} \end{aligned}$$

and $M_{FWT}^1 M_{IWT}^1 = Id$.

We have $M_{FWT}^1 = R \odot M_{FWT}$ and $M_{IWT}^1 = M_{IWT} \odot R^{-1}$, with \odot the element-wise matrix product operator. M_{FWT} and M_{IWT} are the unnormalized transformation matrix. $R \in \mathcal{M}_{\prod_i K_i}$ is the renormalization matrix that ensures preservation of the signal's energy when computing the Forward or Inverse Wavelet Transform: $R[i] = \frac{1}{\|M_{FWT}[i, :]\|_2}$

where $\|M_{FWT}[i, :]\|_2$ is the L_2 norm of the i_{th} row of M_{FWT} .

Note that this renormalization ensures the orthogonality of the transformations so that they verify $M_{IWT}^{1T} = M_{FWT}^1$.

Furthermore, there exist fast implementations of these transforms that have a linear complexity with respect to the number of coefficients [141]. In this work, we have implemented the FWT and IWT algorithms using the fast lifting scheme described in Mallat et al. [142] (Chapter 7.8). This scheme is strictly equivalent to the previous description when the grid size is a power of 2, with the difference that it handles non-dyadic grids through improved computations at the boundaries and that FWT and IWT remain orthogonal transforms. The pseudocode for the algorithms FWT and IWT can be found in Appendix A.2 (Algorithms 5 and 6, respectively), along with a brief explanation and an illustration of how the FWT algorithm operates on a non-dyadic grid (Figure A.2).

3.3.3.4 Preservation of the RKHS structure of the velocity fields

We transposed the classical Haar description of continuous functions to discrete functions defined on a grid. To implement a coarse-to-fine initialization approach for atlas estimation, we will describe the initial momentum vectors not through their values on the grid but by the decomposition of these values in the discrete Haar basis. The multiscale structure will be used to obtain a smooth initial field by setting fine-scale coefficients to 0. By construction, this does not change the fact that the initial vector fields $v_{0,i}$, and thus all the vector fields $v_{t,i}$, are finite combination of the RKHS kernel K_g .

We recall that in the algorithm of Durrleman et al. [54], the initial vector field for a subject i is defined as a finite linear combination of identical Gaussian kernels that are evaluated at an initial set of points:

$$v_{0,i}(x) = \sum_k K_g(x, c_{k,i}(0)) \alpha_{k,i}(0)$$

where K_g is a Gaussian kernel defining a RKHS and $\alpha_{k,i}(0)$ is a vector attached to $c_{k,i}(0)$. In our scheme, we set the initial points $(c_{k,i}(0))_k$ on a grid so that the equation becomes:

$$v_{0,i}(x) = \sum_{k \in \llbracket 0, K_1 \rrbracket \times \dots \times \llbracket 0, K_d \rrbracket} K_g(x, c_k(0)) \alpha_{k,i}(0)$$

where $c_k(0) = k$. Instead of optimizing $v_{0,i}$ by optimizing directly its momentum vectors $\alpha_{k,i}(0)$, we will optimize them in the wavelet domain. More precisely, we define

$$((a_{S,k,i})_k, (d_{s,k,i}^o)_{1 \leq s \leq S,o,k}) = FWT((\alpha_{k,i})_k)$$

where the Wavelet Transform has been applied to the momentum vectors component by component. Note that by construction,

$$(\alpha_{k,i})_k = IWT((a_{S,k,i})_k, (d_{s,k,i}^o)_{1 \leq s \leq S,o,k}).$$

Using the wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{1 \leq s \leq S, o, k}$ instead of the momentum vectors $(\alpha_{k,i})_k$ amounts to switch from the description

$$v_{0,i}(x) = \sum_{k \in [0, K_1] \times \dots \times [0, K_d]} K_g(x, c_k(0)) \alpha_{k,i}(0)$$

$$\text{to } v_{0,i}(x) = \sum_k a_{S,k,i} \tilde{\Phi}_{S,k}(x) + \sum_{s'=s+1}^S \sum_{o,k} d_{s',k,i}^o \tilde{\psi}_{s',k}^o(x)$$

where $\tilde{\Phi}_{s,k}$ (respectively $\tilde{\psi}_{s,k}^o$) is a function defined by replacing $\Phi_{0,k'}$ by $K_g(\cdot, c_{k'}(0))$ in the decomposition of $\Phi_{s,k}$ (resp. of $\psi_{s,k}^o$) in Equation (3.3) (resp. Equation (3.4)).

Even if $v_{0,i}$ is now defined through the vectorial wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{s,o,k}$ instead of the momentum vectors $\alpha_{k,i}$, it remains a linear combination of the $K_g(x, c_k(0))$ so that we are still in the setting of Durrleman et al. [54] and can rely on

$$v_{t,i}(x) = \sum_k K_g(x, c_{k,i}(t)) \alpha_{k,i}(t).$$

Note that we do not use the Haar parameterization outside the initialization. Indeed, the initial grid is deformed under the action of the diffeomorphism $\phi_{t,i}$ when the time evolves, so that the Haar wavelet coefficients would be hard to interpret for $t > 0$.

3.3.4 Multiscale optimization algorithm

The key difference between our scheme and the one of Durrleman et al. [54] is the use of the Haar parameterization in a fixed grid for the initial velocity fields. Thus, we optimize the following cost function:

$$\sum_{i=1}^N \left(\frac{d(O_i, O_{ref} \circ \phi_{1,i}^{-1})^2}{2\sigma^2} + \|v_{0,i}\|_V^2 \right),$$

where $\beta_{0,i}$ is a set of wavelet coefficients $(a_{S,k,i})_k$ and $(d_{s,k,i}^o)_{s,k,o}$ related to the momentum vectors $\alpha_{0,i}$ by $\alpha_{0,i} = IWT(\beta_{0,i}) = M_{IWT}^1 \beta_{0,i}$.

This implies a relationship between the gradient of the cost function with respect to the wavelet coefficients $\nabla_{\beta_{0,i}} E$ and the gradient with respect to the momenta $\nabla_{\alpha_{0,i}} E$:

$$\begin{aligned} \nabla_{\beta_{0,i}} E &= M_{IWT}^{T1} \nabla_{\alpha_{0,i}} E \\ &= M_{FWT}^1 \nabla_{\alpha_{0,i}} E \\ &= FWT(\nabla_{\alpha_{0,i}} E) \end{aligned}$$

where we have used the fact that $M_{IWT}^{T1} = M_{FWT}^1$ because the transform is orthonormal. The gradient $\nabla_{\beta_{0,i}} E$ can thereby be computed for almost the same cost as the one of

$\nabla_{\alpha_{0,i}} E$. If we perform the same gradient descent as in the original algorithm but in the wavelet domain, we obtain exactly the same results as the original algorithm.

To obtain hopefully better results, we enforce some constraints on the wavelet coefficients of the initial velocity fields. Namely, we use a coarse-to-fine **initialization** strategy by optimizing first the initial velocity fields whose wavelets coefficients are null at the finest scales and adding progressively these fine scale coefficients. In the following, we describe in detail our procedure which is summarized in Figure 3.2 and Algorithm 1 and illustrated on a simple registration example in Figure 3.5.

Algorithm 1 Multiscale initialization algorithm.

```

1: Input
2: Set of objects  $(O_i)_{1 \leq i \leq N}$  of dimension  $d$ , template object  $O_{ref}$ , geometric kernel width  $\sigma_g$ ,
   trade-off regularity/fidelity-to-data  $\sigma$ , initial step size  $h$ 
3: Initialization
4:  $j \leftarrow 0$ 
5:  $c_0 \leftarrow$  Regular grid of control points with spacing  $\sigma_g$ 
6: Template object  $O_{ref}(j) \leftarrow O_{ref}$ 
7: Momentum vectors  $\alpha_{0,i}(j) \leftarrow 0$  for each subject  $i$ 
8: Initialize the gradients  $\nabla_{\alpha_{0,i}} E \leftarrow 0$  for each subject  $i$  and  $\nabla_{O_{ref}} E \leftarrow 0$ 
9:  $\beta_{0,i}(j) \leftarrow FWT(\alpha_{0,i}(j))$  for each subject  $i$ 
10: Current scale  $S_j \leftarrow$  maximum scale of  $\beta_{0,0}(j)$ 
11: repeat
12:    $j \leftarrow j + 1$ 
13:   for each subject  $i$  do
14:     Compute the evolution of  $\alpha_i(t)$  using Equation (2.3) {Gradients computation}
15:     Compute  $\phi_i$  by solving the flow equation (Equation (2.1))
16:     Deform the template object  $O_{ref}$  with  $\phi_i$ 
17:     Compute the cost function  $E$  as in Equation (3.2)
18:     Compute the gradients  $\nabla_{\alpha_{0,i}} E$  and  $\nabla_{O_{ref}} E$ 
19:      $\nabla_{\beta_{0,i}} E \leftarrow FWT(\nabla_{\alpha_{0,i}} E) = (a_{S_{max},k}^i)_k \cup (d_{s,k}^{i,o})_{1 \leq s \leq S_{max},k,o}$  {Algorithm 5}
20:     for each detail coefficient  $d_{s,k}^{i,o}$  of  $\nabla_{\beta_{0,i}} E$  do
21:       if  $s < S_j$  then
22:          $d_{s,k}^{i,o} \leftarrow 0$  {Finer scale silencing}
23:       end if
24:     end for
25:   end for
26:    $\beta_{0,i}(j) \leftarrow \beta_{0,i}(j-1) - h \times \nabla_{\beta_{0,i}} E$  for all subjects  $i$  {Parameter update}
27:    $\alpha_{0,i}(j) \leftarrow IWT(\beta_{0,i}(j))$  for all subjects  $i$  {Algorithm 6}
28:    $O_{ref}(j) \leftarrow O_{ref}(j-1) - h \times \nabla_{O_{ref}} E$ 
29:   Compute the total residual value  $\Delta_j$  according to Equation (3.5)
30:   if  $\frac{\Delta_{j-1} - \Delta_j}{\Delta_{j-1}} < 0.01$  and  $S_j > 1$  then
31:      $S_j \leftarrow S_{j-1} - 1$  {Scale refinement step}
32:   end if
33: until Convergence
34: return Template object  $O_{ref}$  and momentum vectors  $\alpha_{0,i}$ 

```

The multiscale optimization can be seen as an initialization of each new scale with the

optimal template-to-subject deformations of the previous coarser scale. More precisely, at iteration j , we only optimize the wavelet coefficients of the vector fields whose scales are above or equal to a current scale S_j . Since the original RKHS setting is preserved, this can easily be done by computing the gradient $\nabla_{\alpha_{0,i}} E$ with the efficient numerical scheme of Durrleman et al. [54] (lines 14-18 in Algorithm 1), applying FWT to $\nabla_{\alpha_{0,i}} E$ to derive the gradient with respect to $\beta_{0,i}$ (line 19) and then setting to 0 the wavelet coefficients whose scale is strictly smaller than S_j (line 22). We then update the coefficients $\beta_{0,i}$ with the modified gradient and recover the updated $\alpha_{0,i}$ using the IWT function (line 26). If we iterate without modifying the current scale, we optimize the cost function in a subspace of functions that are simpler than in the original algorithm, the wavelet transform scale limitation acting as a regularizer.

As we want to optimize on the full set of functions defined by the momentum vectors, we progressively decrease the current scale S_j . We propose to decrease the scale when we are close to convergence at the current scale (line 31). This is measured by computing the mean residual value over subjects at iteration j :

$$\Delta_j = \frac{1}{N} \sum_{i=1}^N d(O_{ref,j} \circ \phi_{j,i}^{-1}, O_i). \quad (3.5)$$

If the residual decrease with respect to the previous iteration is below a threshold of 1%, we decide that the algorithm is close to convergence. In the case of Figure 3.5, the algorithm starts at $S_0 = 3$, performs optimization until (almost) convergence at this scale, goes to scale 2 and performs the subsequent scale transitions in the same manner.

Our optimization procedure ensures that the momenta belonging to the same area are updated with identical values. At a given scale S_j , the velocity fields can vary spatially only at scales coarser than $S_j - 1$. In other words, at scale S_j , the i^{th} initial velocity field implicitly writes as follows:

$$v_{0,i}(x) = \sum_k a_{S_{max},k,i} \widetilde{\Phi}_{S_{max},k}(x) + \sum_{s'=S_j}^{S_{max}} \sum_{k,o} d_{s',k,i}^o \widetilde{\psi}_{s',k}^o(x)$$

where $\widetilde{\Phi}_{S_{max},k}$ and $\widetilde{\psi}_{s',k}^o$ are linear combinations of localized Gaussian kernels $K_g(x, c_k(0))$.

When the algorithm reaches scale 1, the momenta are updated independently of each other. Importantly, unlike previous approaches that represented deformations in a wavelet basis (see Section 3.2.1), when the algorithm reaches this finest scale, the momentum vectors are free of constraints and the parameterization of the velocity fields is equivalent to its original definition, i.e. Equation (3.1). Thus, in theory, our coarse-to-fine algorithm could reach the same solutions as the original one, but as we will see in Section 3.4, the multiscale scheme converges to better solutions.

Contrary to the previous multiscale algorithms developed in the LDDMM framework, our strategy does not add any complexity to the mathematical model. The parameterization of the velocity fields remains identical to that of Durrleman et al. [54]. Computation of the gradients and subsequent cost also remains identical. The only additional complexity arises from the algorithms FWT and IWT, which, following the im-

plementation of Mallat [141], are of linear complexity.

Our code is publicly available in a Git repository¹.

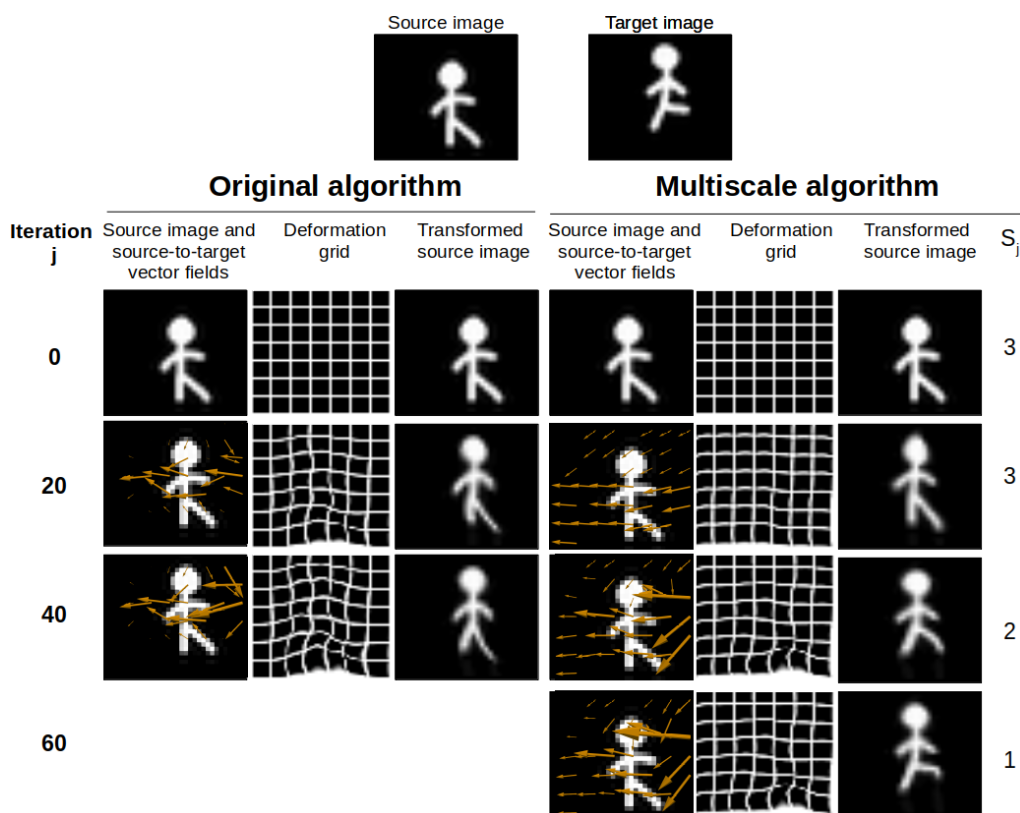


Figure 3.5: Original and multiscale strategies applied to a registration example with $\sigma_g = 4$ ($k_g = 49$ control points). For each algorithm, we display the source-to-target vector fields (orange arrows, scaling factor = 5), the corresponding deformation grid and the transformed source image every 20 iterations until convergence. S_j is the scale constraining the momentum vectors at iteration j . Notice how the multiscale strategy first estimates coarse displacements to move the character to the left, then finer transformations to adjust the position of the arms. In contrast, the original algorithm applies fine-scale deformations to the entire image (except the borders), while the multiscale algorithm estimates fine deformations localized around the character’s right leg, and smoother deformations elsewhere.

3.3.5 Multiscale optimization algorithm with local adaptation

The term *multiscale* can imply either that deformations are estimated in a coarse-to-fine manner, or that deformations of different scales coexist. Here, we enhance our algorithm with a *true* multiscale flavor by enforcing the coexistence of velocity fields of different scales. Since wavelet functions are characterized by their scale, location and orientation, we take advantage of the location component so that we can force the velocity

¹https://github.com/fleurgaudferneau/Deformetrica_multiscale/

field of each subject to be *unevenly* smooth. Namely, the spatial constraints we impose through the wavelet-based multiscale representation can vary according to the amount of variability in the images, so that template-to-subject deformations may remain coarse in areas that contain few information. For the sake of clarity, in the following the multiscale algorithm *without* local adaptation will be considered as the default version and referred to as *multiscale* or *naive multiscale*.

The multiscale algorithm with local adaptation is described in more detail in Appendix A.3.

3.3.6 Alternative multiscale algorithm

We introduce an alternative to our multiscale algorithm in which we combine our naive multiscale procedure with a different renormalization step in the FWT algorithm. Independently of the coarse-to-fine procedure, this allows us to give more weight to coarse scale or fine scale coefficients during optimization. The interested reader is referred to Appendix A.4 for more details.

3.4 Experiments: cross-sectional atlas estimation

In this section, we compare our multiscale algorithm to the original version from Durrleman et al. [54] on three different image datasets. Note that the version of the multiscale algorithm evaluated here is the *naive* one: the added value of the local adaptation extension will be evaluated in Section 3.5.1. The L^2 norm is used to compute distances between objects. The training phase consists in atlas estimation and the test phase consists in registering the estimated template image to a set of new images. Performance of the algorithms is assessed by computing the relative residual error: $R = \frac{\Delta_{j_{end}}}{\Delta_0}$ with $\Delta_{j_{end}}$ the mean residual value over subjects at convergence, and Δ_0 the mean residual value at iteration 0.

Optimization relies on a gradient descent algorithm in which the step sizes h are first scaled by the gradients norm and then diminished by a backtracking algorithm to guaranty a descent. The following parameters are used: trade-off between regularity and fidelity to data $\sigma = 0.1$; initial step size $h = 0.01$; convergence threshold = 0.0001. The initial template image for atlas estimation is given by the mean of the intensities of the training images.

The experiments were run on an Ubuntu 18.04.5 machine equipped with a NVIDIA GPU driver with 12 GB memory.




Figure 3.6: Estimation of the template image by the original and multiscale algorithms on the dataset of handwritten digits with different values of σ_g . For each experiment, five template images (estimated with non-intersecting training sets) are presented on the left, along with the template image from the first training set warped to the first five training images on the right. σ_g : width of the Gaussian kernel; k_g : corresponding number of control points.

3.4.1 Handwritten digits

In this section, we use images of the digit two extracted from the well-known United States postal database of handwritten digits [89]. The size of the images is 28 by 28 pixels. We test the original algorithm against the multiscale strategy using an experimental procedure similar to that of Durrleman et al. [54]: atlas estimation is performed using 20 randomly-chosen training images and the estimated template image is registered to 10 randomly chosen test images with the same parameters as those used during training. The experiment is repeated five times with different training and test sets. There is no intersection between any of the training and test sets. This cross-validation procedure is performed for each version of the algorithm and reproduced with different kernel widths σ_g . Since the five experiments are performed on independent datasets, we use paired Student t-tests to compare the performance of the two algorithms. Table 3.1 shows the mean relative residual error yielded by the algorithms over the experiments. Figure 3.6 presents the five template images estimated by each algorithm with different sets of parameters, along with the template image estimated from the first training set warped

to five of the training images.

Table 3.1 shows that the multiscale algorithm reaches lower residual error than the original algorithm during atlas estimation and registration, with differences that reach significance for $\sigma_g = 2$. Consistent with these results, we observe in Figure 3.6 that the original algorithm yields highly irregular template images, a trend which worsens when the number of control points increases, indicating overfitting. The first template image warped towards the training images yields images that are close to the original ones for $\sigma_g = 3$, but these deformations become unrealistic for lower values of σ_g . These observations belie the quantitative evaluation, which shows that the performance increases with the number of parameters. This discrepancy demonstrates that residual error alone is not sufficient to evaluate the accuracy of the algorithms, as irregular vector fields might be able to match any target image at the cost of producing unnatural deformations.

Unlike the original algorithm, the multiscale procedure produces realistic template images, whose quality is preserved when σ_g decreases. Moreover, all reconstructed images are very close to the original ones. Their quality slightly increases with the number of control points: this is most evident for the third and fifth reconstructed subjects, which become more accurate for lower values of σ_g .

Table 3.1: Performance of the original and multiscale algorithms on the dataset of handwritten digits during the training and test phases. Data are mean \pm standard deviation of relative residual error over five experiments. Bold style indicates statistically significant differences between the algorithms (p-value < 0.05).

σ_g	k_g	Original	Multiscale
Atlas estimation			
3	100	15.8 \pm 12.2	7.1 \pm 0.5
2	196	10.1 \pm 1.7	4.8 \pm 0.4
1.5	361	7.3 \pm 1.7	6.4 \pm 3.4
Registration			
3	100	10.0 \pm 6.1	6.4 \pm 5.3
2	196	8.5 \pm 5.2	4.2 \pm 3.2
1.5	361	9.0 \pm 5.1	4.0 \pm 5.9

3.4.2 Toy data

In the previous experiment, one can remark that the performance of the algorithms diverge most when a high number of parameters is used. Therefore, one might simply be tempted to employ the original algorithm with a lower number of parameters, as in Durleman et al. [54] whose experiments were performed with 36 control points. However, datasets that present a higher amount of details and inter-subject variability may benefit from our multiscale strategy even when a lower number of control points is used. To confront our algorithm with a more difficult task, we designed a dataset of 30 characters.

The size of the images is 28 by 28 pixels. The dataset is presented in Appendix Figure A.1 and can be downloaded at the author’s website².

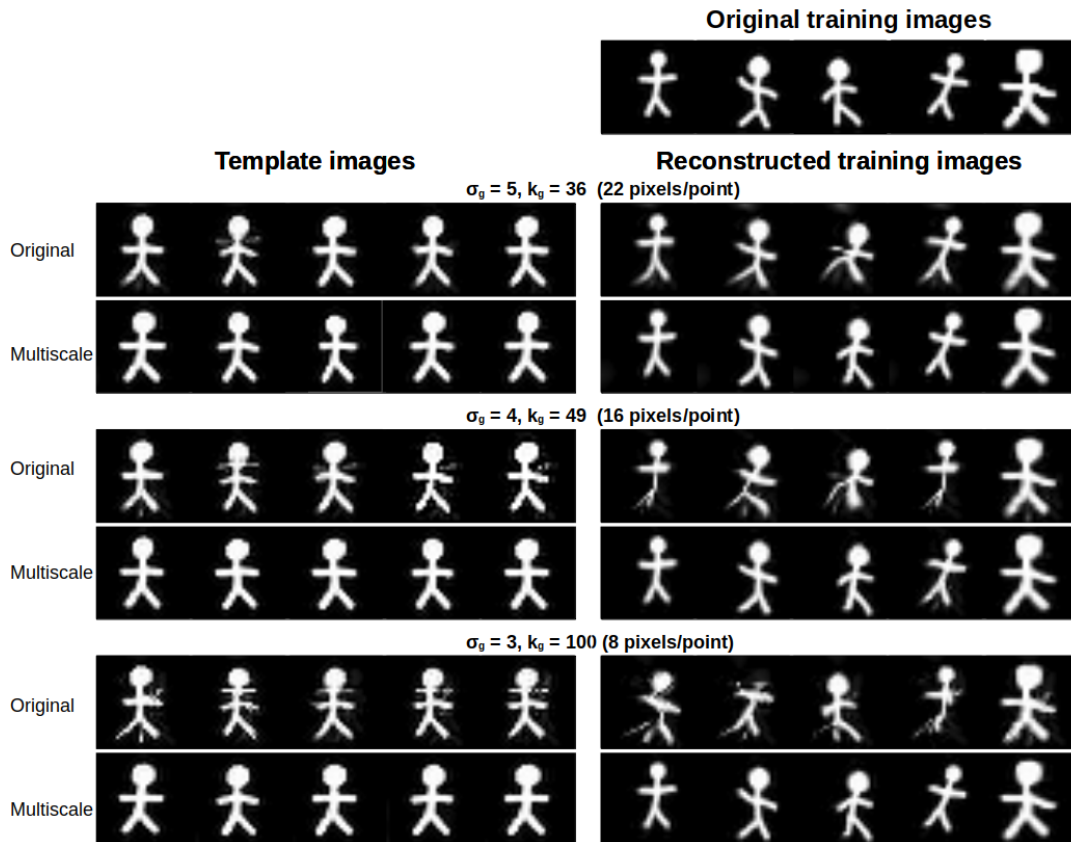


Figure 3.7: Estimation of the template image by the original and multiscale algorithms on the toy dataset with different parameters σ_g . For each experiment, five estimated template images (for each fold of cross-validation) are presented, along with the template image estimated from the first training set warped to the first five training images.

We compare our algorithm to the original version using cross-validation: the dataset is randomly split into a training set (24 images) and a test set (6 images). Each algorithm independently estimates a template image from the training set, and then registers the template to each image in the test set with the same parameters as those used during training. This procedure is repeated five times and reproduced with different kernel widths. No statistical tests are performed because of the overlap between the training sets and between the test sets. Section 3.4.2 displays the mean relative residual error after atlas estimation and registration and Figure 3.7 shows the five template images estimated by each algorithm with different kernels, along with five reconstructed training images.

Section 3.4.2 shows that the multiscale algorithm reaches lower residual error than the original algorithm. The performance of the multiscale strategy increases with the number of control points during training and test, while the original algorithm demonstrates the opposite trend. These results are supported by the qualitative evaluation of the template images. In Figure 3.7, for $\sigma_g = 5$, the two algorithms generate template

²https://fleurgaudfernau.github.io/Multiscale_atlas_estimation/

Table 3.2: Performance of the original and multiscale algorithms on the dataset of artificial characters during the training and test phases. Data are mean \pm standard deviation of relative residual error over five folds of cross-validation.

σ_g	k_g	Original	Multiscale
Atlas estimation			
5	36	14.0 ± 4.5	5.5 ± 0.8
4	49	16.4 ± 3.8	4.4 ± 0.4
3	100	18.4 ± 4.5	3.3 ± 0.5
Registration			
5	36	9.8 ± 6.2	6.0 ± 6.3
4	49	10.7 ± 6.0	4.6 ± 4.2
3	100	13.4 ± 7.5	3.6 ± 3.8

images that present discrete but noticeable differences. With the original version, the arms and legs of the characters appear slightly fuzzier, and the second template image is noisy. The reconstructed images yielded by the original algorithm are blurry (and even erroneous in case of the third subject), while the template and reconstructed images yielded by the multiscale algorithm seem sharp and accurate. As in the previous experiment, the quality of the template images estimated by the original algorithm deteriorates when the number of control points increases: with $\sigma_g = 4$ and $\sigma_g = 3$, images become fuzzier and display more erroneous features such as additional arms below the character’s neck. The morphology of all but one reconstructed characters is also completely erroneous. In contrast, the multiscale algorithm is able to produce stable, sharp and correct template images for all parameters. Similar observations can be made regarding the transformation of the template image towards the five training images: the multiscale strategy succeeds in generating images that are nearly identical to the original ones.

These differences have a simple explanation: the original version simultaneously estimates the overall shape of the characters and details such as the location and orientation of the arms and legs, making it more dependent on the initial template image and leading to the selection of erroneous features, while the multiscale strategy first focuses on estimating the characters main features, which are then refined during the finer scales. This phenomenon is illustrated by movies that show the evolution of the estimated templates across iterations, available at the first author’s webpage³.

3.4.3 Fetal brain images

To evaluate the performance of our multiscale approach on a dataset of clinical images, we use 30 fetal brain images with agenesis of the corpus callosum acquired in hôpital Trousseau, France. Gestational ages are comprised between 32 and 34 weeks of gestation (mean = 32.9 ± 0.6). A more detailed presentation of the dataset is available in Chapter 6.

³https://fleurgaudferneau.github.io/Multiscale_atlas_estimation

Agenesis of the corpus callosum is a developmental anomaly characterized by the total or partial absence of the corpus callosum. It is often associated to anatomical features such as widening of the lateral ventricles (see Section 6.1 for more details). Atlas estimation can help better understand congenital anomalies by providing an insight into how these anatomical characteristics vary together. However, as abnormal fetal brains may present a wide range of defects, this makes atlas estimation more difficult and prone to errors than with datasets of healthy fetuses. Thus, it is crucial to develop algorithms that are able to estimate realistic templates on both healthy and abnormal subjects.

The brain images are volume reconstructed and rigidly aligned according to the procedure described in Section 6.3.2. The final images are of size 105x100x120 voxels. Cross-validation is performed in the same manner as in Section 3.4.2, with 24 images used for atlas estimation and 6 images for testing. Figure 3.8 presents one example of estimated template images for $\sigma_g = 10$ and $\sigma_g = 5$, and Section 3.4.3 displays the mean relative residual error for each algorithm. Results for $\sigma_g = 7$ can be visualized in Appendix Figure A.4. Visual examination of the templates is performed by an expert radiologist.

Section 3.4.3 shows the importance of using a high number of parameters when performing atlas estimation on clinical images with high variability: the performance of both algorithms rises with the number of control points during training and test, which is reflected by the increasing sharpness of the template images in Figure 3.8. During training, the multiscale algorithm achieves lower residual error than the original algorithm for all values of σ_g and estimates more stable templates (as indicated by lower values of standard deviation). Registration results demonstrate that the multiscale algorithm has a higher ability to generalize to new images.

Table 3.3: Performance of the original and multiscale algorithms on the dataset of fetal brains images during the training and test phases. Data are mean \pm standard deviation of relative residual error over five folds of cross-validation.

σ_g	k_g	Original	Multiscale
Atlas estimation			
10	1,320	49.8 \pm 3.4	47.7 \pm 3.4
7	4,050	43.6 \pm 3.7	37.8 \pm 0.9
5	10,080	38.5 \pm 4.0	32.5 \pm 1.6
Registration			
10	1,320	41.1 \pm 6.9	39.9 \pm 6.7
7	4,050	35.0 \pm 6.2	31.1 \pm 5.7
5	10,080	31.4 \pm 5.0	25.6 \pm 5.2

In terms of template quality, the original and multiscale algorithms yield distinct results. The brain images estimated with the latter display sharper features and enhanced contrast between structures, especially for higher values of σ_g . The template images estimated by the original algorithm display more fuzzy areas, especially in the cingulum bundle (white arrows), the frontal horns (blue arrows) and the posterior area of the lateral ventricles (orange circles). On the contrary, these structures are more clearly delineated

in the template images generated by the multiscale algorithm. In addition, this method leads to more pronounced gyration patterns (e.g. the superior temporal sulcus, see red circles) and more visible subcortical brain structures such as the basal ganglia (red arrows). The most salient differences between the estimated templates are located at the medial surface of the brain (sagittal view), which presents high inter-subject variability: the surface estimated by the original algorithm is close to the mean of the image intensities, while the multiscale strategy produces a realistic anatomy, with a visible cingulate sulcus (green arrows). Some differences can also be noted between the reconstructed images yielded by the algorithms. In addition to increased sharpness, the multiscale template image warped to the subjects reveals more abnormal features, such as abnormally shaped corpus callosum in the case of the first subject (yellow arrows, middle column).

Altogether, these results indicate that our multiscale algorithm can be successfully applied to real-world, complex clinical data.

3.5 Additional experiments

3.5.1 Local adaptation extension

To evaluate the performance of the multiscale algorithm with local adaptation compared to the naive multiscale algorithm, we replicate our previous experiment on the handwritten digits (see Section 3.4.1). We recall that during training, atlas estimation is performed using 20 randomly-chosen images; then, the estimated template image is registered to 10 randomly chosen test images with the same parameters as those used during training. The experiment is repeated five times with different training and test sets. Results are available in Appendix A.3.

Appendix Table A.1 shows that there are no statistically significant performance differences between the two multiscale algorithms. Surprisingly, however, the naive multiscale version has lower residual error values in almost all experiments. In Appendix Figure A.3, the template images estimated by the two algorithms appear of comparable quality. The same observation can be made regarding the way the template images are warped to the first five training images. The only noteworthy differences are that the naive multiscale algorithm estimates template images that are more stable between experiments and seem slightly more natural.

3.5.2 Alternative multiscale scheme

In Appendix A.4, we introduced a renormalization factor ρ and the associated variant of the renormalization matrix: R^ρ . In the normal case (i.e. in the rest of this chapter), we have $\rho = 1$ and $R^1 = R$. As detailed in Appendix A.4, if we modify the value of ρ , we introduce some imbalance between coarse scale and fine scale coefficients. In other words, when computing the wavelet coefficients β of the gradient Δ_{E_α} , we can artificially increase coarse scale coefficients (if $\rho < 1$) or fine scale coefficients (if $\rho > 1$).




Figure 3.8: Atlas estimation by the original and multiscale algorithms on the dataset of fetal brain images. For each experiment, the estimated template volume from the first fold of cross-validation is presented in the left column. The first row displays the first two training images; the middle and right columns display the corresponding reconstructed images (i.e. the template image warped to the training images). Salient differences between images are indicated by specific markers. White arrows indicate the cingulum bundle, blue arrows the left frontal horn, red arrows the left basal ganglia, orange circles the posterior part of the right lateral ventricle, red circles the right superior temporal sulcus, yellow arrows the theoretical location of the corpus callosum, and green arrows the cingulate sulcus.

When applying the Inverse Wavelet Transform to β , we recover a gradient with more high or low frequencies than in the input Δ_{E_α} .

We use this different renormalization step to introduce a variant of the coarse-to-fine algorithm: we combine the multiscale gradient descent with different renormalization factors. This modified renormalization step is not performed in a coarse-to-fine fashion: the chosen value of ρ is kept constant during optimization. This can be seen as an additional low-pass or high-pass filter applied to the vector fields.

We study the effect of ρ on the algorithm performance by replicating our experiment on the toy dataset (Section 3.4.2) with different values of ρ and different kernel widths σ_g . Results are presented in Appendix A.4.2 and compared with the baseline results, i.e. the error rate when $\rho = 1$, corresponding to the regular multiscale algorithm. Briefly, we

observe that error rates increase when ρ moves away from 1 (i.e. towards 0 or 2), which is unsurprising because this means giving too much weight to high or low frequency components. With large kernels, i.e. $\sigma_g = 7$ or $\sigma_g = 5$, giving more weight to fine scale components slightly increases the performance. This is interesting as this means that one can choose a sparser parameterization of the vector fields, e.g. to save computational time while increasing the performance by selecting higher frequency components.

3.6 Discussion

In this chapter, we introduced a novel, wavelet-based reparameterization of the initial velocity fields in the LDDMM framework. Taking advantage of the hierarchical property of the wavelet decomposition, we implemented a multiscale optimization strategy for cross-sectional atlas estimation and registration. The transfer of information from coarse to fine scales ensures smarter initialization of the deformations at each level, leading the algorithm to favor smoother solutions while avoiding unrealistic local minima. Contrary to previous multiscale algorithms introduced in the LDDMM framework [215, 85, 156, 154], our approach adds no complexity to the mathematical model. The reparameterization of the velocity fields can be seen as an additional layer of spatial regularization, which preserves the RKHS structure of the vector fields and the efficient numerical scheme used to compute the gradients.

We performed experiments on three datasets of increasing difficulty:

- On the classical dataset of handwritten digits, the multiscale algorithm efficiently avoids unrealistic local minima, especially when using fine parameterizations.
- On a more challenging dataset of hand-drawn characters, the original algorithm is unable to estimate templates that are satisfying with respect to quantitative and qualitative criteria. The multiscale algorithm outperforms the former in both regardless of the number of parameters.
- Finally, the experiment on real-world, complex clinical data, i.e. pathological fetal brain images, showed more subtle differences between the multiscale and the original algorithms, in particular a better ability of the multiscale strategy to preserve anatomical details, including unusual or abnormal ones.

Eventually, these experiments demonstrate that the multiscale algorithm yields higher quality templates with better stability, that are able to generalize to unseen images. Not only does our strategy produce images that have a realistic anatomy, but it leads to enhanced preservation of anatomical details. This makes it particularly promising for tasks involving high inter-subject variability, specifically clinical images.

We also introduced a local adaptation extension of the multiscale algorithm to enable the coexistence of flows of difference scales. Experiments showed that this addition brings no improvement to the multiscale algorithm, suggesting that, in fact, the so-called "naive" coarse-to-fine strategy implicitly performs some kind of local adaptation and favors truly multiscale flows.

Some limitations of this algorithm have to be highlighted. Unlike approaches based on a mixture of kernels, our deformation model relies on a single Gaussian kernel. While this provides the advantage of introducing no additional parameters, the results of our algorithm still depend on the choice of the kernel width σ_g - but to a lesser extent than the original version. Some methodological improvements could be made to the FWT algorithm. In particular, the border treatment of non dyadic grids define scaling and wavelet functions of small support in some image borders. As illustrated in Appendix Figure A.2, this creates areas of uneven size at coarse scales, and therefore vector fields with locally uneven smoothness constraints. To avoid this, one could replace the Haar wavelet by smoother wavelet functions such as the Daubechies wavelet. However, since image borders often contain little information, this issue is unlikely to arise.

To conclude, it is important to remark that the multiscale strategy is versatile in the sense that it is not tailored to a given type of object, a given task or a given deformation model. It could be extended to other mathematical frameworks like those that compose flows of different scales [156, 154] or, as we shall see in the next chapter, to other Computational Anatomy models.

VERSATILE MULTISCALE STRATEGIES

Contents

4.1	Motivation	50
4.2	Methods	52
4.2.1	Overview	52
4.2.2	Model of diffeomorphic deformations	52
4.2.3	Dual multiscale optimization	53
4.3	Experiments - cross-sectional atlasing	57
4.3.1	Comparison of the different multiscale strategies	57
4.3.2	Computational requirements	58
4.4	Experiments - temporal modelling	58
4.4.1	Regression on toy data	61
4.4.2	Regression of fetal brain images	63
4.4.3	Regression of cortical surface meshes	65
4.5	Discussion	67

This chapter is an extension of a paper published in the Proceedings of SPIE Medical Imaging [66]: *A multiscale algorithm for computing realistic image transformations: application to the modelling of fetal brain growth* by Fleur Gaudfernau, Stéphanie Allassonière and Erwan Le Pennec, 2023. This work was also published at the French Colloquium of Signal and Image Processing 2023 [67] as *Un algorithme multiéchelle pour déformer les objets de façon réaliste - application à la modélisation de la croissance du cerveau foetal* by Fleur Gaudfernau, Stéphanie Allassonière and Erwan Le Pennec, September 2023. An experiment on fetal brains using the techniques presented here was subject to a poster presentation at the Fetal, Infant and Toddler Neuroimaging Group 2022.

4.1 Motivation

In Chapter 3, we introduced a coarse-to-fine strategy based on a multiscale representation of diffeomorphisms for cross-sectional atlas estimation. In this chapter, our goal is threefold:

- Improve the previously introduced strategy by combining the multiscale representation of diffeomorphisms with a multiscale representation of the objects.
- Extend our coarse-to-fine strategy to a more challenging task: time-series atlasing, i.e. geodesic regression.
- Finally, evaluate the ability of geodesic regression models to accurately model brain changes during pregnancy.

Cross-sectional atlas estimation is an interesting tool to translate the classical mean-variance analysis to datasets of high dimensional clinical images. However, it is limited in that the reference shape and the anatomical variability are computed at a given age or at the given stage of a disease: for instance, in Section 3.4.3, a template MRI was computed for fetal brains with corpus callosum agenesis between 32-34 gestational weeks and it cannot be considered a representative object for other gestational ages. It is thus crucial to develop models to describe the evolution of a phenomenon over time. In Computational Anatomy, this means modelling the relationship between manifold-valued data (e.g. brain images, meshes...) and a real-valued independent parameter.

The concept of geodesic regression (see Section 2.4.3) was introduced by Fletcher et al. [65] to estimate the geodesic curve that best fits manifold-valued data. A geodesic regression model [64] was later developed in the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework and combined with the sparse parameterization of diffeomorphisms [54] that we use in this thesis. Geodesic regression has also been the basis for more complex models such as the spatio-temporal mixed-effects model [201, 200], where estimation of the average trajectory is combined with estimation of random effects representing the temporal and geometrical individual variations of this trajectory. Although appealing, spatio-temporal models require longitudinal datasets. In this thesis, as we do not work with such data, we restrict ourselves to models that handle time-series data such as geodesic regression and its variants.

As pointed out in Chapter 3, a crucial issue with Computational Anatomy models is the degree of realism of the underlying deformation model: an important feature of geodesic regression is that it is based on the hypothesis that the average trajectory is geodesic. This assumes that the time-dependent changes of the representative shape are modelled by a *single*, unchanging dynamic. In the case of fetal brain growth, this assumption is not accurate since brain development is known to be a complex, heterogeneous and non-linear phenomenon [34, 73, 206]. To the best of our knowledge, geodesic regression has only been applied once to the modelling of the fetal brain growth [134]. In this paper, fetuses were grouped into three age range spanning no more than 4 gestational weeks and separate regressions were applied to each group. Needless to say, it is more desirable to model brain development using a single, *continuous* trajectory to enable inter-subject comparisons. An interesting extension of geodesic regression has been introduced by Chevalier et al. [30, 29], who proposed to model the average trajectory not as a geodesic but as a *piecewise* geodesic. This model has the potential to overcome the limitations of geodesic regression and accurately model the fetal brain growth.

In this chapter, we are interested in applying our versatile multiscale optimization strategy to geodesic regression and piecewise geodesic regression, and evaluate how

these models, combined with our new optimization strategy can enhance the challenging task of modelling fetal brain growth. Further, to improve the modelling of complex trajectories with changing dynamics, we combine the multiscale reparameterization of the initial velocity fields with a multiscale representation of the objects. This upgraded optimization strategy is applied to several models (i.e. cross-sectional atlas and geodesic regression) and to several types of objects (i.e. images and meshes).

4.2 Methods

4.2.1 Overview

In the LDDMM setting, diffeomorphisms are constructed by integrating time-dependant velocity fields [233]. To reduce the search space, we impose that the vector fields belong to a Reproducing Kernel Hilbert Space (RKHS) defined by a kernel K_g [54]. This restrictive spatial constraint forces deformations to occur at a scale close to the kernel width σ_g . This leaves the user to find the right balance between a large kernel, which generates smooth but inaccurate transformations, and a smaller one, which increases the number of parameters to optimize and thus the risk of converging towards unrealistic local minima. To alleviate this problem, we introduce a multiscale optimization strategy based on a multi-resolution representation of both of the deformations and the objects. This additional regularization, easily transferable, is applied to several computational anatomy models.

4.2.2 Model of diffeomorphic deformations

In the LDDMM framework, objects are deformed through flows of diffeomorphisms $(\phi_t)_t$ defined by time-dependant velocity fields $(v_t)_t$. As stated in Section 2.3.2, we impose that the initial velocity fields at $t = 0$ belong to a RKHS V [54]:

$$v_0(x) = \sum_{k=1}^{k_g} K_g(x, c_k) \alpha_k(0) \quad (4.1)$$

where K_g is a Gaussian kernel of width σ_g and the $(\alpha_k)_k$ are momentum vectors attached to control points $(c_k)_k$. The control points are set on a regular grid of spacing σ_g .

The regularity term in Equation (2.4) constrains possible solutions $(v_t)_t$ to define geodesic paths of the ambient space according to the V norm. The $(v_t)_t$ satisfying this condition keep their RKHS structure along time. It follows that a geodesic vector field v_t is fully parameterized by its initial velocity field v_0 : hence, to estimate an optimal transformation ϕ_t , only v_0 has to be optimized. In other words, optimization boils down to a geodesic shooting problem, where only the initial conditions determine the end point of the trajectory. The RKHS structure of the velocity fields allows us to rely on an efficient numerical scheme to compute the gradients [54].

This framework can be used to define standard Computational Anatomy models. Here, we seek to improve the estimation of atlas estimation and geodesic regression through a dual multiscale optimization strategy. We recall the generic formulation of the cost function:

$$E(O_{ref}, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N \left[\sum_{j=1}^{M_i} \frac{d(O_i(t_{i,j}), O_{ref} \circ \phi_i^{-1}(t_{i,j}))}{\sigma_\epsilon} + \|v_{0,i}\|_V^2 \right] \quad (4.2)$$

where:

- N is the number of subjects
- M_i is the number of observations from subject i
- O_{ref} is the template object
- $\alpha_{0,i}$ is the i^{th} set of momentum vectors parameterizing deformation ϕ_i
- $t_{i,j}$ is the j^{th} observation of the i^{th} subject
- $O_{ref} \circ \phi_i^{-1}(t_{i,j})$ is the template deformed by the i^{th} template-to-subject deformation at time $t_{i,j}$
- d is a distance measure, i.e. the L^2 distance for images and the varifold distance for meshes (see Section 2.3.4)
- $\|v_{0,i}\|_V^2$ is the norm of the i^{th} velocity field in the RKHS V
- σ_ϵ is a parameter controlling the trade-off between attachment and regularity.

For the sake of clarity, we briefly summarize the models of interest in this chapter:

- **Cross-sectional atlas estimation:** given a dataset of N objects $(O_i)_{1 \leq i \leq N}$ with one observation each (i.e. $\forall i M_i = 1$), one seeks to estimate a template object O_{ref} representative of the average anatomy, and N template-to-subject deformations $(\phi_i)_{1 \leq i \leq N}$.
- **Registration** is a particular case of atlas estimation where the template object is fixed and $N = 1$.
- **Geodesic Regression:** given M observations $(O_1(t_1), \dots, O_1(t_M))$, one seeks to estimate a template object O_{ref} at age t_0 , along with the average geodesic trajectory $\phi(t)$ that deforms O_{ref} .
- **Piecewise geodesic regression:** an extension of geodesic regression, it allows changes in trajectory dynamics. It estimates P trajectories $(\phi_l(t))_l$ called geodesic components and associated to rupture times $t_R = [t_1, \dots, t_{P-1}]$.

4.2.3 Dual multiscale optimization

As detailed in Section 3.1, the regularity term in Equation (4.2) puts a significant spatial constraint on the vector fields. This is particularly problematic when one deforms clinical images, which often require a fine parameterization of deformations: flexibility comes at the price of making large displacements expensive and fostering unrealistic mappings. Here, we introduce a coarse-to-fine optimization procedure to guide the algorithm towards more realistic local minima. We combine a multi-resolution representation of the deformations with a multi-resolution representation of the objects. This strategy is detailed below and illustrated in Figure 4.1 on a toy example.

4.2.3.1 Multiscale representation of the velocity fields

In Equation (4.1), v_0 is defined as a sum of RKHS basis elements. Its coefficients in this basis are the initial momenta $(\alpha_k(0))_k$, denoted α_0 for simplicity. As described in Chapter 3, in the multiscale approach the same function is described by its coefficients in a different multiscale basis:

$$v_0(x) = \sum_{s=S_j}^{S_{max}} \sum_k \sum_o \Phi_{s,k}^o(x) \beta_{s,k}^o \quad (4.3)$$

where $\Phi_{s,k}^o$ is a wavelet function of scale s , location k and orientation o , the $\beta_{s,k}^o$ are the coordinates of v_0 in this new basis and S_j is the current scale used to perform coarse-to-fine optimization. The wavelet coefficients $\beta_{s,k}^o$, denoted β_0 in the following, are obtained from the conventional parameterization of v_0 using the Forward Wavelet Transform: $\alpha_0 = FWT(\beta_{s,k}^o)$.

In Chapter 3, this multiscale representation proved a straightforward way to implement a coarse-to-fine optimization strategy. When the algorithm is at scale S_j , the wavelet coefficients β_0 whose scale s is strictly smaller than S_j are set to 0, thus erasing the spatial variations of v_0 defined at scales smaller than S_j . At initialization, S_j is set to S_{max} so that the algorithm optimizes smooth velocity fields. These vector fields are progressively refined by adding finer and finer scale coefficients each time the algorithm is close to convergence.

In practice, the change of basis in Equation (4.3) is applied to the gradient $\nabla_{\alpha_0} E$ rather than α_0 . This enables us to rely on the same numerical scheme as in the original algorithm to compute the gradient of the cost function with regard to the momentum vectors. The silencing of fine scale coefficients occurs only during the parameters update step and can be summarized as follows (see Figure 3.2):

1. The gradient $\nabla_{\alpha_0} E$ goes through the wavelet transform: $\nabla_{\beta_0} E \leftarrow FWT(\nabla_{\alpha_0} E)$
2. The wavelet coefficients in $\nabla_{\beta_0} E$ whose scale is smaller than S_j are set to zero.
3. The coordinates of v_0 in the wavelet basis are updated: $\beta_0(j) \leftarrow \beta_0(j-1) - h \nabla_{\beta_0} E$
4. The coordinates of v_0 in the RKHS basis are recovered with the Inverse Wavelet Transform algorithm to get $\alpha_0(j) \leftarrow IWT(\beta_0(j))$.

Importantly, this optimization scheme preserves the RKHS structure of the vector fields and the computational efficiency of the original algorithm.

4.2.3.2 Multiscale representation of the objects

Irregular velocity fields are likely to foster an unrealistic template object O_{ref} . Conversely, complex objects may generate an erroneous template object, which will in turn foster irregular velocity fields. This is why we combine the multiscale optimization of the velocity fields with a multi-resolution representation of the input objects and the template object O_{ref} . We denote by S'_j the scale of the objects at iteration j . S'_j is computed as a function of S_j .

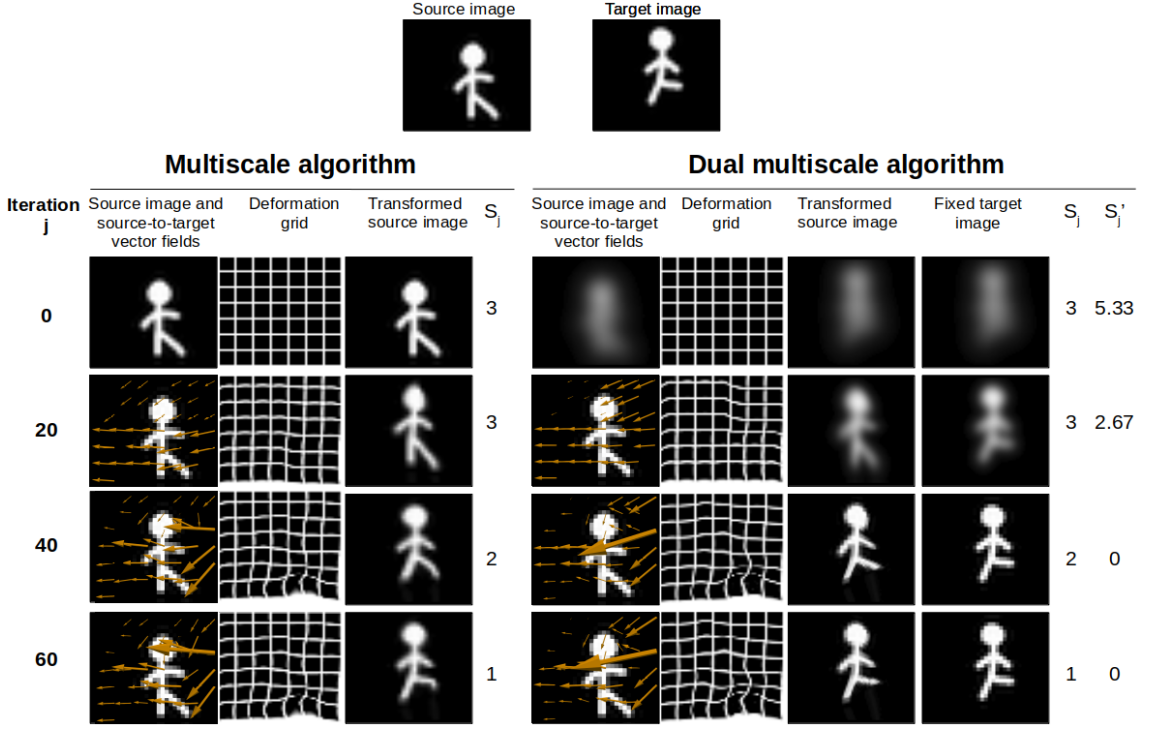


Figure 4.1: Multiscale and dual multiscale strategies applied to a registration example with $\sigma_g = 4$ (i.e. $k_g = 49$ control points). For each algorithm, we display the estimated source-to-target vector fields (orange arrows; scaling factor = 5), the corresponding deformation grid and the transformed source image every 20 iterations until convergence. S_j denotes the scale constraining the momentum vectors at iteration j , S'_j the width of the image filter at iteration j .

Images. In case of images, filtering is performed with Gaussian smoothing. At iteration j , S'_j denotes the width of the Gaussian smoothing kernel and the vector fields are constrained by scale S_j . We set the kernel standard deviation S'_j so that the effective averaging size is proportional to the size of the support of the wavelets at scale S_j :

$$S'_j = \begin{cases} \frac{\max(\text{supp}(\Phi_{S_j})) - \min(\text{supp}(\Phi_{S_j}))}{3} & \text{if } S_j > 1 \\ 0 & \text{otherwise} \end{cases}, \quad (4.4)$$

where $\text{supp}(\Phi_{S_j, (0,0)})$ is the support of the wavelet function of scale S_j located at position $(0, 0)$.

Meshes. With surface meshes, we can choose between two types of filtering:

- *Intrinsic* smoothing, i.e. Laplacian smoothing [90]: for a given number of iterations, each point of the mesh is moved at the average of its neighbors positions. As the distance between objects is computed using the varifold distance, the normal vectors of a mesh are recomputed after each filtering step. The scale of the objects, i.e. the number of iterations during Laplacian smoothing, is computed as

follows:

$$S'_j = \begin{cases} (\max(\text{supp}(\Phi_{S_j})) - \min(\text{supp}(\Phi_{S_j}))) \times 200 & \text{if } S_j > 1 \\ 0 & \text{otherwise} \end{cases}, \quad (4.5)$$

- *Extrinsic* smoothing: instead of modifying the object itself, the idea is to smooth its varifold representation. More details are available in Appendix B.1.

In this chapter, experiments on surface meshes are performed with intrinsic smoothing.

4.2.3.3 Multiscale optimizations

The multiscale representations of the velocity fields and the objects serve as basis for several multiscale strategies:

- The *multiscale momenta* strategy, which was introduced and evaluated in Chapter 3. Coarse-to-fine is performed on the subjects velocity fields from scale $S_0 = S_{max} - 1$ to scale 1.
- The *multiscale objects* strategy, where coarse-to-fine is performed on the objects from scale S'_{max} to scale 0.
- The *dual multiscale* strategy, which starts at scales $S_0 = S_{max} - 1$ and $S'_0 = f(S_0)$ and performs optimization using the filtered subjects images and the reparameterized velocity fields. The two scales are alternatively refined each time the algorithm is close to convergence. When scales $S_j = 1$ and $S'_j = 0$ are reached, both the velocity fields and the objects are free of constraints and optimization proceeds as in the original version of the algorithm.

The dual multiscale strategy is described in Appendix Algorithm 10. For clarity, the pseudo-code focuses on atlas estimation, but all multiscale strategies can also be applied to (piecewise) geodesic regression. The multiscale representation of the objects has only a few variations depending on the model one seeks to estimate. In atlas estimation, the template object O_{ref} is filtered only at initialization and left to be optimized by the algorithm. In registration, the fixed template object is filtered in the same way as the input objects, i.e. at each refinement of scale S'_j (line 39 in Appendix Algorithm 10). In (piecewise) geodesic regression, when fixed, the template object is filtered in the same way as the input objects. When estimated, it is filtered only at initialization and then optimized by the algorithm.

In piecewise geodesic regression, the P velocity fields defining the geodesic components are constrained in the same manner, although in theory one could allow different geodesic components to be constrained by different scales.

After each refinement of scale S'_j , the objects are filtered at a new scale, which leads to a sharp increase of the gradients of the cost function with regard to the momenta and the template object. When the template object is not fixed, optimization of the momentum

vectors is frozen during 3 iterations so that the template object can be updated to match the new filtered objects.

The code is available at a public Git repository¹. As in Chapter 3, the non-multiscale algorithm will be referred to as the *original* algorithm.

4.3 Experiments - cross-sectional atlasing

All experiments were run on an Ubuntu 18.04.5 machine equipped with a NVIDIA GPU driver with 12 GB memory.

4.3.1 Comparison of the different multiscale strategies

In this section, we evaluate how the dual multiscale optimization strategy can enhance the performance of an atlas estimation task - a model which we thoroughly studied in Chapter 3. Here, we evaluate four different optimization strategies: the original algorithm, the multiscale momenta algorithm, the multiscale objects algorithm and the dual multiscale algorithm. We perform the same experiment as in Section 3.4.2 using a dataset of 30 manually designed characters, which can be visualized in Appendix Figure A.1.

For each algorithm, we repeat five experiments as follows: the dataset is randomly split into a training set (24 images) and a test set (6 images); the algorithm performs atlas estimation on the training set and then registers the template to each image in the test set with the same parameters as those used during training. Performance of the algorithms is assessed by computing the relative residual error: $R = \frac{\Delta_{j_{end}}}{\Delta_0}$ with $\Delta_{j_{end}}$ the mean residual value over subjects at convergence, and Δ_0 the mean residual value over subjects at iteration 0. To assess how the performance of each optimization strategy varies depending on the number of parameters of the model, this procedure is repeated with different kernel sizes, i.e. $\sigma_g \in [3, 4, 5, 6, 8, 10]$, corresponding to a number of control points $k_g \in [100, 49, 36, 25, 16, 9]$.

Figure 4.2 shows the evolution of the residual error as a function of k_g for each algorithm and Figure 4.4 shows examples of estimated templates and template-to-subject deformations for the coarser ($\sigma_g = 10$) and finer ($\sigma_g = 3$) parameterizations of the velocity fields. Consistent with the quantitative results, the three multiscale strategies estimate quite similar images. Some subtle differences can however be noticed. With the coarser parameterization, the multiscale momenta and multiscale objects algorithms generate images with fuzzier areas compared to the dual multiscale strategy (e.g. the left arm of the template characters). More noticeably, they also fail to avoid some pitfalls (e.g. the erroneous reconstructions of the third training image) when estimating finer deformations.

¹https://github.com/fleurgaudfernau/Deformetrica_multiscale

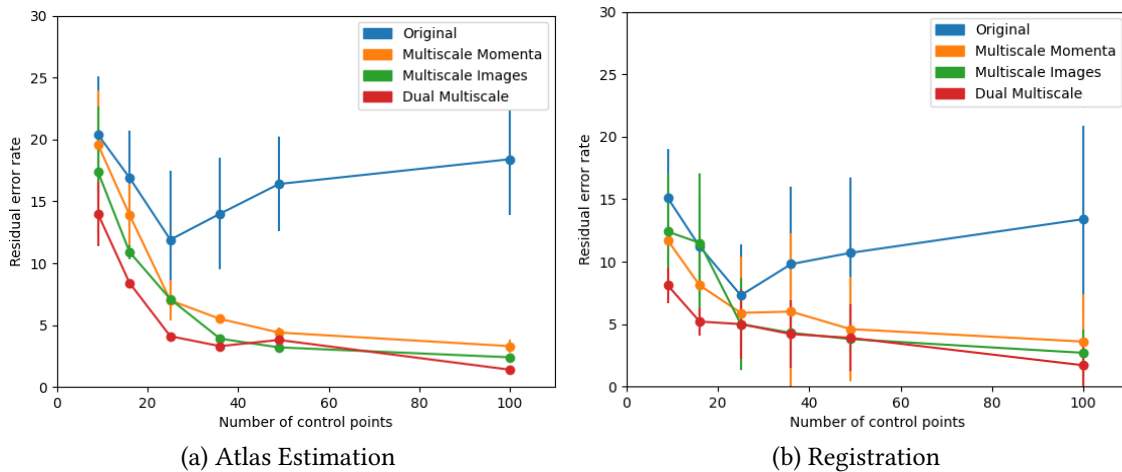


Figure 4.2: Residual error rate (mean and standard deviation) of the training (panel (a)) and test (panel (b)) phases. After atlas estimation, the residual error between the deformed template and each subject is computed. During cross-validation, each algorithm performed atlas estimation 5 times and image registration 30 times.

4.3.2 Computational requirements

To assess the computational requirements associated to each four optimization strategies, we plot in Figure 4.3 the number of iterations until convergence of the previous experiment, for the finer ($\sigma_g = 3$) and coarser ($\sigma_g = 10$) parameterizations. We recall that a single iteration for a given model has a similar computational cost in the four algorithms. Unsurprisingly, the original algorithm requires the least number of iterations and the dual multiscale strategy the most number of iterations. Interestingly, the multiscale objects strategy requires a higher number of iterations than the multiscale momenta strategy, despite equal numbers of coarse-to-fine steps; this is likely due to the fact that following each filtering step, the multiscale objects algorithm needs to adapt to the new target objects. There seems to be a trade-off between quality of estimation and computational speed. During registration, the gradient is computed with regards to one set of momentum vectors, so a high number of iterations is relatively impactless. However, a single iteration has a heavy computational cost in more complex models such as atlas estimation, especially if the deformation kernel is small and the number of objects is high. When running such models, it may be relevant to test first the original algorithm and assess whether or not it is worthy to try the multiscale(s) algorithm(s).

4.4 Experiments - temporal modelling

In the previous section, we demonstrated how the dual multiscale strategy can further enhance cross-sectional atlas estimation compared to the multiscale momenta strategy introduced in Chapter 3. Here, we will no longer compare the different versions of the multiscale algorithm and only assess the performance of the original and *dual multiscale*

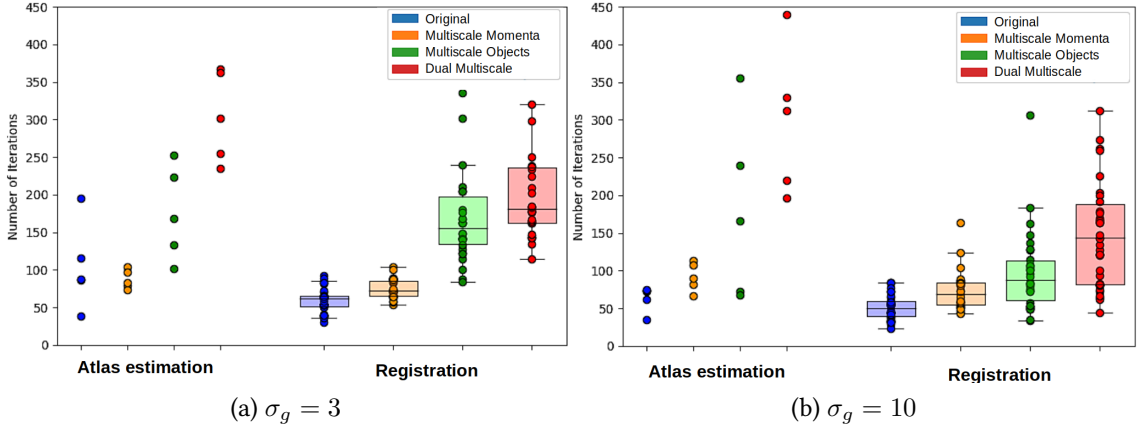


Figure 4.3: Number of iterations until convergence for each atlas estimation (training) and registration (test) experiments with two kernel sizes. During cross-validation, each algorithm performed atlas estimation 5 times and image registration 30 times.

algorithms. We now focus on another model which is of prime importance in the context of this thesis: geodesic regression. Generating a single time-dependent deformation that is able to represent the average changes in a population is a more tedious task than the estimation of a cross-sectional atlas. Moreover, biological phenomena may not always be characterized properly by a single temporal dynamic and fetal brain growth, which is the focus of this thesis, is no exception. This is why we will assess how the combination of the dual multiscale strategy with piecewise regression can enhance the modelling of time-dependent deformations.

Following an approach similar to that of the previous chapter, we evaluate the ability of our algorithm to generate realistic deformations on several time-series datasets. The first one, a toy dataset comprising artificial characters similar to that of the previous experiment, enables for easy interpretation and detection of potential errors. The second and third experiments evaluate the algorithm usefulness in a clinical setting as we use fetal brain images to model time-dependant brain changes during pregnancy.

We design an evaluation procedure that enables us to highlight differences between the original and multiscale strategies in different tasks. As illustrated in Figure 4.5, we implement a 5-fold cross-validation procedure composed of a two-step training phase and a test phase. As optimizing simultaneously the template object O_{ref} and the deformation can be ill-posed, we first perform atlas estimation to provide geodesic regression with an accurate baseline object at time t_0 :

1. Each algorithm estimates a template object O_{ref} from training set subjects close to t_0 in age, along with the template-to-subjects deformations.
2. Using O_{ref} as a starting point, geodesic regression estimates the trajectory $\phi(t)$ that best fits the objects (O_1, \dots, O_N) in the training set.
3. To assess the accuracy of the estimated trajectory, each object in the test set is compared to the age-matched template object from regression using registration: for each test subject i observed at time t_i we compute the deformation $\phi_{1,i}$ between the deformed template $\phi(t_i)(O_{ref})$ and the observation O_i .



Figure 4.4: Estimation of the template image by the original, multiscale momenta, multiscale objects and dual multiscale algorithms on the toy dataset with $\sigma_g = 10$ and $\sigma_g = 3$. The five estimated templates from cross-validation are presented on the left and the first template image estimated wrapped to the first five training images on the right. Red circles indicate differences between the images estimated by the different multiscale strategies. σ_g : width of the Gaussian kernel; k_g : corresponding number of control points.

We compare the performance of two optimization strategies (i.e. original and dual multiscale) associated with two ways of modelling time changes (one-component versus several components geodesic regression). Hence, the evaluation procedure is repeated four times. For each step of the evaluation, performance is assessed by computing the relative residual error.

As the residual distance between each object and the deformed template is part of the objective function (Equation (4.2)), and since we showed in the last chapter that it is not always a reliable indicator of the quality of a transformation, we also compute the Structural Similarity Index Metric [239] (SSIM) at each step of the evaluation:

$$SSIM(I_1, I_2) = l(I_1, I_2) \times c(I_1, I_2) \times s(I_1, I_2)$$

where I_1 and I_2 are the compared images, l is a function comparing the luminance (i.e.




Figure 4.5: Three-step evaluation procedure illustrated on the dataset of cortical surface meshes. Cortices outlined in red belong to the test set. O_{ref} : template object estimated during step (1). $\phi(t)$: (piecewise) geodesic trajectory estimated during step (2).

the mean pixel/voxel intensity) of the images, c compares the contrast (i.e. the standard deviation of the image intensities), and s quantifies the structural similarity between I_1 and I_2 (i.e. the correlation between luminance- and contrast- normalized intensities). SSIM values range between -1 (dissimilarity) and 1 (near-perfect similarity).

Optimization is performed with gradient descent. The algorithms are evaluated using several kernel widths σ_g . Rupture times between components are set at regular intervals. Other parameters are set as follows: trade-off between attachment and regularity $\sigma_\epsilon = 0.1$; initial step size $h = 0.01$; convergence threshold $= 0.0001$.

4.4.1 Regression on toy data

In this experiment, we use a toy dataset of 30 hand-drawn characters. Images are of size 28-by-28 pixels and are defined between times 0 and 19. Despite their simple anatomy, the temporal evolution of the characters is quite complex, with at least 3 different dynamics: a jump phase, a landing phase, and a falling phase. An extract of the dataset can be seen in Figure 4.6a. The cross-validation procedure is repeated with five different kernel widths $\sigma_g \in \{4, 5, 6, 8, 10\}$, corresponding to a number of control points $k_g \in \{49, 36, 25, 16, 9\}$. The template object O_{ref} is estimated using the training set images defined between times 0 and 2 and geodesic regression is computed from $t_0 = 0$. To set the number of components, we evaluated the regression performance of the original algorithm with a number of components ranging from $P = 2$ to $P = 10$. As shown in Appendix Figure B.2, the maximum performance is reached for a number of components between 5 and 8. Thus, in the main experiment, piecewise regression is computed with 5 components and the rupture times are set to $t_R = [3.5, 7, 10.5, 14]$.

Table 4.1 shows the mean SSIM after training (i.e. atlas estimation and regression) and test (i.e. registration). Appendix Table B.1 shows the mean relative residual error after training and test. Figure 4.6 shows 12 example trajectories (one-component or 5-component trajectories associated with the original or dual multiscale optimization) estimated with different values of $\sigma_g \in 4, 6, 10$. Appendix Figure B.3 shows example of

Table 4.1: Mean \pm standard deviation of the SSIM over five folds of cross-validation for atlas estimation, geodesic regression and registration on the toy dataset. At convergence, the SSIM is computed between each train/test subject and the deformed template.

σ_g	k_g	Algorithm	P	Atlas estimation	Regression	Registration
10	9	Original	$\frac{1}{5}$	0.84 ± 0.06	0.30 ± 0.05	0.56 ± 0.13
		Multiscale	$\frac{1}{5}$	0.85 ± 0.05	0.27 ± 0.03	0.58 ± 0.13
8	16	Original	$\frac{1}{5}$	0.79 ± 0.08	0.27 ± 0.03	0.54 ± 0.06
		Multiscale	$\frac{1}{5}$	0.84 ± 0.06	0.26 ± 0.02	0.58 ± 0.07
6	25	Original	$\frac{1}{5}$	0.83 ± 0.06	0.26 ± 0.03	0.54 ± 0.14
		Multiscale	$\frac{1}{5}$	0.90 ± 0.06	0.29 ± 0.05	0.50 ± 0.12
5	36	Original	$\frac{1}{5}$	0.85 ± 0.16	0.29 ± 0.05	0.66 ± 0.09
		Multiscale	$\frac{1}{5}$	0.94 ± 0.06	0.34 ± 0.02	0.77 ± 0.06
4	49	Original	$\frac{1}{5}$	0.84 ± 0.07	0.26 ± 0.04	0.51 ± 0.14
		Multiscale	$\frac{1}{5}$	0.89 ± 0.05	0.28 ± 0.07	0.57 ± 0.10
					0.29 ± 0.06	0.65 ± 0.05
					0.53 ± 0.04	0.80 ± 0.02
					0.27 ± 0.04	0.62 ± 0.06
					0.23 ± 0.02	0.52 ± 0.10
					0.31 ± 0.06	0.72 ± 0.03
					0.53 ± 0.03	0.83 ± 0.01

trajectories estimated with all tested values of σ_g .

After atlas estimation, the multiscale algorithm consistently reaches lower residual error and higher SSIM values than the original algorithm, which is confirmed by the more natural template images in Figure 4.6, especially for lower values of σ_g . The higher error rates achieved by regression compared to atlas estimation are not surprising, since regression has to fit 24 images within a single deformation.

Regardless of which optimization strategy is used, piecewise regression reaches lower residual error than one-component geodesic regression in Appendix Figure B.3. In contrast, Table 4.1 shows slighter differences between one-component and piecewise regression when they are associated to the original optimization strategy. In Figure 4.6, we observe that the original optimization strategy generates irregular and unrealistic trajectories even with large kernels.

In contrast, the multiscale strategy associated to piecewise regression consistently achieves higher SSIM values and lower residual errors than the other methods, highlighting the need to combine a piecewise model with the multiscale strategy to accurately represent complex time-dependent changes. As illustrated Figure 4.6, this association drastically enhances the quality of the characters' trajectories and maintains very good performance even for smaller kernels: fine movements are accurately modelled while preserving the anatomy of the characters.

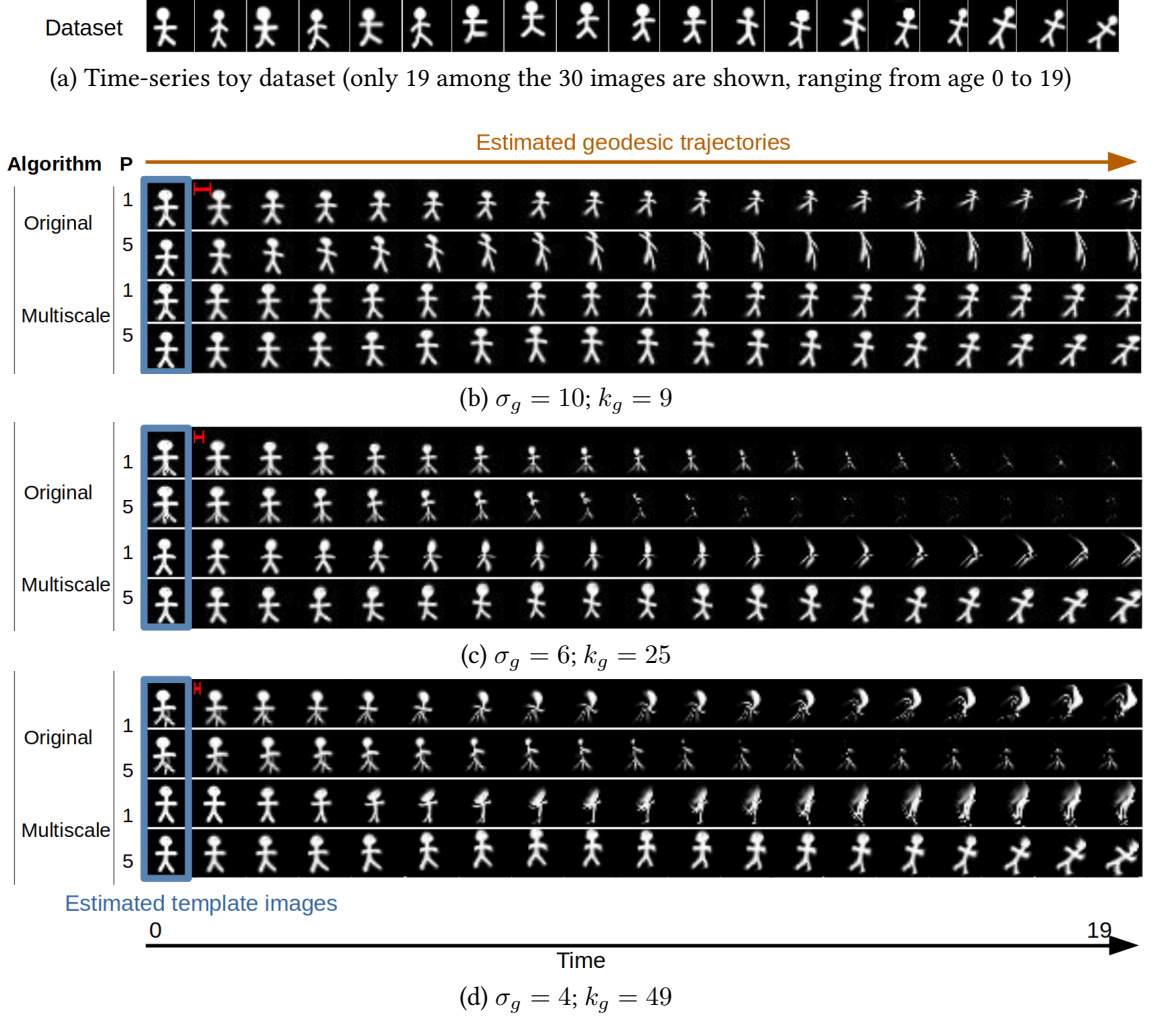


Figure 4.6: Geodesic regression results on the toy dataset for different values of σ_g . We show the results computed on the first fold of cross-validation: images from atlas estimation are outlined in blue and served as starting points to compute the geodesic trajectories from age 0 to 19. Red rulers indicate the kernel width σ_g .

4.4.2 Regression of fetal brain images

Here, we evaluate the efficiency of the multiscale strategy compared to the original algorithm on a more challenging task: estimation of an average trajectory from a cross-sectional dataset of fetal brains imaged at different gestational ages.

We use 42 volume-reconstructed brain MRIs of healthy fetuses from the open-source FeTa dataset [168]. More details about the dataset are available in Chapter 7, Section 7.2. Gestational ages range between 22 and 35 weeks (mean = 29.01 ± 3.96). Images are of size $135 \times 189 \times 155$ voxels. Brains are skull-stripped, rigidly aligned to a common coordinate space and intensity normalized. We perform 5-fold cross-validation procedures with $\sigma_g \in \{8, 10, 20\}$ (i.e. $k_g \in \{8, 160, 4, 256, 560\}$). The starting time t_0 is set at 29 weeks to minimize the amount of deformation the template image undergoes. Atlas estimation is performed on the training set subjects aged between 27 and 31 weeks.

Table 4.2: SSIM values (mean \pm standard deviation) over five folds of cross-validation on the dataset of fetal brain images.

σ_g	k_g	Algorithm	P	Atlas estimation	Regression	Registration
20	560	Original	1	0.90 ± 0.02	0.80 ± 0.06	0.87 ± 0.05
			4		0.80 ± 0.06	0.87 ± 0.05
		Multiscale	1	0.93 ± 0.01	0.81 ± 0.07	0.88 ± 0.06
			4		0.82 ± 0.06	0.88 ± 0.06
10	4,256	Original	1	0.92 ± 0.02	0.80 ± 0.06	0.86 ± 0.04
			4		0.79 ± 0.06	0.88 ± 0.04
		Multiscale	1	0.94 ± 0.01	0.80 ± 0.05	0.86 ± 0.04
			4		0.82 ± 0.06	0.88 ± 0.05
8	8,160	Original	1	0.94 ± 0.01	0.80 ± 0.06	0.84 ± 0.04
			4		0.80 ± 0.06	0.89 ± 0.04
		Multiscale	1	0.94 ± 0.01	0.80 ± 0.06	0.86 ± 0.04
			4		0.83 ± 0.06	0.88 ± 0.05

Piecewise regression is computed with 4 components and rupture times $t_R = [24, 28, 32]$. Each training and test set contains between 33-34 images and 8-9 images (respectively). Images are downsampled by a factor of 5 during optimization.

Quantitative performance. Table 4.2 and Appendix Table B.2 and show the SSIM and relative residual error (respectively) between the subjects and the deformed template brains after atlas estimation, regression and registration. The residual error shows little difference between the original and multiscale strategies during atlas estimation and registration. During regression however, the 4-component model consistently achieves better performance than the one-component model, and the multiscale strategy outperforms the original strategy. In contrast with the previous experiment, the SSIM values in Table 4.2 do not agree with these observations: the mean performance of the multiscale strategy is only slightly superior to that of the original algorithm, and the piecewise-multiscale strategy has slightly higher performance than the one-component multiscale strategy.

Qualitative results. For the sake of brevity, Figure 4.7 provides a visual comparison of the outputs of piecewise geodesic regression for $\sigma_g = 20$ and $\sigma_g = 8$ only. The original algorithm fails to model brain growth in a realistic manner: for $\sigma_g = 20$, the estimated template at 29 weeks has several blurry zones, notably the lateral ventricles (red arrow) and the border between the cortical plate and the cerebrospinal fluid (green arrow), and these errors are propagated to older and younger gestational ages during regression. For $\sigma_g = 8$, changes in brain size are underestimated compared to Gholipour et al. atlas (Figure 4.7e) and the cerebrospinal fluid has exaggerated thickness. These flaws are partially corrected by the multiscale strategy, which generates more natural template images.

Comparison to Gholipour et al. atlas. It is important to remark that compared to the state-of-the-art atlas, the piecewise trajectory generated by the multiscale strategy has many flaws: age-related changes seem mostly limited to an affine transformations. This lack of realism is not surprising given that diffeomorphisms cannot make structures

appear or disappear. Even more striking, already present structures such as the cortical plate (brown arrow), do not undergo the expected gyrification process. These flaws are also present in one-to-one mapping: the template brains wrapped to the test images at 33 and 34 in Figure 4.7b and 4.7d exhibit a smooth cortex, far from the expected convoluted pattern. Appendix Table B.3 displays the average SSIM values between the templates from geodesic regression and Gholipour et al. atlas at each gestational age: compared to the other methods, the multiscale-piecewise strategy achieves higher similarity with the atlas only for younger gestational ages.

4.4.3 Regression of cortical surface meshes

Finally, we assess the performance of the multiscale strategy on cortical surface meshes extracted from the FeTa dataset. The outer cortical surface of 34 subjects is delineated and reconstructed using the dHCP pipeline [140]. Gestational ages range between 22 and 35 weeks (mean= 29.55 ± 5.13 weeks). More details about data preprocessing are available in Section 7.2.




Figure 4.7: Results of the evaluation procedure on the dataset of fetal brains images during the first fold of cross-validation. Panels (a) and (c): template brains in axial view estimated between 23-33 weeks by the original and multiscale algorithms with $P = 4$ components. For comparison, panel (c) shows template images from Gholipour et al. atlas. Panels (b) and (d) shows examples of reconstructed test images during registration.

We perform the usual 5-fold cross-validation procedure with $\sigma_g \in \{5, 4, 3.5\}$ (i.e. $k_g \in \{5, 796, 11, 088, 16, 000\}$). Performance is computed with the residual error rate. We choose $\sigma_w = \sigma_g$ for the kernel K_w of the varifold distance. (Note that this means that different distance metrics are used for each parameterization.) The starting time t_0 is set at 24 weeks. Atlas estimation is performed on the training set subjects aged between 22 and 26 weeks. As in the previous experiment, piecewise regression is computed with 4 components and rupture times $t_R = [24, 28, 32]$.

The residual error rates are presented in Table 4.3. In average, the piecewise multiscale strategy achieves lower error values than the other 3 strategies. Although these quantitative differences are quite small, the gap tends to widen with lower values of σ_g .

Table 4.3: Residual error rate (mean \pm standard deviation) over five folds of cross-validation on the dataset of cortical surface meshes.

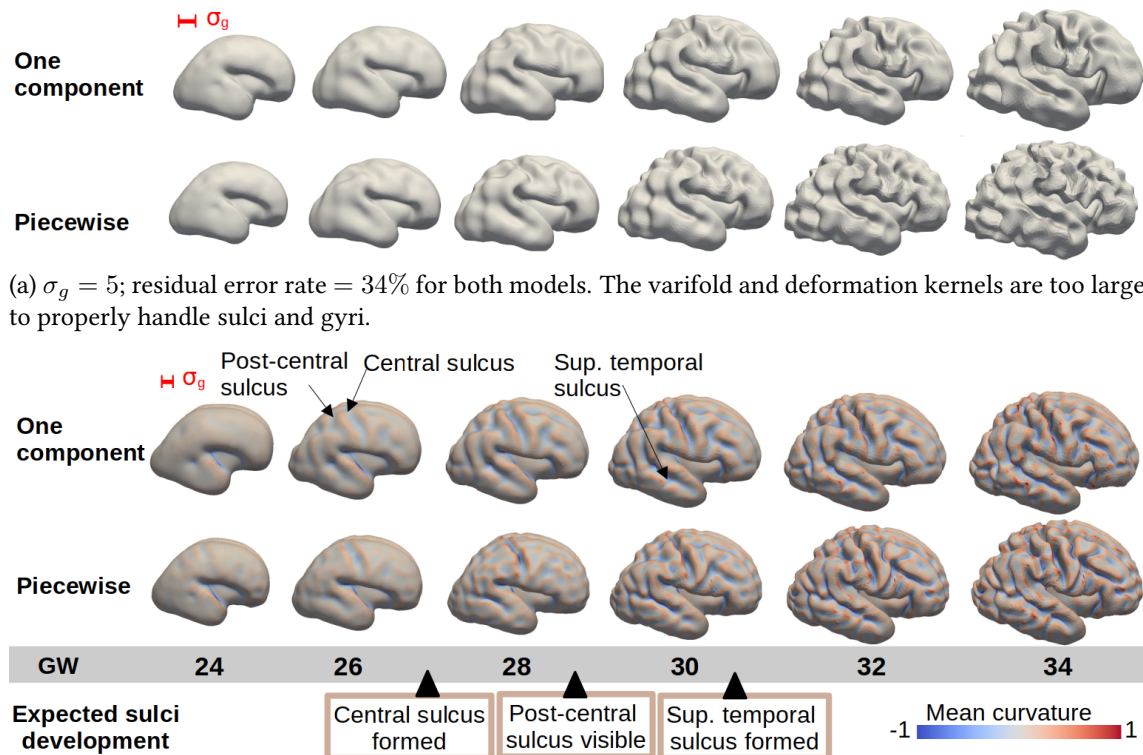
σ_g	k_g	Algorithm	P	Atlas estimation	Regression	Registration
5	5,796	Original	1 4	4.3 ± 1.2	33.8 ± 1.6	8.8 ± 5.4
		Multiscale	1 4		33.3 ± 3.6	7.7 ± 5.0
4	11,088	Original	1 4	2.9 ± 1.2	34.1 ± 1.7	4.8 ± 3.6
		Multiscale	1 4		31.8 ± 1.6	4.2 ± 3.1
3.5	16,000	Original	1 4	5.4 ± 6.1	40.2 ± 2.6	8.6 ± 6.4
		Multiscale	1 4		39.7 ± 2.7	7.8 ± 4.6
		Original	1 4	3.4 ± 1.1	39.8 ± 2.5	6.6 ± 2.8
		Multiscale	1 4		37.8 ± 2.6	5.3 ± 2.7
		Original	1 4	1.4 ± 0.8	49.4 ± 7.6	8.8 ± 9.4
		Multiscale	1 4		44.0 ± 4.6	7.9 ± 8.4
		Original	1 4		42.3 ± 2.6	3.6 ± 2.4
		Multiscale	1 4		40.3 ± 2.6	2.8 ± 2.5

Piecewise versus one-component trajectory. As deformation of surface meshes are easy to visualize, we first compare one-component versus the piecewise trajectories associated to multiscale optimization. Figure 4.8 presents the one-component and piecewise trajectories with the lowest residual error, estimated with $\sigma_g = 5$ (panel (a)) and $\sigma_g = 3.5$ (panel (b)). For the coarsest parameterization, few differences between the deformed templates can be seen, although piecewise regression generates a more precise delineation of sulci and gyri at 34 gestational weeks. With a smaller kernel (i.e. $\sigma_g = 3.5$), gyrification is modelled in a more accurate manner. Comparing the emergence of the major sulci with what is described by Habas et al. [88], we notice that piecewise regression has a more accurate gyrification timing. Indeed, with one-component regression, cortical folding is clearly too advanced, with prominent sulci visible up to 3 weeks ahead their expected emergence .

Multiscale versus original strategy. Since finer kernels seem to model cortical folding more accurately, we compare in Figure 4.9 examples of piecewise trajectories estimated with the original and the multiscale strategies with $\sigma_g = 3.5$. Panels (a) and (b) show the deformations with the highest and lowest performance gaps between the algorithms, respectively. Interestingly, the original optimization strategy generates unnatural deformations of the cortex in both cases, although in panel (b) these barely impact the residual error. The multiscale strategy generates a plausible trajectory of cortical folding. We also notice a more accurate timing of fold emergence, e.g. the central sulcus, supposed to be visible at 24 weeks [88], is absent from the original but not the multiscale templates.

4.5 Discussion

In this chapter, we extended our previously introduced coarse-to-fine optimization strategy based on a multiscale representation of the velocity fields in the LDDMM framework. We presented a new coarse-to-fine strategy based on a multiscale representation of the



(a) $\sigma_g = 5$; residual error rate = 34% for both models. The varifold and deformation kernels are too large to properly handle sulci and gyri.

(b) $\sigma_g = 3.5$; residual error rate = 40% (one-component trajectory), 39% (piecewise trajectory). For the one component trajectory, the post-central sulcus is clearly visible at 26 weeks (3 weeks ahead what is expected); the superior temporal sulcus is very pronounced at 28 weeks, 3 weeks ahead the expected end of its formation.

Figure 4.8: Comparison of one-component and piecewise trajectories associated to multiscale optimization for $\sigma_g = 5$ (panel (a)) and $\sigma_g = 3.5$ (panel (b)). Folds with the lowest residual error difference are presented.

objects, combined it with the previous strategy and demonstrated increased performance on an atlas estimation task.

The presented multiscale strategy is versatile in the sense that it can be applied to several types of objects, combined to different deformation frameworks and to different Computational Anatomy models. Here, we extended its application from cross-sectional to *temporal* modelling. Combining the multiscale strategy to single-dynamic and piecewise trajectories, we demonstrated on toy data that multiscale optimization can model complex time-dependent changes in a more natural way.

We then focused on the complex task of modelling brain development during pregnancy using fetal brain volumes and cortical surface meshes. These experiments can be considered as preliminary findings of Chapter 7, where we will work more in-depth on modelling the fetal brain growth. Both experiments showed that the multiscale strategy combined to piecewise geodesic regression can enhance the modelling of the fetal brain growth, notably by estimating more accurate brain deformations and avoiding unrealistic local minima.

However, Section 4.4.2 also highlighted some limitations of the LDDMM framework,

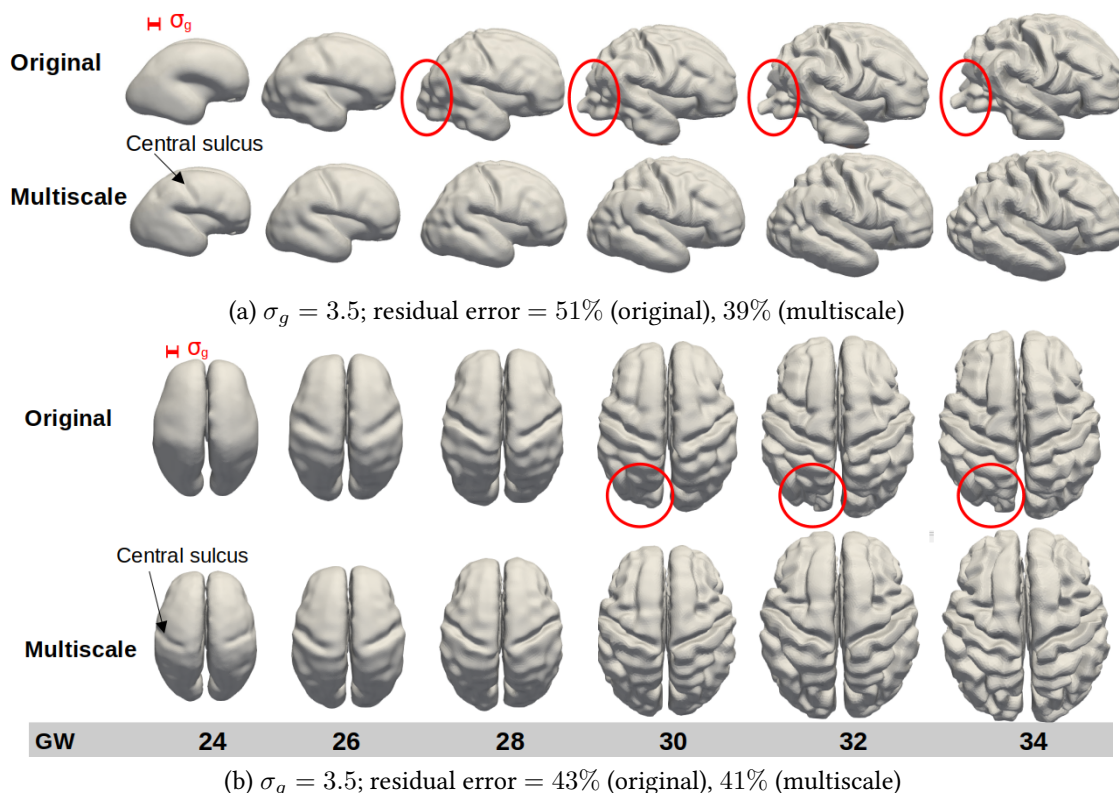


Figure 4.9: Comparison of the original and multiscale strategies associated to piecewise regression for $\sigma_g = 3.5$. Folds with the highest (panel (a)) and lowest (panel (b)) residual error differences are presented. Red circles indicate unnatural deformations.


namely the known inability of diffeomorphisms to generate new structures and the failure to properly model the cortical folding process. As regards the second point, whilst diffeomorphisms can in theory deform the cortical plate, a possible explanation is the use of the L^2 distance between images: since the L^2 norm grants the same weight to all voxels, the matching of large structures with little intensity differences penalizes more the distance function than the matching of thin structures with large intensity differences (e.g. incorrectly convoluted cortical plate). In the future, we could associate our multiscale strategy with algorithms that take into account structure-specific information, for example multi-structure LDDMM [120], which adds penalty terms for mismatched segmented structures in the cost function. Another solution could be to use the metamorphosis model, a derivative of the LDDMM framework which allows variations of the template image intensities [234].

In contrast, Section 4.4.3 demonstrated that our deformation model combined with multiscale optimization is more suited to the modelling of cortical folding than to the modelling of whole-brain intensity changes. This result was to be expected since the LDDMM framework has been successfully applied to surface meshes, notably subcortical structures, numerous times [177, 176, 45, 175, 82]. However, it must be noted that the cortex is a more complex structure than say, for instance, the hippocampi. As of the emergence of tertiary sulci in late gestation, cortices display high inter-subject variability, which makes estimating an average object more challenging. As illustrated in

Figure 4.8, the model parameters, in particular the regularizing kernel K_g , are likely to influence greatly the accuracy of the cortical folds. Further, a prominent question is that of model evaluation: like the L^2 distance between images, the varifold distance is an insufficient estimate of how well two cortices are matched, or how accurate the gyrification process is represented. The questions as to how properly model and evaluate fetal cortical development will be the subject of Chapter 7.

Part II

Exploring healthy and pathological fetal brain growth



THE QUANTITATIVE ANALYSIS OF FETAL BRAIN IMAGES - AN OVERVIEW

Contents

5.1	The growing field of fetal MRI	73
5.1.1	Basic principles of MRI	73
5.1.2	A brief history of fetal MRI	73
5.1.3	Use in a clinical setting	75
5.1.4	Challenges	76
5.2	Fetal MRI Preprocessing	77
5.2.1	Brain extraction	77
5.2.2	Volume reconstruction	78
5.2.3	Tissue segmentation	79
5.3	Quantitative analysis	80
5.3.1	Volumetric analyses	80
5.3.2	Cortical folding	81
5.3.3	Spatiotemporal atlases	83

In this chapter, we will introduce fetal brain imaging by reviewing the basic principles of MRI, the history of fetal brain MRI and its current clinical usage. Then, we will focus on the quantitative analysis of fetal brain MRI. We are not interested in providing an exhaustive overview of the research centered around fetal brain MRI. Rather, we seek to adopt the perspective of a computational imaging researcher and focus on questions of data availability, methodology, technical challenges, and discuss how these elements have shaped the current state of the field.

5.1 The growing field of fetal MRI

5.1.1 Basic principles of MRI

MRI is a non invasive and non irradiating imaging technique. The principle is to measure, localize and plot a nuclear magnetic resonance signal originating from body tissues. The human body is made of multiple tissues with different chemical compositions, many of which containing hydrogen. The hydrogen atom has one proton, or nucleus, which rotates around its axis, generating a small magnetic field called a **spin**. Spins in the body have a random orientation and cancel each other out, resulting in a null overall magnetic field.

During MRI, a strong external magnetic field B_0 (of strength 1.5 or 3 Tesla in fetal MRI) makes the spins line up in the direction of the magnetic field, creating a longitudinal magnetization. Secondly, another magnetic field B_1 formed of radio frequency waves is generated perpendicularly to the first one. Excited by the radio frequency pulse, the protons absorb energy, synchronise and spin in phase with each other. The resulting magnetic field tilts in the direction of B_1 . Finally, the radio frequency signal is stopped and the protons resume to their normal state under the magnetic field B_0 . This phenomenon, known as **relaxation**, causes a radio wave to be emitted. Receiver coils capture this signal, which is used to generate an image by means of a Fourier transformation.

The relaxation is made of two simultaneous realignment processes:

- the T1 relaxation along the longitudinal axis when the high energy protons fall back into a low energy state and realign with B_0 ;
- the T2 relaxation along the transverse axis, based on energy transfer between spins as the precessing protons dephase out and move apart.

The speed of the T1 and T2 relaxations differs based the local environment of the protons. For instance, protons in water continue to spin together longer than protons associated with fatty acid and have a longer T2 relaxation time. The protons' relaxation behavior can be modified by varying the parameters of the radiofrequency pulse sequence, namely the repetition time (TR) between two consecutive pulses and the echo time (TE) between a radio frequency pulse and signal acquisition. Different time settings lead to different types of contrast:

- **T1-weighted MRI**, characterized by low TR and low TE, has low signal for tissues with high water content (for instance, see Figure 5.1, year 1993);
- **T2-weighted MRI**, characterized by high TR and high TE, has high signal for tissues with high water content (for instance, see Figure 5.1, year 1998 to 2018).

5.1.2 A brief history of fetal MRI

Prior to the introduction of medical ultrasonography in the 1960s and MRI in the 1980s, knowledge of fetal anatomy came from *ex vivo* dissections of fetuses acquired from miscarriages, stillbirths or deceased pregnant women [257]. Compared to post-natal MRI,

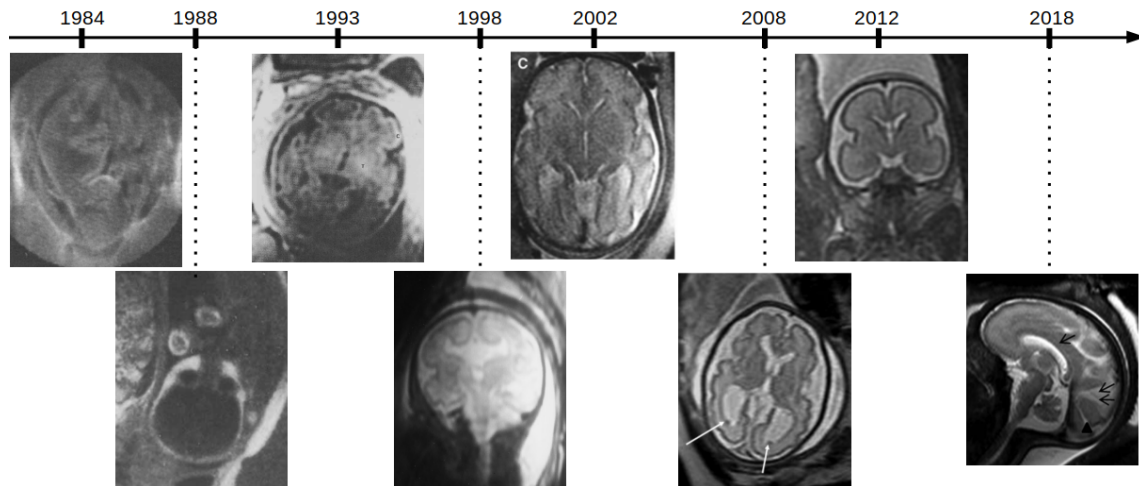


Figure 5.1: Evolution of fetal MRI over the years. **1984**: proton-weighted MRI of a 40 weeks fetus (whole body) from Johnson et al. [110]. **1988**: transverse section of a fetal brain, T1-weighted, from Powell et al. [172]. **1993**: axial view of a fetal brain at 35 weeks, T1-weighted, from Revel et al. [188]. **1998**: coronal view of a fetal brain with agenesis of the corpus callosum at 22 weeks, fast spin echo T2-weighted sequence, from Sonigo et al. [216]. **2002**: Axial image of a healthy fetal brain at 33 weeks, T2-weighted HASTE sequence, from Huisman et al. [100]. **2008**: axial slice of a fetal brain at 27 weeks with bilateral ventriculomegaly, T2-weighted from Rutherford et al. [196]. **2012**: coronal image of a 25-week-old normal fetus, single-shot fast spin-echo T2 sequence from Glenn et al. [81]. **2018**: 3 Tesla sagittal image of a healthy fetus at 33 weeks, turbo-spin echo T2-weighted sequence, from da Silva et al. [213]

the development of fetal MRI has been relatively slow (see Figure 5.1) and strewn with obstacles. The first clinical MR images date back to 1973 [129]. A decade later, Smith et al. [214] conducted fetal MRI for the first time in 6 women during early pregnancy. Long acquisition times (about 2 minutes for a single image), incompatible with fetal motion, meant degraded image quality, especially before the third trimester [110]. At the time, fetal MRI had no diagnostic quality -at least in terms of fetal health- and scans were performed mostly under research protocols [149, 150, 244, 172]. To reduce fetal motion, fetal sedation and curarization were often employed [42, 94, 216] but carried a risk of complications [188].

To make fetal MRI useful in clinical routine, the development of rapid sequences was mandatory. The first to be tested on fetuses was a real-time snapshot MR sequence referred to as Echo Planar imaging (EPI) [146, 147]. However, EPI could not be generalized due to lower resolution, high susceptibility to motion and field distortions [147], which limited its applications to volumetric measurements.

In the 1990s, the field was revolutionized by single shot fast spin echo sequences, i.e. high-quality sequences with a slice acquisition time of less than a second that mitigates in-plane maternal and fetal movements [37, 107]. The most popular of these sequences, Half-Fourier Acquisition Single-shot Turbo Spin Echo (HASTE), reached an acquisition time of 400 ms per slice [99]. This marked the introduction of fast multislice 2D images

which, combined, formed 3D-like images for the study of the fetal anatomy. Whilst early studies had targeted several regions of the fetal body (e.g. central nervous system, lungs, abdomen) [149, 150, 172], the potential of fetal MRI to refine or confirm US findings of brain anomalies soon became a central topic of interest [99, 216].

The field of fetal MRI has grown steadily ever since, but its clinical usage remained limited for a long time due to safety concerns and its elevated cost [217, 246]. It was only in the 2000s that clinical referrals for fetal MRI started to increase [187]. In 2005, a survey conducted reported that only 16 hospitals performed fetal MRI in the United Kingdom, at an average frequency of 17 MRI per year [246]. Only recently did scientific societies published guidelines and recommendations for fetal MRI [232]. Nowadays, fetal MRI is relatively more frequent (of the order of hundreds to a few thousands a year in European countries), but remains circumscribed to specialized centers [25, 26]. Of note, the recently developed high-field (3 Tesla) MRI scanners, which provide increased signal-to-noise ratio compared to 1.5 Tesla machines, are becoming more and more common [145].

5.1.3 Use in a clinical setting

An integral part of pregnancy monitoring, ultrasonography remains the primary evaluation tool of the fetus health [245, 107]. Fetal MRI is used as a complementary tool to confirm or refine a sonographic diagnosis and help parents make informed decisions. Compared to ultrasonography, MRI is less limited by maternal adipose tissue, amniotic fluid volume or fetal position, and visualization of the fetus intracranial compartment is not restricted by the ossified skull [197, 245, 237]. It also benefits from a larger field of view than sonography and multiple image orientations [245, 197]. Most importantly, ultrafast MR sequences capture images of higher quality, i.e. with elevated spatial and tissue contrast resolution, which allows to distinguish individual fetal structures as thin as the corpus callosum [197, 181, 144]. Examination is recommended from 18 gestational weeks [173]. Before this stage, excessive fetal movements, the limited size of anatomical structures, and limited development of the brain pose serious challenges [220, 257].

Fetal MRI is generally targeted at a specific anatomical region, namely the brain in 80% cases [144]. Following abnormal sonographic findings, common indications for prescribing a fetal brain MRI encompass ventriculomegaly, commissural abnormalities (e.g. agenesis of the corpus callosum), cortical malformations (e.g. lissencephaly), posterior fossa anomalies, infections of the central nervous system or other acquired pathologies like haemorrhage and tumors [144, 197, 245, 145]. It has high sensitivity and specificity for detecting cortical development anomalies, especially after 24 gestational weeks [81].

During a typical acquisition protocol, an initial localizer sequence is obtained to assess the fetal position and, if need be, reposition the fetal brain in the center of the coil [173]. Each acquisition serves as a localizer for the next sequences to avoid misregistration due to fetal movement [144, 148]. T2 images are acquired in the three orthogonal planes in an interleaved fashion to reduce signal loss [17]. Ultrafast imaging acquires 3-4 mm thick slices in less than a second, but the whole acquisition usually lasts from 30

minutes to an hour.

Ultrafast T2-weighted sequences such as HASTE, fast spin echo and single-shot fast spin echo are the most commonly used to examine the fetal brain [173]. Due to the elevated water content of the fetal tissues, T2 sequences provide high contrast between brain tissues and the cerebrospinal fluid and excellent visualization of the intracranial anatomy [148, 245]. T1-weighted sequences are primarily used to detect haemorrhage, fat deposition, and calcification, and EPI to visualize bony structures and calcification [173]. More advanced MR-based scanning techniques (e.g. diffusion weighted imaging, diffusion tensor imaging, functional MRI, MR spectroscopy) are out of the scope of this thesis. For a recent review about the different fetal MRI sequences, see [145].

5.1.4 Challenges

Fetal brain MRI differs drastically from neonatal or adult brain MRI and therefore carries its own challenges. Technical difficulties have hampered the use of automated, computer-assisted methods. As detailed in Section 5.2, an entire field of research is now dedicated solely to image preprocessing. The very nature of the developing brain brings its own difficulties, namely the presence of maternal tissues around the fetal brain, its unknown orientation, the limited size of fetal organs, which lowers image resolution, and the rapid and complex brain development, which limits inter subject comparison [79]. Artifacts present in postnatal MRI are also found -and often enhanced- in fetal ones [192]. Partial volume effects (i.e. the contribution of different tissues to a single voxel) are more common due to the small brain size and may vary during gestation as tissue composition changes [17, 192]. Bias field, or intensity non uniformity, occurs when the same tissue have different intensities in different slices [17, 217]. Aliasing artifacts, that appear when the field of view is smaller than the body part being imaged, are frequent in fetal MRI and require adjustment of the field of view [80]. Maternal and especially fetal movements are one of the greatest challenges in the field [181, 19]. Despite the fast acquisition of 2D images, in-plane motion artifacts cannot be entirely avoided, which creates intensity and spin artifacts such as ghosting and blurring effects [217]. Inter-slice motion artifacts and limited through-plane resolution impair the geometrical consistency of the whole image and prevent the direct visualization of a 3D anatomical volume [192]. This also makes the entire acquisition process highly operator-dependent and thus all the more costly [145].

Non-technical challenges are also worthy of mention, although they are not directly related to scientific research. There are worldwide disparities in the use of fetal MRI [3]. Even among European countries, the use of fetal MRI varies according to many factors such as national recommendations, health insurance coverage, judiciary pressure, and the presence of research institutions [26]. Unsurprisingly, the gestational age at which the exam is conducted strongly depends on time limits for pregnancy termination [26]. In Europe, it ranges from 18 weeks (in Sweden) to the end of pregnancy (in France, *inter alia*) [26]. In low and middle income countries, the lack of healthcare services and infrastructure, the limited training opportunities in radiology, along with more restrictive legislation about pregnancy termination restrict the number of performed scans [3].

Therefore, research in the field is also very restrained: only one paper about fetal MRI was published in low and middle income countries between 2000 and 2020 [209]. The topic of Caucasian-centered research was recently addressed by Wu et al. [249], who published the first fetal brain atlas from a Chinese population.

The above-mentioned difficulties, in particular elevated cost and technical issues, along with the relative novelty of fetal MRI and the sensitive nature of the data, explain why open-source data is lacking. On a positive note, the first publicly available fetal MRI dataset was published in 2021, providing researchers with 280 T2-weighted fetal brain reconstructions from four different institutions [168]. This dataset will be used in Chapter 7.

5.2 Fetal MRI Preprocessing

To perform quantitative analysis of any clinical images, some prior preprocessing steps are inevitable to obtain comparable data. This is all the more true when such data comes from prenatal brain MRI, as several obstacles specific to fetal MRI impede the direct analysis of the brain anatomy:

- the unknown orientation of the fetus with regard to the scanner;
- the presence of the placenta and maternal tissues around the fetal brain;
- as detailed in Section 5.1.4, the acquisition of 2-dimensional slices hampers the direct visualization of a volumetric image.

These specificities also prevent the application of sophisticated tools developed for postnatal MRI such as Freesurfer [63] and FSL [108] to fetal images. Further, postnatal preprocessing tools are mostly designed for T1-weighted images. This encouraged researchers in prenatal screening to develop their own techniques in order to obtain a high resolution volumetric brain image from motion corrupted stacks of 2D slices acquired in orthogonal orientations [34] Although these preliminary steps are not the subject of this thesis, they represent a requisite step for every study in the field. Thus, in the following sections we will present the different preprocessing methods for fetal brain MRI in a non-exhaustive manner.

Brain extraction, i.e. the task of delineating the fetal brain from the surrounding tissues such as the placenta, is often a prerequisite for volume reconstruction. As numerous studies have focused separately on either task, we shall review first brain extraction techniques, then volume reconstruction strategies.

5.2.1 Brain extraction

Template-based methods. Template matching was a rather popular method in the early days of fetal brain extraction. The first algorithm for this task selected the mid-sagittal slice by detecting the fetus eye with template matching, then isolated the brain on this 2D slice and subsequently on the 3D volume using a graph cut approach [9]. In Taleb et al. [222], each 2D slice was mapped to an age-specific template to estimate a

brain mask. Quality of the 2D brain extraction was assessed by computing a similarity measure between the estimated masks. A related approach [221] formulated the 2D rigid registration problem to a brain template as a block-matching problem. In Tourbier et al. [231], after manual brain localization, several brain templates were registered to each 2D slice and brain extraction was performed with a global voting strategy.

Learning-based methods. Other methods favored supervised approaches. A random forest classifier was trained to distinguish maternal from fetal tissues and the fetus orientation was recovered by estimating the position of the centroid of several tissues [104]. Keraudren et al. [118] localized the fetal brain by describing candidate regions with scale and rotation invariant features and trained a Support Vector Machine to delineate a bounding box around the brain. Then, slice-by-slice patch-based extraction was performed inside the bounding box to extract the brain [117]. Kainz et al. [111] computed rotation invariant descriptors of the 3D volume and trained a random forest classifier to produce a probability map of brain voxels, which was refined for final brain extraction. Alansary et al. [4] decomposed each 2D slice into superpixels, computed descriptors for each superpixel, and trained a random forest to generate a probability map of the brain. Recently, deep-learning methods have entered the field and shown constant performance increase [33]. Rajchl et al. [184] employed a Convolutional Neural Network (CNN) trained with weak annotations by non-expert raters. Khalili et al. [119] predicted brain masks from input 2D patches with a multi-scale CNN. Salehi et al. [198] reused a U-net CNN which showed high performance on 3D adult brain extraction, and adapted it to 2D fetal MRI slices. A two-step procedure whereby a first CNN performs coarse localization of the brain region, followed by fine brain delineation by a second CNN, was introduced in a state-of-the pipeline for brain reconstruction [56]. More recently, another brain extraction algorithm based on a U-net CNN trained on pathological cases [185] was integrated to the pipeline.

5.2.2 Volume reconstruction

Typical volume reconstruction techniques comprise two main steps [34]. First, motion correction is performed through Slice-to-Volume rigid Registration (SVR) in order to correct the discrepancy between positions of the 2D slices. The best alignment between each image and an arbitrarily chosen target slice is achieved by minimizing a similarity metric (e.g. Normalized Mutual Information, Cross Correlation, Mean Square Intensity) with gradient-descent optimization. Additional steps may include outlier rejection to identify and exclude highly motion corrupted slices. Then, Super Resolution Reconstruction (SRR) restores a 3D volume either by performing scattered data interpolation or solving an inverse problem with spatial regularization.

The majority of reconstruction algorithms perform SVR registration and SRR in an iterative manner. Pioneer strategies iterated between SVR registration and scattered data interpolation [195, 109]. As an alternative to SVR, Kim et al. [121] introduced a motion correction technique in which slices are collectively aligned by matching their intersections, followed by Gaussian-weighted interpolation. The following algorithms formulated the SRR inverse problem in a variational framework. Rousseau et al. [194]

expressed the SRR task as an inverse problem solved with Total Variation regularization. Similarly, Gholipour et al. [72] formulated the SRR task as a maximum likelihood error norm minimization problem and performed outlier rejection to reduce the weight of motion-corrupted voxels and slices. Building on this idea, another algorithm [47] iterated between a SVR approach similar to that of Rousseau et al. [194] and a Bayesian formulation of the SRR problem with a complete outlier rejection scheme based on the Expectation-Minimization algorithm. A GPU support of this algorithm was developed and complemented with an automatic detection of the slice with least motion to serve as reference during SVR [112]. Another study [229] further improved the SRR step by minimizing the Total Variation energy with convex optimization. Ebner et al. [56] estimated an initial high resolution volume using scattered data approximation, followed by a classical iterative registration-reconstruction procedure with robust outlier detection.

Recent years have also witnessed the emergence of machine learning approaches. A CNN was trained to predict a good initial slice alignment prior to the SVR registration, which proved robust for scans with extreme slice motion corruption [95]. A 3D CNN was used to upsample each low-resolution stack during the SRR step prior to SVR motion correction [151].

To perform the quantitative analysis of fetal brains, the availability and ease of use of the aforementioned preprocessing algorithms is of prime importance. Several authors have proposed fully automated reconstruction pipelines, mostly in recent years. The Baby Brain Toolkit [193] reconstructs fetal brains with non local denoising and previously introduced SRR algorithms [195, 194]. However, brain extraction is optional and performed manually. The open-source toolkit [112] performs volume reconstruction but does not include extraction nor reorientation of the fetal brain. The first complete pipeline [230] includes template-based brain localisation and extraction, along with reorientation in the standard radiological anatomical planes. Finally, the recently introduced state-of-the-art reconstruction pipeline NiftyMIC [56] offers brain localization and extraction, iterative SVR and outlier-robust SRR, and registration of the reconstructed brain to a template space.

5.2.3 Tissue segmentation

As we shall see in the following, segmentation of the brain into different tissue classes is almost always an initial step for quantifying the fetal brain volumetry and morphology. Automatic segmentation is essential to avoid time-consuming manual segmentation, but it is prone to errors due to the elevated age- and health-related anatomical variability in populations of fetuses.

Multi-atlas segmentation. Multi-atlas segmentation [73] registers several labelled template images to the target images and subsequently combine the warped labels, considered as probabilistic priors, using some label fusion procedure (e.g. the STAPLE algorithm [240]). The construction of atlases, i.e. images representative of the fetal brain at different gestational ages, will be the focus of Section 5.3.3. Weigl et al. [243, 242] introduced a semi-supervised method that simultaneously learns a spatiotemporal latent

atlas from a small number of segmented examples along with the subjects segmentations. Gholipour et al. [76] introduced a multi atlas multi-shape segmentation procedure tailored to segment similar tissues (e.g. the right and left ventricles). The algorithm used STAPLE to obtain shape priors from label fusion and a probabilistic shape model that penalizes intersections among structures. To tackle the low efficiency of multi-atlas segmentation on dissimilar subjects, a semi-supervised learning approach was introduced to propagate the atlas labels to the most similar image regions based on a voxel-wise graph interconnecting similar regions in all images [124].

Deep learning-based approaches. In recent years, many studies have addressed the challenge of fetal brain segmentation with deep learning methods [33]. Only a few [61, 98, 167, 169] explicitly focused on the segmentation of pathological brains. Some developed methods for the segmentation of specific structures, such as the cortical plate [51, 93, 59, 52] or the ventricles [169]. The popular U-Net architecture, which performs segmentation on 2-dimensional images, was used in several of these works [169, 93, 52] with adaptations such as a multi-view aggregation approach [93] and a topological loss function [52] to enforce morphological consistency of the 3D segmented structure. Other studies segmented multiple tissue classes using once again the U-Net architecture [167], its 3D extensions [114, 264, 62, 61] or a CNN-based encoder-decoder [98]. Novel training approaches were proposed, for example to handle missing annotations [62] or noisy (i.e. partially erroneous) labels [114] in the segmented images.

5.3 Quantitative analysis

In the last decade, the postprocessing techniques developed for fetal MRI have made possible to study quantitatively the growth of brain tissues. Our aim here is to provide the reader with an overview of the techniques employed, the structures of interest and the limitations of the studies that have quantified the evolution of healthy and abnormal brains in volume-reconstructed fetal images. For extensive reviews about the findings of such studies, see [17, 218, 181, 19, 192, 34].

The analysis of healthy fetal brains has mostly focused on measurements of specific brain volumes, quantification of the cortical folding process, and longitudinal analyses of brain evolution throughout pregnancy. We will also detail how these analyses have been applied to pathological fetal brains. So far, they have been restricted to few abnormalities and pathologies, including mainly ventriculomegaly, Chronic Heart Disease (CHD), and spina bifida. Studies focusing on agenesis of the corpus callosum have been reviewed in Section 6.2.

5.3.1 Volumetric analyses

As many congenital disorders can alter several brain compartments [179, 49, 126, 180], it is clinically relevant to analyze healthy and pathological volume changes of brain tissues. Volumetric studies typically perform manual or automatic brain segmentation in order

to compute tissue volumes, which are then incorporated in temporal models. Gholipour et al. [75] performed linear and quadratic fits of the total brain volume with age. The relationship between gestational age and several tissue volumes such as gray matter, white matter, and ventricles were tested using linear, quadratic and exponential models [206, 38]. Relying on a cohort of 166 fetuses, normative growth curves of several structures were built with quantile regression and used to compare hemispheric growth rates [8]. Some studies focused on a single structure such as the hippocampi [105] and the cerebellum [205], which were manually segmented to carry out inter-hemispheric volume comparisons and statistical modelling with age.

Numerous volumetric analyses have been performed on pathological fetal brain. Following manual or semi-automatic brain segmentation, regression models were employed to compare the temporal evolution of several tissue volumes between fetuses with CHD and healthy controls [35, 191] and between two subtypes of CHD [180]. Gholipour et al. [76] developed a multi-atlas based segmentation method robust to ventricular abnormalities and compared the reliability of ventriculomegaly diagnosis using a measure of ventricular volume versus atrial diameter. Two studies [126, 207] extracted the ventricles and several brain structures to perform volumetric comparisons between fetuses with ventriculomegaly and normal fetuses, with contrasted results. Rajagopalan et al. [179] assessed alterations of several brain tissue volume in fetal donors after laser surgery for twin-twin transfusion syndrome. The growth trajectories of several regional volumes and the cortical plate were compared between fetuses with Down Syndrome and healthy controls [226].

5.3.2 Cortical folding

5.3.2.1 Healthy fetal brains

The development of the cortical plate and the gyrification of the healthy fetal brain have been widely explored using quantitative measures of the cortical surface curvature, such as cortical thickness, mean curvature, gyrification index, and sulcal depth. These analyses usually rely on segmentation and reconstruction of the cortical plate, derivation of cortical folding measures at the global, regional or vertex/triangle level, and temporal modelling.

Global and regional analyses. Two studies [247, 250] averaged several folding measures over the whole brain and found several of these measures to accurately predict gestational age using linear [250] and Gompertz models [247]. Clouchoux et al. [36] computed the mean gyrification index and cortical surface area of each hemisphere and studied their association with gestational age. Hemispheric-level cortical folding measures were compared between fetuses and preterm newborns, revealing inter-group differences in cortical geometry and brain folding dynamics [130]. The gyrification index was used to investigate the emergence of sulcal and gyral patterns at the global and regional level [96]. The authors then used a pair of new measures to further quantify the evolution of gyral nodes and sulcal pits across pregnancy [97].

Local analyses. Habas et al. [88] computed a vertex-wise mean curvature index, built quantitative maps of cortical folding and showed a linear evolution of cortical folding between 20 and 28 gestational weeks. Schwartz et al. 2016 [203] modelled the evolution of cortical surface area at the triangle level using Gompertz models. Clouchoux et al. [36] introduced an algorithm to establish individual probability maps of sulci location based on mean and Gaussian curvatures.

Analysis of sulcal pattern. Sulcal pattern analysis, a method originally developed for postnatal brains, was used to identify sulcal basins in fetal brains [102]. Individual sulcal patterns were compared to template brains with spectral matching. The method was then used to derive a template-based automatic sulcus labelling pipeline, where the contribution of each template depends on the similarity of gyrification with the individual cortical surface [258]. The authors then studied the temporal emergence and spatial distribution of sulcal pits in the fetal period [260].

Deformation-based analyses. The cortical folding process was also explored using tensor-based morphometry, which computes the deformation of images with regard to a reference anatomy in order to capture local shape differences [182, 183]. This enabled to compute growth maps of the cerebrum and the cortical plate and study the directionality of volume changes at the voxel level. This work was extended by Pontabry et al. [171], who performed feature selection to extract sparse local deformation fields that provide information about which zones of the cortical plate undergo major changes across gestation.

5.3.2.2 Pathological fetal brains

Traditional curvature analyses were performed to identify differences in gyrification dynamics in healthy fetuses versus fetuses with CHD and ventriculomegaly. Linear regressions were used to compare the evolution of curvature measures between controls and fetuses with a severe form of CHD [35], highlighting gyrification delays in the pathological group. The volumes and local curvatures of the ventricles and cortical plate were computed in groups healthy fetuses and fetuses with ventriculomegaly [207]. Another study [16] identified several regions of delayed cortical folding related to ventriculomegaly. They later adopted a less typical approach based on manifold learning in order to find associations between the abnormal growths of the cortical and ventricular surfaces [15]. A common underlying representation of vertices belonging to both anatomies provided information about which cortical areas are related to specific dilated ventricular regions.

Sulcal pattern analysis was also applied to fetuses with brain malformations [103], agenesis of the corpus callosum [225], CHD [165], Down Syndrome [259] and ventriculomegaly [224]. Each study computed the sulcal pattern similarity between abnormal fetuses and healthy templates and assessed differences in position, depth and area of sulcal basins. Compared to healthy controls, all groups had reduced similarity to the reference patterns and the method showed higher sensitivity compared to traditional gyrification measures.

5.3.3 Spatiotemporal atlases

Several studies have modelled the evolution of normal brain anatomy across pregnancy by providing intensity brain templates and tissue probability maps for each gestational age [86, 162]. These models, termed atlases, have a large scope of application as they provide an insight into healthy brain growth and may serve as reference to segment new subjects and characterize abnormal brains. They differ in several regards: screening modality (only T2-weighted MRI is considered here), represented pregnancy stage(s), construction methodology, and regions of interest (cortical surface versus whole-brain intensity atlases). As of now, most spatiotemporal atlases have represented healthy brain growth, with the exception of Zhan et al. [262], who used post-mortem fetuses, and Fidon et al. [60], who modelled the evolution of brains with spina bifida.

5.3.3.1 Intensity atlases

The existing spatiotemporal intensity atlases of fetal brain growth are summarized in Table 5.2. The prevailing method for atlas building is made of two components: at each gestational age, a deformable template brain is estimated ("Spatial normalization" method) with a specific kernel-weighting of the subjects ("Temporal regression" method). Two main ways of estimating a template brain exist:

- *Groupwise registration*: given an initial template image, the template is iteratively registered to each subject image to obtain an average template-to-subject deformation, and this mean deformation is subsequently applied to the template, resulting in a new template image, which is refined by repeating the previous steps until convergence.
- *Pairwise registration* iteratively selects a reference subject in the dataset and registers each of the remaining $N - 1$ subjects to the reference. The $N - 1$ deformations characterizing a given subject are averaged to obtain a mean brain image for the reference subject, and the final template is obtained by averaging the mean shapes.

These methods have been utilized within different deformation frameworks, notably affine [87], free-form [208] and diffeomorphic [74, 73, 249, 262] deformations. Most atlas construction methods [73, 208, 242, 74, 249, 60, 254] combine each averaging step of the spatial normalization with a kernel-weighting of the subjects: the more a subject is close to the template gestational age, the more it will contribute to its estimation. Thus, while kernel regression can in theory compute a template at any gestational age, in practice the output atlas is discrete and can be temporally inconsistent [202].

Alternative methods for atlas construction have also been experimented. In an approach akin those developed in these thesis, Licandro et al. [134] built an atlas of the fetal brain using geodesic regression in a large deformation framework. Pei et al. [170] used unsupervised learning and relied on an atlas synthesis network and a deformable registration network to simultaneously learn the age-dependent template and register the individual images to the age-matched atlases.

Of note, only a few of these atlases are publicly available [60, 87, 208, 73, 262, 249, 235]. The links to download the templates are listed on the author’s website¹.

5.3.3.2 Cortical surface atlases

There are fewer cortical surface atlases of the fetal brain than intensity atlases. They are summarized in Table 5.1. In several papers [36, 88, 254], atlas building was not the primary goal of the study and served to perform quantitative analysis of the cortical folding process. Thus, only a few studies [248, 253] also produced segmentation maps of the developing cortex. The only publicly available cortical surface atlas is the one from Zhan et al. [262], which limits possible comparisons between different construction methods. Other open-source spatiotemporal atlases were built from postnatal images and are also listed on the author’s website.

Table 5.1: Fetal cortical surface spatiotemporal atlases in the literature. N: number of images used to build the atlas. Age range: gestational weeks spanned by the atlas. Bold style indicates atlases built from abnormal fetal brains.

Year	Authors	N	Age range	Spatial normalization	Temporal regression	Atlas
2011	Habas et al. [88]	40	20-28	Groupwise registration	None	Inner cortical surface templates
2011	Clouchoux et al. [36]	12	25-35	Groupwise registration	None	Inner cortical surface templates
2013	Zhan et al. [262]	34	15-22	Groupwise registration	None	Outer cortical surface templates
2015	Wright et al. [248]	80	23-37	Pairwise registration	Kernel regression	Inner cortical surface templates, parcellation maps
2019	Xia et al. [253]	25	26-29	Groupwise registration	None	Inner cortical surface templates, parcellation maps
2022	Xu et al. [254]	90	23-38	Extracted from intensity templates (see Table 5.2)		Inner cortical surface templates

¹<https://fleurgaudferneau.github.io/Datasets/>

Table 5.2: Fetal brain spatiotemporal atlases in the literature. N: number of images used to build the atlas. Age range: gestational weeks spanned by the atlas. Bold style indicates atlases built from abnormal fetal brains.

Year	Authors	N	Age range	Spatial normalization	Temporal regression	Atlas
2010	Habas et al. [87]	20	21-25	Groupwise registration	Temporal polynomial model	Intensity templates, probability maps
2012	Serag et al. [208]	80	23-37	Pairwise registration	Adaptative kernel regression	Intensity templates, segmentation maps
2013	Zhan et al. [262]	34	15-22	Groupwise registration	None	Intensity templates, cortical surfaces
2013	Weigl et al. [242]	32	20-30	Registration	Kernel regression	Segmentation maps
2014	Gholipour et al. [74]	40	26-36	Groupwise registration	Kernel regression	Intensity templates, segmentation maps
2016	Licandro et al. [134]	45	18-30	Geodesic regression		Intensity templates, segmentation maps
2017	Gholipour et al. [73]	81	21-38	Groupwise registration	Adaptative kernel regression	Intensity templates, segmentation maps
2021	Li et al. [133]	35	23-36	Pairwise registration	None	Intensity templates
2021	Wu et al. [249]	115	22-34	Groupwise registration	Adaptative kernel regression	Intensity templates, segmentation maps
2021	Pei et al. [170]	82	22-32	Registration network	Atlas generation network	Intensity templates, probability maps
2021	Fidon et al. [60]	37	21-34	Weighted average	Kernel weighting	Intensity templates, segmentation maps
2022	Urru et al. [235]	20	32-39	Groupwise registration	None	Intensity templates, segmentation maps
2022	Xu et al. [254]	90	23-38	Pairwise and groupwise registration	Adaptative kernel regression	Intensity templates, cortical surfaces

ANALYSIS OF THE ANATOMICAL VARIABILITY OF PATHOLOGICAL FETAL BRAINS

Contents

6.1	Motivation	86
6.2	Related work	88
6.3	Dataset	89
6.3.1	Data description	89
6.3.2	Image preprocessing pipeline	91
6.3.3	Final dataset	93
6.4	Shape analysis pipeline	94
6.4.1	Geometrical tools	94
6.4.2	Shape analysis pipeline	96
6.5	Results	100
6.5.1	Healthy brains versus brains with CCA	100
6.5.2	Complete CCA versus Partial CCA	101
6.6	Discussion	103
6.6.1	Fetal MRI preprocessing	105
6.6.2	Applying deformation models to fetal brains	105
6.6.3	The anatomical variability of fetal brains with CCA	106
6.6.4	Partial and complete CCA differentially affect brain anatomy	107
6.7	Conclusion	107

This chapter is an extended version of a paper published in the Perinatal, Preterm and Paediatric Image Analysis Workshop at MICCAI 2021 [68]: *Analysis of the Anatomical Variability of Fetal Brains with Corpus Callosum Agenesis* by Fleur Gaudfernau, Eléonore Blondiaux and Stéphanie Allassonnière at MICCAI 2021 - Perinatal, Preterm and Paediatric Image Analysis Workshop in October 2021.

6.1 Motivation

In this chapter, we are interested in applying diffeomorphic mapping to pathological fetal brains. Corpus callosum agenesis (CCA) is one of the most common congenital

brain anomalies, with a prevalence at birth of 0.02% [131]. It is characterized by the total or partial absence of the largest commissure of the brain, responsible for the transmission of sensory, motor and cognitive information between hemispheres [131]. Diagnosis is usually suspected during the second-trimester routine ultrasound, and confirmed by a MRI scan [131]. In complement with genetic screening, fetal MRI is valuable to provide clinicians with additional information, since the presence of other anomalies is the only consensual prognosis factor for neurodevelopmental delays [199]. In the presence of associated defects, accounting for 45% cases [199], the outcome is usually poor, with impairments affecting motor control, coordination and language [55]. Predicting the outcome is challenging in isolated CCA, where 20-30% children demonstrate a broad spectrum of cognitive deficits [199, 55], resulting in heterogenous medical counselling across hospitals and countries [48]. Further, CCA has a diverse clinical picture in the morphology of the corpus callosum: partial CCA, diagnosed when only some parts of the corpus callosum are missing, is considered less common, more delicate to detect and it is thus harder to make a prognosis in the absence of associated anomalies [92, 143]. To provide parents with informed counselling, it is crucial to identify anatomical markers linked to the neurodevelopmental outcome of children with complete and partial CCA as early as possible during pregnancy.

Quantitative *in vivo* analysis of fetal brains has long been limited by the scarcity of fetal MRI data and its restriction to 2D slices [34]. As explained in more details in Chapter 5, unpredictable fetal and maternal motion make the acquisition of 3D images challenging. With the advent of fast single shot multi-slice MRI sequences, combined with postprocessing techniques, it is now possible to acquire stacks of 2D images with reasonable in-plane motion, perform inter-slice motion correction, and reconstruct a high resolution volumetric image of the fetal brain [17]. Taking advantage of these recent developments, the quantitative assessment of normal and pathological brain development has attracted growing interest [17]. However, to this day, only few studies have investigated quantitatively anatomical alterations in fetuses with CCA [123, 225, 204], and their focus was on specific brain structures rather than global trends. Another limitation is the difficulty to compare fetal brains of different gestational ages, since they undergo rapid and drastic changes across pregnancy [73].

Whole brain shape analysis can provide information about which structures are impaired along with corpus callosum. To perform such global analysis, one can think of image registration, which maps a population average brain template onto individual images in order to measure a distance from normality. In a clinical setting, functions called diffeomorphisms are an appropriate choice for computing shape changes, as they are high dimensional, topology-preserving, and sensitive to small anatomical variations. The previously introduced LDDMM setting [233, 31] (see Chapter 2) is a powerful method for computing such functions, which are seen as geodesics on a Riemannian manifold. Diffeomorphisms can be efficiently computed through a discrete parameterization [54]. The LDDMM framework also provides geometrical tools such as parallel transport, which enables the comparison of subjects that are at different developmental stages. To the best of our knowledge, deformation models have never been applied to abnormal fetal brains.

In this chapter, we do not seek to develop or enhance mathematical models, but rather

to exploit already existing Computational Anatomy tools and show that they harbor much potential for analysing clinical images presenting high anatomical heterogeneity. We propose to explore the anatomical variability of fetal brains diagnosed with CCA by introducing a novel shape analysis pipeline based on diffeomorphic brain mapping and specifically tailored to the specificities of fetal MRI. Following data preprocessing, a healthy template is registered to each subject, and age-related differences between fetuses are corrected by transporting deformations to a common space. Finally, using Principal Component Analysis (PCA) and classification, we will identify deformations specific to fetal brains with CCA and study alterations differentiating brains with partial CCA from those with complete CCA.

6.2 Related work

Unsurprisingly, the brain anatomy of individuals with CCA has been mostly investigated in adults [12, 159, 166] and children [14, 91]. As far as studies focusing on the fetal brain are concerned, quantitative analyses remain scarce. Indeed, the field is geared towards clinical practice. Studies mostly aim at identifying and quantifying anatomical alterations related to CCA [71, 210, 115], or relating anatomical features to postnatal outcome [49, 78] with methods overwhelmingly based on visual assessment of stacks of two-dimensional MRIs by radiologists.

Quantitative studies of fetal brains with CCA (i.e. involving at least some degree of automation during analysis) are listed in Table 6.1. To the best of our knowledge, quantitative analyses of volumetric fetal MRI have been attempted in only three studies [123, 225, 204], with a focus on specific brain structures. Knezovic et al. [123] measured the hippocampal volume by means of manual delineation on volume reconstructed brains from 39 healthy fetuses and 46 fetuses with CCA. Statistical tests showed reduced hippocampal volume in impaired fetuses. Tarui et al. [225] investigated cortical folding in 17 controls and 7 fetuses with CCA. Following surface reconstruction of the inner cortical plate, sulcal developing basins were identified and matched to healthy sulcal templates, revealing abnormalities in sulcal positions. Schwartz et al. [204] also explored the cortical morphology of 46 fetuses with CCA and 22 healthy fetuses. Measures of cortical surface area, thickness, asymmetry and gyrification level were automatically computed in several areas to assess group differences, showing reduced cerebral wall thickness and structural asymmetry in several brain regions in fetuses with CCA.

Of note, two studies [116, 106] have analysed Diffusion Tensor Imaging (DTI) data, which maps the diffusion of water molecules within brain tissues and allows to reconstruct the 3D route of white matter fibers. Kasprian et al. [116] analysed DTI data from 20 fetuses with CCA and 20 age-matched healthy fetuses between 20 and 37 GW. Trajectories of abnormal white matter tracts, namely the Probst bundles and sigmoid bundles, were visualized in subjects with complete and partial CCA, respectively. Alterations in somatosensory and motor pathways were also highlighted. Using the same method, Jakab et al. [106] further showed that fetuses with CCA have aberrant organization of the brain connectome, notably decreased interhemispheric structural connectivity and




Figure 6.1: Brain T2-weighted MR images of fetuses from Hôpital Trousseau in sagittal view. In panels (a) and (c), arrows indicate the expected location of the corpus callosum. For fetuses with partial CCA (b), arrows indicate the missing part of the corpus callosum.

increased connectivity in intrahemispheric alternative information pathways.

Table 6.1: Quantitative analyses of fetal brains with CCA. Data: type of data used (either T2-weighted MRI or DTI). N: number of subjects included in the study (subjects with CCA / healthy controls)

Reference	Data	N	Study	Method
[116]	DTI	(20/20)	White matter connectivity	Tractography
[106]	DTI	(20/20)	Brain connectome organization	Tractography
[225]	MRI	(7/17)	Cortical folding	Semi-automated segmentation, sulcal pattern analysis
[123]	MRI	(39/46)	Hippocampal volume	Manual segmentation, volume measurement
[204]	MRI	(46/22)	Cortical folding, brain asymmetry	Semi-automated segmentation, quantitative measures

6.3 Dataset

6.3.1 Data description

Data consist of retrospectively selected fetal MRIs acquired between 2006 and 2019 at Hôpital Trousseau, in Paris, France, which is a nationwide reference center for CCA. The database contains 61 healthy fetuses scanned between 29 and 37 GW (mean = 32.40 ± 1.69) and 118 fetuses diagnosed with CCA scanned between 24 and 37 GW (mean = 31.88 ± 2.45), including 81 and 37 fetuses with partial and complete CCA, respectively. A histogram of the subjects gestational ages is presented in Figure 6.2a.

Healthy fetuses had normal central nervous system findings at MRI examination. In the group with CCA, the pathology was detected either at the second (between 22-24 GW) or third trimester (between 31- 33 GW) screening ultrasound examination, which

consistently checks for the presence of the corpus callosum. After findings of abnormality of the corpus callosum, the routine ultrasound was followed by an expert ultrasound assessment including neurosonography to investigate other associated fetal anomalies before MRI scanning. Patients had a consultation with a fetal medicine specialist, a pediatric neurologist and a clinical geneticist in the referral prenatal center and were offered amniocentesis for fetal karyotyping and chromosomal microarray. Since pregnancy termination is possible until the end of pregnancy in France and given that associated anomalies such as gyrfication delays are not visible at earlier stages of pregnancy, the fetal MRI is programmed around 30-34 GW following the pre-diagnosis step. Rather than examining solely the corpus callosum, radiologists look for indirect features of CCA, in particular the bundles of Probst, i.e. abnormal rearrangement of the callosal fibers visible only in complete CCA or large ventricles.

Corpus callosum anomalies were defined as: (1) *complete CCA*, namely the complete absence of the corpus callosum and (2) *partial CCA*, i.e. the absence of one or more of the five segments of the corpus callosum resulting in an abnormally shaped corpus callosum. Selection of the MRIs was performed by an expert radiologist. Inclusion criteria were as follows: fetuses affected by isolated or associated partial or complete corpus callosum abnormalities. Fetuses with a short corpus callosum, defined as a complete corpus callosum with an antero-posterior diameter below the third percentile, were not included in this study. Examples of fetal brain images with normal corpus callosum, partial CCA and complete CCA are presented in Figure 6.1.




Figure 6.2: Histograms of the subjects gestational ages at the date of MRI acquisition for the full dataset (a) and the volume-reconstructed images used for the analysis (b).

6.3.1.1 Image acquisition protocol

Fetal brain MRI was performed using repeated T2 half-Fourier Single Shot Fast Spin Echo (HASTE), or Single-Shot half-Fourier Turbo Spin Echo (SSFSE). MRIs were acquired on a 1.5 Tesla MRI system Achieva Philips (Best, the Netherlands) before 2016 and Optima

MR450w General Electric (Waukesha, WI, USA) after 2016. Maternal sedation was systematically offered to reduce fetal motion artefacts. The scanning parameters are specified in Table 6.2. To ensure proper volumetric reconstruction, only images with at least 3 stacks of slices in the three standard radiological planes are included in the analysis.

Table 6.2: MR acquisition parameters of the dataset.

In-plane resolution	0.8 or 1.6 mm
Field of view	256x256 or 512x512 mm
Echo time	150-200 ms
Repetition time	3,500-4,000 ms
Slice thickness	4 mm
Acquisition matrix	320 × 320
Flip angle	90°

6.3.2 Image preprocessing pipeline

In order to perform quantitative analyses, isotropic brain volumes are reconstructed from the 2D fetal images using NiftyMIC software version 0.7.5 [56]. As our dataset comprises routinely acquired clinical images, the brain extraction step preceding volumetric reconstruction often had a high rate of false positive voxels. To obtain accurate 3D brain volumes without resorting to manual correction, we designed a two-step preprocessing pipeline. As the newest version of NiftyMIC includes a more efficient brain extraction algorithm [185], it is important to note that our image processing pipeline is now obsolete. In the following, a brief overview of the pipeline will be given, and the interested reader is referred to Appendix C.1 for more details.

6.3.2.1 Semi-automated brain extraction and volume reconstruction pipeline

Brain extraction and volume reconstruction. Isotropic high resolution 3D volume reconstruction of fetal brains is performed using the open-source state-of-the-art NiftyMIC software that takes as input stacks of low resolution 2D slices. First, brain extraction is performed in each 2D stack with a coarse-to-fine approach that localizes and then extracts the fetal brain using two CNNs [57]. Intensity non-uniformity is corrected using N4 bias field correction. Finally, high resolution reconstruction is performed by iterating between SVR registration for motion correction and SRR with robust rejection of slices that are misregistered or artifact-corrupted.

Correction of erroneous brain extraction. Following these fully automated steps, a visual assessment of the reconstructed images is performed: as specified in Figure 6.3, 38% of the dataset was discarded because of image quality issues. Examples of excluded volumetric images are presented in Appendix Figure C.1. The remaining data comprises images with good reconstruction quality. However, the fetal brain was overdetected in 46% of brains with CCA and 24% of healthy brains. To avoid a tedious and time-consuming manual correction of the erroneous 3D brain masks, the following steps

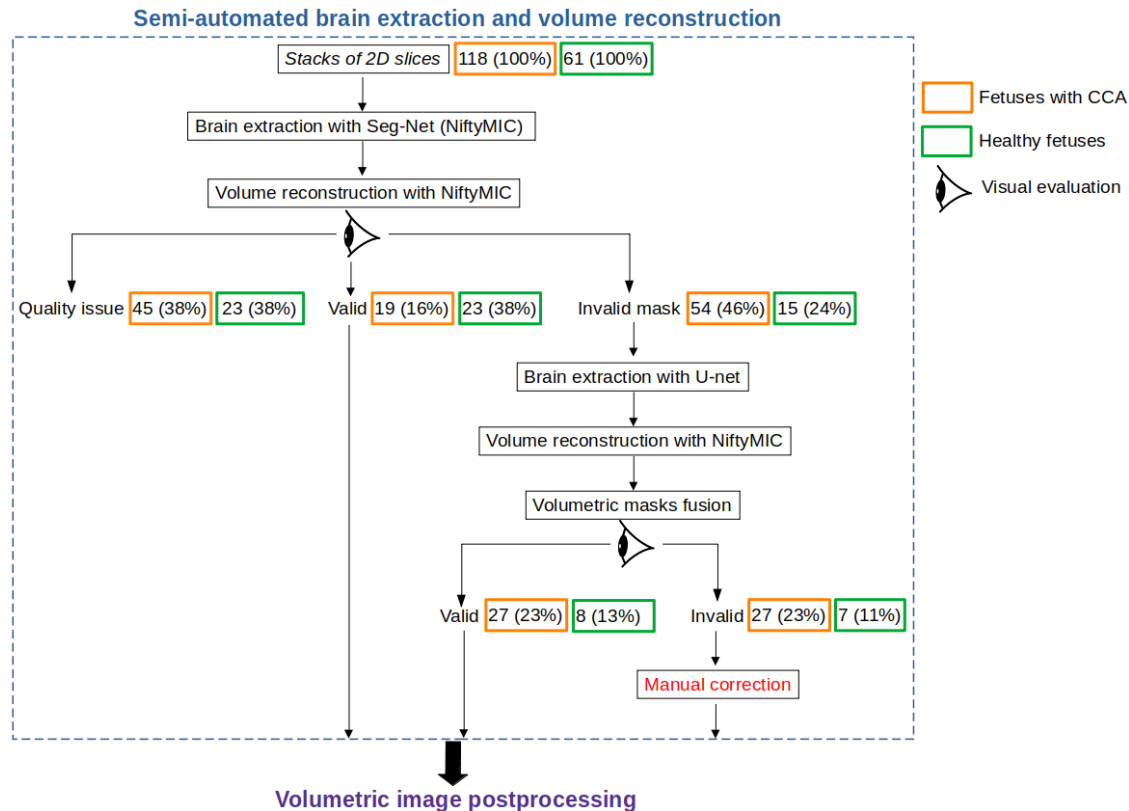


Figure 6.3: Steps of the semi-automated brain extraction and volume reconstruction pipeline. Values in orange and green rectangles denote the number and percentage of subjects going through each step of the pipeline in the group with CCA and the control group, respectively. Eye icons indicate visual evaluation of the quality of brain extraction and volume reconstruction.

are performed: brain extraction with a U-net CNN [198], volume reconstruction with NiftyMIC, fusion of the 3D masks obtained from NiftyMIC and U-net, and possible manual correction of the resulting brain mask.

The semi-automated volume reconstruction pipeline is further detailed in Appendix C.1.1.

6.3.2.2 Volumetric image postprocessing pipeline

As described in Appendix C.1.2, the volume-reconstructed brain image is post-processed to enable inter-subject comparison. As the brain mask of the reconstructed image may still contain voxels belonging to the skull, we first reorient the fetal brain: the coronal, sagittal and axial planes are automatically identified based on length and symmetry measurements and flipped in the right direction by minimizing the sum of squared differences between all possible orientations and an age-matched reference fetal brain from a state-of-the-art spatiotemporal atlas [73]. A correct brain mask is extracted from the reference brain, rigidly registered to each erroneous fetal brain mask, and used to re-mask the fetal brain. To enable inter-subject comparisons, fetal brains are transported

to a common anatomical space by performing affine registration to the previously used reference template. Finally, intensity normalization and histogram matching to the template brain are performed.

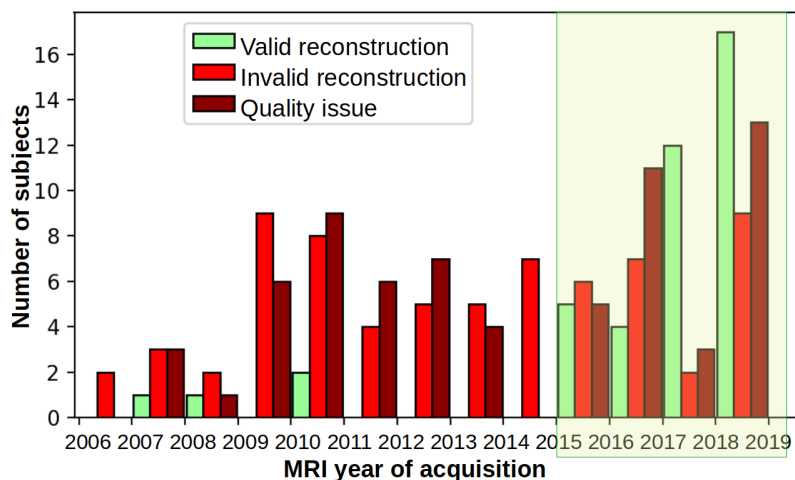


Figure 6.4: Histogram of the year of MRI acquisition for images with different reconstruction outcomes: valid brain extraction and reconstruction, invalid brain extraction (which was later corrected) and quality issue.

6.3.3 Final dataset

Among the 61 healthy fetal brains which underwent the preprocessing pipeline, 24% had incorrect brain masks (later corrected), 38% were discarded because of significant inter-slice motion and 38% had valid reconstructed images. 46% of fetuses with CCA had incorrect brain masks, 38% were discarded because of significant inter-slice motion and only 16% had valid reconstructed images. After semi-automated correction of erroneous brain extraction, 62% of the images in both groups had valid brain volumes which were used in the analyses described in the next section.

Gestational age at the time of acquisition did not seem to impact the reconstruction outcome: fetuses with valid brain volumes had a mean age of 32.35 ± 2.02 weeks, fetuses with invalid brain extraction a mean age of 31.87 ± 1.91 weeks and those with poor quality volumes a mean age of 32.23 ± 2.22 weeks. On the contrary, the MRI year of acquisition seemed to play a significant role in the success of volume reconstruction, with valid reconstructions having been imaged mostly between 2015 and 2019 (see Figure 6.4).

The final dataset comprises 111 MRIs from 38 healthy fetuses scanned between 29 and 37 GW (mean = 32.39 ± 1.69 GW) and 73 fetuses diagnosed with CCA scanned between 26 and 37 GW (mean = 31.87 ± 2.06 GW). In the group with abnormal corpus callosum, 51 fetuses have partial CCA and 22 complete CCA. A histogram of the subjects gestational ages is presented in Figure 6.2b.

6.4 Shape analysis pipeline

6.4.1 Geometrical tools

Our image processing and shape analysis pipeline are based on specific shape analysis tools developed in the LDDMM framework. Here, we provide a description of the theoretical setting behind these geometrical tools. More details are available in Chapter 2.

6.4.1.1 The LDDMM framework

As described in Section 2.3, the LDDMM framework [233, 31, 153] is a mathematical setting to compute diffeomorphic shape transformations. A flow of diffeomorphisms is considered as a Riemannian Manifold of infinite dimension, and shapes are seen as objects on that manifold, transformed through deformations of the whole ambient space.

Diffeomorphisms are constructed by integrating time-dependent vector fields, considered as infinitesimal linearized deformations. Namely, to build a flow of diffeomorphisms ϕ_t , one integrates the flow equation, which describes the motion of a particle x along the curve $x(t)$ under the effect of a time-varying vector field v_t :

$$\begin{cases} \frac{dx(t)}{dt} = v_t(x(t)) \\ x(0) = x_0. \end{cases} \quad (6.1)$$

This model builds a flow of diffeomorphisms $\phi_t : x_0 \rightarrow x(t) \forall t \in [0, 1]$. The diffeomorphism of interest ϕ_1 is the end point of the path $x(t) : \forall x_0 \in D, \phi_1(x_0) = x(1)$.

Here, we restrict ourselves to a finite parameterization of the velocity field v [54]. Namely, we impose that v belongs to a finite dimensional subspace of a Reproducing Kernel Hilbert Space (RKHS) V , i.e. the set of square integrable functions convolved with an interpolation kernel K_g :

$$v_t(x) = \sum_{k=1}^{k_g} K_g(x, c_k(t)) \alpha_k(t), \quad (6.2)$$

where $(\alpha_k)_k$ is a set of momentum vectors attached to k_g control points $(c_k)_k$, and K_g is a Gaussian kernel of width σ_g . K_g acts as a spatial regularizer restricting the range of deformations the model is able to express.

The Riemannian manifold is provided with a right invariant metric, defined as the total kinetic energy needed between the identity map I_d and the diffeomorphism ϕ_1 , i.e.:

$$d(I_d, \phi_1) = \int_0^1 \|v_t\|_V^2 dt$$

This means that distances are computed as the length of the minimal geodesic path connecting two elements. To ensure that the transformation ϕ_1 is diffeomorphic, we will impose regularity conditions so that the vector fields v_t define geodesic paths, i.e. the shortest paths between ϕ_0 and ϕ_1 according to the norm $\int_0^1 \|v_t\|_V^2 dt$.

It has been shown that if an initial velocity field v_0 is written in the following manner:

$$v_0(x) = \sum_{k=1}^{k_g} K_g(x, c_k(0)) \alpha_k(0),$$

then the vector fields v_t along geodesic paths of direction v_0 remain defined as a linear combination of RKHS basis elements [152]. In other words, Equation (6.2) is true for any time t .

Further, the kinetic energy along geodesic paths is preserved over time, i.e. $\forall t \in [0, 1], \|v_t\|_V = \|v_0\|_V$. Moreover, the evolution of the control point positions $(c_k(t))_k$ and momentum vectors $(\alpha_k(t))_k$ satisfy Hamiltonian equations that describe the movement of a set of particles, with K_g modeling the forces of repulsion between particles [152].

It follows that the vector fields along geodesics paths are fully parameterized by the initial velocity field v_0 . This has two main consequences:

(1) v_0 is the tangent space representation of the diffeomorphism ϕ_1 at the identity map I_d , which enables one to define tangent-space statistics of the diffeomorphism ϕ_1 and to characterize ϕ_1 with standard statistical tools such as PCA.

(2) to estimate ϕ_1 , one only needs to estimate the system's initial conditions $\alpha_0 = (\alpha_k(0))$ and $c_0 = (c_k(0))$. In other words, to compute shape changes, one solves a geodesic shooting problem: given a set of initial momentum vectors and control points characterizing a flow of diffeomorphisms, we compute the trajectory of a given shape. The end point of this trajectory, i.e. the final deformed image, is then compared to the target shape, and the initial conditions can be modified accordingly.

In this framework, optimization is performed by minimizing a cost function whose formulation depends on the chosen model. It is typically composed of a data fidelity term, i.e. the Euclidean L^2 distance between the images to be registered, plus a term penalizing the kinetic energy of the deformation.

6.4.1.2 Computational Anatomy models of interest

In the following, we briefly detail the geometrical tools used in this work, namely registration, geodesic regression and parallel transport (see Chapter 2).

Registration. Registration seeks the transformation ϕ_1 that best warps a source image I_{ref} onto a target image I_2 . The control points $c_{0,1}$ and momenta $\alpha_{0,1}$ that define

ϕ_1 are optimized by minimizing a cost function E :

$$E(c_0, \alpha_0) = \frac{\|I_2 - I_{ref} \circ \phi_1^{-1}\|^2}{\sigma_\epsilon^2} + \int_{t=0}^1 \|v_t\|_V^2, \quad (6.3)$$

where σ_ϵ controls the trade-off between the two terms. The first term is the sum of squared differences between the deformed source image and the target image. The second term penalizes the kinetic energy of the transformation, ensuring that only geodesic vector fields are considered as potential solutions.

Geodesic regression. Geodesic regression can be seen as the generalization of linear regression to shapes. Given a set of N images $(I_i)_{1 \leq i \leq N}$ observed at times $(t_i)_{1 \leq i \leq N}$, we seek the geodesic trajectory γ_t , defined by control points c_0 and momentum vectors α_0 that best fit the images. The cost function is defined as follows:

$$E(c_0, \alpha_0) = \sum_{i=1}^N \left(\frac{\|I_i - I_{ref} \circ \phi_{t_i}^{-1}\|^2}{\sigma_\epsilon^2} \right) + \int_{t=0}^1 \|v_t\|_V^2, \quad (6.4)$$

We refer to I_{ref} as the template image: it is the point from which the trajectory is computed. Here, I_{ref} is user-defined and can be any of the images in the set $(I_i)_{1 \leq i \leq N}$.

Parallel transport. Parallel transport is a differential geometry notion that considers two known transformations γ_t and ϕ_2 defined by the sets of parameters $(c_{0,1}, \alpha_{0,1})$ and $(c_{0,2}, \alpha_{0,2})$, respectively. Typically, γ_t describes the known evolution of a reference shape, and ϕ_2 describes the transformation that maps the reference subject at a given time to a new subject. Parallel transport enables one to transport the diffeomorphism ϕ_2 at any time point along the reference trajectory. We denote by $P_{\gamma, t_1, t_2}(\alpha_{0,2})$ the Parallel Transport of momentum vectors $\alpha_{0,2}$ along the trajectory γ_t from time t_1 to time t_2 . The geodesic shooting of the transported momenta can be used to predict the shape of a new subject at any time t_i , defining a trajectory that is parallel to the reference one. For details about the computation of Parallel Transport, the reader is referred to Louis et al. [135].

6.4.2 Shape analysis pipeline

Registering a reference average brain, called template, to healthy or pathological brains yields transformations that encode subject-specific anatomical deviations from normality. As brains undergo important structural changes during gestation, we compare each fetal brain to a healthy template of the *same age* using registration. To enable inter-subject comparisons, deformations are transported to a common space using parallel transport. PCA is applied to the transported subject deformations to reduce dimension and extract relevant features. Finally, these features are fed to a Support Vector Machine (SVM) to perform patient classification. The steps of our shape analysis pipeline are summarized in Figure 6.5 and detailed below.

Shape transformations, namely geodesic regression, registration and parallel trans-




Figure 6.5: Shape analysis pipeline. Red curve: average trajectory $\gamma(t)$ estimated by geodesic regression. Green curves: trajectory between a subject i and an age-matched template from regression, parameterized by the momentum vectors $\alpha_{0,i}$ (green arrows). Blue arrows: parallel transport of $\alpha_{0,i}$ from t_i to t_0 along the average trajectory.

port, are computed using the open-source software Deformetrica, version 4.3.0 [23]. Regression and registration are estimated with a kernel width and a spacing between control points $\sigma_g = 5$, corresponding to a number of points $k_g = 10,080$. Movies of the different steps of the shape analysis pipeline are available at the author’s webpage¹.

6.4.2.1 Geodesic regression of template brains

To take into account the anatomical changes that occur during gestation, each fetal brain is compared to an age-matched healthy brain. We use as reference a spatiotemporal atlas defined at each week of gestation, constructed from 81 healthy fetuses scanned between 19 and 39 GW [73]. We extract the 13 template brains between 26 to 38 GW and spatially normalize them to the space of the template at age 31 GW. From this discrete set of templates, we construct a *continuous* trajectory of normal brain changes from 26 to 38 GW by performing geodesic regression. This trajectory $\gamma(t)$ (red curve in Figure 6.5) is described by a pair of vectors, namely the control points c_0 and momenta α_0 defined at time $t_0 = 31$ GW. The point from which the geodesic is computed, i.e. the template brain at age 31 GW, has been chosen to minimize the amount of deformation that the template image undergoes. It will be referred to as T_{ref} in the following.

¹https://fleurgaudferneau.github.io/Shape_analysis/

6.4.2.2 Registration to an age-matched template.

Our dataset comprises $N = 111$ MRIs from 38 healthy fetuses and 73 fetuses with CCA. We denote by I_i the brain volume of the i^{th} subject and by t_i its gestational age. To characterize subject i , we extract the template brain from the geodesic trajectory whose age is closest to t_i and register it to I_i using geodesic shooting. Given an initial set of control points $c_{0,i}$ and momenta $\alpha_{0,i}$, geodesic shooting computes the trajectory of a shape under the flow of diffeomorphisms defined by $c_{0,i}$ and $\alpha_{0,i}$ (green paths). By comparing the deformed template image and the subject image, registration optimizes the $c_{0,i}, \alpha_{0,i}$ that best warp the template image to match the subject image. Note that the same set of $k_g = 10,080$ control points is used to estimate the geodesic regression and each registration.

6.4.2.3 Parallel transport

The diffeomorphism computed by registration encodes, for each subject, the distance between its anatomy and that of an age-matched healthy brain template. However, to enable comparisons between subjects, transformations need to exist in the same space. The momenta parameterizing each deformation are parallel transported to the tangent space of T_{ref} . In brief, parallel transport translates the deformation from the template brain towards subject i , defined by $c_{0,i}$ and $\alpha_{0,i}$, along the trajectory $\gamma(t)$ from the time point t_i (age at MRI acquisition) to $t_0 = 31$ GW (blue arrows). It adjusts for anatomical differences related to gestational age while preserving components of the transformation non-related to age. In other words, parallel transport provides a way of artificially transporting the subjects anatomies to the *same gestational stage*. We denote by $P_{\gamma, t_i, t_0}(\alpha_{0,i})$ the momentum vectors transports from t_i to t_0 along γ .

Note that we chose to transport the subjects momentum vectors to the space of T_{ref} in order to explore the anatomical variability of fetal brains with CCA at this specific stage of pregnancy, but we could have chosen any time point along the trajectory $\gamma(t)$.

6.4.2.4 Principal Components Analysis

After parallel transport, we obtain $N = 111$ high-dimensional transformations ($d \times k_g = 30,240$) living in the same Euclidean space. From this dataset, we now want to extract *directions* of deformation that can discriminate between healthy fetuses and fetuses with CCA. Given the high dimension of the transformations and the low sample size, the momenta cannot directly be used as features to predict whether or not a fetal brain has CCA. To reduce feature space and extract interpretable deformation modes, PCA is applied to the subjects deformation fields.

We denote by β_i the transported momentum vectors of subject i of dimension $3k_g$ and we introduce the N by $3k_g$ matrix of transported momenta:

$$X = (\beta_1, \dots, \beta_N)^T.$$

The empirical mean of the transported momenta is given by $\bar{X} = \sum_{i=1}^N \frac{\beta_i}{N}$. We compute the mean-centered matrix of transported momenta Z , defined as $Z = X - \bar{X}$.

The $3k_g$ by $3k_g$ empirical covariance matrix is given by

$$\Sigma = Z^T Z.$$

Then, eigendecomposition of Σ is performed in the form of

$$\Sigma = U \Lambda U^{-1}$$

in which U is a matrix of size $3k_g$ by N , whose columns (U_1, \dots, U_n) are the eigenvectors of Σ , and Λ a diagonal matrix of size $N \times N$, whose diagonal elements ($\lambda_1, \dots, \lambda_n$) are the eigenvalues of Σ .

Each eigenvector U_j is associated to an eigenvalue λ_j , representing the amount of variability that is explained by U_j . We extract the first M components that characterize 90% of the sample shape variability.

6.4.2.5 Deformation modes

Being a linear combination of momentum vectors, each eigenvector can generate a diffeomorphism, called deformation mode, which represents how the template brain anatomy varies within the population.

The i^{th} mode of deformation is given by:

$$m_i = \bar{X} + c\sigma_i U_i$$

with $c \in [-4, -2, 0, 2, 4]$, $\sigma_i = \sqrt{\lambda_i}$, and U_i the i^{th} eigenvector.

In order to visualize the deformation mode m_i , we apply the generated diffeomorphism to the template brain T_{ref} using geodesic shooting. At 0σ , T_{ref} is deformed by the average deformation \bar{X} and this deformed brain is interpreted as the average anatomy of the population.

In T2-weighted fetal MRI, thinness and hypointensity of the corpus callosum make it difficult to discern. To make corpus callosum deformations visible, geodesic shooting is also applied to the template parcellation image as provided by Gholipour et al. [73]. In the original segmentation, 126 different anatomical regions are labeled. To help interpretation, we simplify the segmentation image by fusing labels together, resulting in a total of 21 anatomical labels: motor cortex, frontal cortex, temporal cortex, insula, cingulum, hippocampi, amygdala, occipital cortex, sensory cortex, parietal cortex, basal ganglia, thalami, lateral ventricles, brainstem, cerebellum, subthalamic nuclei, hippocampal commissure, cortical plate, fornix, white matter and cerebrospinal fluid.

The projection of the momenta of subject j on deformation i is computed as follows: $P_{\beta_j} = \beta_j^T U_i$. P_{β_j} can be seen as a score quantifying how much β_j is represented by the

i^{th} deformation mode.

6.4.2.6 Classification

The deformation modes associated to the highest eigenvalues are those that best explain the anatomical variability of the data. However, they do not only encode shape variations related to CCA, but also components of rigid registration correction and normal inter-subject variability. To identify which modes of deformation can best discriminate between healthy fetuses and fetuses with CCA, we perform classification with a SVM equipped with a radial basis function kernel, that receives as input the subjects scores on the deformation modes. The SVM parameters (width of the gaussian kernel and penalty) are tuned using grid-search. The dataset is randomly split into a training (70% of the data) and a test set (30% of the data) to perform 5-fold cross validation. To extract deformation modes specific to CCA, we perform forward feature selection: starting from an initial model with no input features, we train the model with each of the M principal deformations independently and keep the one that best enhances the model accuracy. This process is repeated iteratively until the addition of a new deformation does not augment the accuracy.

As we also want to extract shape deformations that differentiate between complete and partial CCA, the PCA is replicated with only the transported momentum vectors of subjects with CCA ($N = 73$). To identify the modes of deformation that best discriminate between these subtypes of CCA, the classification and feature selection procedure is replicated.

6.5 Results

6.5.1 Healthy brains versus brains with CCA

As described in Figure 6.6a, we extracted from PCA the first 67 components that characterize 90% of the sample shape variability. The final SVM-based classification model reaches a 86% ($\pm 7\%$) accuracy. As indicated in Figure 6.6b, feature selection selected 4 deformation modes. Interestingly, feature selection did not retain the first component of PCA, which accounts for 12% of the sample shape variability. Visual inspection of the related deformation mode, presented in Appendix C.2.1, indicates it corrects for brain misalignment and characterizes subjects with large ventricles.

We present in Figure 6.7 the second component, which drives most of the model accuracy (see Figure 6.6b) and captures the anatomical variability of the dataset. The segmentation image of the template brain is transformed by the second mode of deformation in directions -4σ and -2σ , on which healthy subjects generally score higher, and in directions $+2\sigma$ and $+4\sigma$, on which subjects with CCA generally score higher. Appendix Figure C.5 presents the same mode of deformation applied to the template brain intensity image. Complete movies of these deformation modes are available at the

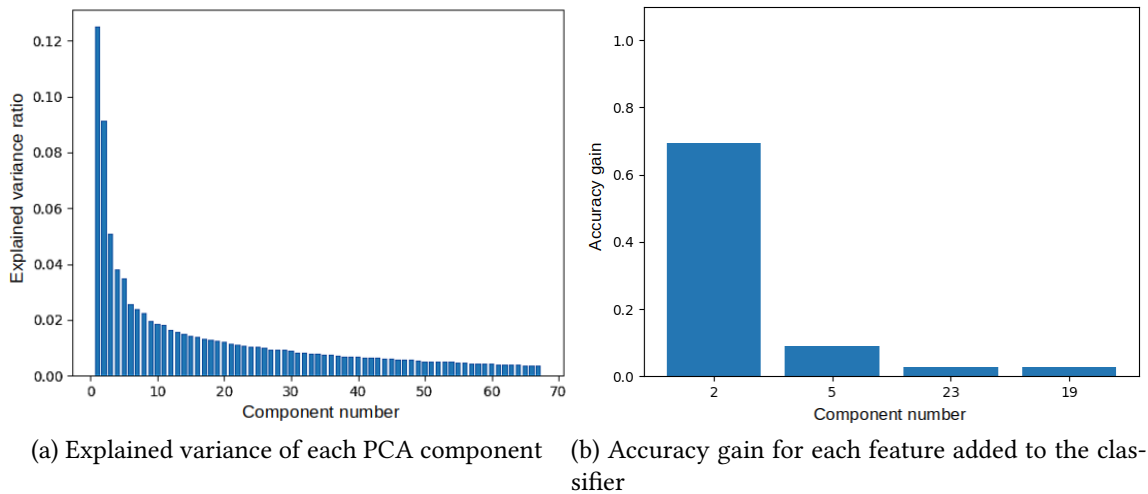


Figure 6.6: Results from PCA (a) and SVM-based classification (healthy fetuses versus fetuses with CCA) (b).

author’s webpage². As shown in Figure 6.7b, the score distributions of subjects with complete CCA is more spread out than that of control subjects. While healthy fetuses are mostly characterized by negative scores, fetuses with complete CCA reach a wider range of values.

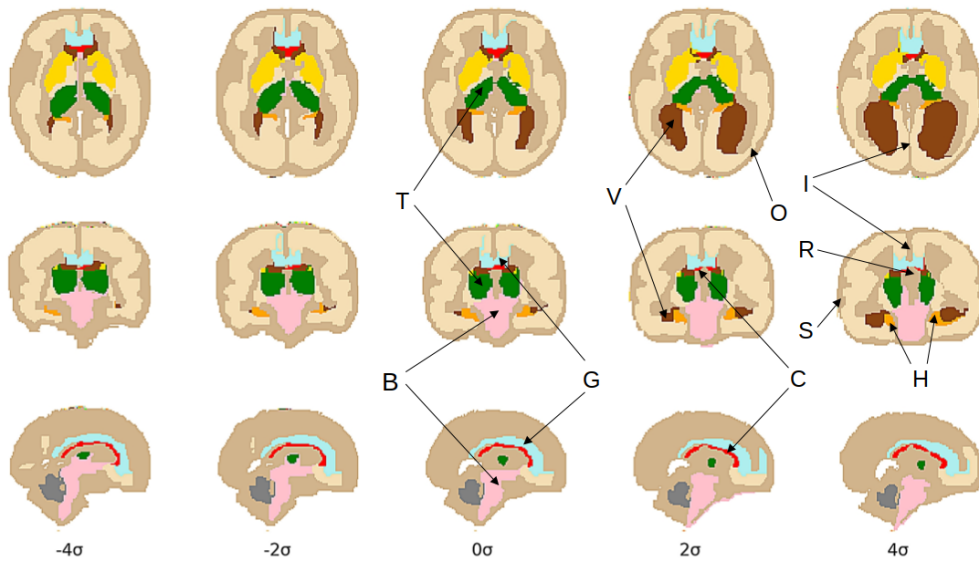
The direction of deformation that mostly characterizes subjects with CCA reveals a thinning and a shortening of corpus callosum (C) on sagittal view - despite a negligible volume change according to Figure 6.7c. It is folded into a V-like shape, with a stronger distortion towards its posterior part. The volume of the cingulate gyrus (G) is also reduced. Compared to the reference brain at 0σ , the volume of the lateral ventricles (LV) is increased by 74% (Figure 6.7c). They are widely spaced and parallel, with prominent occipital horns and atrium, corresponding to colpocephaly. Dilation is slightly stronger in the right ventricle. The volume of the occipital cortical and subcortical region (O) is reduced, especially in the right hemisphere. On axial and coronal views, the hippocampi (H) appear thinner and verticalized. The superior temporal sulci (S) seem less pronounced. On coronal view, the thalami (T) are parallelized and displaced away from the interhemispheric fissure. The shape of the brainstem (B) is abnormal on sagittal view, with prominent pons and midbrain.

6.5.2 Complete CCA versus Partial CCA


As described in Figure 6.8a, we extracted from PCA the first 48 components that characterize 90% of the sample shape variability. The final classification model reaches a 90% ($\pm 7\%$) accuracy. The first component of PCA, which accounts for 13% of the sample shape variability, drives most of the classification accuracy (Figure 6.8b).

Figure 6.9a shows the first deformation mode applied to the segmentation of the


²https://fleurgaudferneau.github.io/Shape_analysis/



(a) Second mode of deformation applied to the segmentation of the template at age 31 GW



(b) Distribution of the subjects scores on component 2.



(c) Volume changes of several structures between 0σ and -2σ and 0σ and $+2\sigma$

Figure 6.7: Second mode of deformation extracted from PCA at -4σ , -2σ , 0σ , $+2\sigma$, and $+4\sigma$ (a), distribution of the subjects scores on the related component (b), volume changes of several structures at -2σ and $+2\sigma$ (c). In panel (a), the deformed template is presented in axial, coronal and sagittal views. B: brainstem. C: corpus callosum. G: cingulate gyrus. H: hippocampi. I: interhemispheric fissure. O: occipital cortex. R: roof of the third ventricle. S: superior temporal sulcus. T: thalami. V: lateral ventricles.

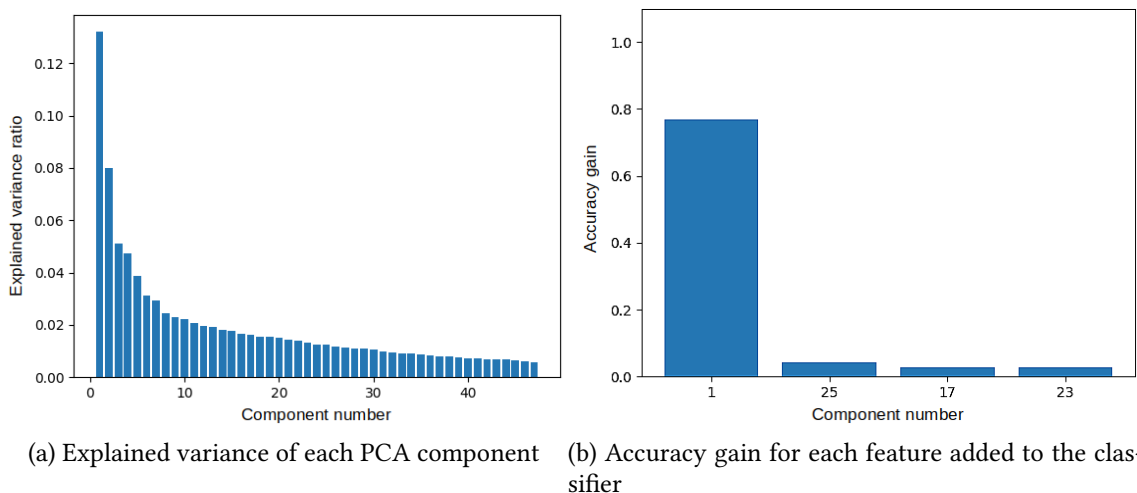



Figure 6.8: Results from PCA (a) and SVM-based classification (b) (fetuses versus with complete CCA versus fetuses with partial CCA).

template brain and Appendix Figure C.6 shows the same deformation applied to the MR image. According to Figure 6.9b, the direction $+\sigma$ characterizes subjects with complete CCA while the direction $-\sigma$ characterizes subjects with partial CCA. Once again, the distribution of the scores of subjects with complete CCA is more widespread.


The direction of deformation that mostly characterizes subjects with complete CCA reveals an important thinning of the corpus callosum (C) on sagittal view, along with a volume reduction of 8% according to Figure 6.9c. Similar to the "healthy subjects versus CCA" contrast, the lateral ventricles (V) are enlarged -though to a lesser extent, with a volume increase of 66%- indicating that subjects with complete CCA are more affected by colpocephaly. Once again, the volume of the occipital cortical and subcortical region (O) is reduced in the right hemisphere and the superior temporal sulci (S) is less pronounced. On coronal view, the thalami (T) appear thinner. Figure 6.9c shows that the whole brain, cingulum, hippocampi, basal ganglia, amygdala and thalami of subjects with complete CCA have reduced volumes compared to subjects with partial CCA.

6.6 Discussion


In this chapter, we addressed the challenge of exploring quantitatively alterations in abnormal fetal brains. We developed a shape analysis pipeline adapted to the specificities of clinical fetal MRI. Geometrical models based on diffeomorphisms, that were originally designed for postnatal imaging, enabled us to compare fetuses of different ages and investigate brain alterations *globally*, without requiring any prior assumption. Such models are adapted to the scarcity of medical data and to the need for interpretable results. This preliminary work opens new perspectives for the quantitative analysis of fetal brains with developmental alterations.



(a) First mode of deformation applied to the segmentation of the template brain at 31 GW



(b) Distribution of the subjects scores on component 1



(c) Volume changes of several structures between 0 and -2σ and 0 and $+2\sigma$

Figure 6.9: First mode of deformation extracted from PCA at -4σ , -2σ , 0σ , $+2\sigma$, and $+4\sigma$ (a), distribution of the subjects scores on the related component (b), volume changes of several structures at -2σ and $+2\sigma$ (c). In panel (a), the deformed template is presented in axial, coronal and sagittal views. B: brainstem. C: corpus callosum. G: cingulate gyrus. H: hippocampi. I: interhemispheric fissure. O: occipital cortex. R: roof of the third ventricle. S: superior temporal sulcus. T: thalami. V: lateral ventricles.

6.6.1 Fetal MRI preprocessing

In this chapter, particular attention was given to data preprocessing. As our analysis draws on whole-brain shape comparisons, the accuracy of brain extraction and alignment can impact the results and is thus of prime importance.

We developed a semi-automated volume reconstruction and image postprocessing pipeline in order to correct erroneous brain extraction and enable inter-subject comparison. While this pipeline is now obsolete, some lessons can be drawn from the problems encountered. This analysis exploited retrospectively selected fetal MRIs acquired during clinical routine. Although we used a state-of-the-art processing pipeline [56] for brain extraction and volume reconstruction, the brain extraction algorithm showed poor results on our dataset. The reconstruction task was less challenged by healthy fetal brains, with 60% of correct reconstruction following brain extraction by NiftyMIC, and only 18% of manual correction required. This may suggest that brain extraction algorithms are less robust to developmental defects. However, it should also be noted that healthy fetuses in our dataset were imaged more recently (i.e. after 2015) than fetuses with CCA, which likely led to higher image quality. It would be interesting to reiterate volume reconstruction with the newest version of NiftyMIC in order to increase the number of subjects and evaluate how the quality of preprocessing might impact the results.

6.6.2 Applying deformation models to fetal brains

In this work, we applied for the first time deformation models based on diffeomorphisms to abnormal fetal brains. One of the main advantages of this approach is that it provides a novel and practical way of dealing with the gestational age heterogeneity in datasets of fetal images by transporting subjects-specific deformations to a common space.

Another benefit of deformation models is they enable to target the whole brain. Hence, our method does not require tedious manual segmentations -nor automated ones, which are less reliable on abnormal brains [61]. This is in contrast with previous studies on brains with CCA [241, 123, 225, 159, 204], which often had fewer data and focused on specific brain structures. It is also important to note that unlike most papers, our analysis is not restricted to the study of volumetric changes. While our approach is related to tensor-based morphometry [182, 183], the latter only reflects *local* volume changes, while our pipeline also includes *global* transformations and therefore provides richer and more complex information about anatomical alterations.

The Computational Anatomy tools we employed come with several limitations. As registration was computed in the space of the healthy template brain using topology-preserving deformations, structures specific to brains with CCA such as Probst's bundles could not be studied. Moreover, parallel transport assumes that the speed of growth of impaired fetal brains is similar to that of healthy brains, which is in contradiction with reported growth delays for fetal brains with CCA [225]. Furthermore, the small spacing between control points yielded irregular deformations, that can be anatomically inaccurate. This issue has already been addressed in Chapter 3 and 4. It would be interesting

to replicate our analysis using the multiscale optimization strategy and study how the extracted deformation modes vary with the width of the regularizing kernel K_g .

Of note, the number of subjects in the control group was lower than in the group with CCA. Because the control group only included 38 subjects within a limited age range (29-37 GW), the reference trajectory for normal brain development was built using open-source template brains [73]. In the future, efforts could be made towards increasing the sample size in order to extract more robust features and define our own reference trajectory.

6.6.3 The anatomical variability of fetal brains with CCA

Our shape analysis pipeline extracted a mode of deformation that depicts the anatomical variability related to the health status of fetuses, i.e. normality of the corpus callosum. Replicating the analysis on fetuses with CCA only, we also identified a mode of deformation that characterizes complete versus partial CCA. It is important to note that global deformations that belong to the same mode of deformation *correlate* together, enabling us to identify sets of alterations that often co-occur.

The distribution of the scores of subjects with CCA on the second component of PCA was more widespread than that of healthy subjects. This might reflect the greater anatomical variability of abnormal fetuses compared to healthy ones.

The deformation mode characterizing healthy fetuses versus fetuses with CCA highlighted well-known defects of brains with CCA. As expected, the corpus callosum had abnormal shape and size. It was especially distorted in its posterior segment, which is usually the missing part in partial CCA [186]. The cingulate gyrus, commonly absent in CCA [14], was also reduced. CCA is often accompanied by the development of a pair of aberrant callosal fibers, called Probst bundles, that run parallel to the midline, and a rearrangement of the midline cerebral structures [131]. The most common alterations include colpocephaly [131, 14], which was clearly visible in the second mode of deformation. Ventricles dilation and volume reduction of the occipital cortical and subcortical brain matter were uneven across hemispheres, which may reflect a tendency for abnormal brain asymmetry [78, 204]. The observed volume reduction of the occipital region coincides with findings of decreased thickness of the cerebral wall in the lateral occipital region [204]. Consistent with findings of abnormal shape and rotation of the hippocampi in fetuses with CCA [78, 123], we observed verticalized hippocampi, probably because of the extension of the temporal ventricular horns into the parahippocampal gyri. Both observations might be related to reduced volume of the ventral cingulum bundle, the fibers of which normally have an initial course below the body of the corpus callosum and then course within the parahippocampal gyrus in the inferior and medial temporal lobe [159]. We also observed underdeveloped superior temporal sulcus, which might be related to delayed sulcation [241] or altered cortical folding [225]. Verticalization and displacement of the thalami, which are not reported in the literature, probably result from the widening of the interhemispheric fissure. It has been suggested that in CCA other interhemispheric connections, such as indirect thalamic nuclei connections, sup-

ply the absence of callosal fibers [14]. Understanding whether the displacement of the thalami is a marker of the absence or presence of such indirect connections and related to neurodevelopmental outcome could help understand the differences in outcome of patients with apparently isolated CCA. Surprisingly, we observed a strong deformation of the brainstem, which is not a typical feature of CCA. This result likely originates from inaccurate segmentation of the brainstem during image processing, which tended to exclude the medulla.

6.6.4 Partial and complete CCA differentially affect brain anatomy

Partial CCA, considered less common than complete CCA, is also more likely to be underdiagnosed [143, 178]. Unsurprisingly, this form of the disease is also less studied [212]. The papers exploring the morphological anomalies linked to CCA often have a larger group of subjects with complete CCA than with partial CCA [159, 225, 123, 116, 106, 204], and sometimes study solely complete CCA [14, 241]. Among the quantitative analyses of fetuses with CCA, only Schwartz et al. [204] compared complete CCA with partial CCA.

In this study, we retrospectively collected a large dataset in which fetuses with partial CCA were predominant, enabling us to compare the two main forms of the disease. The findings of the previous analysis can be refined by examining the deformation mode characterizing fetuses with complete versus partial CCA. The deformation in the $+\sigma$ direction, characterizing the complete absence of the corpus callosum, revealed abnormal features, many of which consistent with the previous analysis. This suggests that common brain alterations known as indirect signs of CCA (e.g. colpocephaly) are more likely to be present in complete CCA. The distribution of the scores of subjects with complete CCA was also more widespread compared to those with partial CCA, indicating that they present a wider range of anatomical defects. These results are consistent with a recent study that found cortical morphological alterations to be more widespread and common in fetuses with complete CCA than in those with partial CCA [204].

This analysis reinforces previous findings, but also illustrates why complete CCA is easier to characterize than partial CCA. According to the two deformation modes we presented, there is a morphological gradient spanning from healthy fetal brains to brains affected by complete CCA. With less frequent and salient anatomical alterations, fetal brains with partial CCA stand in the middle. However, they do not benefit from a better clinical outcome [55, 157], hence the need to further characterize them.

6.7 Conclusion

In this chapter, we presented a novel shape analysis pipeline to characterize the anatomical variability of fetuses with abnormal corpus callosum. Together, our findings draw a typical profile of brains with partial and complete CCA, which is in agreement with the results of more local methods, validating our approach. Our method could help un-

derstand the mechanisms of the rearrangements linked to CCA, and, above all, identify the anatomical defects related to poor clinical outcome in isolated CCA. Indeed, recent studies suggest that children with isolated complete or partial CCA have some degree of subtle neurodevelopmental impairments [39, 157, 127]. The next step would be to correlate anatomical deformations during the fetal period to postnatal neurodevelopmental outcomes. However, this is a difficult task since follow-up mostly revolves around the more severely impacted individuals.

From a broader perspective, the tools we introduced here are promising for the depiction of healthy and abnormal fetuses and can be generalized to *any* dataset of fetal brain MRIs acquired at different time points.

A SPATIO-TEMPORAL ATLAS OF THE DEVELOPING CORTEX

Contents

7.1	Motivation	110
7.2	Data	111
7.2.1	FeTa dataset	112
7.2.2	Fondation Lumière	113
7.2.3	Image preprocessing	113
7.3	Methods	115
7.3.1	Deformation model	115
7.3.2	Model of cortical surface development	117
7.3.3	Statistical model	119
7.3.4	Optimization procedure	120
7.3.5	Kernel regression model	122
7.3.6	Performance assessment	123
7.4	Preliminary results	123
7.4.1	Atlases visual comparison	124
7.4.2	Atlas evaluation	125
7.4.3	Influence of the hyper-parameters	127
7.5	Discussion	128

This chapter describes ongoing analyses and will be converted into an article for submission. Preliminary results have been the subject of poster presentations at the French Colloquium of Artificial Intelligence in Biomedical Imaging 2023 and AI4Health Summer School 2023.

7.1 Motivation

Gyrification is the process by which the cerebral cortex evolves from a smooth surface into an arrangement of folds made of sulci and gyri. The increase in surface area without an equivalent increase in volume allows the cortex to pack a large number of neurons and connections, which is believed to be essential to many human-specific cognitive processes [137]. The folding process starts at the second trimester of pregnancy and accelerates during the third trimester [34]. The major gyri seen in adult brains are already present at birth [132].

As described in Section 5.3.2, investigating and modelling the gyrification process in pregnancy has been the focus of many studies. However, there is still a need for fetal and neonatal cortical surface atlases, i.e. reference surfaces that describe the average healthy cortical folding process and provide a common space for quantitative analysis. Whilst intensity atlases describing the appearance of the fetal brain in T2 MRI have become quite common [73, 249, 88, 36, 133, 60], fetal cortical surface atlases remain quite scarce and are not available to the research community (see Section 5.3.3.2). A noteworthy example of such an atlas is the one from Wright et al. [248], who built cortical surface templates between 22 and 37 gestational weeks.

In the neonatal period, more work on the subject has been reported [252, 20, 1], probably because of the higher data availability and the existence of open-source preprocessing tools such as the dHCP pipeline for extraction and reconstruction of the cortical surface in neonatal T1 MR images [140]. Wu et al. [252] built a spatiotemporal cortical surface atlas between 39 and 44 gestational weeks from 764 term-born neonates using group-wise spherical surface registration. Another atlas [20] was published in the same year, ranging between 36-44 weeks and later extended to 28-44 weeks. It was constructed using multimodal surface matching for spatial normalization and adaptive kernel regression for temporal regression. Adamson et al. [1] published a single-point cortical surface atlas at 40 gestational weeks generated using multimodal surface matching. Importantly, all these atlases are publicly available, which is not the case of their antenatal counterparts.

One striking characteristic of *spatiotemporal* brain atlases is that the time component is not explicitly integrated into the model. Some authors [252] simply compute independent template objects at different gestational ages. Others [248, 20] combine the template estimation (i.e. spatial normalization) step with kernel regression [43], which weights the contribution of each subject to the template based on its age. Kernel regression is a simple and highly popular temporal "regression" method, able to generate accurate template objects at a *given time*. Unfortunately, the method completely overlooks the process by which a template object at time n transforms into the next template at time $n + 1$. As a result, the so-called spatiotemporal atlas is just a mere collection of unrelated, discrete templates.

In this chapter, we propose to use a new method for spatiotemporal atlas building, which relies on complex Computational Anatomy models and on tools introduced earlier in this thesis. Namely, we propose to switch the focus from the estimation of age-

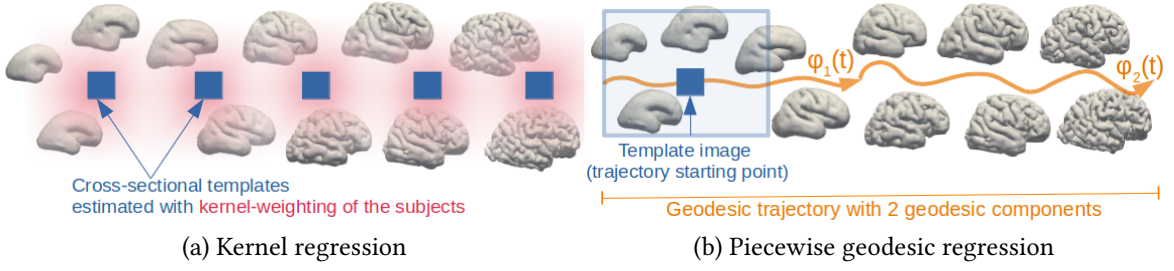


Figure 7.1: Comparison of the main atlas building method (a) with our proposed method (b). The cortical meshes represent training subjects of different gestational ages. Blue squares: estimated template shapes. Orange curve: geodesic trajectory.

dependent cortical surfaces to the estimation of how these surfaces evolve. This approach is much like building the growth curve of the height and weight of children, where we are more interested in the average growth dynamic and the positioning of each child with respect to this dynamic than in the mean value at a given time. The proposed modeling framework has the following characteristics:

- An average cortical surface and its developmental trajectory are estimated with geodesic regression, ensuring the temporal continuity of the atlas;
- The mean geodesic trajectory is *piecewise*, i.e. it can model changes in growth dynamics;
- The *spatial variability* around the average trajectory is also estimated, enabling us to study the time-varying anatomical variability of cortical folding and detect gyration anomalies;
- Estimation of the different model components (i.e. template shape, piecewise trajectory, and anatomical variability) is performed *simultaneously*, and not sequentially as in Chapter 6;
- Estimation of the model parameters is performed with the multiscale strategy introduced in Chapters 3 and 4.

A comparison of kernel regression versus our proposed model is available in Figure 7.1. Note that the spatial variability around the mean trajectory is not represented.

We use the proposed model to build a spatiotemporal atlas of the developing cortex between 20 and 37 gestational weeks. In Section 7.2, we present the two fetal MRI datasets used in this study and detail the extraction of cortical surfaces. In Section 7.3, we introduce the mathematical framework on which our model relies on, along with the statistical model and the optimization procedure. Finally, we present and evaluate the spatiotemporal atlas in Section 7.4 and discuss our results in Section 7.5.

7.2 Data

The data used in this study originate from two distinct datasets: the open-source Fetal Tissue Annotation and Segmentation Dataset (FeTa) dataset [168] and a French clinical

research platform (Fondation Lumière). A histogram of the gestational ages for the FeTa and Lumière datasets can be seen in Figure 7.2a.




Figure 7.2: Histograms of the subjects gestational ages at the date of MRI acquisition for the full dataset (a) and the extracted cortical surfaces used for the analysis (b).

7.2.1 FeTa dataset

The FeTA dataset¹ [168] consists of 120 high-resolution fetal brain volumes reconstructed from multiple Single Shot Fast Spin Echo (SSFSE) scans and manual brain segmentations into seven different tissues. The dataset comprises both normal and pathological brains. Images were acquired during clinical routine in two different institutions.

University Children’s Hospital Zurich. 80 acquisitions were performed between 2016 and 2019. Data was acquired using 1.5 Tesla and 3 Tesla GE whole-body scanners (Signa Discovery MR450 and MR750). T2-weighted SSFSE sequences were acquired with the following scanning parameters: in plane resolution 0.5 mm, field of view 200-240 mm, echo time 120ms minimum, repetition time 2000-3500 ms, slice thickness 3-5 mm, acquisition matrix 256x224 (1 Tesla) or 320x224 (3 Tesla) and flip angle 90°. Brain extraction was performed using a semi-automated atlas-based custom module. High-resolution reconstruction was performed with Tourber et al. algorithm [229].

University of Vienna. The data was acquired using 1.5 Tesla (Philips Ingenia/Intera, Best, the Netherlands) and 3 Tesla magnets (Philips Achieva, Best, the Netherlands). T2-weighted SSFSE sequences were acquired with the following scanning parameters: in plane resolution 0.65-1.17mm, echo time 80-140ms, slice thickness 3-5 mm. The preprocessing pipeline consisted of data denoising, in-plane super resolution, automatic brain masking step, isotropic motion correction and volumetric super-resolution reconstruction with NiftyMIC [56] and rigid alignment to a common reference space.

We selected 65 MR brain volumes from healthy fetuses in the FeTa dataset (32 from the University Children’s Hospital Zurich and 33 from the University of Vienna). The dataset does not comprise longitudinal data. Gestational ages range between 19 and 35 GW (mean = 27.96 ± 4.24).

¹<https://feta.grand-challenge.org/>

7.2.2 Fondation Lumière

Data consist of 87 MR brain images from healthy fetuses. The dataset does not comprise longitudinal data. Gestational ages range between 17 and 37 GW (mean = 27.48 ± 4.82).

Images were acquired in 2021 at the Plateforme Lumière in Paris, the first French clinical research platform dedicated to fetal imaging. Pregnant women were recruited as part of the research protocol aiming to assess the feasibility of advanced MR techniques in different clinical situations: normal pregnancy between 16 and 36 gestational weeks, normal pregnancy with screening difficulties, pathological pregnancy outside common indication for MRI and fetal virtopsy after pregnancy termination or *in utero* death. Women were usually screened only one time during pregnancy, although some of them underwent 2, 3 or even 4 MRI. The images selected for this study came from the group "normal pregnancy between 16 and 36 gestational weeks".

Acquisition protocol. Fetal brain MRI was performed using repeated T2 SSFSE sequences on Optima MR450w General Electric (Waukesha, WI, USA) at 1.5 Tesla. Acquisition parameters are indicated in Table 7.1. No maternal or fetal sedation was performed. As screening was part of a research protocol, image acquisition along a given axis could be repeated several times to increase image quality. The whole screening procedure could take up to 1 hour.

Table 7.1: MR acquisition parameters of the dataset.

In-plane resolution	1 mm
Field of view	512x512 mm
Echo time	100-150 ms
Repetition time	1,500-2,000 ms
Slice thickness	3.5-4 mm
Acquisition matrix	512 × 512
Flip angle	90°

7.2.3 Image preprocessing

The image preprocessing pipeline is summarized in Figure 7.3. It consists of high-resolution volume reconstruction and cortical surface extraction and reconstruction. Since images from the FeTa dataset are already volumetric, only acquisitions from the Fondation Lumière went through the reconstruction pipeline.

7.2.3.1 Volume reconstruction (Lumière dataset)

Isotropic high resolution 3D volume reconstruction of fetal brains is performed using the open-source state-of-the-art NiftyMIC software version 0.9.0. [56] (and not version 0.7.5 as in the previous chapter). NiftyMIC takes as input stacks of low resolution 2D




Figure 7.3: Steps of the preprocessing pipeline. Values in orange and green rectangles denote the number and percentage of subjects going through each step of the pipeline in the FeTa and Lumière datasets, respectively. Eye icons indicate visual evaluation of the quality of volume reconstruction and cortical surface extraction.

slices. First, brain extraction is performed using a U-net neural network [185]. Note that, unlike in Chapter 6, the brain extraction step led to mostly correct results. Intensity non-uniformity is corrected using N4 bias field correction. Finally, high resolution reconstruction is performed by iterating between SVR registration for motion correction and SRR with robust rejection of slices that are misregistered or corrupted by artifacts. This outlier rejection step depends on a threshold β (set to 0.8 here): slices whose similarity value with the estimated volumetric image is lower than β are considered as outliers. Then, we perform visual evaluation of the reconstructed volumetric images: when the number of outlier slices is too great, resulting in a brain volume corrupted with black strips, we reiterate the volume reconstruction step with $\beta = 0.5$. As indicated in Figure 7.3, only 23 (26%) of fetal brain images had correct volumes with $\beta = 0.8$. Lowering β to 0.5 allowed us to recover another 25 (29% of the dataset) valid brain volumes. A comparison of brain volumes obtained with $\beta = 0.8$ and $\beta = 0.5$ is available in Appendix D.1.

The output brain volumes were validated by an expert radiologist. Of the initial 87 acquisitions from Fondation Lumière, we obtained 48 (55%) valid volume-reconstructed

fetal brains. In Chapter 6, 62% of the data from Hôpital Trousseau had valid brain volumes after reconstruction with an earlier version of NiftyMIC.

7.2.3.2 Cortical surface extraction

Segmentation and cortical surface reconstruction. Brain segmentation is performed with the automatic segmentation DrawEM algorithm [139]. Then, the outer cortical surface (i.e. between the developmental cortical gray matter and the cerebrospinal fluid) and inner cortical surface (i.e. between cortical gray matter and white matter) are reconstructed using the dHCP structural pipeline [140], which was designed for neonates.

Manual correction. Validation of the reconstructed cortical surfaces is performed by an expert radiologist. In the majority of cases (77% and 65% of the Lumière et FeTa brain volumes, respectively), the reconstructed cortical surfaces are erroneous. This is not surprising since the dHCP pipeline is not designed for fetal brains; in addition, Figure 7.2 shows that younger fetuses are more likely to fail at cortical surface extraction. Whenever possible, erroneous meshes are manually corrected using ParaView [2] version 5.10.1 by flattening abnormal cortical protrusions with Laplacian smoothing. Details and examples of the performed corrections are available in Appendix D.2. For the remaining 48 invalid cortical surfaces, we will use another brain segmentation algorithm or manually correct the brain segmentation maps prior to cortical surface extraction.

7.2.3.3 Final dataset

The final dataset comprises in total 65 inner cortical surface meshes, 36 from FeTa (mean = 28.55 ± 6.11 GW) and 29 from Fondation Lumière (mean = 31.02 ± 3.61 GW). A histogram of the gestational ages of the extracted cortical surfaces can be seen in Figure 7.2b.

7.3 Methods

In this section, we present the geometrical model that allows us to compute the representative trajectory of cortical surface development.

7.3.1 Deformation model

The LDDMM framework. To compute shape transformations, we choose to work in the LDDMM framework [233, 31, 153], which was described in Section 2.3. In this framework, a shape can be warped by a diffeomorphic transformation constructed by integrating time-dependent vector fields $(v_t)_{t \in [0,1]}$ belonging to a Hilbert space V . More

specifically, a flow of diffeomorphisms ϕ_t is obtained by integrating the flow equation:

$$\begin{cases} \frac{\partial \phi_t}{\partial t} = v_t \circ \phi_t \\ \phi_0 = I_d. \end{cases} \quad (7.1)$$

The diffeomorphisms ϕ_t defined by the velocity fields v_t belong to a group \mathcal{G} which has the structure of a Riemannian manifold. Distances in \mathcal{G} are computed as the length of the minimal geodesic path connecting two elements via the metric:

$$d(I_d, \phi_1) = \int_0^1 \|v_t\|_V^2 dt$$

where I_d is the identity map and ϕ_1 the end point of the flow of the diffeomorphisms.

When trying to match two shapes O_1 and O_2 on the manifold, we will restrict the search space to the vector fields v_t that define geodesic paths in \mathcal{G} , i.e. the shortest paths between ϕ_0 and ϕ_1 according to the norm $\|\cdot\|_V$.

Finite parameterization of deformations. To simplify computations, we place ourselves in a finite dimensional setting by relying on a finite parameterization of the velocity field v_t [54]. Namely, we impose that v_t belongs to a finite dimensional subspace of a RKHS V :

$$v_t(x) = \sum_{k=1}^{k_g} K_g(x, c_k(t)) \alpha_k(t), \quad (7.2)$$

where $(\alpha_k)_k$ is a set of momentum vectors attached to k_g control points $(c_k)_k$, and the regularizing kernel K_g is a Gaussian kernel of width σ_g .

If an initial velocity v_0 at time $t = 0$ writes according to Equation (7.2), i.e. as a linear combination of initial momentum vectors $\alpha_k(0)$ and control points $c_k(0)$, then the vector fields v_t along geodesic paths of direction v_0 remain defined as a linear combination of RKHS basis elements and thus write according to Equation (7.2) [152].

Moreover, the evolution of the control point positions $(c_k(t))_k$ and momentum vectors $(\alpha_k(t))_k$ along geodesic paths satisfies Hamiltonian equations. Thus, to compute a time-dependent trajectory of cortical surfaces, we only need to define a shape O_{ref} , the initial set of momenta $\alpha(0)$ and the control points c_0 (fixed in our case). By integrating the Hamiltonian equations, we obtain the evolution of the momentum vectors over time and then the velocity field at any time t using Equation (7.2). Integration of the flow equation yields a flow of diffeomorphisms $(\phi_t)_{t \in [0,1]}$. Finally, ϕ_t can be applied to the object O_{ref} . We recall that these steps of the optimization procedure are illustrated in Chapter 2, Figure 2.3.

Metric between objects. In this chapter, we work with meshes with different numbers of vertices. To compute the fidelity term in the cost function, we must define a distance metric between surface meshes. To this end, we rely on the varifold frame-

work, in which meshes are embedded into a Hilbert space and algebraic operations and distances are defined. A varifold O is represented as the distribution of its points with unit normal vectors attached. Let W and its dual W' be Hilbert spaces. The varifold distance between two meshes O_1 and O_2 is defined as

$$\begin{aligned} d_W(O_1, O_2)^2 &= \|O_1 - O_2\|_{W'}^2 \\ &= \langle O_1, O_1 \rangle_{W'} + \langle O_2, O_2 \rangle_{W'} - 2 \times \langle O_1, O_2 \rangle_{W'} \end{aligned}$$

and the inner product between two meshes is given as:

$$\langle O_1, O_2 \rangle_{W'} = \sum_{k=1}^{m_1} \sum_{l=1}^{m_2} K_w(x_k^1, x_l^2) \frac{(n_k^{1T} n_l^2)^2}{\|n_k^1\| \|n_l^2\|}$$

where K_w is a Gaussian kernel of width σ_w , m_1 and m_2 are the number of cells (triangles) in the meshes O_1 and O_2 (respectively), $(x_k^1)_k$ and $(x_k^2)_k$ the cell *centers* of O_1 and O_2 and $(n_k^1)_k$ and $(n_k^2)_k$ the cell *normals* of O_1 and O_2 .

7.3.2 Model of cortical surface development

In this section, we introduce the mathematical model representing the average cortical development process during pregnancy. First, we propose to rely on a piecewise geodesic regression model, illustrated in Figure 7.4.




Figure 7.4: Example of a piecewise trajectory $\gamma(t)$ with $P = 3$ components. t_1 and t_2 denote the rupture times between the geodesic components ϕ_1 , ϕ_2 , and ϕ_3 that belong to the manifold of diffeomorphisms \mathcal{G} . $\alpha_{0,1}$, $\alpha_{0,2}$ and $\alpha_{0,3}$ are the initial momentum vectors generating the diffeomorphisms ϕ_1 , ϕ_2 , and ϕ_3 (respectively). Blue points denote the observed cortical surfaces, seen as points on a shape manifold \mathcal{S} .

Piecewise trajectory. In the following, we consider a data set of N subjects, where each subject is associated to a single observation O_j at time t_j : $(O_1(t_1), \dots, O_N(t_N))$. We seek to estimate a trajectory $\gamma(t)$ representing the average cortical surface development.

$\gamma(t)$ is formed of P geodesic components $(\phi^l(t))_{1 \leq l \leq P}$, separated by a set of rupture times $t_R = [t_1, t_2, \dots, t_{P-1}]$, with $t_1, < \dots, < t_{P-1}$.

$$\gamma(t)(O_{ref}^1) = \phi^1(t) \star O_{ref}^1 \mathbb{1}_{[t_{min}, t_1]}(t) + \sum_{l=2}^{P-1} \phi^l(t) \star O_{ref}^l \mathbb{1}_{[t_{l-1}, t_l]}(t) + \phi^P(t) \star O_{ref}^P \mathbb{1}_{[t_P, t_{max}]}(t) \quad (7.3)$$

where :

- t_{min} and t_{max} are the minimum and maximum of the observation times (t_1, \dots, t_N) , respectively
- the l^{th} diffeomorphism ϕ^l is parameterized by the set of momentum vectors $\alpha_{0,l}$ and the fixed control points c_0 :

$$\forall l \in [1, P] \quad \phi^l(t) = \mathcal{Exp}_{c_0, t_l, t-t_l}(\alpha_{0,l})$$

- $\forall l \in [1, P]$, O_{ref}^l is the representative shape at the rupture time t_l :

$$O_{ref}^l = \phi^{l-1}(t_l) \star O_{ref}^{l-1}$$

Note that the mean trajectory $\gamma(t)$ and thus the shape changes $\gamma(t)(O_{ref})$ are *continuous*: there is no need to define P representative shapes $(O_{ref}^1, \dots, O_{ref}^P)$. We only need to define a *single* reference shape O_{ref} . For the sake of simplicity, we decide that O_{ref} is defined at some time $t_0 \in t_R$: if $l \in [1, P]$ is the index of the rupture time at which the template shape is defined, then $t_l = t_0$ and $O_{ref}^l = O_{ref}$.

Space shifts. From the average trajectory $\gamma(t)$, we want to generate *individual* cortical surface trajectories. To achieve this goal, we borrow the idea of *space reparameterization* from spatio-temporal models [44, 22]. More precisely, we introduce for each subject $i \in [1, N]$ a space-shift momentum ω_i encoding the anatomical differences between the subject's observation $O_i(t_i)$ and the representative shape changes $\gamma(t_i)(O_{ref})$. In the same spirit as the shape analysis pipeline of Chapter 6, this enables us to:

- Obtain subject-specific trajectories of cortical development;
- Analyze the anatomical variability of the dataset at each gestational age.

To derive an individual trajectory given the representative trajectory $\gamma(t)$ and a space shift ω_i , we rely on Parallel Transport, a tool introduced in Section 2.2. We note $P_{\gamma, t_i, t}(\omega_i)$ the parallel transport of ω_i along the trajectory $\gamma(t)$ from time t_i to a given time t . We then compute the exp-parallelization of γ by ω_i (i.e. the geodesic shooting of the parallel transported vectors):

$$\eta_{\gamma}^{\omega_i}(t) = \mathcal{Exp}_{\gamma(t)}(P_{\gamma, t_i, t}(\omega_i)) \quad (7.4)$$

This procedure is illustrated in Section 7.3.2. Finally, the subject-specific cortical development is obtained by applying the deformation of the representative curve by the space shift ω_i to the template object: $\gamma_i(t) \circ O_{ref}$, with $\gamma_i(t) = \eta_{\gamma}^{\omega_i}(t)$.




Figure 7.5: Construction of an individual trajectory $\eta_\gamma^{\omega_i}$ from the representative trajectory $\gamma(t)$. **Left:** the space shift ω_i encoding the differences between $O_i(t_i)$ and $\gamma(t_i)(O_{ref})$ is orthogonalized with respect to the initial momentum vectors $\alpha_{0,2}$ and subsequently parallel transported along the curve γ . **Right:** exp-parallelization yields the trajectory of the subject i .

Further, we impose orthogonality between the space shift ω_i and the momentum vectors $(\alpha_{0,1}, \dots, \alpha_{0,P})$ defining the representative trajectory. This condition ensures that the space shifts encode only space-related reparameterizations and no time reparameterizations [200]: the individual trajectories γ_i cannot be accelerated/decelerated, nor temporally shifted compared to γ . Since we are studying a population of healthy fetuses, it makes sense not to allow abnormal development timelines such as growth delays.

Parallel Transport numerical scheme. In this chapter, we had to adapt the parallel transport numerical scheme of Louis et al. [135] to high dimensional parallel transport (i.e. transport of large numbers of momentum vectors). Details are available in Appendix D.3 and Appendix Figure D.3 illustrates parallel transport on an example.

7.3.3 Statistical model

Mixed effects model. We place ourselves within the framework of mixed effects models and we distinguish the *fixed* effects, or *population* parameters θ^{pop} and the *random* effects, or *individual* parameters $(\theta^i)_{i \in [1, N]}$.

We suppose that each observation from a subject i is a noisy deformation of the individual curve γ_i :

$$O_i(t_i) | \theta^{pop}, \theta^i \sim \mathcal{N}(\gamma_i(t_i)(O_{ref}), \sigma_\epsilon^2 I_d)$$

where σ_ϵ is the noise standard deviation and $\gamma_i(t_i) = \mathcal{E}xp_{\gamma(t)}(\omega_i)$.

Further, we model the space shift ω_i as a linear combination of q sources $S_i = (S_i^1, \dots, S_i^q)$ in the spirit of Independent Component Analysis (ICA):

$$\omega_i = A_{t_0} S_i$$

with $S_i \in \mathbb{R}^q$ the sources, A_{t_0} a $(k_g \times d) \times q$ modulation matrix, k_g the number of control points and d the dimension of the ambient space.

This helps to reduce the dimension by identifying the principal sources of anatomical variation. Further, the aforementioned orthogonality condition is enforced by projecting each column of matrix A_{t_0} onto $(\alpha_{0,1}, \dots, \alpha_{0,P})^\perp$ for the metric K_g . The projected modulation matrix is denoted $A_{t_0}^\perp$.

Finally, we can introduce the population parameters, which define the representative trajectory of the template shape: $\theta^{pop} = ((\alpha_{0,l})_{l \in [1,P]}, (t_l)_{l \in [1,P]}, O_{ref}, A_{t_0}, \sigma_\epsilon)$. The individual parameters are modeled as independent samples from a normal distribution: $\forall i \in [1, N] S_i \sim \mathcal{N}(0, 1)$. However, during optimization they will be treated as regular parameters since we will use a deterministic algorithm.

Bayesian framework. We also place ourselves within a Bayesian framework: making prior assumptions about the model parameters regularizes and guides the estimation procedure. We set the following standard conjugate distributions on the parameters:

$$\begin{cases} t_l & \sim \mathcal{N}(\bar{t}_l, s_t^2) \quad \forall l \in [1, P] \\ \alpha_{0,l} & \sim \mathcal{N}(\bar{\alpha}_0, s_\alpha^2) \quad \forall l \in [1, P] \\ O_{ref} & \sim \mathcal{N}(\bar{O}_{ref}, s_r^2) \\ A_{t_0} & \sim \mathcal{N}(\bar{A}_{t_0}, s_a^2) \\ \sigma_\epsilon & \sim \mathcal{W}^{-1}(v_\epsilon, n_\epsilon^2) \end{cases}$$

where \mathcal{N} is the normal distribution (multiscalar except for t_l), \mathcal{W}^{-1} the inverse Wishart distribution, and $\bar{t}_l, s_t, \bar{\alpha}_0, s_\alpha, \bar{O}_{ref}, s_r, \bar{A}_{t_0}, s_a, v_\epsilon$ and n_ϵ are the model hyper-parameters. The setting of the hyper-parameters is detailed in Appendix D.4.1.

7.3.4 Optimization procedure

Gradient descent. We seek to maximize the log-likelihood of the model, i.e. $\log(p((O_i(t_i))_{1 \leq i \leq N}, \theta))$:

$$\log(p((O_i(t_i))_{1 \leq i \leq N}, \theta)) = \sum_{i=1}^N \left(\log p(O_i(t_i) | S_i, \theta) + \log(p(S_i)) \right) + \log(p(\theta^{pop})) \quad (7.5)$$

Which is equivalent to minimizing the following cost function:

$$E((S_i)_{1 \leq i \leq N}, \theta) = \sum_{i=1}^N \left(\frac{d(O_i(t_i), \gamma_i(t_i) \circ O_{ref})^2}{2\sigma_\epsilon^2} - \log(p(S_i)) \right) - \log(p(\theta^{pop})) \quad (7.6)$$

where

- S_i denotes the q sources of subject i
- d is the Varifold distance (see Section 7.3.1)
- $\gamma_i(t_i) = \mathcal{E}xp_{\gamma(t)}(\omega_i)$
- The usual regularity terms are the Bayesian prior log likelihoods, specified in Appendix D.4.2.

Multiscale optimization. To enhance the estimation results, we rely on the coarse-to-fine strategy introduced in Chapter 3 and Chapter 4. Briefly, this multiscale strategy relies on two different multiscale parameterizations:

- The geodesic components of the main trajectory are constrained at a given scale S_j by expressing the gradient $\Delta E \alpha_{0,l}$ in a wavelet basis and setting to 0 some details coefficients. Updating the momentum vectors $\alpha_{0,l}$ with this smoothed gradient naturally yields smoother transformations;
- The observations and the template O_{ref} are filtered by Laplacian smoothing at a scale S'_j , which depends on the value of S_j .

The scales S_j and S'_j are progressively refined in a coarse-to-fine fashion, thus alleviating the constraints on the deformations and the objects.

Initialization. As detailed in Chapter 3, the outcome of gradient descent strongly depends on the initial parameters $\theta^{[0]}$. Plus, the $\theta^{[0]}$ are also used to set the hyper-parameters values in this Bayesian model. It is thus important to provide a good set of initial values for the parameters. The number of components P , and the time t_0 at which the template shape is defined are fixed and provided by the user, as are the initial rupture times $(t_l^{[0]})_{l \in [1,P]}$. The initialization procedure for the rest of the parameters is described in Appendix D.4.3.

Tuning of the hyper-parameters. To model cortical development in an accurate manner, several choices need to be made with regard to the model fixed parameters:

- the width σ_g of the kernel K_g controlling the amplitude of the deformations. Following preliminary experiments, the following values will be tested: $\sigma_g \in [3, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, 3.8, 3.9, 4]$. This corresponds to a number of control points k_g between 11, 000 (for $\sigma_g = 4$) and 25, 000 (for $\sigma_g = 3$).
- the width σ_w of the kernel K_w used to compute the varifold distance: since decoupling too much σ_g and σ_w might be irrelevant, for each value of σ_g , the following values of σ_w will be evaluated: $\sigma_w \in [\sigma_g - 0.2, \sigma_g - 0.1, \sigma_g, \sigma_g + 0.1, \sigma_g + 0.2]$.
- the number of components P and the initial rupture times $(t_l)_{1 \leq l \leq P}$: we will test piecewise models with $P \in [2, 3, 4, 5, 6]$ and evenly distributed rupture times

Note that the parameters $\sigma_g, \sigma_w, P, (t_l)_{1 \leq l \leq P}$ and q cannot be changed during optimization as this would modify the dimensionality of the problem. By default, the time at which the template object O_{ref} is defined will always be set to the first rupture time t_1 .

Unfortunately, estimation of the main statistical model, combined with the prior initialization procedure, has a heavy computational cost. As running it for dozens of instances would not be realistic, tuning of the hyper-parameters is performed on a simplified model, i.e. piecewise geodesic regression without space shifts in a frequentist setting: this model is equivalent to the piecewise regression model of Chapter 4.

7.3.5 Kernel regression model

To assess the performance of the proposed model against the popular kernel regression method, we implement our own kernel regression algorithm, which is a simple modification of the cross-sectional atlas model already used in this thesis.

The problem is as follows: from the dataset of N subjects $(O_1(t_1), \dots, O_N(t_N))$, we want to estimate M reference shapes $(O_{ref}^1, \dots, O_{ref}^M)$ at discrete time points (t_1, \dots, t_M) . Each template shape (O_{ref}^m) is estimated independently from the others using the classical deformable template model [6], where we assume that each observation $O_i(t_i)$ is a smooth deformation of O_{ref}^m , plus an additive random white noise ϵ_i :

$$O_i(t_i) = O_{ref}^m \circ \phi_{1,i}^{-1} + \epsilon_i, \quad \forall i \in [1, n] \quad (7.7)$$

where $\phi_{1,i} = \mathcal{Exp}_{c_0}(\alpha_{0,i})$ is the i^{th} template-to-subject deformation.

We then modify the classical cost function, optimized through gradient descent, so that the contribution of the i^{th} subject to the m^{th} template is weighted by the difference between t_i and t_m :

$$E(O_{ref}^m, (\alpha_{0,i})_{1 \leq i \leq N}) = \sum_{i=1}^N p(t_i, t_l) \left(\frac{d(O_i(t_i), O_{ref}^m \circ \phi_{1,i}^{-1})^2}{\sigma_\epsilon^2} + \|v_{0,i}\|_V^2 \right), \quad (7.8)$$

where $p(t_i, t_l) = \frac{K(t_i - t_l)}{\sum_{j=1}^N K(t_j - t_l)}$ and K is the Gaussian kernel with standard deviation 1.

To save computational time, subjects with weights lower than 0.1 are not included in the model. Optimization is performed using the same parameters σ_g and σ_w as our proposed piecewise regression model. The codes of kernel regression and our method will be made available in Deformetrica ².

²https://github.com/fleurgaudfernau/Deformetrica_multiscale

7.3.6 Performance assessment

The question of model performance is a crucial one, especially when introducing a new model for cortical surface development. One can think of several different ways of assessing the performance of a model:

- The obvious choice would be to compute the varifold distance between the observations and the deformed template, but its value depends on the kernel width σ_w , and we showed in Chapters 3 and 4 that it is not a reliable metric to evaluate the accuracy of cortical deformations. Since the varifold metric is not a point-wise distance, values are hard to interpret;
- Computing another distance metric independent of the cost function, e.g. the average Hausdorff distance [101], defined as the average of all the distances from a vertex in O_1 to the closest vertex in O_2 .
- From an anatomical perspective, it is important to assess how well the gyrification process is modelled: to this end, we can compute standard curvature measures;
- The resulting atlas can be compared to existing atlases: unfortunately, cortical surface atlases from healthy fetuses are not publicly available. Hence, we will only compare the performance of our proposed method against the state-of-the-art method (Section 7.3.5);
- Finally, from a clinical perspective, it is important to evaluate the ability of our model to discriminate healthy versus pathological brain cortices.

At each gestational week, we will compute the curvature measures presented in Table 7.2 in order to compare the gyrification level of our atlas to that of the subjects. The surface area (SA) and gyrification index (GI) are global measures, while the mean curvature (MC) and the shape index (SI) are local measures computed at the vertex level and then averaged over the cortex.

7.4 Preliminary results

In this section, we report preliminary results from our atlas building method and kernel regression. Both models are used to estimate a spatiotemporal atlas between 20 and 37 GW from the 65 inner cortical surface meshes extracted from the Lumière and FeTa datasets. The experiments were run on an Ubuntu 18.04.5 machine equipped with a NVIDIA GPU driver with 12 GB memory.

We use simple piecewise geodesic regression on the dataset, i.e. piecewise regression without space shifts (similar to the model of Chapter 4). This yields an average trajectory close to what we could expect from our final model, with the difference that the spatial variability has not been estimated. Parameters are set as follows: deformation kernel width $\sigma_g = 3.3$ ($k_g = 18,876$ control points), varifold kernel width $\sigma_w = 3.3$, noise standard deviation $\sigma_\epsilon = 0.1$, $t_0 = 24$ GW, $P = 4$ and rupture times $t_R = [24, 28, 32]$. Prior to geodesic regression, a template O_{ref} is estimated at age $t_0 = 24$ GW with kernel-weighting of the subjects. O_{ref} is subsequently used as the geodesic regression starting

Table 7.2: Curvature measures.

Name	Formula	Meaning
Surface area	$SA = \sum_{k=1}^m a_k$ with m the number of mesh triangles, a_k the area of the k^{th} triangle.	Cortical surface
Gyrification index [136]	$GI = \frac{SA}{SAC}$ with SAC the surface of the mesh convex hull	Amount of cortex buried within the sulcal folds
Mean curvature [189]	$MC = \frac{1}{n} \sum_{v=1}^n \frac{k_1(v)+k_2(v)}{2}$ with n the number of vertices, k_1 and k_2 the point-wise maximum and minimum curvatures.	Level of convexity (positive values) or concavity (negative values)
Positive and negative shape index [125]	$SI(v) = \frac{2}{\pi} \arctan\left(\frac{k_1(v)+k_2(v)}{k_1(v)-k_2(v)}\right)$ PSI is the average positive SI. NSI is the average negative SI.	PSI characterizes gyral shapes, from smooth (0) to angular (1). NSI characterizes sulcal shapes, from smooth (0) to angular (-1).

point. Note that we set t_0 to a young age because backward modelling of the cortical changes from t_0 to the minimum age (i.e. 20 GW) is challenging as it requires smoothing the surface and eliminating folds. Optimization is performed by gradient descent with adaptive step size and combined with multiscale optimization, with an initial step size $h = 0.01$ and a convergence threshold of 0.0001. Estimation took 9 hours and 17 minutes.

For kernel regression, discrete template images are independently estimated at each gestational week using the same parameters σ_g , σ_w and σ_ϵ and the same optimization procedure. For each age, the initial template shape is a high-resolution inflated surface mesh scaled to the mean size of the subjects cortices at that age. For each atlas estimation task, computation time ranges between 1 hours and 3 hours and estimation of the 18 templates took approximately 24 hours.

7.4.1 Atlases visual comparison

An interactive atlas viewer is currently being developed at the first author's website³. The goal is to enable visualization of the atlas at any gestational age and display local curvature measures at the vertex level.

Figure 7.6 presents the spatiotemporal atlases built with the two methods every two weeks of gestation between 22 and 36 GW. Both methods yield relatively similar results, with concurrent emergence of the main sulci (e.g., the central sulcus). However, some of the main sulci are better delineated by kernel regression: the postcentral sulcus is clearly visible at 28 GW in kernel regression but not in geodesic regression and the superior temporal sulcus is better defined. We also notice that the cortical template

³<https://fleurgaudfernau.github.io/Atlas/>

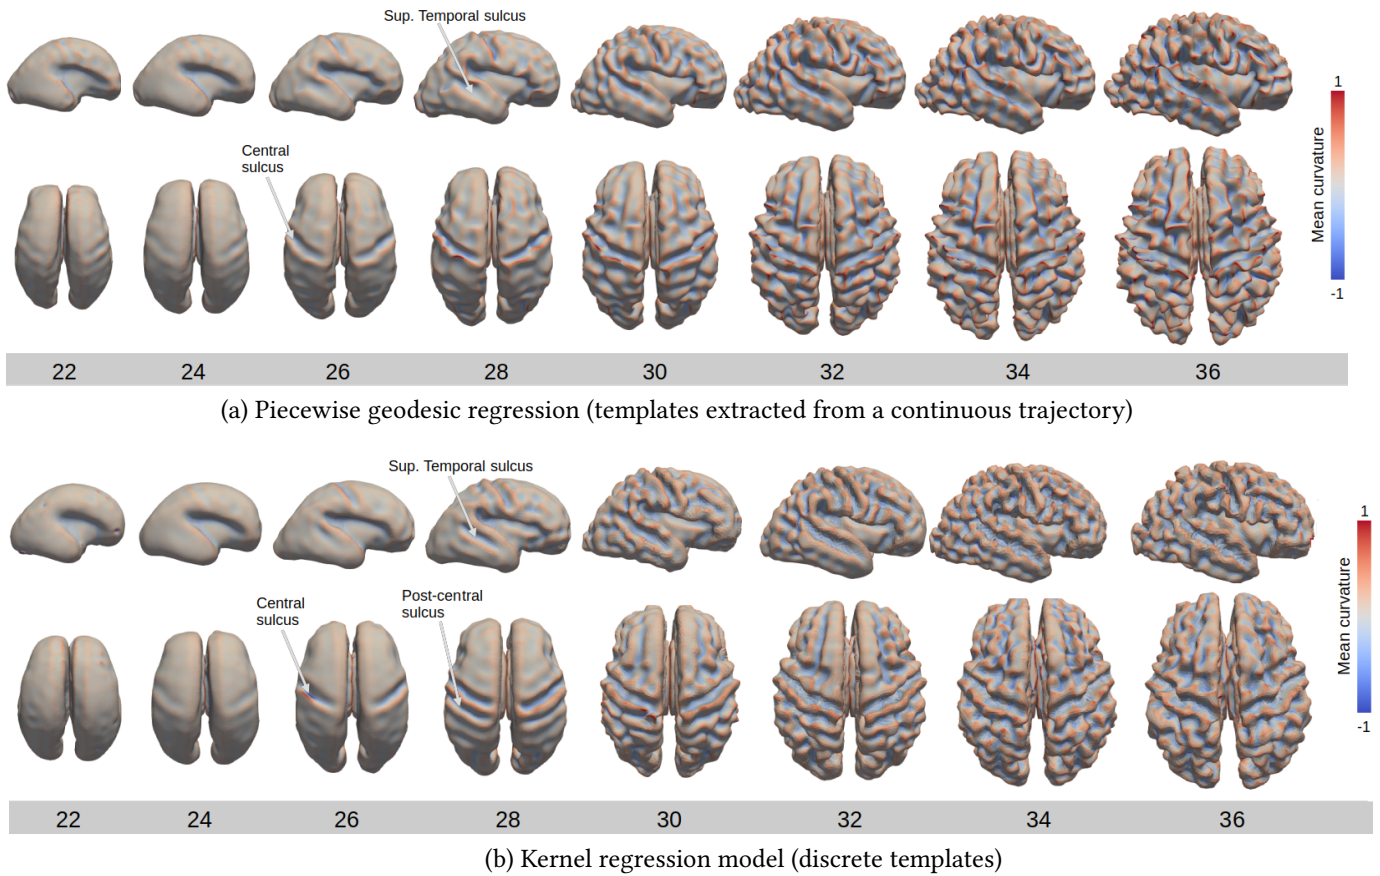


Figure 7.6: Visual comparison of the spatiotemporal atlases build with our method (panel (a)) and kernel regression (panel (b)) in right lateral (top row) and dorsal (bottom row) views between 22 and 36 GW. Colors indicate the point-wise mean curvature. Positive values correspond to the gyri and negative values to the sulci.

extracted from geodesic regression at 22 GW has slightly abnormal folding, with a too visible central sulcus.

7.4.2 Atlas evaluation

7.4.2.1 Distance metrics

Table 7.3 presents the mean varifold distance at each gestational week between the subjects and the templates constructed by our method or kernel regression. As expected, the distances increase with gestational age as the cortex becomes more complex. Little differences can be seen between the algorithms, although our method has a slightly lower average distance. Table 7.4 shows the mean Hausdorff distance at each gestational week for the two methods. This metric mostly agrees with the previous one, although some of the differences highlighted by the varifold distances are less salient.

Table 7.3: Residual error for our method and kernel regression, computed as the varifold distance with $\sigma_w = 3.3$ between each subject and the closest template. The mean value (top row) and standard deviation (bottom row) are displayed for each gestational age. Values are divided by 10^6 . μ denotes the average distance over subjects.

GW	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	μ
Our method	3.67	\	5.84	6.83	6.78	8.07	6.70	6.46	9.23	8.02	11.22	10.91	11.58	15.43	16.40	18.75	17.67	10.85
	0	\	0	0.28	1.5	2.36	1.49	1.23	4.11	0.33	4.12	1.68	1.67	2.39	2.32	2.35	0	4.56
Kernel regression	3.77	\	5.44	6.98	6.84	8.36	7.03	6.59	9.76	9.09	11.56	10.55	11.44	15.47	16.31	18.57	17.23	10.93
	0	\	0	0.05	1.52	2.44	1.59	1.26	4.32	0.7	4.14	1.68	1.68	2.48	2.27	2.33	0	4.45

Table 7.4: Average Hausdorff distance (mm) between each subject and the closest template for our method and kernel regression. The mean value (top row) and standard deviation (bottom row) are displayed for each gestational age. μ is the average distance over subjects.

GW	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	μ
Our method	0.82	\	1.41	1.41	1.05	1.45	1.14	1.27	1.70	1.27	1.22	1.19	1.21	1.37	1.36	1.19	1.10	1.27
	0.0	\	0.0	0.14	0.14	0.42	0.22	0.22	1.18	0.09	0.27	0.18	0.15	0.16	0.27	0.12	0.0	0.40
Kernel regression	1.06	\	1.40	1.49	1.06	1.47	1.13	1.27	1.75	1.46	1.21	1.10	1.16	1.34	1.34	1.21	1.09	1.28
	0.0	\	0.0	0.16	0.16	0.44	0.24	0.16	0.06	0.28	0.16	0.15	0.17	0.27	0.13	0.12	0.0	0.41

7.4.2.2 Cortical folding accuracy

In Figure 7.7, we compute curvature measures for each cortical template and compare them to the ground truth, i.e. the curvature measures of each subject used to build the atlases. Both our method and kernel regression yield atlases with surface area and gyri-fication index close to the expected values. However, at advanced gestational ages, when the variability of cortical geometry increases, both indices are slightly underestimated. In Figure 7.7c, the evolution of mean curvature is also close to that of the subjects, although geodesic regression produces templates with overestimated mean curvature at younger and older gestational ages. This is consistent with Figure 7.6a, where templates around 22-26 GW lack smoothness and those around 32-36 weeks have elevated mean curvature values around the gyri. Finally, evolution of the positive and negative shape index indicates that the gyri become smoother while the sulci become more angular with gestational age, a tendency already observed in Hu. et al [96]. For these values, both atlases are mostly correct but two issues can be raised: geodesic regression builds gyri that are not convex enough around 20-24 GW and too convex around 32-36 weeks, and sulci that are too concave around 20-24 GW and not angular enough around 32-36 weeks. This is consistent with Figure 7.6a: in middle pregnancy, sulcal shapes are a bit too marked and in late gestation, gyri have too angular shapes, while sulci are too wide. A comparison of the atlases sulci and gyri at 36 GW with subjects at the same age is presented in Figure 7.8.

Whilst kernel regression has a more accurate cortical geometry, some temporal dis-




Figure 7.7: Comparison of five curvature measures between the subjects (ground-truth) and two spatiotemporal atlases at each gestational week.

crepancy between consecutive ages can be seen, e.g. abnormally elevated mean curvature at 23 GW and discrepant positive and negative shape indices around 20-24 GW. This phenomenon is illustrated in Appendix Figure D.4.

7.4.3 Influence of the hyper-parameters

Here, we report preliminary results on the tuning of the geodesic regression model parameters σ_g , σ_w , P and t_R . We compare the evolution of five curvature measures for different values of these parameters.

Size of the deformation and varifold kernels. We test piecewise geodesic regression with $\sigma_g = \sigma_w \in [3.3, 3.5, 3.7]$, i.e. $k_g \in [18, 876, 16, 275, 13, 800]$ and the same parameters as in the previously presented model. Results are presented in Appendix Figure D.5. Higher values of σ_g lead to underestimated surface area, gyrfication level, and smoother sulci at advanced gestational ages. This indicates that kernels with standard deviations above 3.5 are too large to accurately represent cortical folds.

Number of components and rupture times. The previously introduced model with $\sigma_g = \sigma_w = 3.3$ is tested with the following dynamics: $P = 2$ and $t_R = 24$, $P = 3$ and $t_R = [24, 30]$, $P = 4$ and $t_R = [24, 28, 32]$, $P = 5$ and $t_R = [24, 27, 30, 33]$, $P = 6$ and $t_R = [24, 26.5, 29, 31.5, 33]$. Results are presented in Appendix Figure D.6. Models with a higher number of components seem to underestimate more the increase in surface area and gyrfication at older gestational ages, which is quite counter-intuitive.




Figure 7.8: Comparison of the atlas templates at age 36 GW (panels (a) and (b)) with two subjects of age 36 GW (panels (c) and (d)). Cortices are presented in right lateral view. Arrows indicate the central sulcus. Colors indicate local shape index values, ranging from -1 (blue) to 1 (red). In panel (a), notice the slightly exaggerated sharpness of the gyri compared to panels (c) and (d). The sulci are also slightly too wide.

Expectedly, the 2-component model also achieves a lower level of gyrification and generates smoother gyri and sulci in late gestation. In contrast, the model with 3 geodesic components achieves the highest level of cortical folding (i.e. highest gyrification index and mean curvature) and generates smoother gyri (i.e. lower positive shape index).

7.5 Discussion

In this chapter, we introduced a new method for spatiotemporal atlas construction and applied it to healthy cortical development during the fetal period. Relying on two large datasets of fetal brain images, we compared our method against state-of-the-art kernel regression.

Data: A strength of this study is that it relies on data coming from different institutions and countries. As noted in previous work [249], data origin and representativeness matter. As usual with prenatal MRI, high-resolution volume reconstruction and cortical surface extraction were challenging steps. Contrary to volume reconstruction [57], there does not exist a fully automated pipeline for the reconstruction of cortical surface in fetuses and we used a pipeline designed for T2-weighted postnatal MRI [140]. Only 23% and 35% of brain volumes from Lumière and FeTa had successful cortical surface extraction. The discrepancy probably reflects the age difference between the two datasets and the fact that the dHCP pipeline is more efficient on older fetal brains. In the future, we will modify the dHCP pipeline to adapt it to fetal brains, in the spirit of what was done by Xu et al. [254], and hopefully recover more cortical surfaces from the dataset. This key issue is a recurrent one in the field of fetal MRI, where the lack of fully automated tools and open-source datasets hampers the development of quantitative analyses. Indeed, we note the involvement of manual segmentation of the fetal cortex [38, 34, 97, 250] or the use of homemade tools for cortical surface extraction [130, 254] in a number of papers.

Model of cortical surface development. Arguing that kernel regression does not account for the spatial variability of the fetal anatomy, we introduced a new model to build a *truly* spatio-temporal atlas. This work is built upon methods and experiments developed in the previous chapters: space shifts were integrated onto a piecewise geodesic regression model to estimate normal anatomical variability during cortical folding, and optimization was performed using a multiscale strategy whose efficiency was demonstrated in Chapter 4. As in Chapter 6, we combined different tools from Computational Anatomy, namely atlas estimation, registration and parallel transport. This time however, optimization of the main trajectory and the space shifts is not sequential but *simultaneous*, which hopefully will lead to better results. Unlike kernel regression, our model could explicitly take into account longitudinal data, e.g. fetuses that underwent more than one MRI during pregnancy [22].

Atlas evaluation. In this chapter, we paid careful attention to the question of model evaluation. As visual evaluation and distance measures between the templates and the subjects are insufficient proxy of atlas quality, we computed curvature measures to assess the accuracy of the modelled cortical folding. Standard measures, i.e. surface area, gyrification index and mean curvature, indicated good levels of cortical folding, i.e. burial of the cortex inside the sulci. At older gestational ages, gyrification levels were a bit lower than that of the input subjects, but an *average* template cannot be expected to reach the same level of complexity as an individual cortex. The shape index, which characterizes gyral and sulcal shapes, indicated some shape inaccuracies in middle and late pregnancy. The overly pronounced folds in middle pregnancy likely result from the choice of t_0 , i.e. the time at which the template is defined, at 24 GW. Additional experiments will be performed to test alternative values of t_0 . In late pregnancy, increased gyri sharpness might be corrected by adding local smoothness constraints to the vector fields defining the geodesic components. A possibility could be to use the alternative multiscale algorithm presented in Appendix A.4 and grant more weight to coarse scale coefficients during optimization.

Kernel regression modelled cortical folding in a slightly more accurate manner, which was expected since the model is designed to perform well at a given gestational age. However, unlike kernel regression, our method guarantees temporal consistency of the cortical templates. Both models can only be compared so far since they are based on different philosophies. The true advantage of our model is its ability to classify new subjects.

Limitations. One may wonder about the relevance of using a more complex model. It is no wonder that kernel regression is popular as it is a simple and efficient method to depict brain anatomy at a given time. Our model is derived from spatio-temporal models [200], which have proven useful, *inter alia*, for longitudinal modelling of hippocampal shape changes in Alzheimer’s disease [177, 44]. In neurodegenerative diseases, a patient is likely to undergo several exams and brain scans, which is rarely the case during prenatal monitoring. In the absence of such longitudinal data, characterizing the growth dynamic of the fetal brain is a more difficult task, hence the need to carefully evaluate this new model. Another important pitfall is the need to set a high number of parameters and the long initialization procedure to initialize the parameters, which can be prohibitive.

Future work. As mentioned above, future work will include improving the cortical surface extraction and reconstruction step. We will pursue the tuning of the model parameters and finally run our final model to estimate both an average trajectory and the normal anatomical variability around the mean. The final model will be tested on an independent dataset to classify healthy versus abnormal cortical anatomy.

GENERAL CONCLUSION

In this thesis, we began to bridge the gap between the quantitative analysis of the fetal brain and geometrical deformation models in the LDDMM framework. To this end, we needed to address challenges specific to the analysis of fetal brain MRI, namely the lack of publicly available data, heterogeneous data quality, the complexity of brain growth and the difficulty of comparing subjects at different gestational ages. Fortunately, the LDDMM framework has properties that are very compatible with these needs, namely the ability to handle small datasets, to transport subjects in space and time and to provide explicability and interpretability.

In this work, we proposed both methodological and practical approaches: in a first part, we provided new multiscale optimization strategies to enhance the estimation of cross-sectional reference shapes and time-dependent trajectories. In a second part, we started from a clinical question and developed a shape analysis pipeline to extract anatomical deformations specific to fetal brains with corpus callosum agenesis. We also characterized partial versus complete agenesis. Finally, we introduced a new model to accurately depict prenatal cortical development.

The innovative nature of the subject required the use of a variety of tools and methods from different disciplines. In Chapter 3, we manipulated mathematical objects, e.g. the Haar wavelet and addressed an optimization problem. In Chapter 4, we continued down the same road, this time attributing more importance to clinical applications. In Chapter 6 and 7, our goals were both clinical and methodological: we relied on the previously introduced mathematical tools and experiments to answer specific research questions and presented new methods that can be of use outside these specific questions. The contribution of this thesis spans from new algorithms in the field of Computational Anatomy to new approaches in medical image analysis.

8.1 Multiscale optimization(s)

As noted above, the LDDMM framework has a variety of properties and tools suitable to the analysis of different types of data. Whilst it requires few assumptions, in practice the deformation of high-dimensional, highly complex medical images can easily fall into unrealistic local minima. One often has to resort to fine-tuning of the parameters, specifically the width of the kernel controlling the spatial regularity of the transformations. To alleviate this problem, we introduced in Chapters 3 and 4 a coarse-to-fine optimization strategy based on a wavelet-like reparameterization of the velocity fields

and a filtering of the objects. The strength of this method is that it modifies a single step of the optimization procedure while preserving the deformation model. This makes the multiscale strategy highly versatile since it can be applied to different types of objects, combined to other deformation models or even associated to other wavelet functions such as Daubechies. As often with coarse-to-fine strategies [190], ours only *alleviate* the local minima issue: in practice, the final scale of the deformation still depends on the regularizing kernel size. This is well illustrated by our experiments: the advantages of the multiscale strategy are undeniable on simple data such as handwritten digits and hand-drawn characters but more subtle on clinical objects.

This work is also intertwined with the question of *model evaluation*. The residual data term, e.g. L^2 norm or varifold distance, is of course quite limited. For images, we used another similarity metric independent of the cost function, namely the Structural Similarity Index Metric. However, similarity between source and target objects is not the only desirable feature: deformation regularity is also very important, especially when modelling anatomical changes, which requires visual evaluation. Of note, we could have compared our multiscale strategy to other algorithms, but this was outside this thesis objectives: we were more interested in adapting tools from the LDDMM framework to specific analyses than reaching state-of-the art performance.

8.2 Analysis of fetal MRI

As mentioned in Chapter 5, the quantitative analysis of fetal brain MRI is a recent field of research, which raises exciting challenges as well as technical difficulties. In Chapters 6 and 7, data preprocessing was a rather problematic step. In Chapter 6, volume reconstruction was performed with NiftyMIC 0.7.5 [56] and brain extraction required semi-automated corrections. In Chapter 7, volume reconstruction was performed with NiftyMIC 0.9.0, which includes a novel brain extraction algorithm [185], and no corrections were needed, illustrating the fast technical advances in the field. The reconstruction success rate was only 62% for hôpital Trousseau and 55% for the Lumière dataset. In the first case, this number seemed linked to the year of MRI acquisition and not to the fetus age (likely because sedation was offered). In the second case, younger fetuses were less likely to have a valid brain volume, probably because of motion artefacts. Of note, such difficulties are rarely reported in the literature. This indicates that current state-of-the-art reconstruction techniques lack robustness regarding lower quality data, which can be problematic because of poor data availability. In Chapter 7, extraction of the cortical surface was an even more challenging step because of the lack of fully automated pipelines for segmentation and structure reconstruction. This situation will hopefully evolve soon as the subject has recently been drawing interest [235].

8.3 Fetal brain analysis pipeline

In Chapter 6, we combined existing Computational Anatomy tools to design a pipeline for characterizing pathological fetal brains: an average trajectory of healthy brain changes

is built, subjects are registered to an age-matched template and transported to a common space to extract anatomical deformations through PCA. Compared to existing analysis frameworks, our pipeline has the advantage of enabling whole-brain analysis (although it could also be applied to single structures) and smoothing out age-related differences between fetal brains. In the future, we aim at making our code publicly available and hopefully apply the pipeline to other pathologies. In cases where a postnatal outcome, e.g. intellectual disability is known, directions of anatomical variability could be correlated to the outcome in order to identify new prenatal predictors of disability. While the tools used in this chapter are similar to that of Chapter 7, the spirit of the work is a bit different: the building of the reference trajectory is not the main focus of the analysis and in fact, one may use existing spatiotemporal atlases. Different tools such as registration and parallel transport are used in a *sequential* manner, making the pipeline more adapted to the characterization of a population of abnormal fetal brains. In contrast, the model of Chapter 7 is geared to the *simultaneous* estimation of the mean and variability of healthy brain changes and could be used for detecting abnormal subjects rather than characterizing them.

8.4 Prenatal brain development atlas

In Chapter 7, we introduced an alternative method to the widely-used kernel regression to build an atlas of the developing brain. Kernel regression is a rather simple method where a template object at a given gestational age is estimated independently of the others. Our more complex model focuses on the average growth dynamic. Preliminary experiments were performed in Chapter 4, where we showed that diffeomorphic deformations have difficulty modelling MR intensity changes across pregnancy. As evoked earlier, this task would require a high degree of model adaptation, or even the use of more elaborate models [234].

On the contrary, the LDDMM framework is more adapted to the wrapping of anatomical meshes. In Chapter 7, we designed an integrative model that estimates simultaneously a single template object, the piecewise geodesic trajectory that deforms the template with time and the anatomical variability around the mean. This model harbors great potential to widen the scope of spatio-temporal atlases: not only building representative anatomical objects and their parcellations, our method enables the analysis of time-dependent growth dynamics, the positioning of any subject with respect to the average development, the displacement of said subject in space and time and finally the characterization of the anatomical variability of healthy or pathological populations. To summarize, what we propose is merely a change of perspective and a focus on the dynamic aspect of prenatal brain growth. An interesting path of research would be to include postnatal data into our model to span longer developmental periods. This is actually doable as neonatal cortical surface meshes are already publicly available [5].

8.5 Perspective

Computational Anatomy and fetal brain research are dynamic research fields. In addition to the numerous research directions that we already evoked, one major objective stands out: to make the tools developed here publicly available and *easily* operable. This is no simple task as their use requires at the bare minimum a good understanding of the presented models.

The methodological and clinical questions tackled in this thesis are complex and multifaceted ones. By adapting deformation models, we sought to provide novel ways of characterizing fetal brains and widen the field of fetal brain MRI. We believe that our approach of combining deformation models with fetal brain imaging is promising to advance the field of fetal neuroimaging.

APPENDIX **A**

SUPPLEMENTARY MATERIAL OF CHAPTER 3

A.1 Characters dataset




Figure A.1: Dataset of manually designed characters.

A.2 Algorithmic implementations of the transforms FWT and IWT

To compute the Forward and Inverse Wavelet Transforms, explicit computation of the related matrices M_{FWT} and M_{IWT} is not necessary since the algorithms rely on local operations on basis coefficients. To compute the wavelet coefficients of an array X of dimension d , the FWT algorithm (Algorithm 5) relies on a one-dimensional Haar Forward algorithm (Algorithm 2), which computes local means and differences along one axis. In Algorithm 5, lines 4-12 amount to computing $M_{FWT}X$. This process is illustrated on a simple example in Fig. A.2. Uneven numbers of rows/columns are handled by computing *weighted* averages and differences so that the boundaries are given the same importance as the rest of the array.

Finally, the output wavelet coefficients have to be normalized: this is done by computing explicitly the matrix M_{FWT} in Algorithm 4 and computing the renormalization matrix R from M_{FWT} :

$$R[i] = \frac{1}{\|M_{FWT}[i, :]\|_2}$$

where $\|M_{FWT}[i, :]\|_2$ is the L_2 norm of the i_{th} row of M_{FWT} .

Another way of considering R is by reminding ourselves that the embedded spaces V_s and W_s were constructed by normalizing scaling and wavelet functions with a factor $2^{-sd/2}$. Hence, when computing the transformation between scales $s - 1$ and s , R is such that $R[i] = 2^{sd/2}$ (except at the boundaries).

Note that, since the normalization factor when going from scale $s - 1$ to scale s is known, computing M_{FWT} should not be necessary. However, when the algorithm operates on non-dyadic grids, computations at the boundaries are modified, and so are the expressions of M_{FWT} and R . To avoid errors, it is thus simpler to compute R from M_{FWT} .

The IWT algorithm (Algorithm 6) runs in a manner that is symmetrical to the FWT algorithm, by using the one-dimensional Haar Backward algorithm (Algorithm 3) to compute finer-scale coefficients from coarse scale coefficients one axis at a time.

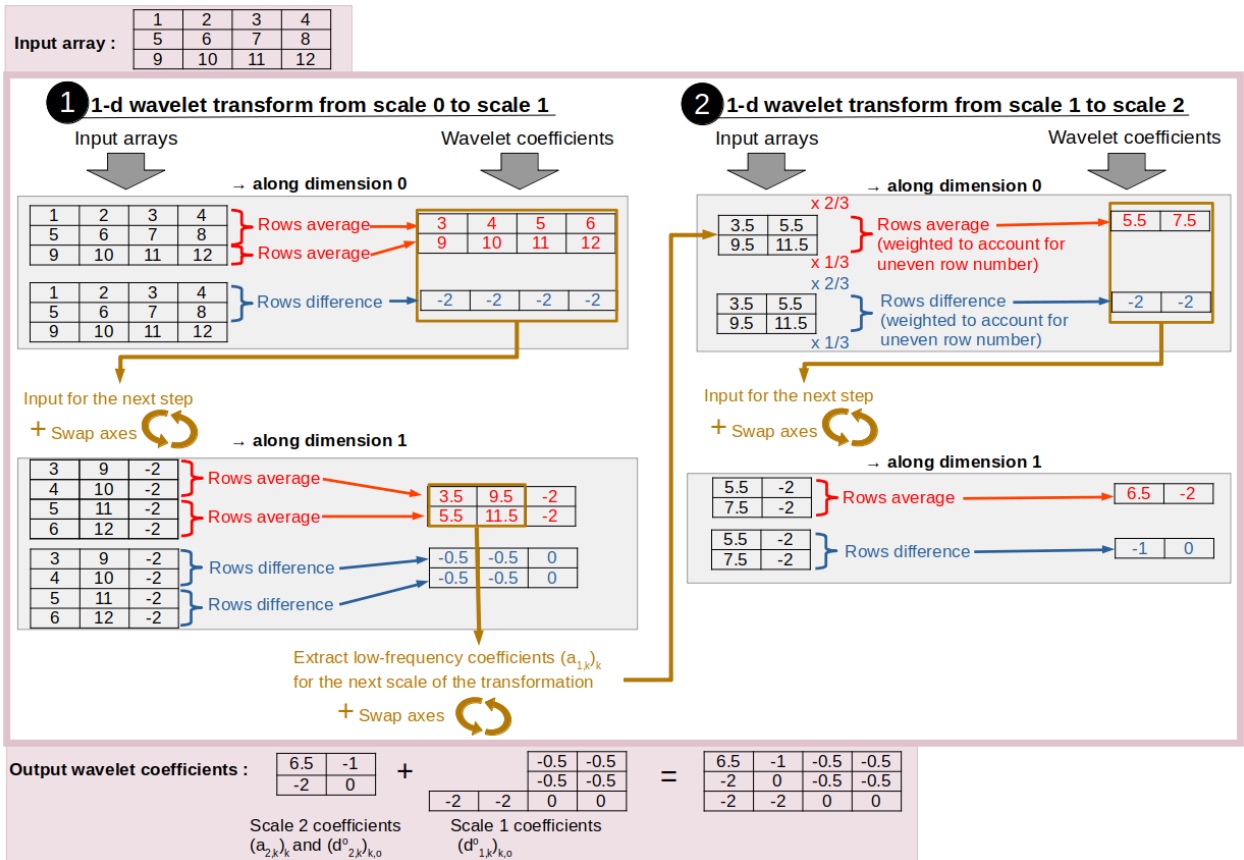


Figure A.2: The FWT Algorithm (Algorithm 5) applied to a 3-by-4 array. Steps in grey squares illustrate the 1-dimensional Haar Forward step (Algorithm 2). Running the IWT Algorithm (Algorithm 6) amounts to running the illustrated steps backward. For the sake of simplicity, the renormalization step (line 16 in Algorithm 5) is not featured.

Algorithm 2 1-dimensional Haar Forward step

```
1: Input
2:  $\beta$ : array of shape  $(k_1, \dots, k_D)$ ,  $d$ : axis along which to compute the transform,  $K_d$ :
   size of the original array along axis  $d$ ,  $s$ : current scale,  $w_s$ : list of scales for each axis
    $[1, \dots, D]$ , stored in ascending order
3: Initialization
4: Swap axes in  $\beta$  to put axis  $d$  in position 0
5: if  $k_d > 1$  then
6:   if  $k_d$  is even then
7:      $\beta_a \leftarrow (\beta[0 :: 2] + \beta[1 :: 2])/2$  {Average the consecutive rows of  $\beta$ }
8:     if  $K_d \neq 2 \times k_d^s$  then
9:        $\delta \leftarrow K_d/2^s - (k_d - 1)$  {Weighting of the border}
10:       $\beta_a[-1] \leftarrow (\beta[-2] + \delta * \beta[-1])/(1 + \delta)$ 
11:     end if
12:      $\beta_d \leftarrow \beta[0 :: 2] - \beta_a$  {Difference between the consecutive rows of  $\beta$ }
13:   else
14:      $\beta_a \leftarrow (\beta[0 : -1 : 2] + \beta[1 :: 2])/2$  {Average the consecutive rows of  $\beta$ }
15:      $\beta_d \leftarrow \beta[0 : -1 : 2] - \beta_a$  {Difference between the consecutive rows of  $\beta$ }
16:      $\beta_a \leftarrow \text{concatenate}([\beta_a, \beta[-1]])$  {Add the last unpaired row}
17:   end if
18:    $\beta \leftarrow \text{concatenate}([\beta_a, \beta_d])$ 
19: end if
20:  $w_s[d] \leftarrow \lceil k_d/2 \rceil + w_s[d]$ 
21: return  $\beta, w_s$ 
```

Algorithm 3 1-dimensional Haar Backward step

```

1: Input
2:  $\beta$ : array of wavelet coefficients of shape  $(k_1, \dots, k_D)$ ,  $d$ : axis of  $\beta$  along which to
   apply the backward transform,  $K_d$ : size of the original array along axis  $d$ ,  $s$ : current
   scale,  $w_s$ : list of scales for each axis  $[1, \dots, D]$ , stored in ascending order
3: Initialization
4:  $\beta \leftarrow$  Swap axes in  $\beta$  to put axis  $d$  in position 0
5:  $X \leftarrow$  zero array of the shape of  $\beta$ 
6:  $n_{low} \leftarrow w_s[d][0]$ 
7:  $\beta_a \leftarrow \beta[: n_{low}]$  {Low-frequency coefficients}
8:  $\beta_d \leftarrow \beta[n_{low} : ]$  {High-frequency coefficients}
9: if  $w_s[d][1] > 1$  then
10:  if  $w_s[d][1]$  is even then
11:     $X[0 :: 2] \leftarrow \beta_a + \beta_d$ 
12:     $X[1 :: 2] \leftarrow 2 \times \beta_a - X[0 :: 2]$ 
13:    if  $K_d \neq 2 \times k_d^s$  then
14:       $\delta \leftarrow K_d/2^s - (k_d - 1)$ 
15:       $X[-1] \leftarrow ((1 + \delta) \times \beta_a[-1] - X[-2])/\delta$ 
16:    end if
17:  else
18:     $X[0 : -1 : 2] \leftarrow \beta_a[: -1] + \beta_d$ 
19:     $X[-1] \leftarrow \beta_a[-1]$ 
20:     $X[1 :: 2] \leftarrow 2 \times \beta_a[: -1] - X[0 : -1 : 2]$ 
21:  end if
22: end if
23:  $w_s[d] \leftarrow w_s[d][1 : ]$ 
24: return  $X, w_s$ 

```

Algorithm 4 Renormalization

```
1: {Compute the Haar Forward Matrix  $M_{FWT}$  and the renormalization matrix  $R$ }
2: Input
3:  $(K_1, \dots, K_D)$ : shape of the array to wavelet transform
4: Initialization
5:  $n \leftarrow \prod_{i=1}^D K_i$ 
6:  $M_{FWT} \leftarrow$  zero array of shape  $(n, n)$ 
7: for  $i$  in range( $n$ ) do
8:    $Z \leftarrow$  zero array of shape  $(K_1, \dots, K_D)$ 
9:    $i_{th}$  element of  $Z \leftarrow 1$ 
10:   $\beta, w_s, R \leftarrow$  Haar_Forward( $Z, \rho = 0$ ) {Algorithm 5}
11:   $M_{FWT}[:, i] \leftarrow$  flatten( $\beta$ )
12: end for
13:  $R \leftarrow$  zero array of shape  $(K_1, \dots, K_D)$ 
14: for  $i$  in range( $n$ ) do
15:   $i^{th}$  element of  $R \leftarrow \frac{1}{\|M_{FWT}[i,:]\|_2}$ 
16: end for
17: return  $R$ 
```

Algorithm 5 Forward Haar Wavelet Transform (FWT algorithm)

```
1: {Compute the wavelet coefficients  $\beta$  of an array  $X$ }
2: Input
3:  $X$ : array of shape  $(K_1, \dots, K_D)$ ,  $\rho$ : renormalization factor (default: 1)
4: Initialization
5:  $\beta \leftarrow X$  {Array to store the wavelet coefficients}
6:  $w_s \leftarrow [[K_1], \dots, [K_D]]$  {For each axis, store the scales for which the wavelet coefficients have been computed in ascending order}
7:  $S_{max} \leftarrow \lceil \log_2(\max(K_1, \dots, K_D)) \rceil$ 
8: for  $s$  in range( $S_{max}$ ) do
9:   $\beta_{current} \leftarrow \beta[:, w_s[0][0], \dots, : w_s[D][0], ]$  {Low-frequency coefficients at scale  $s - 1$ }
10:  for  $d$  in range( $D$ ) do
11:     $\beta_{current}, w_s, R \leftarrow$  Haar_forward_1d_step( $\beta_{current}, d, K_d, s, w_s$ ) {Algorithm 2}
12:     $\beta[:, w_s[0][0], \dots, : w_s[D][0], ] \leftarrow \beta_{current}$ 
13:  end for
14: end for
15: if  $\rho \neq 0$  then
16:   $R \leftarrow$  renormalize(( $K_1, \dots, K_D$ )) {Algorithm 4}
17:   $\beta \leftarrow \beta \times R^\rho$ 
18: end if
19: return  $\beta, w_s, R$ 
```

Algorithm 6 Inverse Haar Wavelet Transform (IWT algorithm)

```

1: {Given its wavelet coefficients  $\beta$ , compute array  $X$  (i.e. coefficients of scale 0)}
2: Input
3:  $\beta$ : array of wavelet coefficients of shape  $(K_1, \dots, K_D)$ ,  $\rho$ : renormalization factor
   (default: 1),  $w_s$ : list of scales for each axis  $[1, \dots, D]$ , stored in ascending order
4: Initialization
5:  $X \leftarrow \beta$ 
6:  $S_{max} \leftarrow \lceil \log_2(\max(K_1, \dots, K_D)) \rceil$ 
7: if  $\rho \neq 0$  then
8:    $R \leftarrow \mathbf{renormalize}((K_1, \dots, K_D))$  {Algorithm 4}
9:    $X \leftarrow X/R^\rho$ 
10: end if
11: for  $s$  in range( $S_{max} - 1, -1, -1$ ) do
12:    $X_{current} \leftarrow X[:, w_s[0][1], \dots, : w_s[D][1]]$ 
13:   for  $d$  in range( $D$ ) do
14:      $X_{current} \leftarrow \mathbf{Haar\_Backward\_1d\_step}(X_{current}, d, K_d, s, w_s)$  {Algorithm 3}
15:      $X[:, w_s[0][1], \dots, : w_s[D][1]] \leftarrow X_{current}$ 
16:   end for
17: end for
18: return  $X$ 

```

A.3 Multiscale algorithm with local adaptation

A.3.1 Local adaptation extension

As highlighted in Section 3.2, the term *multiscale* can imply either that deformations are estimated in a coarse-to-fine manner, or that deformations of different scales coexist. The multiscale strategy we presented is closer to the former case. Here, we seek to enhance our algorithm with a *true* multiscale flavor by enforcing the coexistence of velocity fields of different scales. Since wavelet functions are characterized by their scale, location and orientation, we take advantage of the location component so that we can force the velocity field of each subject to be *unevenly* smooth. Namely, the spatial constraints we impose through the wavelet-based multiscale representation can vary according to the amount of variability in the images, so that template-to-subject deformations may remain coarse in areas that contain few information. Our motivation is twofold: alleviate the computational cost in areas of low variability and force the optimization to focus on challenging areas.

The multiscale algorithm with local adaptation is described in Algorithm 8. It only requires the addition of two steps to the the naive multiscale algorithm. Of note, the local adaptation can only handle images at the moment.

Screening of smooth areas. The local adaptation step (Algorithm 7) takes place after each coarse-to-fine step. We recall that the current scale S_j is related to wavelet functions whose support is of size $(2^{S_j})^d$ control points (or less if the functions are located at the image borders). Hence, each scale S_j is related to a list of unique wavelet positions on the grid of control points:

$$P = [k \in \mathbb{Z}^d, k_i \in [0, 2^{S_j}, 2 \times 2^{S_j}, 3 \times 2^{S_j} \dots, L]]$$

where $L = k_g(y) - 2$ if the grid is dyadic and $k_g(y) - 1$ otherwise, with $k_g(y) + 1$ the number of control points along axis y .

The position of a wavelet function $\widetilde{\psi_{S_j, k}^o}$ on the grid of control points can easily be converted to pixel coordinates on an image array, thus it is straightforward to compute the variability inside the related area on the image. Following each coarse-to-fine step, we compute for each subject i the average residual image Δ_j^i at iteration j according to the formula:

$$\Delta_j^i = \|I_{ref, j} \circ \phi_{j, i}^{-1} - I_i\|_2^2, \quad (\text{A.1})$$

Then, we compute the average residual value $\Delta_{j, k}^i$ of every area at position k . If $\Delta_{j, k}^i$ is below a given threshold, the area is considered as smooth and stored in a dictionary.

Update of the template-to-subjects deformations. The gradient descent step is performed as described in the previous section. The only change to the procedure is that we not only set to 0 detail coefficients defined at scales finer than S_j (line 24 in Algorithm 8), but also detail coefficients whose position is stored in Z (line 26). Consequently, spatial constraints over the subjects' velocity fields will be uneven: in areas of low variability, the velocity fields will be constrained to remain smooth, whereas in areas with

high variability, the velocity fields will be able to express finer and finer components as the coarse-to-fine procedure progresses.

Algorithm 7 Local adaptation

```

1: {For each subject image, identify smooth areas and store their position in a dictionary}
2: Input
3: Z: dictionary of smooth areas
4:  $S_j$ : current scale
5: Initialisation
6:  $P \leftarrow$  list of unique wavelet positions at scale  $S_j$ 
7:  $Z[i][S_j] \leftarrow \emptyset$  for each subject  $i$ 
8: for each subject  $i$  do
9:   Compute  $\Delta_j^i$  the residuals image of subject  $i$  according to Eq. (A.1)
10:  Compute  $r_{j,max}^i$  the maximum average residuals value over zones in  $P$ 
11:  Compute the average residuals  $r_{j,k}^i$  inside area position  $k \in P$ 
12:  for each position  $k$  in  $P$  do
13:    if  $k$  is already included in  $Z[i][S_{j-1}]$  then
14:      Add  $k$  to  $Z[i][S_j]$  {The area belongs to an already smoothed larger area}
15:    else
16:      if  $r_{j,k}^i < 0.01 \times r_{j,max}^i$  then
17:        Add  $k$  to  $Z[i][S_j]$ 
18:      end if
19:    end if
20:  end for
21: end for
22: return  $Z$ 
    
```

Algorithm 8 Coarse-to-fine initialization algorithm with local adaptation

-
- 1: {For clarity, steps linked to the local adaptation extension are indicated in blue}
 - 2: **Input**
 - 3: Set of images $(I_i)_{1 \leq i \leq N}$ of dimension d , template image I_0 , geometric kernel width σ_g , trade-off regularity/fidelity-to-data σ , initial step size h
 - 4: **Initialization**
 - 5: $j \leftarrow 0$
 - 6: Regular grid of control points c_0 with spacing σ_g
 - 7: Template image intensities $I_0(j) \leftarrow I_0$
 - 8: Momentum vectors $\alpha_{0,i}(j) \leftarrow 0$ for each subject i
 - 9: $\beta_{0,i}(j) \leftarrow FWT(\alpha_{0,i}(j))$ for each subject i
 - 10: Initialize the gradients $\nabla_{\alpha_{0,i}} E \leftarrow 0$ for each subject i and $\nabla_{I_0} E \leftarrow 0$
 - 11: Current scale $S_j \leftarrow$ maximum scale of $\beta_{0,0}(j)$
 - 12: Dictionary of smooth areas $Z[i][S_j] \leftarrow \emptyset$ for each subject i
 - 13: **repeat**
 - 14: $j \leftarrow j + 1$
 - 15: **for** each subject i **do**
 - 16: Compute the evolution of $\alpha_i(t)$ using Eq. (2.3) {Gradients computation}
 - 17: Compute ϕ_i by solving the flow equation (Eq. (2.1))
 - 18: Deform the template image I_0 with ϕ_i
 - 19: Compute the cost function E as in Eq. (2.7)
 - 20: Compute the gradients $\nabla_{\alpha_{0,i}} E$ and $\nabla_{I_0} E$
 - 21: $\nabla_{\beta_{0,i}} E \leftarrow FWT(\nabla_{\alpha_{0,i}} E) = (a_{S_{max},k}^i)_k \cup (d_{s,k}^{i,o})_{1 \leq s \leq S_{max},k,o}$ {Wavelet transform the gradients with Algorithm 5}
 - 22: **for** each detail coefficient $d_{s,k}^{i,o}$ of $\nabla_{\beta_{0,i}} E$ **do**
 - 23: **if** $s < S_j$ **then**
 - 24: $d_{s,k}^{i,o} \leftarrow 0$ {Finer scale silencing}
 - 25: **else if** $k \in Z[i][s]$ **then**
 - 26: $d_{s,k}^{i,o} \leftarrow 0$ {Smooth areas silencing}
 - 27: **end if**
 - 28: **end for**
 - 29: **end for**
 - 30: $\beta_{0,i}(j) \leftarrow \beta_{0,i}(j-1) - h \times \nabla_{\beta_{0,i}} E$ for all subjects i {Parameter update}
 - 31: $\alpha_{0,i}(j) \leftarrow IWT(\beta_{0,i}(j))$ for all subjects i {Algorithm 6}
 - 32: $I_0(j) \leftarrow I_0(j-1) - h \times \nabla_{I_0} E$
 - 33: Compute the total residual value $r_j = \|\Delta_j\|_1$ according to Eq. (3.5)
 - 34: **if** $\frac{r_{j-1}-r_j}{r_{j-1}} < 0.01$ and $S_j > 1$ **then**
 - 35: $S_j \leftarrow S_{j-1} - 1$ {Scale refinement step}
 - 36: $Z \leftarrow$ **Local_adaptation**(Z, S_j) {Algorithm 7}
 - 37: **end if**
 - 38: **until** Convergence
 - 39: **return** Template image I_0 and momentum vectors $\alpha_{0,i}$
-

A.3.2 Experiments: atlas estimation on handwritten digits with local adaptation

σ_g	k_g	Original	Multiscale	Multiscale & local adaptation
Atlas estimation				
3	100	15.8 ± 12.2	7.1 ± 0.5	9.4 ± 3.2
2	196	10.1 ± 1.9	4.8 ± 0.4	5.3 ± 0.5
1.5	361	7.3 ± 1.7	6.4 ± 3.4	4.3 ± 1.4
Registration				
3	100	10.0 ± 6.1	6.4 ± 5.3	7.5 ± 6.0
2	196	8.5 ± 5.2	4.2 ± 3.2	6.5 ± 6.3
1.5	361	9.0 ± 5.1	4.0 ± 5.9	5.3 ± 7.0

Table A.1: Performance of the original, multiscale and multiscale with local adaptation algorithms on the dataset of handwritten digits during the training and test phases. Data are mean \pm standard deviation of relative residual error over five experiments. No statistically significant differences were found between the two versions of the multiscale algorithm. σ_g : width of the Gaussian kernel; k_g : number of control points.

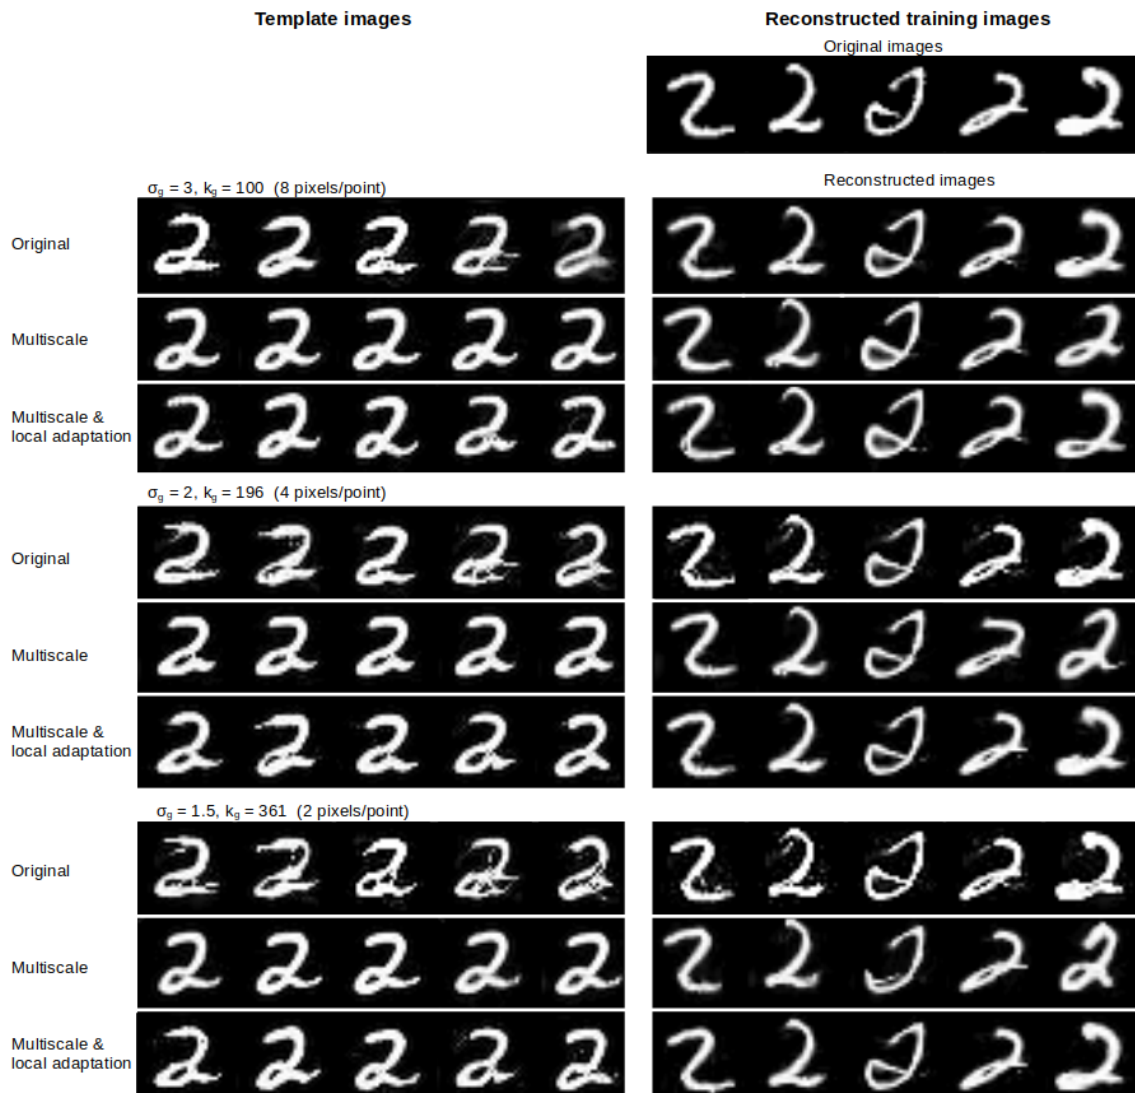


Figure A.3: Estimation of the template image by the original, multiscale and multiscale with local adaptation algorithms on the dataset of handwritten digits with 3 different parameters σ_g . For each experiment, five template images (estimated with non-intersecting training sets) are presented on the left, along with the template image estimated from the first training set wrapped to the first five training images on the right. σ_g : width of the Gaussian kernel; k_g : corresponding number of control points. The number of voxels per control point are indicated in parentheses.

A.4 Alternative multiscale scheme

A.4.1 Modifying the renormalization step

We previously introduced the matrices M_{FWT}^1 and M_{IWT}^1 associated to the transformations FWT and IWT, with $M_{FWT}^1 = R \odot M_{FWT}$ and $M_{IWT}^1 = M_{IWT} \odot R^{-1}$, where \odot denotes the element-by-element matrix product. We recall that IWT and FWT are orthogonal transforms, i.e.: $M_{IWT}^{1T} = M_{FWT}^1$.

We will now try to explore the following question: **how will the optimisation of the momentum vectors be impacted if we modify the renormalization step?**

For the sake of simplicity, in the following we will denote by α and β the coefficients $\alpha_{0,i}$ and $\beta_{0,i}$ of a given subject i , respectively.

We introduce a renormalization factor $\rho \in \mathbb{R}$ and the associated variant of the renormalization matrix: R^ρ , which verifies $R^1 = R$. The matrices associated to this renormalization write as $M_{FWT}^\rho = R^\rho \odot M_{FWT}$ and $M_{IWT}^\rho = M_{IWT} \odot R^{\rho^{-1}}$.

Note that, even if the matrices M_{FWT}^ρ M_{IWT}^ρ are no longer orthonormal when $\rho \neq 1$, this modified renormalization has no impact upon our coarse-to-fine procedure (Algorithm 1), since the gradient of the cost function with respect to the wavelet coefficients is computed using the FWT algorithm (Algorithm 5) as described in section Section 3.3.3.3. Indeed, if we overlook the finer scale silencing step (i.e. the coarse-to-fine step) and simply perform gradient descent in the wavelet domain, we can write:

$$\begin{aligned}
 \beta_{j+1} &= \beta_j + h \times \Delta_\beta E \\
 &= \beta_j + h \times FWT(\Delta_\alpha E) \\
 &= \beta_j + h \times M_{FWT}^\rho \Delta_\alpha E \\
 \alpha_{j+1} &= IWT(\beta_{j+1}) \\
 &= M_{IWT}^\rho \beta_j + h \times M_{IWT}^\rho M_{FWT}^\rho \Delta_\alpha E \\
 &= \alpha_j + h \times M_{IWT} \odot R^{\rho^{-1}} \odot R^\rho \odot M_{FWT} \Delta_\alpha E \\
 &= \alpha_j + h \times M_{IWT} M_{FWT} \Delta_\alpha E \quad \text{and since } M_{IWT} M_{FWT} = I_d \\
 &= \alpha_j + h \times \Delta_\alpha E
 \end{aligned}$$

To assess the influence of ρ , we now compute a variant of the gradient with regard to the wavelet coefficients using the **transpose version** of the IWT algorithm (Algorithm 9) instead of the FWT algorithm, i.e. $\Delta_\beta E' = M_{IWT}^{\rho T} \Delta_\alpha E$, which we recall is equivalent to $\Delta_\beta E' = M_{FWT}^\rho \Delta_\alpha E$ only when $\rho = 1$.

We can now study how this modifies the update of the momentum vectors:

$$\begin{aligned}
 \beta_{j+1} &= \beta_j + h \times M_{IWT}^{\rho^T} \Delta_\alpha E \\
 &= \beta_j + h \times M_{IWT}^T \odot R_\rho^{-1T} \odot \Delta_\alpha E \\
 &= \beta_j + h \times R_\rho^{-1} \odot M_{IWT}^T \Delta_\alpha E \\
 \alpha_{j+1} &= IWT(\beta_{j+1}) \\
 &= M_{IWT}^\rho \beta_j + h \times M_{IWT}^\rho \odot R_\rho^{-1} \odot M_{IWT}^T \Delta_\alpha E \\
 &= \alpha_j + h \times M_{IWT} \odot R^{\rho-1} \odot R^{\rho-1} \odot M_{IWT}^T \Delta_\alpha E \\
 &= \alpha_j + h \times M_{IWT} \odot R^{\rho-2} \odot M_{IWT}^T \Delta_\alpha E
 \end{aligned}$$

Fortunately, we have some information about the unnormalized transform matrices M_{FWT} and M_{IWT} . Since $M_{IWT_1}^T = M_{FWT_1}$, we have $(M_{IWT} \odot R^{-1})^T = R \odot M_{FWT}$ and thus $M_{IWT}^T = R^2 \odot M_{FWT}$. Hence, the update of α becomes:

$$\begin{aligned}
 \alpha_{j+1} &= \alpha_j + h \times M_{IWT} \odot R^{\rho-2} \odot R^2 \odot M_{FWT} \Delta_\alpha E \\
 &= \alpha_j + h \times M_{IWT} \odot R^{\rho-2} \odot R^2 \odot M_{FWT} \Delta_\alpha E \\
 &= \alpha_j + h \times R^{\rho-2} \odot R^2 \odot \Delta_\alpha E
 \end{aligned}$$

And we verify that when $\rho = 1$, we recover $\alpha_{j+1} = \alpha_j + h \times \Delta_\alpha E$. When $\rho \neq 1$ however, the gradient $\Delta_\alpha E$ is multiplied by the matrix $R^{\rho-2} \odot R^2$, which contains coefficients of the form $\sim 2^{sd(1-\rho)}$. Thus, if $\rho > 1$, fine scale coefficients in $\Delta_\alpha E$ will be given more weight than coarse scale coefficients. Conversely, if $\rho < 1$, low frequency components will be given more weight and the velocity fields will be smoother. Hence, by varying the value of ρ , we have another way of performing coarse-to-fine optimisation. If ρ is kept constant, this is equivalent to applying a low pass or high pass filter on the vector fields. In the following experiments, we will combine the modified renormalization step with the regular multiscale strategy.

Algorithm 9 Haar Backward Transpose

- 1: {Compute the wavelet coefficients β of an array X by computing the Inverse wavelet transform matrix M_{IWT} and applying M_{IWT}^T to X }
 - 2: **Input**
 - 3: X : array of shape (K_1, \dots, K_D) , ρ : renormalization factor
 - 4: **Initialization**
 - 5: $\beta, w_s, R \leftarrow \mathbf{Haar_Forward}(X, \rho)$ {Algorithm 5}
 - 6: $n \leftarrow \prod_{i=1}^D K_i$
 - 7: $M_{IWT} \leftarrow$ zero array of shape (K_1, \dots, K_D)
 - 8: **for** i in range(n) **do**
 - 9: $\mathcal{B} \leftarrow$ zero array of shape (K_1, \dots, K_D)
 - 10: i^{th} element of $\mathcal{B} \leftarrow 1$
 - 11: $M_{IWT}[:, i] \leftarrow \text{flatten}(\mathbf{Haar_Backward}(\mathcal{B}, w_s))$ {Algorithm 6}
 - 12: **end for**
 - 13: $\beta \leftarrow$ zero array of shape (K_1, \dots, K_D)
 - 14: **for** i in range(n) **do**
 - 15: i^{th} element of $\beta \leftarrow M_{IWT}^T[i, :] \times \text{flatten}(X)$
 - 16: **end for**
 - 17: **return** β
-

A.4.2 Additional experiments - alternative multiscale algorithm

A.4.2.1 Algorithm

Here, we perform experiments with our regular multiscale algorithm combined with a modified renormalization step. This step can be seen as a (slightly) low pass or high pass filter applied to the velocity fields. Hence, this introduces an imbalance in our coarse to fine procedure, which can be in favor of the coarse or fine scale signal components.

In Appendix A.4, we introduced a renormalization factor ρ and the associated variant of the renormalization matrix: R^ρ . In the normal case (i.e. in the rest of the paper), we have $\rho = 1$ and $R^1 = R$. When computing the wavelet coefficients β of the gradient Δ_{E_α} , we can artificially increase coarse scale coefficients (if $\rho < 1$) or fine scale coefficients (if $\rho > 1$). When applying the Inverse Wavelet Transform to β , we recover a gradient with more high or low frequencies than in the original input Δ_{E_α} . Note that, since we combine this imbalance with coarse-to-fine optimization, the effect of ρ is not always the same: for instance, if the current scale S_j is the coarsest scale (i.e. all fine scale coefficients are silenced by the regular multiscale strategy), the effect of ρ is likely negligible and will increase only as finer and finer scale coefficients are reintegrated.

We can now combine our regular multiscale gradient descent with different renormalization factors ρ . We study the effect of ρ on the algorithm performance by replicating our experiment on the toy dataset with different values of ρ and different kernel width σ_g . Results are presented in Table A.2 (atlas estimation step) and Table A.3 (registration, i.e. test step) and compared with the baseline results, i.e. the error rate when $\rho = 1$, corresponding to the regular multiscale algorithm.

A.4.2.2 Results

We observe that error rates increase when ρ moves away from 1 (i.e. towards 0 or 2), which is unsurprising because this means giving too much weight to high or low frequency components. With larger kernels, i.e. $\sigma_g = 7$ or $\sigma_g = 5$, giving more weight to fine scale components slightly increases the performance. This is interesting as this means that one can choose a sparser parameterization of the vector fields, e.g. to save computational time, while ensuring better performance by selecting higher frequency components.

Table A.2: Performance of the alternative multiscale algorithm with different values of ρ and σ_g during **atlas estimation** on the toy dataset. Data are mean \pm standard deviation of relative residual error over five folds of cross-validation. In bold, we indicate results obtained with $\rho = 1$ (baseline), corresponding to the results of Section 3.4.2. In green, we indicate experiments with similar or higher average performance than the baseline.

$\rho / \sigma_g(k_g)$	7(16)	5 (36)	4 (49)	3 (100)	2.3 (144)	2 (196)
0	16.8 \pm 1.5	13.3 \pm 0.9	11.5 \pm 5.9	12.5 \pm 3.4	13.5 \pm 4.9	15.7 \pm 1.5
0.1	14.1 \pm 1.7	12.3 \pm 1.7	10.3 \pm 4.5	9.5 \pm 1.9	11.5 \pm 2.9	15.4 \pm 3.6
0.2	14.3 \pm 2.7	10.0 \pm 1.2	9.1 \pm 2.0	8.2 \pm 1.0	9.5 \pm 1.9	15.3 \pm 2.6
0.3	14.0 \pm 1.9	9.4 \pm 1.0	8.4 \pm 3.5	6.6 \pm 1.0	12.3 \pm 1.1	12.8 \pm 1.1
0.4	12.7 \pm 1.6	9.5 \pm 1.5	7.1 \pm 1.7	6.8 \pm 1.4	9.7 \pm 0.8	11.7 \pm 1.2
0.5	12.1 \pm 2.0	9.3 \pm 1.4	7.9 \pm 3.7	5.6 \pm 0.8	8.1 \pm 1.6	10.9 \pm 1.1
0.6	12.6 \pm 2.2	7.5 \pm 1.1	5.8 \pm 1.1	5.2 \pm 0.8	7.4 \pm 1.0	8.9 \pm 2.7
0.7	10.6 \pm 1.5	7.3 \pm 1.2	6.6 \pm 2.3	4.3 \pm 0.6	6.2 \pm 0.8	8.2 \pm 2.6
0.8	12.1 \pm 1.7	6.9 \pm 1.0	4.5 \pm 0.1	6.9 \pm 1.6	5.2 \pm 1.8	7.4 \pm 1.2
0.9	10.6 \pm 0.8	5.7 \pm 0.6	4.9 \pm 0.6	4.1 \pm 0.3	5.2 \pm 1.2	6.9 \pm 1.7
1	11.0 \pm 2.0	5.5 \pm 0.8	4.4 \pm 0.4	3.3 \pm 0.5	4.1 \pm 0.9	6.5 \pm 2.7
1.1	10.1 \pm 1.4	5.5 \pm 1.1	4.3 \pm 0.7	2.8 \pm 0.4	4.3 \pm 1.1	4.9 \pm 1.3
1.2	11.5 \pm 2.5	5.3 \pm 1.1	4.0 \pm 0.6	3.0 \pm 0.8	4.2 \pm 1.2	7.1 \pm 4.3
1.3	10.2 \pm 1.4	4.8 \pm 0.8	4.2 \pm 0.6	2.9 \pm 1.0	4.8 \pm 2.6	7.8 \pm 4.2
1.4	10.5 \pm 1.9	4.8 \pm 0.5	4.9 \pm 1.5	4.3 \pm 3.1	7.1 \pm 2.4	7.8 \pm 1.5
1.5	10.7 \pm 1.7	4.8 \pm 0.7	4.5 \pm 1.4	3.6 \pm 2.2	7.6 \pm 3.1	8.0 \pm 2.0
1.6	8.9 \pm 0.7	5.2 \pm 1.0	5.7 \pm 2.3	6.1 \pm 2.1	7.9 \pm 1.8	7.5 \pm 1.2
1.7	10.2 \pm 1.3	6.0 \pm 1.6	6.6 \pm 4.4	6.1 \pm 2.0	11.2 \pm 5.3	7.5 \pm 2.2
1.8	10.8 \pm 1.9	7.8 \pm 2.3	8.4 \pm 5.7	7.7 \pm 3.0	11.7 \pm 4.0	13.5 \pm 5.5
1.9	11.6 \pm 2.9	8.1 \pm 2.0	9.3 \pm 7.7	6.3 \pm 1.7	11.6 \pm 3.4	13.4 \pm 4.4
2.0	13.6 \pm 5.4	9.0 \pm 3.4	10.3 \pm 4.2	9.0 \pm 2.4	12.3 \pm 3.4	13.3 \pm 5.4

Table A.3: Performance of alternative multiscale algorithm with different values of ρ and σ_g during **registration** on the toy dataset. Data are mean \pm standard deviation of relative residual error over five folds of cross-validation. In bold, we indicate results obtained with $\rho = 1$, corresponding to the regular multiscale algorithm and to the results of Section 3.4.2. In green, we indicate experiments with similar or higher average performance than the baseline.

$\rho / \sigma_g(k_g)$	7(16)	5 (36)	4 (49)	3 (100)	2.3 (144)	2 (196)
0	15.5 \pm 6.4	11.8 \pm 8.2	9.4 \pm 10.7	10.4 \pm 10.7	11.4 \pm 9.3	12.8 \pm 10.4
0.1	11.3 \pm 6.0	11.8 \pm 7.6	8.6 \pm 6.8	9.1 \pm 10.6	10.9 \pm 9.6	11.1 \pm 6.8
0.2	10.4 \pm 5.2	9.5 \pm 6.7	7.5 \pm 7.7	6.4 \pm 4.6	7.4 \pm 5.6	9.9 \pm 8.1
0.3	11.0 \pm 6.6	11.6 \pm 11.1	6.4 \pm 5.7	7.3 \pm 8.4	6.6 \pm 4.6	8.5 \pm 7.3
0.4	9.5 \pm 5.8	10.2 \pm 9.4	5.3 \pm 4.3	6.9 \pm 6.7	6.4 \pm 6.0	8.2 \pm 6.4
0.5	16.6 \pm 9.7	7.8 \pm 5.8	6.2 \pm 6.2	5.3 \pm 5.0	7.3 \pm 6.6	8.0 \pm 5.4
0.6	9.7 \pm 6.2	7.7 \pm 6.0	5.9 \pm 6.4	6.2 \pm 6.3	8.1 \pm 7.4	7.7 \pm 6.1
0.7	9.4 \pm 6.6	13.4 \pm 12.3	4.7 \pm 3.8	7.6 \pm 9.1	5.8 \pm 6.8	6.4 \pm 5.8
0.8	8.0 \pm 4.4	5.6 \pm 4.0	4.3 \pm 3.3	4.6 \pm 3.9	6.5 \pm 6.5	5.4 \pm 5.1
0.9	8.2 \pm 5.7	6.5 \pm 5.6	5.7 \pm 5.6	4.8 \pm 6.0	4.4 \pm 3.7	6.4 \pm 6.7
1	9.4 \pm 6.4	6.0 \pm 6.3	3.8 \pm 2.8	3.6 \pm 3.8	3.4 \pm 6.3	3.4 \pm 3.8
1.1	7.2 \pm 3.8	5.6 \pm 4.1	4.1 \pm 3.9	5.1 \pm 6.8	4.2 \pm 5.2	4.9 \pm 5.6
1.2	9.2 \pm 5.9	5.0 \pm 4.2	4.4 \pm 3.9	3.7 \pm 3.7	6.1 \pm 9.9	7.8 \pm 7.3
1.3	7.0 \pm 4.3	4.0 \pm 4.3	4.8 \pm 5.9	3.3 \pm 4.4	5.8 \pm 7.1	8.6 \pm 9.8
1.4	7.6 \pm 5.1	4.9 \pm 4.4	5.3 \pm 6.2	3.9 \pm 4.4	6.2 \pm 7.6	9.1 \pm 7.8
1.5	8.1 \pm 5.5	5.0 \pm 4.0	7.1 \pm 7.2	5.1 \pm 7.1	7.8 \pm 6.1	10.1 \pm 10.5
1.6	7.2 \pm 4.5	6.2 \pm 7.0	5.1 \pm 3.8	6.5 \pm 7.0	8.7 \pm 17.8	10.7 \pm 13.4
1.7	7.2 \pm 4.6	7.6 \pm 13.0	6.5 \pm 8.5	9.8 \pm 10.6	8.1 \pm 5.5	7.1 \pm 7.5
1.8	7.2 \pm 3.4	9.4 \pm 12.6	9.8 \pm 13.7	11.5 \pm 13.7	10.5 \pm 12.7	11.1 \pm 11.9
1.9	12.3 \pm 12.7	13.5 \pm 6.1	9.2 \pm 12.6	10.0 \pm 15.8	11.0 \pm 12.3	12.6 \pm 14.4
2.0	12.2 \pm 11.6	11.7 \pm 16.2	12.2 \pm 22.0	12.2 \pm 13.0	12.4 \pm 12.0	11.8 \pm 15.1

A.5 Additional results - fetal brain MRI




Figure A.4: Atlas estimation by the original and multiscale algorithms on the dataset of fetal brain images with $\sigma_g = 7$. The estimated template volume from the first fold of cross-validation is presented in the left column. The first row displays the first two training images; the middle and right columns display the corresponding reconstructed images (i.e. the template image warped to the training images). Salient differences between images are indicated by specific markers. White arrows indicate the cingulum bundle, blue arrows the left frontal horn, red arrows the left basal ganglia, orange circles the posterior part of the right lateral ventricle, red circles the right superior temporal sulcus, yellow arrows the theoretical location of the corpus callosum, and green arrows the cingulate sulcus.

SUPPLEMENTARY MATERIAL OF CHAPTER 4

B.1 Extrinsic mesh smoothing

Instead of modifying the object itself, the idea of extrinsic smoothing is to smooth its varifold representation. More precisely, we recall that the varifold representation of a mesh O_1 is the distribution of its triangle centers $(c_k^1)_k$ with unit normal vectors $(n_k^1)_k$ attached and that the varifold distance writes as:

$$d_W(O_1, O_2)^2 = \langle O_1, O_1 \rangle_{W'} + \langle O_2, O_2 \rangle_{W'} - 2 \times \langle O_1, O_2 \rangle_{W'}$$

$$\text{with} \quad \langle O_1, O_2 \rangle_{W'} = \sum_{k=1}^{m_1} \sum_{l=1}^{m_2} K_w(x_k^1, x_l^2) \frac{(n_k^1 \cdot n_l^2)^2}{\|n_k^1\| \|n_l^2\|}$$

where K_w is a Gaussian kernel of width σ_w , m_1 and m_2 are the number of triangles in the meshes O_1 and O_2 (respectively), $(x_k^1)_k$ and $(x_k^2)_k$ the cell *centers* of O_1 and O_2 and $(n_k^1)_k$ and $(n_k^2)_k$ the cell *normals* of O_1 and O_2 .

To smooth the mesh O_1 , we replace its normal vectors $(n_k^1)_k$ by their Gaussian convolution:

$$n_k^1 = \sum_{j=1}^{m_1} K_s(x_k^1, x_j^1) n_j^1$$

where K_s is a Gaussian kernel of width S'_j .

S'_j is set in manner similar to the Gaussian image filter (Section 4.2.3.2):

$$S'_j = \begin{cases} \frac{\max(\text{supp}(\Phi_{S_j})) - \min(\text{supp}(\Phi_{S_j}))}{3} \times \sigma_g & \text{if } S_j > 1 \\ 0 & \text{otherwise} \end{cases},$$

where $\text{supp}(\Phi_{S_j, (0,0)})$ is the support of the wavelet function of scale S_j located at position $(0, 0)$.

The effect of varifold smoothing is illustrated in Figure B.1. Note that mesh extrinsic smoothing is likely less efficient than mesh intrinsic smoothing (Section 4.2.3.2), as Laplacian smoothing will smooth the position as well as the orientation of the mesh normal vectors.




Figure B.1: Smooth varifold representation of a cortical surface mesh (age: 24 GW) at scale $S'_j = 8$. The red square indicates the zoomed area. Red arrows indicate the normal vectors, scaled by a factor of 10.

B.2 Dual multiscale Algorithm

Algorithm 10 Multiscale objects optimization algorithm (applied to atlas estimation)

- 1: {For clarity, differences with Algorithm 1 are colored in blue }
 - 2: **Input**
 - 3: Set of objects $(O_i)_{1 \leq i \leq N}$ of dimension d , template object O_{ref} , geometric kernel width σ_g , trade-off regularity/fidelity-to-data σ , initial step size h
 - 4: **Initialization**
 - 5: $j \leftarrow 0$
 - 6: $O_{ref}(j) \leftarrow O_{ref}$ and $\alpha_{0,i}(j) \leftarrow 0$ for each subject i
 - 7: $\beta_{0,i}(j) \leftarrow FWT(\alpha_{0,i}(j))$ for each subject i
 - 8: Current momenta scale $S_j \leftarrow$ maximum scale of $\beta_{0,0}(j)$
 - 9: Current objects scale $S'_j \leftarrow f(S_j)$ {Computation of S'_j using Equation (4.4) or 4.5}
 - 10: Filter O_{ref} and $(O_i)_{1 \leq i \leq N}$ at scale S'_j
 - 11: Order \leftarrow dict('momenta': $[-1, 0] \times (S_{max} - 1)$, 'objects': $[0, -1] \times (S_{max} - 1)$) {Dictionary of the order of the coarse-to-fine steps}
 - 12: **repeat**
 - 13: $j \leftarrow j + 1$
 - 14: **for** each subject i **do**
 - 15: Compute the evolution of $\alpha_i(t)$ using Equation (2.3) {Gradients computation}
 - 16: Compute ϕ_i by solving the flow equation (Equation (2.1))
 - 17: Deform the template object O_{ref} with ϕ_i
 - 18: Compute the cost function E as in Equation (3.2)
 - 19: Compute the gradients $\nabla_{\alpha_{0,i}} E$ and $\nabla_{O_{ref}} E$
 - 20: $\nabla_{\beta_{0,i}} E \leftarrow FWT(\nabla_{\alpha_{0,i}} E) = (a_{S_{max},k}^i)_k \cup (d_{s,k}^{i,o})_{1 \leq s \leq S_{max},k,o}$ {Algorithm 5}
 - 21: **for** each detail coefficient $d_{s,k}^{i,o}$ of $\nabla_{\beta_{0,i}} E$ **do**
 - 22: **if** $s < S_j$ **then**
 - 23: $d_{s,k}^{i,o} \leftarrow 0$ {Finer scale silencing}
 - 24: **end if**
 - 25: **end for**
 - 26: **end for**
 - 27: $\beta_{0,i}(j) \leftarrow \beta_{0,i}(j-1) - h \times \nabla_{\beta_{0,i}} E$ for all subjects i {Parameter update}
 - 28: $\alpha_{0,i}(j) \leftarrow IWT(\beta_{0,i}(j))$ for all subjects i {Algorithm 6}
 - 29: $O_{ref}(j) \leftarrow O_{ref}(j-1) - h \times \nabla_{O_{ref}} E$
 - 30: Compute the total residual value Δ_j according to Equation (3.5)
 - 31: **if** $\frac{\Delta_{j-1} - \Delta_j}{\Delta_{j-1}} < 0.01$ **then**
 - 32: **if** Order['momenta'][0] $\neq 0$ **then**
 - 33: $S_j \leftarrow S_{j-1} - 1$ {Scale refinement step for the momenta}
 - 34: Order['momenta'] = Order['momenta'][1:]
 - 35: **end if**
 - 36: **if** Order['objects'][0] $\neq 0$ **then**
 - 37: $S'_j \leftarrow f(S_j)$ {Scale refinement step for the objects using Equation (4.4) or 4.5}
 - 38: Order['objects'] = Order['objects'][1:]
 - 39: Filter the objects $(O_i)_{1 \leq i \leq N}$ at scale S'_j
 - 40: **end if**
 - 41: **end if**
 - 42: **until** Convergence
 - 43: **return** Template object O_{ref} and momentum vectors $\alpha_{0,i}$
-

B.3 Tuning the number of components in geodesic regression




Figure B.2: Residual error rate versus the number of components on a piecewise geodesic regression task with the original algorithm. Mean and standard deviation are indicated for a 5-fold cross-validation on a dataset of 30 images.

B.4 Additional results - regression on toy data

Table B.1: Mean \pm standard deviation of relative residual error over five folds of cross-validation for atlas estimation, geodesic regression and registration on the toy dataset. σ_g : width of the Gaussian kernel; k_g : number of control points. P : number of components in geodesic regression.

σ_g	k_g	Algorithm	P	Atlas estimation	Regression	Registration
10	9	Original	1	17.0 ± 7.2	58.2 ± 8.7	42.8 ± 19.9
			5		48.3 ± 14.6	36.4 ± 16.4
		Multiscale	1	11.4 ± 2.0	66.4 ± 5.1	36.7 ± 16.5
			5		36.6 ± 1.5	16.9 ± 9.6
8	16	Original	1	18 ± 5.7	58.3 ± 6.8	43.9 ± 11.3
			5		52.7 ± 4.5	48.5 ± 12.9
		Multiscale	1	7.5 ± 1.7	56.6 ± 3.0	32.7 ± 4.1
			5		35.2 ± 2.7	11.9 ± 5.6
6	25	Original	1	17.8 ± 6.7	62.3 ± 6.5	39.9 ± 14.0
			5		58.5 ± 4.9	51.6 ± 13.3
		Multiscale	1	3.4 ± 1.6	56.6 ± 3.8	24.9 ± 6.2
			5		34.2 ± 2.7	16.7 ± 8.4
5	36	Original	1	13.8 ± 4.6	61.5 ± 7.5	44.4 ± 14.3
			5		57.4 ± 7.6	35.8 ± 8.3
		Multiscale	1	3.6 ± 2.9	59.9 ± 6.0	24.8 ± 5.1
			5		32.9 ± 2.6	12.4 ± 2.6
4	49	Original	1	15.0 ± 5.6	58.8 ± 8.7	47.3 ± 5.7
			5		62.7 ± 4.7	39.5 ± 9.8
		Multiscale	1	3.4 ± 2.5	58.9 ± 7.6	17.6 ± 5.2
			5		33.2 ± 2.2	11.1 ± 2.5

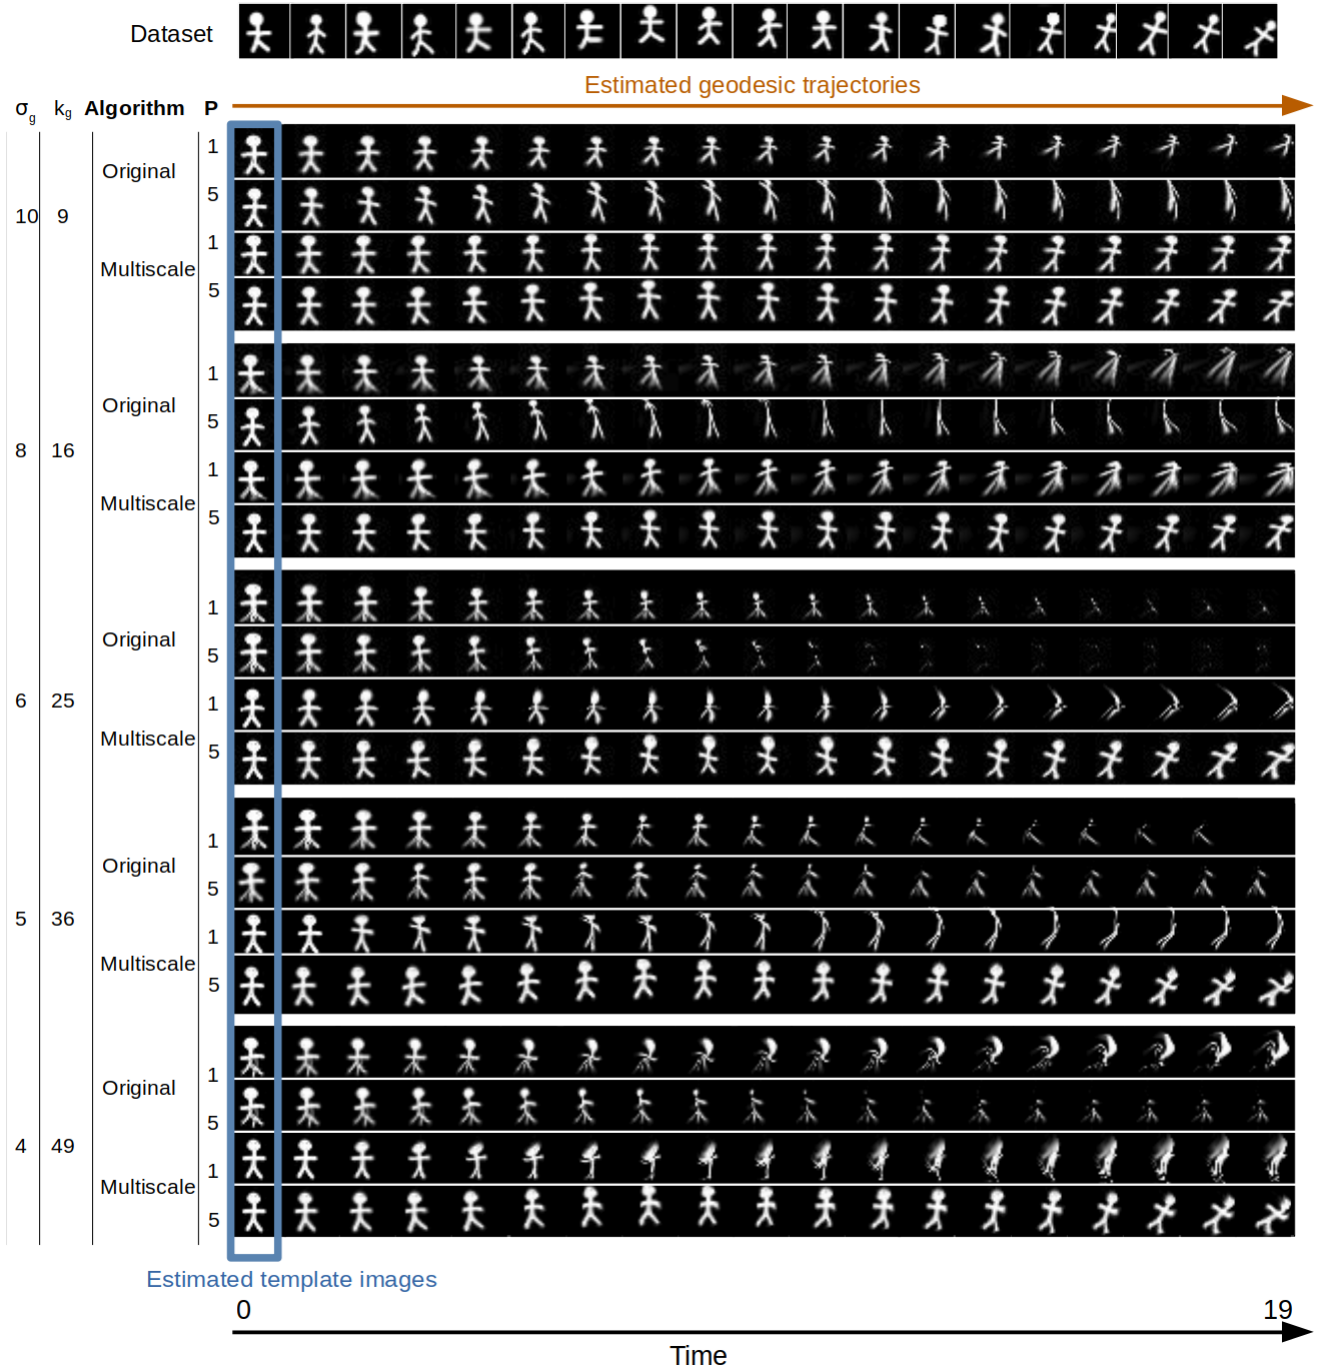


Figure B.3: Geodesic regression results on the toy dataset for $\sigma_g = 10$, $\sigma_g = 8$, $\sigma_g = 6$, $\sigma_g = 5$, and $\sigma_g = 4$. First row: presentation of the dataset (only 19 among the 30 images are shown, ranging from age 0 to 19). We show the results computed on the first fold of cross-validation: images from atlas estimation are outlined in blue and served as starting points to compute the geodesic trajectories from age 0 to 19.

B.5 Additional results - regression on fetal brain images

Table B.2: Mean \pm standard deviation of relative residual error over five folds of cross-validation for atlas estimation, geodesic regression and registration on the dataset of fetal brain images. σ_g : width of the Gaussian kernel; k_g : number of control points. P : number of components in geodesic regression.

σ_g	k_g	Algorithm	P	Atlas estimation	Regression	Registration
20	560	Original	1	26.3 ± 1.3	74.4 ± 21.8	27.3 ± 9.9
			4		70.6 ± 16.0	27.3 ± 10.5
		Multiscale	1	24.2 ± 0.7	57.8 ± 4.6	27.7 ± 12.3
			4		50.8 ± 2.9	28.9 ± 11.9
10	4, 256	Original	1	20.7 ± 2.8	73.7 ± 11.9	23.8 ± 7.9
			4		68.7 ± 11.9	23.9 ± 10.6
		Multiscale	1	23.2 ± 1.4	55.5 ± 3.4	22.3 ± 8.6
			4		51.5 ± 2.4	21.4 ± 11.4
8	8, 160	Original	1	17 ± 1.7	77 ± 19.1	18.3 ± 7.7
			4		71.6 ± 8.0	19.5 ± 8.8
		Multiscale	1	15.0 ± 1.9	57.3 ± 14.4	20.9 ± 8.2
			4		49.3 ± 18.7	20.3 ± 10.4

Table B.3: SSIM values between the templates estimated during geodesic regression and Gholipour et al. atlas [74] at each gestational week. Values are mean (top rows) and standard deviations (bottom rows) over 5 folds of cross-validation. One can notice that smaller regularizing kernels achieve higher SSIM values. Interestingly, the piecewise-multiscale strategy provides benefit only for younger gestational ages and lower values of σ_g . After 29 GW, all models and strategies have equal structural similarity to the state-of-the-art templates.

σ_g	k_g	Algorithm	P	22	23	24	25	26	27	28	29	30	31	32	33	34	35		
20	560	Original	1	0.86	0.84	0.84	0.82	0.81	0.80	0.79	0.78	0.76	0.75	0.73	0.72	0.70	0.68		
				0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	
			4	0.84	0.84	0.85	0.82	0.81	0.80	0.79	0.78	0.76	0.75	0.73	0.72	0.71	0.69		
				0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	
		Multiscale	1	0.89	0.88	0.87	0.85	0.83	0.82	0.80	0.79	0.77	0.75	0.73	0.73	0.71	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.01	0.00	0.01	0.00	0.00		
			4	0.88	0.88	0.87	0.85	0.83	0.82	0.80	0.79	0.78	0.75	0.74	0.73	0.70	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		
		10	4, 256	Original	1	0.84	0.84	0.83	0.83	0.81	0.80	0.79	0.78	0.77	0.75	0.74	0.73	0.71	0.69
						0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
					4	0.84	0.84	0.84	0.83	0.81	0.80	0.79	0.78	0.77	0.75	0.74	0.73	0.71	0.69
						0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00
Multiscale	1			0.84	0.84	0.84	0.83	0.81	0.80	0.79	0.78	0.77	0.75	0.74	0.73	0.71	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		
	4			0.89	0.88	0.88	0.85	0.83	0.82	0.81	0.80	0.78	0.76	0.74	0.73	0.71	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		
8	8, 160			Original	1	0.84	0.84	0.83	0.82	0.81	0.81	0.79	0.78	0.77	0.76	0.74	0.73	0.71	0.70
						0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	
					4	0.83	0.83	0.84	0.83	0.81	0.81	0.79	0.78	0.77	0.76	0.74	0.73	0.71	0.69
						0.02	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	
		Multiscale	1	0.84	0.83	0.84	0.83	0.81	0.81	0.80	0.79	0.77	0.76	0.74	0.73	0.71	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
			4	0.88	0.88	0.87	0.85	0.83	0.82	0.81	0.80	0.78	0.76	0.74	0.73	0.71	0.69		
				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			

SUPPLEMENTARY MATERIAL OF CHAPTER 6

C.1 Fetal MRI processing pipeline

C.1.1 Semi-automated brain extraction and volume reconstruction pipeline

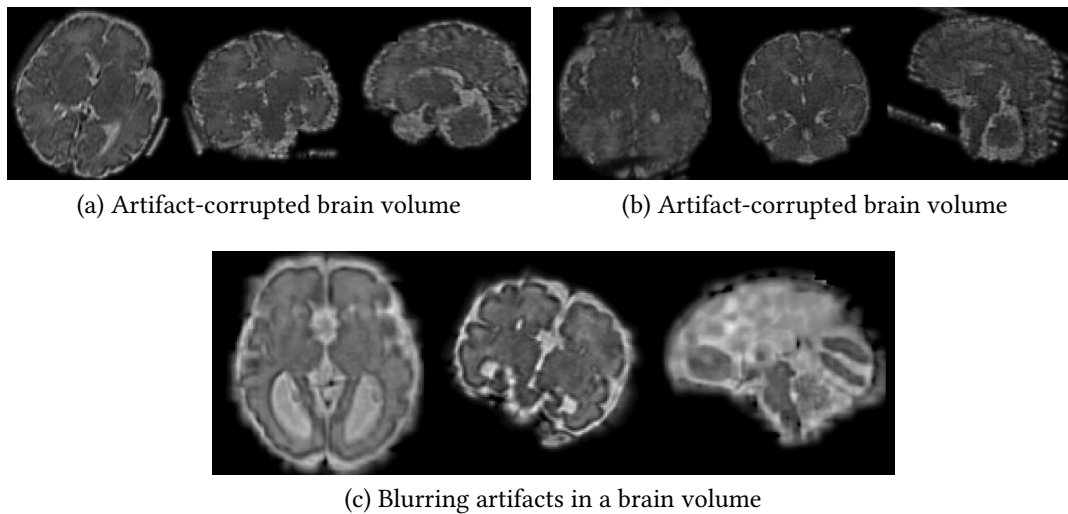


Figure C.1: Volume-reconstructed fetal brain images excluded from the analysis.

Following volume reconstruction with NiftyMIC, 46% of brains with CCA and 24% of healthy brains have erroneously reconstructed images as detailed in Fig. 6.3. Most of the time, this is caused by incorrect delineation of the brain, which has a high rate of false positive voxels (see for example Subjects 2 and 3 in Fig. C.3). Instead of correcting manually the volumetric brain masks, which would be very time consuming, we design a semi-automated volume reconstruction pipeline, which is described in Fig. 6.3. In cases where brain extraction with NiftyMIC is erroneous, we reiterate the extraction step using the U-net CNN from [198]. U-net is affected by the same defect as NiftyMIC as it yields a high proportion of false positive voxels, and was deemed less efficient than NiftyMIC in previous experiments [56]. Thus, our goal is not to extract a more relevant brain mask with the U-net algorithm, but to combine the extraction results from U-net and NiftyMIC. Fusion is carried out in the following manner: brain extraction is performed

with NiftyMIC and U-net, and two brain volumes are reconstructed; we then extract the two volumetric masks yielded by both reconstructions; the intersection of the masks is computed and used to re-mask the two volumetric images. Visual examinations of the re-masked images are performed and the reconstruction with highest quality is selected. This mask fusion procedure is illustrated in the second column of Fig. C.3.

In 23% of fetuses with CCA and 11% of healthy fetuses, the mask fusion yields poor results, characterized by the incomplete elimination of false positive voxels. In such cases, manual refinement of the volumetric masks are performed using ITK-SNAP, Version 3.6 [261], in approximately 10 minutes per subject. Note that the mask fusion step renders the manual correction less time-consuming as it eliminates large amounts of false positive voxels. The third column of Fig. C.3 shows an example of a fetal brain for which manual correction was performed. In this worst-case scenario, the semi-automated reconstruction pipeline takes approximately one hour per subject, compared with 20 minutes in the optimal scenario.




Figure C.2: Volumetric image postprocessing pipeline

C.1.2 Volumetric image postprocessing pipeline

Reorientation. As the fetus orientation is unknown during image acquisition, it is necessary to identify the coronal, sagittal and axial planes of the reconstructed image and rotate them if necessary. First, the inferior-superior, antero-posterior and right-left axis are automatically identified based on length and symmetry measurements: the size of the brain is computed along each of the three axes, and the antero-posterior axis is identified as the axis with the highest length measure. Then, symmetry indices are computed along the two remaining axes, and the right-left axis is identified as the axis with the highest symmetry index. Finally, we must ensure that each of the axes is correctly oriented. To reorient the antero-posterior and inferior-superior axes, we extract the 2D median sagittal plane. Areas of the putative anterior half and posterior half of the brain are compared, and the antero-posterior axis is flipped if the surface of the putative anterior half is larger than that of the putative posterior half. Similar area measurements on the putative superior and inferior halves of the brain are performed in order to reorient the inferior-superior axis.

Mask correction. The subject brains are aligned and cropped to a size of 105x100x120 voxels. To correct small errors (i.e. false positive voxels) during the brain extraction step, a correct brain mask is extracted from the reference brain of [73], rigidly registered to each erroneous fetal brain mask using Deformetrica, Version 4.3.0 [23], and used to re-mask the fetal brain. Note that this automated mask correction step can only be applied to mildly erroneous brain masks (see Subject 1 in Fig. C.3) as the overall shape of the brain has to resemble that of the template brain. Brain masks that are already correct, such as that of Subject 3 in Fig. C.3, go through this step without being affected.

Affine registration to a reference template. To enable inter-subject comparisons and eliminate position and size differences, fetal brains are spatially normalized by performing affine registration to a common anatomical space, namely an age-matched template brain from Gholipour et al. atlas [73]. Finally, intensity normalization and histogram matching to the template are performed.

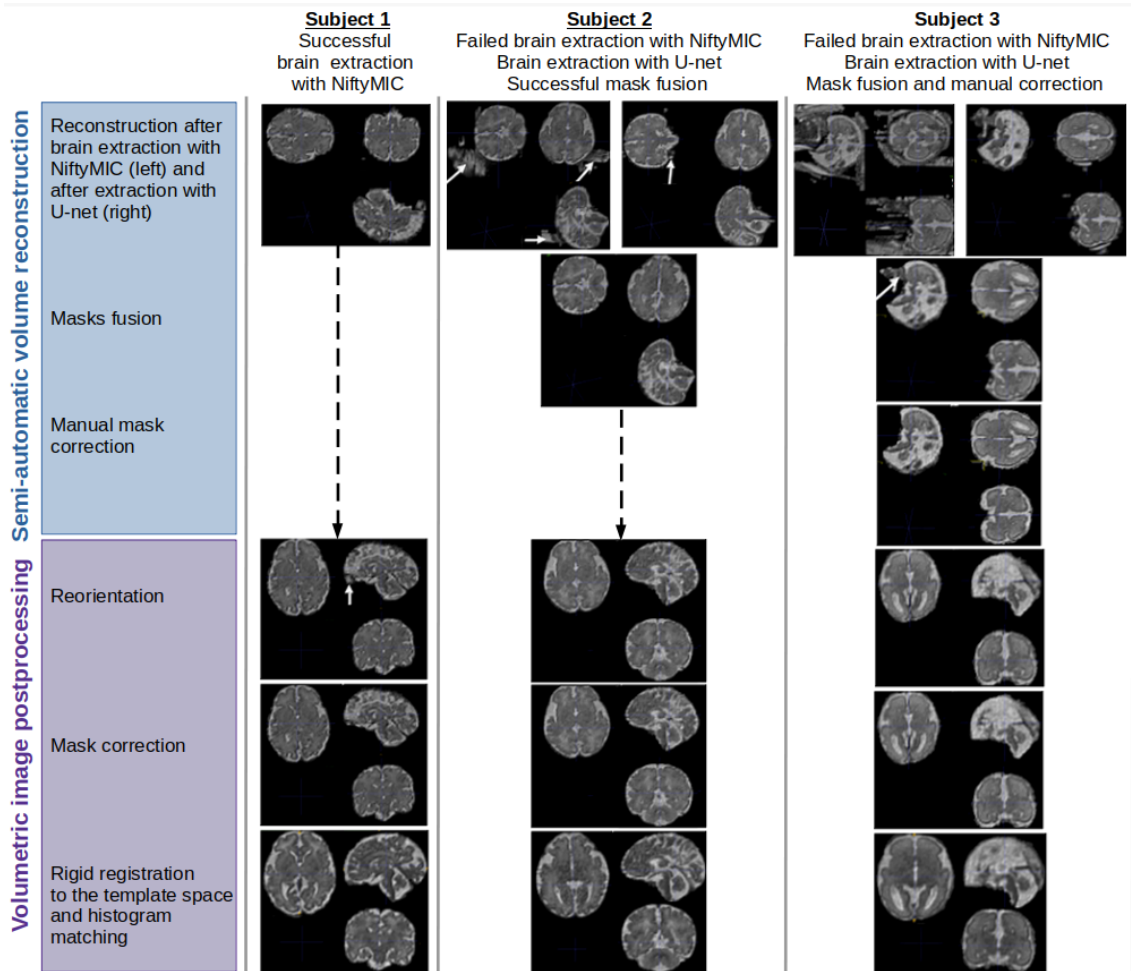
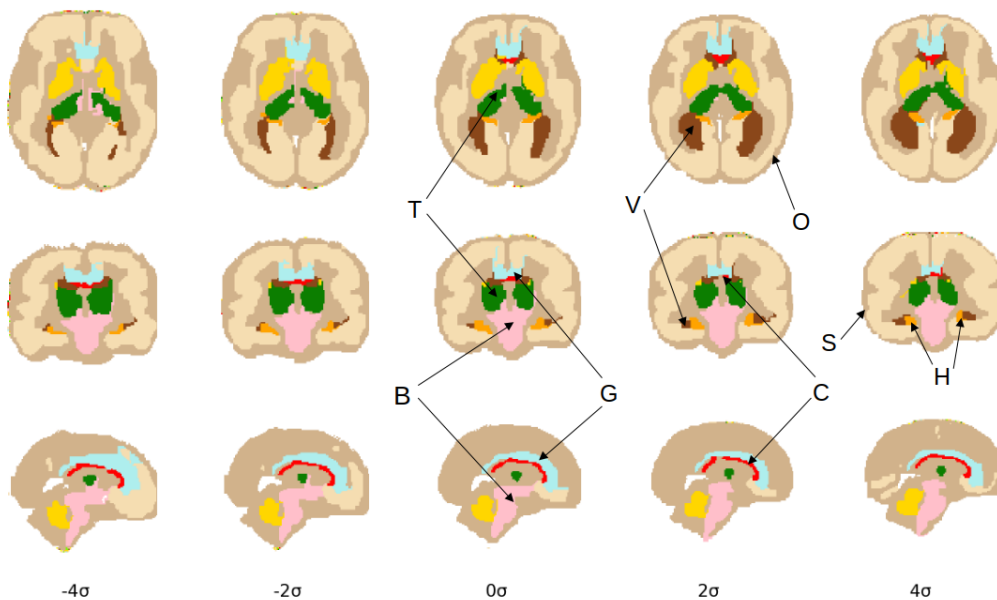



Figure C.3: Illustration of the fetal MRI postprocessing pipeline on three examples. Top three rows: semi-automatic volume reconstruction pipeline. Three bottom rows: volumetric image postprocessing pipeline. First column: reconstruction of Subject 1 after brain extraction with NiftyMIC yields a good quality image and no mask correction is needed. Second column: volumetric images of Subject 2 obtained after brain extraction with NiftyMIC (left) and U-net (right). The intersection of the two erroneous masks is computed, yielding a correct volumetric mask. Third column: volumetric images of Subject 3 obtained after brain extraction with NiftyMIC (left) and U-net (right). Fusion of the two erroneous masks does not eliminate all false positive voxels, hence manual correction is performed. White arrows indicate groups of voxels erroneously classified as fetal brain.

C.2 Additional results - deformation modes

C.2.1 First mode of deformation (healthy brains versus CCA)




(a) First mode of deformation applied to the parcellation image of the template brain at age 31 GW.




(b) Distribution of the subjects scores on component 1.

Figure C.4: First mode of deformation extracted from PCA at -4σ , -2σ , 0σ , $+2\sigma$, and $+4\sigma$ (a) and distribution of the subjects scores on the related component (b). Healthy subjects are more likely to score lower on the related component than subjects with CCA, though their distributions are closer to that of component 2. In panel (a), the deformed template is presented in axial, coronal and sagittal views. B: brainstem. C: corpus callosum. G: cingulate gyrus. H: hippocampi. I: interhemispheric fissure. O: occipital cortex. R: roof of the third ventricle. S: superior temporal sulcus. T: thalami. V: lateral ventricles.

C.2.2 Second mode of deformation in MR intensities (healthy brains versus CCA)




(a) Second mode of deformation applied to the template brain at age 31 GW.




(b) Distribution of the subjects scores on component 2.

Figure C.5: Second mode of deformation extracted from PCA at -4σ , -2σ , 0σ , $+2\sigma$, and $+4\sigma$ (a) and distribution of the subjects scores on the related component (b). In panel (a), the deformed template is presented in axial, coronal and sagittal views. B: brainstem. C: corpus callosum. G: cingulate gyrus. H: hippocampi. I: interhemispheric fissure. O: occipital cortex. R: roof of the third ventricle. S: superior temporal sulcus. T: thalami. V: lateral ventricles.

C.2.3 First mode of deformation in MR intensities (Complete versus partial CCA)



(a) First mode of deformation applied to the segmentation of the template at age 31 GW.



(b) Distribution of the subjects scores on component 1.

Figure C.6: First mode of deformation extracted from PCA at -4σ , -2σ , 0σ , $+2\sigma$, and $+4\sigma$ (a) and distribution of the subjects scores on the related component (b). In panel (a), the deformed template is presented in axial, coronal and sagittal views. B: brainstem. C: corpus callosum. G: cingulate gyrus. H: hippocampi. I: interhemispheric fissure. O: occipital cortex. R: roof of the third ventricle. S: superior temporal sulcus. T: thalami. V: lateral ventricles.

APPENDIX **D**

SUPPLEMENTARY MATERIAL OF CHAPTER 7

D.1 Volume reconstruction pipeline (Lumière dataset)





(a) $\beta = 0.8$; 34/46 rejected slices (22.7 GW)(b) $\beta = 0.5$; 2/46 rejected slices(c) $\beta = 0.8$; 31/33 rejected slices (27.3 GW)(d) $\beta = 0.5$; 3/33 rejected slices(e) $\beta = 0.8$; 50/63 rejected slices (30.3 GW)(f) $\beta = 0.5$; 3/63 rejected slices(g) $\beta = 0.8$; 41/68 rejected slices (35.7 GW)(h) $\beta = 0.5$; 5/68 rejected slices

Figure D.1: Comparison of brain volumes reconstructed with NiftyMIC [56] with an outlier rejection threshold of 0.8 (panels (a), (c), (e) and (g)) and 0.5 (panels (b), (d), (f) and (h)).

D.2 Manual correction of extracted cortical surfaces

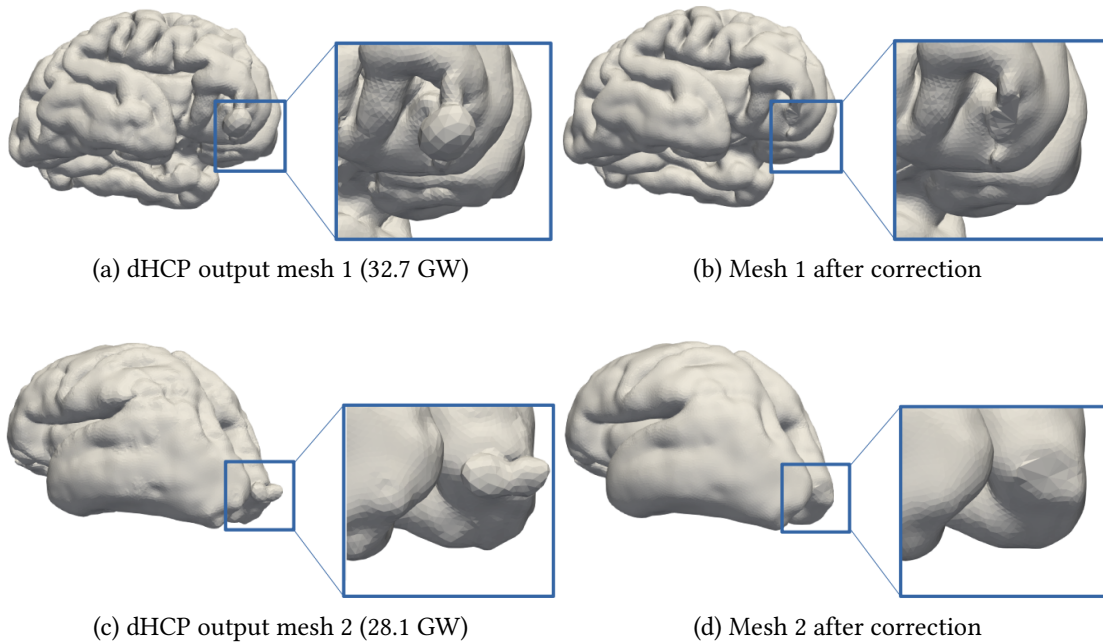


Figure D.2: Cortical surface meshes extracted with the dHCP pipeline before (panels (a) and (c)) and after manual correction and smoothing (panels (b) and (d)). In both cases, vertices belonging to the protrusion on the original mesh were manually selected using Paraview and removed from the mesh. Laplacian smoothing of the protrusion was performed with $n = 5,000$ iterations, resulting in a flat surface, which was subsequently fused to the original mesh. Laplacian smoothing of the entire corrected mesh was performed with $n = 200$ iterations.

D.3 Modification of parallel transport

Here, we summarize the numerical scheme used for parallel transport [135] and describe the changes we made in bold style.

We note $P_{\gamma, t_0, t}(\omega)$ the parallel transport of ω along the trajectory $\gamma(t)$ from time t_0 to a given time t . $P_{\gamma, t_0, t}(\omega)$ is uniquely defined by the integration from $u = t_0$ to t of the differential equation $\nabla \dot{\gamma}(u) P_{\gamma, t_0, u}(\omega) = 0$ with $P_{\gamma, t_0, t_0}(\omega) = \omega$ where ∇ is the Levi-Civita covariant derivative.

We denote by $J_{\gamma(t)}^\omega(h)$ the Jacobi Field emerging from $\gamma(t)$ in the direction $\omega \in T_{\gamma(t)}\mathcal{M}$. The following proposition relates the parallel transport to a Jacobi field:

For all $t > 0$ small enough and $\omega \in T_{\gamma(0)}\mathcal{M}$, we have:

$$P_{\gamma, 0, t}(\omega) = \frac{J_{\gamma(0)}^\omega(t)}{t} + O(t^2) \quad (\text{D.1})$$

To estimate $P_{\gamma, 0, 1}(\omega)$, we can subdivide $[0, 1]$ into N intervals and iteratively compute the Jacobi Fields with an error of $O(\frac{1}{N^2})$ at each step and a total error of $O(\frac{1}{N})$.

D.3.1 Parallel Transport algorithm

Divide $[0, 1]$ into N intervals of length $h = \frac{1}{N}$ where $N \in \mathcal{N}$. We note ω_k the momenta of the transported diffeomorphism, c_k the control points and α_k the momenta of the geodesic γ at time $\frac{k}{N}$.

We iterate the following procedure N times:

1. Compute the controls points c_{k+1} momentum vectors α_{k+1} along the main geodesic at time $\frac{k+1}{N}$ using a Runge-Kutta 2 method to solve the Hamiltonian equations (see Equation (2.3)).
2. Compute the control points $c_{k+1}^{\pm h}$ of the perturbed geodesic $\gamma_{\pm h}$ with initial momenta $\alpha_k \pm h\omega_k$ and control points c_k by solving the Hamiltonian equations
3. Approximate the Jacobi field J_{k+1} by $J_{k+1} = \frac{c_{k+1}^{+h} - c_{k+1}^{-h}}{2h}$
4. Compute the cometric matrix (i.e. RKHS matrix) $K_{c_{k+1}}$ Note that this step is expensive as K_{c_0} has dimension $d \times k_g$ by $d \times k_g$ with k_g the number of control points.
5. Compute the transported momenta $\tilde{\omega}_{k+1}$ according to Equation (D.1).

$$K_{c_{k+1}} \tilde{\omega}_{k+1} = J_{k+1} \quad (\text{D.2})$$

In the initial algorithm, the equation is solved by computing the inverse cometric matrix $K_{c_{k+1}}^{-1}$. But when the number of control points is large, the inversion may have a lot of inadequate solutions. Hence, we invert the regularized matrix $K_{c_{k+1}}^R: K_{c_{k+1}}^R = K_{c_{k+1}} + \lambda I_d$, where λ is a regularizing factor. As $K_{c_{k+1}} = K_{c_0}$, this step is only performed once.

6. Correct the value $\tilde{\omega}_{k+1}$ to ensure preservation of the scalar product $\langle \alpha_{k+1}, \omega_{k+1} \rangle_{c_{k+1}} = \langle \alpha_k, \omega_k \rangle_{c_k}$ and conservation of the norm, i.e. $\|\omega_{k+1}\|_V = \|\omega_k\|_V$.

D.3.2 Parallel Transport example

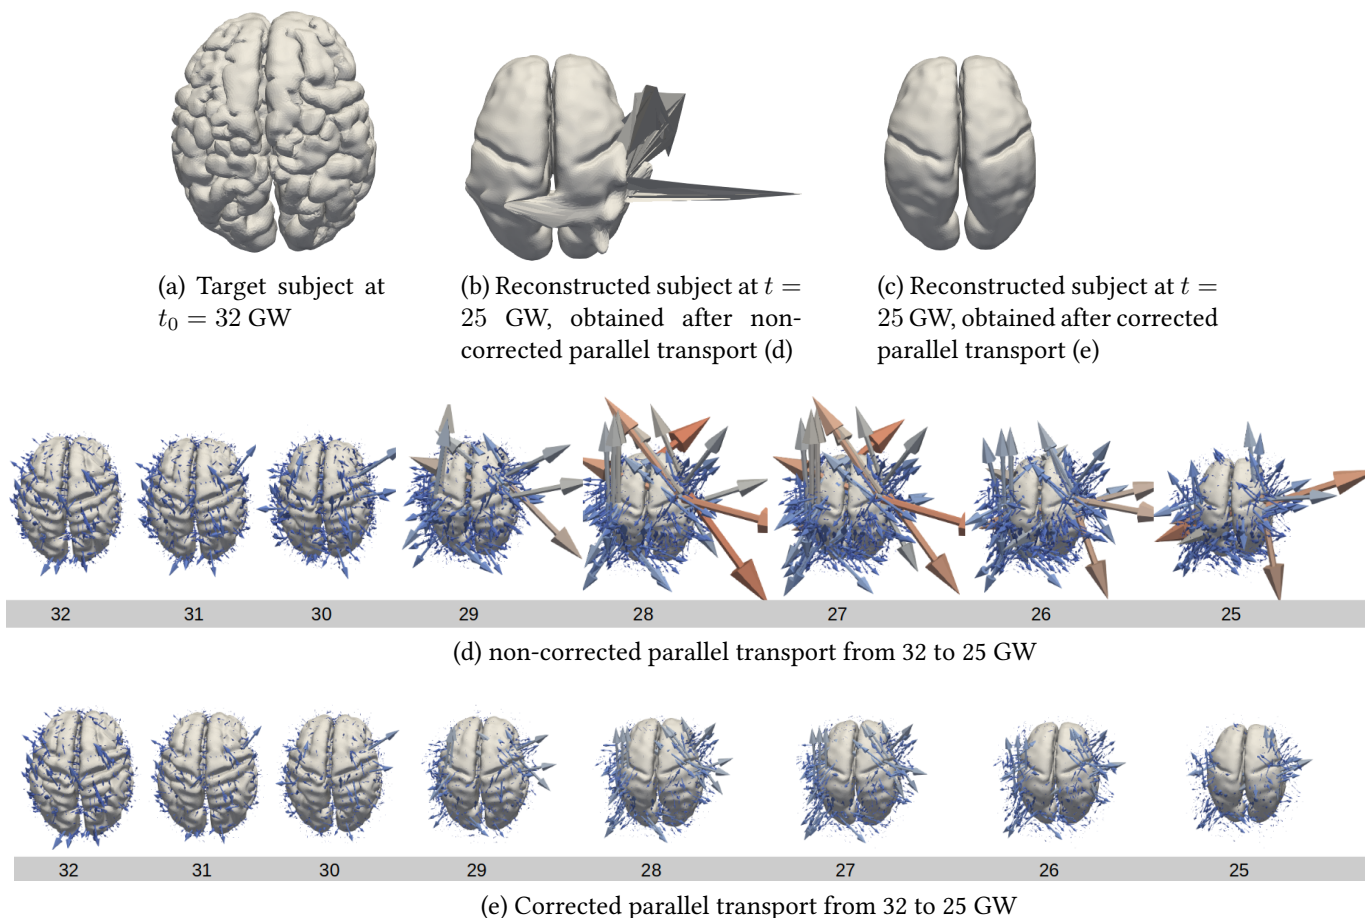


Figure D.3: Parallel transport of a cortical surface mesh (panel (a)) from $t_0 = 32$ to $t = 25$ GW. In panels (d) and (e), we display the reference trajectory of the cortical template warped by the reference geodesic γ . At 32 GW, registration computes the vector field (blue arrows, scaling factor = 10) that optimally warps the age-matched template (i.e. age 32 GW) to the subject. The momenta are then transported from the template space at age 32 GW to template space at age 25 GW using either the non-corrected parallel transport (panel (d)) either the corrected parallel transport (panel (e)). At 25 GW, shooting of the cortical template with the transported momenta yields the reconstructed cortex of the subject (panels (b) and (c)), i.e. the expected appearance of the subject at 25 GW. In panel (d), improper inversion of the cometric matrix yields aberrant transported momenta from 29 GW and thus aberrant reconstruction at 25 GW (panel (b)).

D.4 Optimization procedure

D.4.1 Hyper-parameters setting

The setting of the hyper-parameters depends either on the initial parameters $\theta^{[0]}$ provided by the user (hence the importance of the initialization), or on other fixed parameters that are provided by the user, like the deformation kernel K_g . The prior distributions impose that the parameters cannot deviate too much from the initial values provided by the initialization procedure.

Prior distribution of the rupture times. $\forall 1 \leq l \leq P$ $\bar{t}_l = t_l^{[0]}$; s_t^2 is provided by the user.

Prior distribution of the momentum vectors.

$\forall 1 \leq l \leq P$ $\overline{\alpha_{0,l}} = \alpha_0, l^{[0]}$; $s_\alpha = \|K_c\|_F$ with K_c the $(k_g \times d) \times (k_g \times d)$ RKHS kernel matrix (i.e. the matrix of distances between control points as defined by the kernel K_g) and $\|\cdot\|_F$ the Frobenius norm.

Prior distribution of the template shape.

$\overline{O_{ref}} = O_{ref}^{[0]}$; $s_r = \sigma_g$, with σ_g the standard deviation of the deformation kernel K_g .

Prior distribution of the modulation matrix.

$\overline{A_{t_0}} = A_{t_0}^{[0]}$; $s_a = \sigma_g$, with σ_g the standard deviation of the deformation kernel K_g .

Prior distribution of the noise parameter.

The degree of freedom hyper-parameter n_ϵ is calculated by computing the degree of freedom of the problem.

Scale parameters: $v_\epsilon = 0.01 \frac{r_0}{n_\epsilon}$, where r_0 is the initial residual value $r_0 = \sum_{i=1}^N d(O_i(t_i), \gamma_i(t_i) \star O_{ref})$

The initial noise standard deviation parameter is also set to: $\sigma_\epsilon^{[0]} = v_\epsilon$

D.4.2 Bayesian prior log likelihoods

Sources: $\log(p(s_i)) = -\frac{1}{2}s_i^2$

Population parameters: $\log(p(\theta^{pop})) = \sum_{l=1}^P \log(p(t_l)) + \sum_{l=1}^P \log(p(\alpha_{0,l})) + \log(p(O_{ref})) + \log(p(A_{t_0})) + \log(p(\sigma_\epsilon))$

where:

- $\forall l \in [1, P]$ $\log(p(t_l)) = -\frac{1}{2} \frac{(t_l - \bar{t}_l)^2}{s_t^2}$
- $\forall l \in [1, P]$ $\log(p(\alpha_{0,l})) = -\frac{1}{2} \frac{\|\alpha_{0,l} - \overline{\alpha_{0,l}}\|_2^2}{s_\alpha^2}$

- $\log(p(O_{ref})) = -\frac{1}{2} \frac{\|O_{ref} - \overline{O_{ref}}\|_2^2}{s_r^2}$
- $\log(p(A_{t_0})) = -\frac{1}{2} \frac{\|A_{t_0} - \overline{A_{t_0}}\|_2^2}{s_a^2}$
- $\log(\sigma_\epsilon) = \nu_\epsilon \times (\log(\sigma_\epsilon^2) + \frac{n_\epsilon}{\sigma_\epsilon^2})$

D.4.3 Initialization of the parameters

The initialization procedure essentially solves the same problem but in a *sequential* manner, while the estimation of the main statistical model optimizes these parameters all together.

Given the dataset $(O_i(t_i))_{1 \leq i \leq N}$, the initialization procedure for the of the parameters is as follows:

1. An initial template shape $O_{ref}^{[0]}$ is estimated at time t_0 from the subjects close to t_0 in age using cross-sectional atlas estimation. Note that we use here the traditional kernel weighting of the subjects as in Section 7.3.5
2. Piecewise geodesic regression computes the average trajectory γ that deforms O_{ref} and best fits the observations $(O_i(t_i))_{1 \leq i \leq N}$, yielding the initial set of momentum vectors for the geodesic components $(\alpha_{0,l})_{l \in [1,P]}$
3. For each observation i , the template shape at time t_i , i.e. $\gamma(t_i)(O_{ref})$ is registered to the observation $(O_i(t_i))$, yielding the initial momentum vectors ω_i
4. ω_i is parallel transported to the tangent space at time t_0 , yielding the transported momentum $P_{\gamma,t_i,t_0}(\omega_i)$
5. ICA is performed on the set of transported momenta $(P_{\gamma,t_i,t_0}(\omega_i))_{1 \leq i \leq N}$, yielding the initial modulation matrix $A_{t_0}^{[0]}$ and the sources $(s_i^{[0]})_{1 \leq i \leq N}$.

D.5 Temporal discrepancy in kernel regression




Figure D.4: Temporal discrepancy in kernel regression. Colors indicate local shape index values, ranging from -1 (blue) to 1 (red). Some folds are visible on the template at 22 GW and disappear at 23 GW.

D.6 Hyper-parameters and curvature measures

D.6.1 Deformation and varifold kernels

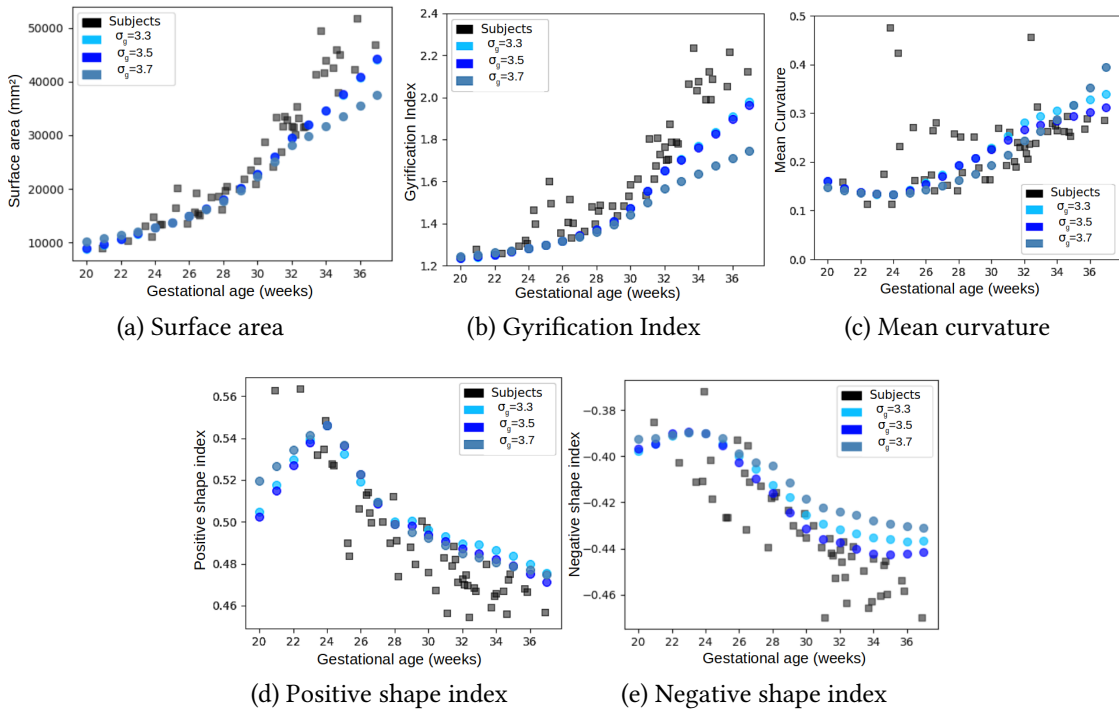


Figure D.5: Comparison of five curvature measures between the subjects and three spatiotemporal atlases build with different values of σ_g and σ_w at each gestational week.

D.6.2 Number of components




Figure D.6: Comparison of five curvature measures between the subjects and four spatiotemporal atlases build with different number of components at each gestational week.

BIBLIOGRAPHY

- [1] Christopher Adamson et al. “Parcellation of the neonatal cortex using Surface-based Melbourne Children’s Regional Infant Brain atlases (M-CRIB-S)”. In: Scientific Reports 10 (Mar. 2020). DOI: [10.1038/s41598-020-61326-2](https://doi.org/10.1038/s41598-020-61326-2).
- [2] James Paul Ahrens, Berk Geveci, and C. Charles Law. “ParaView: An End-User Tool for Large Data Visualization”. In: Visualization Handbook (Jan. 2005).
- [3] Janet Akinmoladun and Dilly O. C. Anumba. “Fetal imaging and diagnosis services in developing countries ; A call to action”. In: Tropical Journal of Obstetrics and Gynaecology 36.1 (2019), pp. 1–7. DOI: [10.4103/TJOG.TJOG_59_18](https://doi.org/10.4103/TJOG.TJOG_59_18).
- [4] Amir Alansary et al. “Automatic Brain Localization in Fetal MRI Using Superpixel Graphs”. In: July 2015, pp. 13–22. ISBN: 3319279289. DOI: [10.1007/978-3-319-27929-9_2](https://doi.org/10.1007/978-3-319-27929-9_2).
- [5] Bonnie Alexander et al. “A new neonatal cortical and subcortical brain atlas: The Melbourne Children’s Regional Infant Brain (M-CRIB) atlas”. In: NeuroImage 147 (Oct. 2016). DOI: [10.1016/j.neuroimage.2016.09.068](https://doi.org/10.1016/j.neuroimage.2016.09.068).
- [6] Stéphanie Allasonnière, Yali Amit, and Alain Trouvé. “Toward a coherent statistical framework for dense deformable template estimation”. In: The Journal of the Royal Statistical Society 69 (Jan. 2007).
- [7] Yali Amit. “A Nonlinear Variational Problem for Image Matching”. In: Siam Journal on Scientific Computing 15 (Jan. 1994). DOI: [10.1137/0915014](https://doi.org/10.1137/0915014).
- [8] Nickie Andescavage et al. “Complex Trajectories of Brain Development in the Healthy Human Fetus”. In: Cerebral cortex (New York, N.Y. : 1991) 27 (Oct. 2016). DOI: [10.1093/cercor/bhw306](https://doi.org/10.1093/cercor/bhw306).
- [9] Jérémie Anquez, Elsa Angelini, and Isabelle Bloch. “Automatic Segmentation of Head Structures on Fetal MRI”. In: Aug. 2009, pp. 109 –112. DOI: [10.1109/ISBI.2009.5192995](https://doi.org/10.1109/ISBI.2009.5192995).
- [10] John Ashburner and Karl J. Friston. “Nonlinear spatial normalization using basis functions”. In: Human Brain Mapping 7.4 (1999), pp. 254–266. DOI: [https://doi.org/10.1002/\(SICI\)1097-0193\(1999\)7:4<254::AID-HBM4>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1097-0193(1999)7:4<254::AID-HBM4>3.0.CO;2-G).
- [11] Ruzena Bajcsy and Chaim Broit. “Matching of deformed images.” In: 1982, 351–353.
- [12] Vincent Beaulé et al. “Cortical thickness in adults with agenesis of the corpus callosum”. In: Neuropsychologia 77 (2015), pp. 359–365. ISSN: 0028-3932. DOI: <https://doi.org/10.1016/j.neuropsychologia.2015.09.020>.

- [13] 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).
- [14] Audrey Bénézit et al. “Organising white matter in a brain without corpus callosum fibres”. In: Cortex 63 (Feb. 2015), pp. 155–171. DOI: [10.1016/j.cortex.2014.08.022](https://doi.org/10.1016/j.cortex.2014.08.022).
- [15] Oualid Benkarim et al. “A novel approach to multiple anatomical shape analysis: Application to fetal ventriculomegaly”. In: Medical Image Analysis 64 (June 2020), p. 101750. DOI: [10.1016/j.media.2020.101750](https://doi.org/10.1016/j.media.2020.101750).
- [16] Oualid Benkarim et al. “Cortical folding alterations in fetuses with isolated non-severe ventriculomegaly”. In: NeuroImage: Clinical 18 (Jan. 2018), pp. 103–114. DOI: [10.1016/j.nicl.2018.01.006](https://doi.org/10.1016/j.nicl.2018.01.006).
- [17] Oualid Benkarim et al. “Toward the automatic quantification of in utero brain development in 3D structural MRI: A review: Quantification of Fetal Brain Development”. In: Human Brain Mapping 38 (Feb. 2017). DOI: [10.1002/hbm.23536](https://doi.org/10.1002/hbm.23536).
- [18] Alain Berlinet and Christine Thomas-Agnan. Reproducing Kernel Hilbert Space in Probability and Statistics. Springer New York, NY, Jan. 2004. ISBN: 978-1-4613-4792-7. DOI: [10.1007/978-1-4419-9096-9](https://doi.org/10.1007/978-1-4419-9096-9).
- [19] Anat Biegon and Chen Hoffmann. “Quantitative magnetic resonance imaging of the fetal brain in utero: Methods and applications”. In: World journal of radiology 6 (Aug. 2014), pp. 523–9. DOI: [10.4329/wjr.v6.i8.523](https://doi.org/10.4329/wjr.v6.i8.523).
- [20] Jelena Bozek et al. “Construction of a neonatal cortical surface atlas using Multimodal Surface Matching in the Developing Human Connectome Project”. In: NeuroImage 179 (June 2018). DOI: [10.1016/j.neuroimage.2018.06.018](https://doi.org/10.1016/j.neuroimage.2018.06.018).
- [21] Martins Bruveris, Laurent Risser, and François-Xavier Vialard. “Mixture Of Kernels And Iterated Semidirect Product Of Diffeomorphisms Groups”. In: SIAM Journal on Multiscale Modeling and Simulation 10 (Aug. 2011). DOI: [10.1137/110846324](https://doi.org/10.1137/110846324).
- [22] Alexandre Bône, Olivier Colliot, and Stanley Durrleman. “Learning Distributions of Shape Trajectories from Longitudinal Datasets: A Hierarchical Model on a Manifold of Diffeomorphisms”. In: June 2018. DOI: [10.1109/CVPR.2018.00966](https://doi.org/10.1109/CVPR.2018.00966).
- [23] Alexandre Bône et al. “Deformetrica 4: An Open-Source Software for Statistical Shape Analysis”. In: Sept. 2018. ISBN: 978-3-030-04746-7. DOI: [10.1007/978-3-030-04747-4_1](https://doi.org/10.1007/978-3-030-04747-4_1).
- [24] Alexandre Bône et al. Deformetrica Version 4.3.0. 2018.
- [25] P. Caro-Domínguez, Leidy García-Díaz, and Monica Rebollo Polo. “Survey about the current use of fetal MRI in Spain”. In: Radiología (English Edition) 65.3 (2023), pp. 195–199. ISSN: 2173-5107. DOI: <https://doi.org/10.1016/j.rxeng.2021.01.004>.

- [26] Marie Cassart and Catherine Garel. “European overview of current practice of fetal imaging by pediatric radiologists: a new task force is launched.” In: Pediatr Radiol. (Nov. 2020), pp. 1794–1798. DOI: [10.1007/s00247-020-04710-4](https://doi.org/10.1007/s00247-020-04710-4).
- [27] Pascal Cathier. “Iconic Feature Registration with Sparse Wavelet Coefficients”. In: vol. 9. Feb. 2006, pp. 694–701. ISBN: 978-3-540-44727-6. DOI: [10.1007/11866763_85](https://doi.org/10.1007/11866763_85).
- [28] Nicolas Charon and Alain Trouvé. “The Varifold Representation of Nonoriented Shapes for Diffeomorphic Registration”. In: SIAM Journal on Imaging Sciences 6 (Dec. 2013), 2547–2580. DOI: [10.1137/130918885](https://doi.org/10.1137/130918885).
- [29] Juliette Chevallier, Vianney Debavelaere, and Stéphanie Allasonnière. “A Coherent Framework for Learning Spatiotemporal Piecewise-Geodesic Trajectories from Longitudinal Manifold-Valued Data”. In: SIAM Journal on Imaging Sciences 14 (Mar. 2021), pp. 349–388. DOI: [10.1137/20M1328026](https://doi.org/10.1137/20M1328026).
- [30] Juliette Chevallier, Stéphane Oudard, and Stéphanie Allasonnière. “Learning spatiotemporal piecewise-geodesic trajectories from longitudinal manifold-valued data”. In: 31st Conference on Neural Information Processing Systems (NIPS 2017). Long Beach, United States, Dec. 2017.
- [31] Gary Christensen, Richard Rabbitt, and Michael Miller. “Deformable template using large deformation kinematics”. In: IEEE transactions on image processing: a publication of the IEEE Signal Processing Society 5 (Feb. 1996), pp. 1435–47. DOI: [10.1109/83.536892](https://doi.org/10.1109/83.536892).
- [32] G.E. Christensen and H.J. Johnson. “Consistent image registration”. In: IEEE Transactions on Medical Imaging 20.7 (2001), pp. 568–582. DOI: [10.1109/42.932742](https://doi.org/10.1109/42.932742).
- [33] Tommaso Ciceri et al. “Review on deep learning fetal brain segmentation from Magnetic Resonance images”. In: Artificial Intelligence in Medicine 143 (June 2023), p. 102608. DOI: [10.1016/j.artmed.2023.102608](https://doi.org/10.1016/j.artmed.2023.102608).
- [34] Cedric Clouchoux and Catherine Limperopoulos. “Novel applications of quantitative MRI for the fetal brain”. In: Pediatr Radiol. 42 (Jan. 2012), pp. 24–32. DOI: [10.1007/s00247-011-2178-0](https://doi.org/10.1007/s00247-011-2178-0).
- [35] Cedric Clouchoux et al. “Delayed Cortical Development in Fetuses with Complex Congenital Heart Disease”. In: Cerebral cortex (New York, N.Y. : 1991) 23 (Sept. 2012). DOI: [10.1093/cercor/bhs281](https://doi.org/10.1093/cercor/bhs281).
- [36] Cedric Clouchoux et al. “Quantitative in vivo MRI measurement of cortical development in the fetus”. In: Brain structure & function 217 (May 2011), pp. 127–39. DOI: [10.1007/s00429-011-0325-x](https://doi.org/10.1007/s00429-011-0325-x).
- [37] Fergus V. Coakley et al. “Fetal MRI: A Developing Technique for the Developing Patient”. In: American Journal of Roentgenology 182.1 (2004). PMID: 14684546, pp. 243–252. DOI: [10.2214/ajr.182.1.1820243](https://doi.org/10.2214/ajr.182.1.1820243).

- [38] James M. Corbett-Detig et al. “3D global and regional patterns of human fetal subplate growth determined in utero”. In: Brain structure and function 215 (Oct. 2010), pp. 255–263. DOI: [10.1007/s00429-010-0286-5](https://doi.org/10.1007/s00429-010-0286-5).
- [39] Sara Bernardes da Cunha et al. “Neurodevelopmental Outcomes following Prenatal Diagnosis of Isolated Corpus Callosum Agenesis: A Systematic Review”. In: Fetal Diagnosis and Therapy 48.2 (Jan. 2021), pp. 88–95. ISSN: 1015-3837. DOI: [10.1159/000512534](https://doi.org/10.1159/000512534).
- [40] Claire Cury et al. “Spatiotemporal analysis for detection of pre-symptomatic shape changes in neurodegenerative diseases: Initial application to the GENFI cohort”. In: NeuroImage 188 (2019), pp. 282–290. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2018.11.063>.
- [41] Sarvamangala D. R. and Raghavendra Kulkarni. “Convolutional neural networks in medical image understanding: a survey”. In: Evolutionary Intelligence 15 (Mar. 2022), pp. 1–22. DOI: [10.1007/s12065-020-00540-3](https://doi.org/10.1007/s12065-020-00540-3).
- [42] Fernand Daffos et al. “Fetal curarization for prenatal magnetic resonance imaging.” In: Prenat Diagn. (May 1988). DOI: [10.1002/pd.1970080412](https://doi.org/10.1002/pd.1970080412).
- [43] Bradley Davis et al. “Population Shape Regression From Random Design Data.” In: vol. 90. Jan. 2007, pp. 1–7. DOI: [10.1007/s11263-010-0367-1](https://doi.org/10.1007/s11263-010-0367-1).
- [44] Vianney Debavelaere, Stanley Durrleman, and Stéphanie Allasonnière. “Learning the Clustering of Longitudinal Shape Data Sets into a Mixture of Independent or Branching Trajectories”. In: International Journal of Computer Vision 128 (Dec. 2020). DOI: [10.1007/s11263-020-01337-8](https://doi.org/10.1007/s11263-020-01337-8).
- [45] Vianney Debavelaere et al. “Clustering of Longitudinal Shape Data Sets Using Mixture of Separate or Branching Trajectories”. In: Oct. 2019, pp. 66–74. ISBN: 978-3-030-32250-2. DOI: [10.1007/978-3-030-32251-9_8](https://doi.org/10.1007/978-3-030-32251-9_8).
- [46] Noémie Debroux, Carole Guyader, and Luminita Vese. “Multiscale Registration”. In: Apr. 2021, pp. 115–127. ISBN: 978-3-030-75548-5. DOI: [10.1007/978-3-030-75549-2_10](https://doi.org/10.1007/978-3-030-75549-2_10).
- [47] Maria Deprez et al. “Reconstruction of fetal brain MRI with intensity matching and complete outlier removal”. In: Medical image analysis 16 (Aug. 2012). DOI: [10.1016/j.media.2012.07.004](https://doi.org/10.1016/j.media.2012.07.004).
- [48] Vincent Des Portes et al. “Outcome of isolated agenesis of the corpus callosum: A population-based prospective study”. In: European Journal of Paediatric Neurology 22.1 (Jan. 2018), pp. 82–92. DOI: [10.1016/j.ejpn.2017.08.003](https://doi.org/10.1016/j.ejpn.2017.08.003).
- [49] Mariana Diogo et al. “Improved prognostication in isolated callosal agenesis: fetal magnetic resonance imaging-based scoring system”. In: Ultrasound in Obstetrics & Gynecology 58 (June 2020). DOI: [10.1002/uog.22102](https://doi.org/10.1002/uog.22102).

- [50] Jenna Dittmar and Piers Mitchell. “From cradle to grave via the dissection room: the role of foetal and infant bodies in anatomical education from the late 1700s to early 1900s”. In: (Dec. 2016). DOI: [10.17863/CAM.10407](https://doi.org/10.17863/CAM.10407).
- [51] Haoran Dou et al. “A Deep Attentive Convolutional Neural Network for Automatic Cortical Plate Segmentation in Fetal MRI”. In: ().
- [52] Priscille de Dumast et al. Segmentation of the cortical plate in fetal brain MRI with a topological loss 2020. arXiv: [2010.12391 \[eess.IV\]](https://arxiv.org/abs/2010.12391).
- [53] Paul Dupuis and Ulf Grenander. “Variational Problems on Flows of Diffeomorphisms for Image Matching”. In: Q. Appl. Math. LVI.3 (Sept. 1998), 587–600. ISSN: 0033-569X.
- [54] Stanley Durrleman, Stéphanie Allasonnière, and Sarang Joshi. “Sparse Adaptive Parameterization of Variability in Image Ensembles”. In: International Journal of Computer Vision 101 (Nov. 2012), pp. 1–23. DOI: [10.1007/s11263-012-0556-1](https://doi.org/10.1007/s11263-012-0556-1).
- [55] Francesco D’Antonio et al. “Outcomes Associated With Isolated Agenesis of the Corpus Callosum: A Meta-analysis”. In: 138.3 (Sept. 2016). e20160445. ISSN: 0031-4005. DOI: [10.1542/peds.2016-0445](https://doi.org/10.1542/peds.2016-0445).
- [56] Michael Ebner et al. “An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI”. In: Neuroimage 206 (2020).
- [57] Michael Ebner et al. “An Automated Localization, Segmentation and Reconstruction Framework for Fetal Brain MRI”. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2018. Ed. by Alejandro F. Frangi et al. Cham: Springer International Publishing, 2018, pp. 313–320. ISBN: 978-3-030-00928-1.
- [58] Lei Feng et al. “Population-averaged macaque brain atlas with high-resolution ex vivo DTI integrated into in vivo space”. In: Brain structure & function 222 (Dec. 2017). DOI: [10.1007/s00429-017-1463-6](https://doi.org/10.1007/s00429-017-1463-6).
- [59] Ahmed E. Fetit et al. “A deep learning approach to segmentation of the developing cortex in fetal brain MRI with minimal manual labeling”. In: Proceedings of the Third Conference on Medical Imaging with Deep Learning. Ed. by Tal Arbel et al. Vol. 121. Proceedings of Machine Learning Research. PMLR, July 2020, pp. 241–261.
- [60] Lucas Fidon et al. “A spatio-temporal atlas of the developing fetal brain with spina bifida aperta”. In: Open Research Europe 1 (Oct. 2021), p. 123. DOI: [10.12688/openreseurope.13914.1](https://doi.org/10.12688/openreseurope.13914.1).
- [61] Lucas Fidon et al. “Distributionally Robust Segmentation of Abnormal Fetal Brain 3D MRI”. In: CoRR abs/2108.04175 (2021). URL: <https://arxiv.org/abs/2108.04175>.
- [62] Lucas Fidon et al. “Label-Set Loss Functions for Partial Supervision: Application to Fetal Brain 3D MRI Parcellation”. In: Sept. 2021.

- [63] Bruce Fischl. “FreeSurfer”. In: *NeuroImage* 62 (Jan. 2012), pp. 774–81. DOI: [10.1016/j.neuroimage.2012.01.021](https://doi.org/10.1016/j.neuroimage.2012.01.021).
- [64] James Fishbaugh et al. “Geodesic shape regression with multiple geometries and sparse parameters”. In: *Medical Image Analysis* 39 (July 2017), pp. 1–17. DOI: [10.1016/j.media.2017.03.008](https://doi.org/10.1016/j.media.2017.03.008).
- [65] Thomas Fletcher. “Geodesic Regression on Riemannian Manifolds”. In: *Proceedings of the Third International Workshop on Mathematical Foundations of Computational Anatomy*. Ed. by Xavier Pennec, Sarang Joshi, and Mads Nielsen. Toronto, Canada, Sept. 2011, pp. 75–86.
- [66] Fleur Gaudfernau, Stéphanie Allassonnière, and Erwan Le Pennec. “A multiscale algorithm for computing realistic image transformations: application to the modelling of fetal brain growth”. In: *Medical Imaging 2023: Image Processing*. Ed. by Olivier Colliot and Ivana Išgum. Vol. 12464. International Society for Optics and Photonics. SPIE, 2023, p. 1246404. DOI: [10.1117/12.2654259](https://doi.org/10.1117/12.2654259).
- [67] Fleur Gaudfernau, Stéphanie Allassonnière, and Erwan Le Pennec. “Un algorithme multi-échelle pour déformer les objets de façon réaliste - application à la modélisation de la croissance du cerveau foetal”. In: *29° Colloque sur le traitement du signal et des images*. 2023-1358. Grenoble: GRETSI - Groupe de Recherche en Traitement du Signal et des Images, Sept. 2023, p. 1093–1096.
- [68] Fleur Gaudfernau, Eléonore Blondiaux, and Stéphanie Allassonnière. “Analysis of the Anatomical Variability of Fetal Brains with Corpus Callosum Agenesis”. In: *MICCAI 2021 - Perinatal, Preterm and Paediatric Image Analysis Workshop*. Ed. by Carole H. Sudre et al. Vol. 12959. LNCS - Lecture Notes in Computer Science. Strasbourg, France: Springer, Oct. 2021. DOI: [10.1007/978-3-030-87735-4_26](https://doi.org/10.1007/978-3-030-87735-4_26).
- [69] Fleur Gaudfernau et al. “Wavelet-Based Multiscale Initial Flow For Improved Atlas Estimation in the Large Diffeomorphic Deformation Model Framework”. Preprint. 2022.
- [70] Smadar Gefen et al. “Surface Alignment of an Elastic Body Using a Multiresolution Wavelet Representation”. In: *IEEE transactions on bio-medical engineering* 51 (Aug. 2004), pp. 1230–41. DOI: [10.1109/TBME.2004.827258](https://doi.org/10.1109/TBME.2004.827258).
- [71] Neda Ghassemi et al. “Ultrasound and Magnetic Resonance Imaging of Agenesis of the Corpus Callosum in Fetuses: Frontal Horns and Cavum Septi Pellucidi Are Clues to Earlier Diagnosis”. In: *Journal of Ultrasound in Medicine* 39 (June 2020). DOI: [10.1002/jum.15348](https://doi.org/10.1002/jum.15348).
- [72] Ali Gholipour, Judy Estroff, and Simon Warfield. “Robust Super-Resolution Volume Reconstruction From Slice Acquisitions: Application to Fetal Brain MRI”. In: *IEEE transactions on medical imaging* 29 (Oct. 2010), pp. 1739–58. DOI: [10.1109/TMI.2010.2051680](https://doi.org/10.1109/TMI.2010.2051680).

- [73] Ali Gholipour et al. “A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth”. In: Scientific Reports 7 (Dec. 2017). DOI: [10.1038/s41598-017-00525-w](https://doi.org/10.1038/s41598-017-00525-w).
- [74] Ali Gholipour et al. “Construction of a Deformable Spatiotemporal MRI Atlas of the Fetal Brain: Evaluation of Similarity Metrics and Deformation Models”. In: vol. 17. Sept. 2014, pp. 292–9. ISBN: 978-3-319-10469-0. DOI: [10.1007/978-3-319-10470-6_37](https://doi.org/10.1007/978-3-319-10470-6_37).
- [75] Ali Gholipour et al. “Fetal brain volumetry through MRI volumetric reconstruction and segmentation”. In: International journal of computer assisted radiology and surgery 6 (May 2011), pp. 329–39. DOI: [10.1007/s11548-010-0512-x](https://doi.org/10.1007/s11548-010-0512-x).
- [76] Ali Gholipour et al. “Multi-Atlas Multi-Shape Segmentation of Fetal Brain MRI for Volumetric and Morphometric Analysis of Ventriculomegaly”. In: NeuroImage 60 (Apr. 2012), pp. 1819–31. DOI: [10.1016/j.neuroimage.2012.01.128](https://doi.org/10.1016/j.neuroimage.2012.01.128).
- [77] Fabrici Girolamo. De Formato Foetu. 1600.
- [78] Sarah Glatter et al. “Beyond Isolated and Associated: A Novel Fetal MR Imaging–Based Scoring System Helps in the Prenatal Prognostication of Callosal Agenesis”. In: American Journal of Neuroradiology 42.4 (2021), pp. 782–786. ISSN: 0195-6108. DOI: [10.3174/ajnr.A7064](https://doi.org/10.3174/ajnr.A7064).
- [79] Orit A. Glenn. “MR imaging of the fetal brain.” In: Pediatric Radiology 40 (2010), 68–81. DOI: [10.1007/s00247-009-1459-3](https://doi.org/10.1007/s00247-009-1459-3).
- [80] Orit A. Glenn and Anthony J Barkovich. “Magnetic Resonance Imaging of the Fetal Brain and Spine: An Increasingly Important Tool in Prenatal Diagnosis, Part 1”. In: American Journal of Neuroradiology 27.8 (2006), pp. 1604–1611. ISSN: 0195-6108.
- [81] Orit A. Glenn et al. “Malformations of Cortical Development: Diagnostic Accuracy of Fetal MR Imaging”. In: Radiology 263.3 (2012). PMID: 22495681, pp. 843–855. DOI: [10.1148/radiol.12102492](https://doi.org/10.1148/radiol.12102492).
- [82] Pietro Gori et al. “A Bayesian Framework for Joint Morphometry of Surface and Curve meshes in Multi-Object Complexes”. In: Medical Image Analysis 35 (Aug. 2016). DOI: [10.1016/j.media.2016.08.011](https://doi.org/10.1016/j.media.2016.08.011).
- [83] Ulf Grenander. General Pattern Theory: A Mathematical Study of Regular Structures. Oxford Mathematical Monographs. Clarendon Press, 1993. ISBN: 9780198536710. URL: <https://books.google.fr/books?id=Z-8YAQAAlAAJ>.
- [84] Ulf Grenander and Michael Miller. “Computational anatomy: An emerging discipline”. In: Quarterly of Applied Mathematics LVI (Dec. 1998), pp. 617–694. DOI: [10.1090/qam/1668732](https://doi.org/10.1090/qam/1668732).
- [85] Barbara Gris, Stanley Durrleman, and Alain Trounev. “A Sub-Riemannian Modular Framework for Diffeomorphism-Based Analysis of Shape Ensembles”. In: SIAM Journal on Imaging Sciences 11 (May 2016). DOI: [10.1137/16M1076733](https://doi.org/10.1137/16M1076733).

- [86] Laura Gui et al. “Developmental Brain Atlases”. In: Brain Mapping: An Encyclopedic Reference 3 (Dec. 2015), pp. 589–602. DOI: [10 . 1016/B978-0-12-397025-1.00067-1](https://doi.org/10.1016/B978-0-12-397025-1.00067-1).
- [87] Piotr Habas et al. “A spatiotemporal atlas of MR intensity, tissue probability and shape of the fetal brain with application to segmentation”. In: NeuroImage 53 (Nov. 2010), pp. 460–70. DOI: [10.1016/j.neuroimage.2010.06.054](https://doi.org/10.1016/j.neuroimage.2010.06.054).
- [88] Piotr Habas et al. “Early Folding Patterns and Asymmetries of the Normal Human Brain Detected from in Utero MRI”. In: Cerebral cortex (New York, N.Y. : 1991) 22 (May 2011), pp. 13–25. DOI: [10.1093/cercor/bhr053](https://doi.org/10.1093/cercor/bhr053).
- [89] Trevor Hastie et al. “The Elements of Statistical Learning: Data Mining, Inference, and Prediction”. In: Math. Intell. 27 (Nov. 2004), pp. 83–85. DOI: [10.1007/BF02985802](https://doi.org/10.1007/BF02985802).
- [90] Leonard R. Herrmann. “Laplacian-Isoparametric Grid Generation Scheme”. In: Journal of Engineering Mechanics-asce 102 (1976), pp. 749–907. URL: [https : / / api.semanticscholar.org/CorpusID:123716570](https://api.semanticscholar.org/CorpusID:123716570).
- [91] Steven W. Hetts et al. “Anomalies of the Corpus Callosum: An MR Analysis of the Phenotypic Spectrum of Associated Malformations”. In: American Journal of Roentgenology 187.5 (2006). PMID: 17056927, pp. 1343–1348. DOI: [10.2214/AJR.05.0146](https://doi.org/10.2214/AJR.05.0146).
- [92] Jagoda Hofman-Hutna et al. “Corpus Callosum Agenesis: An Insight into the Etiology and Spectrum of Symptoms”. In: Brain sciences 10 (Sept. 2020). DOI: [10 . 3390 / brainsci10090625](https://doi.org/10.3390/brainsci10090625).
- [93] Jinwoo Hong et al. “Fetal Cortical Plate Segmentation Using Fully Convolutional Networks With Multiple Plane Aggregation”. In: Frontiers in Neuroscience 14 (2020). ISSN: 1662-453X. DOI: [10.3389/fnins.2020.591683](https://doi.org/10.3389/fnins.2020.591683). URL: <https://www.frontiersin.org/articles/10.3389/fnins.2020.591683>.
- [94] Laura Horvath and John W. Seeds. “Arrest of Fetal Movement With Pancuronium Bromide to Enable Antenatal Magnetic Resonance Imaging of Holoprosencephaly.” In: Obstetric Anesthesia Digest. (Apr. 1990). DOI: [10.1055/s-2007-999629](https://doi.org/10.1055/s-2007-999629).
- [95] Benjamin Hou et al. “3-D Reconstruction in Canonical Co-Ordinate Space From Arbitrarily Oriented 2-D Images”. In: IEEE Transactions on Medical Imaging PP (Sept. 2017). DOI: [10 . 1109/TMI.2018.2798801](https://doi.org/10.1109/TMI.2018.2798801).
- [96] Hui-Hsin Hu et al. “Regional quantification of developing human cortical shape with a three-dimensional surface-based magnetic resonance imaging analysis in utero”. In: The European journal of neuroscience 34 (Oct. 2011), pp. 1310–9. DOI: [10 . 1111 / j . 1460-9568.2011.07855.x](https://doi.org/10.1111/j.1460-9568.2011.07855.x).
- [97] Hui-Hsin Hu et al. “Shape and curvedness analysis of brain morphology using human fetal magnetic resonance images in utero”. In: Brain structure & function 218 (Nov. 2012). DOI: [10.1007/s00429-012-0469-3](https://doi.org/10.1007/s00429-012-0469-3).
- [98] Xiaona Huang et al. “Deep Learning-Based Multiclass Brain Tissue Segmentation in Fetal MRIs”. In: Sensors 23.2 (2023). ISSN: 1424-8220. DOI: [10.3390/s23020655](https://doi.org/10.3390/s23020655).

- [99] Thierry Huisman et al. “Fetal magnetic resonance imaging of the brain: Technical considerations and normal brain development”. In: European radiology 12 (Sept. 2002), pp. 1941–51. DOI: [10.1007/s00330-001-1209-x](https://doi.org/10.1007/s00330-001-1209-x).
- [100] Thierry Huisman et al. “Fetal magnetic resonance imaging of the central nervous system: a pictorial essay”. In: European radiology 12 (Sept. 2002), pp. 1941–51. DOI: [10.1007/s00330-001-1281-2](https://doi.org/10.1007/s00330-001-1281-2).
- [101] Daniel P. Huttenlocher, Gregory A. Klanderman, and William Rucklidge. “Comparing images using the Hausdorff distance”. In: IEEE Transactions on Pattern Analysis and Machine Intelligence 15.9 (1993), pp. 850–863. DOI: [10.1109/34.232073](https://doi.org/10.1109/34.232073).
- [102] Kiho Im and Patricia Grant. “Sulcal pits and patterns in developing human brains”. In: NeuroImage 185 (Mar. 2018). DOI: [10.1016/j.neuroimage.2018.03.057](https://doi.org/10.1016/j.neuroimage.2018.03.057).
- [103] Kiho Im et al. “Quantitative Folding Pattern Analysis of Early Primary Sulci in Human Fetuses with Brain Abnormalities”. In: AJNR. American journal of neuroradiology 38 (May 2017). DOI: [10.3174/ajnr.A5217](https://doi.org/10.3174/ajnr.A5217).
- [104] Mark Ison et al. “Fully Automated Brain Extraction and Orientation in Raw Fetal MRI”. In: Oct. 2012. DOI: [10.13140/2.1.2312.8966](https://doi.org/10.13140/2.1.2312.8966).
- [105] Francois Jacob et al. “Fetal Hippocampal Development: Analysis by Magnetic Resonance Imaging Volumetry”. In: Pediatric research 69 (May 2011), pp. 425–9. DOI: [10.1203/PDR.0b013e318211dd7f](https://doi.org/10.1203/PDR.0b013e318211dd7f).
- [106] András Jakab et al. “Disrupted developmental organization of the structural connectome in fetuses with corpus callosum agenesis”. In: NeuroImage (Feb. 2015). DOI: [10.1016/j.neuroimage.2015.02.038](https://doi.org/10.1016/j.neuroimage.2015.02.038).
- [107] Deborah A. Jarvis and Paul D. Griffiths. “Current state of MRI of the fetal brain in utero”. In: Journal of Magnetic Resonance Imaging 49.3 (2019), pp. 632–646. DOI: <https://doi.org/10.1002/jmri.26316>.
- [108] Mark Jenkinson et al. “FSL”. In: NeuroImage 62.2 (2012). 20 YEARS OF fMRI, pp. 782–790. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2011.09.015>.
- [109] Shuzhou Jiang et al. “MRI of Moving Subjects Using Multislice Snapshot Images With Volume Reconstruction (SVR): Application to Fetal, Neonatal, and Adult Brain Studies”. In: IEEE transactions on medical imaging 26 (Aug. 2007), pp. 967–80. DOI: [10.1109/TMI.2007.895456](https://doi.org/10.1109/TMI.2007.895456).
- [110] Ian R. Johnson et al. “Imaging the pregnant human uterus with nuclear magnetic resonance.” In: American journal of obstetrics and gynecology 148 (May 1984), pp. 1136–9. DOI: [10.1016/0002-9378\(84\)90642-2](https://doi.org/10.1016/0002-9378(84)90642-2). URL: <https://api.semanticscholar.org/CorpusID:7804073>.

- [111] Bernhard Kainz et al. “Fast Fully Automatic Brain Detection In Fetal Mri Using Dense Rotation Invariant Image Descriptors”. In: Apr. 2014. DOI: [10 . 1109 / ISBI . 2014 . 6868098](https://doi.org/10.1109/ISBI.2014.6868098).
- [112] Bernhard Kainz et al. “Fast Volume Reconstruction From Motion Corrupted Stacks of 2D Slices”. In: IEEE Transactions on Medical Imaging 34 (Mar. 2015). DOI: [10 . 1109 / TMI . 2015 . 2415453](https://doi.org/10.1109/TMI.2015.2415453).
- [113] Jonathan Kaplan and David Donoho. “The Morphlet Transform: A Multiscale Representation for Diffeomorphisms”. PhD thesis. Nov. 2021.
- [114] Davood Karimi et al. “Learning to segment fetal brain tissue from noisy annotations”. In: Medical Image Analysis 85 (2023), p. 102731. ISSN: 1361-8415. DOI: <https://doi.org/10.1016/j.media.2022.102731>.
- [115] Katrin Karl et al. “Cavum septi pellucidi (CSP) ratio: a marker for partial agenesis of the fetal corpus callosum”. In: Ultrasound in Obstetrics & Gynecology 50.3 (2017), pp. 336–341. DOI: <https://doi.org/10.1002/uog.17409>. URL: <https://api.semanticscholar.org/CorpusID:206540039>.
- [116] Gregor Kasprian et al. “Assessing prenatal white matter connectivity in commissural agenesis”. In: Brain : a journal of neurology 136 (Jan. 2013), pp. 168–79. DOI: [10 . 1093 / brain/aws332](https://doi.org/10.1093/brain/aws332).
- [117] Kevin Keraudren et al. “Automated Fetal Brain Segmentation from 2D MRI Slices for Motion Correction”. In: NeuroImage (July 2014), pp. –. DOI: [10 . 1016 / j . neuroimage . 2014 . 07 . 023](https://doi.org/10.1016/j.neuroimage.2014.07.023).
- [118] Kevin Keraudren et al. “Localisation of the brain in fetal MRI using bundled SIFT features”. In: vol. 16. Sept. 2013, pp. 582–9. ISBN: 978-3-642-38708-1. DOI: [10 . 1007 / 978 - 3 - 642 - 40811 - 3 _ 73](https://doi.org/10.1007/978-3-642-40811-3_73).
- [119] Nadieh Khalili et al. “Automatic Segmentation of the Intracranial Volume in Fetal MR Images”. In: Sept. 2017, pp. 42–51. ISBN: 978-3-319-67560-2. DOI: [10 . 1007 / 978 - 3 - 319 - 67561 - 9 _ 5](https://doi.org/10.1007/978-3-319-67561-9_5).
- [120] Ali Khan, Lei Wang, and Mirza Faisal Beg. “Multistrukture Large Deformation Diffeomorphic Brain Registration”. In: IEEE Trans Biomed Eng 60 (Feb. 2013), pp. 544–53. DOI: [10 . 1109 / TBME . 2012 . 2230262](https://doi.org/10.1109/TBME.2012.2230262).
- [121] Kio Kim et al. “Intersection Based Motion Correction of Multislice MRI for 3-D in Utero Fetal Brain Image Formation”. In: IEEE transactions on medical imaging 29 (Oct. 2009), pp. 146–58. DOI: [10 . 1109 / TMI . 2009 . 2030679](https://doi.org/10.1109/TMI.2009.2030679).
- [122] Christian Klingenberg. “Walking on Kendall’s Shape Space: Understanding Shape Spaces and Their Coordinate Systems”. In: Evolutionary Biology 47 (Dec. 2020). DOI: [10 . 1007 / s11692 - 020 - 09513 - x](https://doi.org/10.1007/s11692-020-09513-x).
- [123] Vinka Knezović et al. “Underdevelopment of the Human Hippocampus in Callosal Agensis: An In Vivo Fetal MRI Study”. In: American Journal of Neuroradiology 40 (Feb. 2019). DOI: [10 . 3174 / ajnr . A5986](https://doi.org/10.3174/ajnr.A5986).

- [124] Lisa Koch et al. “Graph-Based Label Propagation in Fetal Brain MR Images”. In: Sept. 2014, pp. 9–16. ISBN: 978-3-319-10580-2. DOI: [10.1007/978-3-319-10581-9_2](https://doi.org/10.1007/978-3-319-10581-9_2).
- [125] Jan J. Koenderink and Andrea J. van Doorn. “Surface shape and curvature scales”. In: *Image and Vision Computing* 10.8 (1992), pp. 557–564. ISSN: 0262-8856. DOI: [https://doi.org/10.1016/0262-8856\(92\)90076-F](https://doi.org/10.1016/0262-8856(92)90076-F).
- [126] Vanessa Kyriakopoulou et al. “Cortical Overgrowth in Fetuses With Isolated Ventriculomegaly”. In: *Cerebral cortex (New York, N.Y. : 1991)* 24 (Mar. 2013). DOI: [10.1093/cercor/bht062](https://doi.org/10.1093/cercor/bht062).
- [127] Beatrix Labadi and Anna Beke. “Mental State Understanding in Children with Agenesis of the Corpus Callosum”. In: *Frontiers in Psychology* 8 (Feb. 2017). DOI: [10.3389/fpsyg.2017.00094](https://doi.org/10.3389/fpsyg.2017.00094).
- [128] Ka Chun Lam, Tsz Ching Ng, and Lok Ming Lui. “Multiscale Representation of Deformation via Beltrami Coefficients”. In: *Multiscale Modeling & Simulation* 15 (Jan. 2017), pp. 864–891. DOI: [10.1137/16M1056614](https://doi.org/10.1137/16M1056614).
- [129] Paul C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242.5394 (Mar. 1973), pp. 190–191. DOI: [10.1038/242190a0](https://doi.org/10.1038/242190a0).
- [130] Julien Lefèvre et al. “Are Developmental Trajectories of Cortical Folding Comparable Between Cross-sectional Datasets of Fetuses and Preterm Newborns?” In: *Cerebral cortex (New York, N.Y. : 1991)* 26 (June 2015). DOI: [10.1093/cercor/bhv123](https://doi.org/10.1093/cercor/bhv123).
- [131] Martina Leombroni et al. “Fetal midline anomalies: Diagnosis and counselling Part 1: Corpus callosum anomalies.” In: *European journal of paediatric neurology* 22 6 (2018), pp. 951–962.
- [132] Gang Li et al. “Construction of 4D High-definition Cortical Surface Atlases of Infants: Methods and Applications”. In: *Medical Image Analysis* 25 (Apr. 2015). DOI: [10.1016/j.media.2015.04.005](https://doi.org/10.1016/j.media.2015.04.005).
- [133] Haotian Li et al. “Mapping fetal brain development based on automated segmentation and 4D brain atlasing”. In: *Brain Structure and Function* 226 (July 2021). DOI: [10.1007/s00429-021-02303-x](https://doi.org/10.1007/s00429-021-02303-x).
- [134] Roxane Licandro et al. “A Longitudinal Diffeomorphic Atlas-Based Tissue Labeling Framework for Fetal Brains using Geodesic Regression”. In: Feb. 2016.
- [135] Maxime Louis et al. “Parallel transport in shape analysis: a scalable numerical scheme”. In: Nov. 2017. DOI: [10.1007/978-3-319-68445-1_4](https://doi.org/10.1007/978-3-319-68445-1_4).
- [136] Eileen Luders et al. “A curvature-based approach to estimate local gyrification on the cortical surface”. In: *NeuroImage* 29 (Mar. 2006), pp. 1224–30. DOI: [10.1016/j.neuroimage.2005.08.049](https://doi.org/10.1016/j.neuroimage.2005.08.049).

- [137] Jan Lui, David Hansen, and Arnold Kriegstein. “Development and Evolution of the Human Neocortex”. In: Cell 146 (July 2011), pp. 18–36. DOI: [10.1016/j.cell.2011.06.030](https://doi.org/10.1016/j.cell.2011.06.030).
- [138] Abeyratne M. K., Willi Freeden, and Carsten Mayer. “Multiscale deformation analysis by Cauchy-Navier wavelets”. In: Journal of Applied Mathematics 2003 (Dec. 2003). DOI: [10.1155/S1110757X03206033](https://doi.org/10.1155/S1110757X03206033).
- [139] Antonios Makropoulos et al. “Automatic Whole Brain MRI Segmentation of the Developing Neonatal Brain”. In: IEEE transactions on medical imaging 33 (May 2014). DOI: [10.1109/TMI.2014.2322280](https://doi.org/10.1109/TMI.2014.2322280).
- [140] Antonios Makropoulos et al. “The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction”. In: NeuroImage 173 (Jan. 2018). DOI: [10.1016/j.neuroimage.2018.01.054](https://doi.org/10.1016/j.neuroimage.2018.01.054).
- [141] Stéphane Mallat. A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way. 3rd. USA: Academic Press, Inc., 2008. ISBN: 0123743702.
- [142] Stéphane Mallat. “Theory for multi-resolution signal decomposition - The wavelet representation”. In: IEEE Transactions on Pattern Analysis and Machine Intelligence - PAMI 11 (Jan. 1989).
- [143] Mirra Manevich-Mazor et al. “Added value of fetal MRI in the evaluation of fetal anomalies of the corpus callosum: a retrospective analysis of 78 cases”. In: Ultrasound in Obstetrics & Gynecology 48.S1 (2016), pp. 73–73. DOI: <https://doi.org/10.1002/uog.16228>.
- [144] Lucia Manganaro et al. “Fetal MRI of the central nervous system: State-of-the-art”. In: European Journal of Radiology 93 (2017), pp. 273–283. ISSN: 0720-048X. DOI: <https://doi.org/10.1016/j.ejrad.2017.06.004>.
- [145] Lucia Manganaro et al. “Fetal MRI: what’s new? A short review”. In: European Radiology Experimental 7 (Aug. 2023), pp. 337–50. DOI: [10.1186/s41747-023-00358-5](https://doi.org/10.1186/s41747-023-00358-5).
- [146] Peter Mansfield. “Multi-planar image formation using NMR spin echoes”. In: Journal of Physics C: Solid State Physics 10.3 (Feb. 1977), p. L55. DOI: [10.1088/0022-3719/10/3/004](https://doi.org/10.1088/0022-3719/10/3/004).
- [147] Peter Mansfield et al. “Echo planar imaging of the human fetus in utero at 0.5 T”. In: The British journal of radiology 63 (Dec. 1990), pp. 833–41. DOI: [10.1259/0007-1285-63-755-833](https://doi.org/10.1259/0007-1285-63-755-833).
- [148] Gabriele Masselli et al. “Fetal MRI of CNS abnormalities”. In: Clinical Radiology 75.8 (2020), 640.e1–640.e11. ISSN: 0009-9260. DOI: <https://doi.org/10.1016/j.crad.2020.03.035>.
- [149] Shirley McCarthy et al. “Magnetic resonance imaging of fetal anomalies in utero: early experience”. In: American Journal of Roentgenology 145.4 (1985). PMID: 3875982, pp. 677–682. DOI: [10.2214/ajr.145.4.677](https://doi.org/10.2214/ajr.145.4.677).

- [150] Shirley Mccarthy et al. “Obstetrical magnetic resonance imaging: fetal anatomy.” In: Radiology 154.2 (1985). PMID: 3966129, pp. 427–432. DOI: [10.1148/radiology.154.2.3966129](https://doi.org/10.1148/radiology.154.2.3966129).
- [151] Steven McDonagh et al. “Context-Sensitive Super-Resolution for Fast Fetal Magnetic Resonance Imaging”. In: (Feb. 2017).
- [152] Michael Miller, Alain Trouvé, and Laurent Younes. “Geodesic Shooting for Computational Anatomy”. In: Journal of mathematical imaging and vision 24 (Jan. 2006), pp. 209–228. DOI: [10.1007/s10851-005-3624-0](https://doi.org/10.1007/s10851-005-3624-0).
- [153] Michael Miller, Alain Trouvé, and Laurent Younes. “On the Metrics and Euler-Lagrange Equations of Computational Anatomy”. In: Annual review of biomedical engineering 4 (Feb. 2002), pp. 375–405. DOI: [10.1146/annurev.bioeng.4.092101.125733](https://doi.org/10.1146/annurev.bioeng.4.092101.125733).
- [154] Michael I. Miller, Daniel J. Tward, and Alain Trouvé. “Coarse-to-Fine Hamiltonian Dynamics of Hierarchical Flows in Computational Anatomy”. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 2020, pp. 3760–3765. DOI: [10.1109/CVPRW50498.2020.00438](https://doi.org/10.1109/CVPRW50498.2020.00438).
- [155] Jan Modersitzki. FAIR: Flexible Algorithms for Image Registration. Society for Industrial and Applied Mathematics, Sept. 2008. DOI: [10.1137/1.9780898718843](https://doi.org/10.1137/1.9780898718843).
- [156] Klas Modin, Adrian Nachman, and Luca Rondi. “A Multiscale Theory for Image Registration and Nonlinear Inverse Problems”. In: Advances in Mathematics 346 (Mar. 2018). DOI: [10.1016/j.aim.2019.02.014](https://doi.org/10.1016/j.aim.2019.02.014).
- [157] Marie-Laure Moutard et al. “Isolated corpus callosum agenesis: a ten-year follow-up after prenatal diagnosis (How are the children without corpus callosum at 10 years of age?)” In: Prenatal Diagnosis 32.3 (2012), pp. 277–283. DOI: <https://doi.org/10.1002/pd.3824>.
- [158] Oliver Musse, Fabrice Heitz, and Jp Armspach. “Topology preserving deformable image matching using constrained hierarchical parametric models”. In: Image Processing, IEEE Transactions on 10 (Aug. 2001), pp. 1081–1093. DOI: [10.1109/83.931102](https://doi.org/10.1109/83.931102).
- [159] Yasuhiro Nakata et al. “Diffusion Abnormalities and Reduced Volume of the Ventral Cingulum Bundle in Agensis of the Corpus Callosum: A 3T Imaging Study”. In: AJNR. American journal of neuroradiology 30 (Feb. 2009), pp. 1142–8. DOI: [10.3174/ajnr.A1527](https://doi.org/10.3174/ajnr.A1527).
- [160] Marc Niethammer, Roland Kwitt, and François-Xavier Vialard. “Metric Learning for Image Registration”. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019), pp. 8455–8464. DOI: [10.1109/CVPR.2019.00866](https://doi.org/10.1109/CVPR.2019.00866).

- [161] Vincent Noblet et al. “3-D deformable image registration: A topology preservation scheme based on hierarchical deformation models and interval analysis optimization”. In: IEEE transactions on image processing: a publication of the IEEE Signal Processing Society 14 (June 2005), pp. 553–66. DOI: [10.1109/TIP.2005.846026](https://doi.org/10.1109/TIP.2005.846026).
- [162] Kenichi Oishi, Linda Chang, and Hao Huang. “Baby brain atlases”. In: NeuroImage 185 (Apr. 2018). DOI: [10.1016/j.neuroimage.2018.04.003](https://doi.org/10.1016/j.neuroimage.2018.04.003).
- [163] Kenichi Oishi et al. “Multi-Contrast Human Neonatal Brain Atlas: Application to Normal Neonate Development Analysis”. In: NeuroImage 56 (May 2011), pp. 8–20. DOI: [10.1016/j.neuroimage.2011.01.051](https://doi.org/10.1016/j.neuroimage.2011.01.051).
- [164] Francisco Oliveira and Joao Tavares. “Medical image registration: A review”. In: Computer methods in biomechanics and biomedical engineering 17 (Jan. 2014), pp. 73–93. DOI: [10.1080/10255842.2012.670855](https://doi.org/10.1080/10255842.2012.670855).
- [165] Cynthia Ortinau et al. “Early-Emerging Sulcal Patterns Are Atypical in Fetuses with Congenital Heart Disease”. In: Cerebral cortex (New York, N.Y. : 1991) 29 (Oct. 2018). DOI: [10.1093/cercor/bhy235](https://doi.org/10.1093/cercor/bhy235).
- [166] Julia P. Owen et al. “Resting-State Networks and the Functional Connectome of the Human Brain in Agenesis of the Corpus Callosum”. In: Brain Connectivity 3.6 (2013). PMID: 24063289, pp. 547–562. DOI: [10.1089/brain.2013.0175](https://doi.org/10.1089/brain.2013.0175).
- [167] Kelly Payette, Raimund Kottke, and András Jakab. “Efficient Multi-class Fetal Brain Segmentation in High Resolution MRI Reconstructions with Noisy Labels”. In: Oct. 2020, pp. 295–304. ISBN: 978-3-030-60333-5. DOI: [10.1007/978-3-030-60334-2_29](https://doi.org/10.1007/978-3-030-60334-2_29).
- [168] Kelly Payette et al. “An automatic multi-tissue human fetal brain segmentation benchmark using the Fetal Tissue Annotation Dataset”. In: Scientific Data 8 (July 2021). DOI: [10.1038/s41597-021-00946-3](https://doi.org/10.1038/s41597-021-00946-3).
- [169] Kelly Payette et al. “Longitudinal Analysis of Fetal MRI in Patients with Prenatal Spina Bifida Repair”. In: Oct. 2019, pp. 161–170. ISBN: 978-3-030-32874-0. DOI: [10.1007/978-3-030-32875-7_18](https://doi.org/10.1007/978-3-030-32875-7_18).
- [170] Yuchen Pei et al. “Learning Spatiotemporal Probabilistic Atlas of Fetal Brains with Anatomically Constrained Registration Network”. In: vol. 12907. Sept. 2021, pp. 239–248. ISBN: 978-3-030-87233-5. DOI: [10.1007/978-3-030-87234-2_23](https://doi.org/10.1007/978-3-030-87234-2_23).
- [171] Julien Pontabry et al. “A discriminative feature selection approach for shape analysis: application to fetal brain cortical folding study”. In: Medical Image Analysis 35 (Jan. 2015). DOI: [10.1016/j.media.2016.07.005](https://doi.org/10.1016/j.media.2016.07.005).
- [172] Martin C. POWELL et al. “Magnetic resonance imaging (MRI) in obstetrics. II. Fetal anatomy”. In: BJOG: An International Journal of Obstetrics & Gynaecology 95.1 (1988), pp. 38–46. DOI: <https://doi.org/10.1111/j.1471-0528.1988.tb06478.x>. eprint: <https://obgyn.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1471-0528.1988.tb06478.x>.

- [173] Daniela Prayer et al. "ISUOG Practice Guidelines (updated): performance of fetal magnetic resonance imaging". In: Ultrasound in Obstetrics & Gynecology 61.2 (2023), pp. 278–287. DOI: <https://doi.org/10.1002/uog.26129>.
- [174] Anqi Qiu and Michael I. Miller. "Multi-structure network shape analysis via normal surface momentum maps". In: NeuroImage 42.4 (2008), pp. 1430–1438. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2008.04.257>.
- [175] Anqi Qiu, Laurent Younes, and Michael Miller. "Principal component based diffeomorphic surface mapping". In: IEEE Trans Med Imaging (Feb. 2012). DOI: [10.1109/TMI.2011.2168567](https://doi.org/10.1109/TMI.2011.2168567).
- [176] Anqi Qiu et al. "Atlas Generation for Subcortical and Ventricular Structures With Its Applications in Shape Analysis". In: IEEE transactions on image processing : a publication of the IEEE Signal Processing Society 19 (June 2010), pp. 1539–47. DOI: [10.1109/TIP.2010.2042099](https://doi.org/10.1109/TIP.2010.2042099).
- [177] Anqi Qiu et al. "Parallel Transport in Diffeomorphisms Distinguishes the Time-Dependent Pattern of Hippocampal Surface Deformation due to Healthy Aging and the Dementia of the Alzheimer's Type". In: NeuroImage 40 (Apr. 2008), pp. 68–76. DOI: [10.1016/j.neuroimage.2007.11.041](https://doi.org/10.1016/j.neuroimage.2007.11.041).
- [178] Rehab Raafat, Tamer Abdelrahman, and Mona Fouad. "The prevalence and the adding value of fetal MRI imaging in midline cerebral anomalies". In: Egyptian Journal of Radiology and Nuclear Medicine 51 (Dec. 2020). DOI: [10.1186/s43055-020-0146-0](https://doi.org/10.1186/s43055-020-0146-0).
- [179] Vidya Rajagopalan et al. "Fetal neurodevelopmental recovery in donors after laser surgery for twin-twin transfusion syndrome". In: Prenatal Diagnosis 41 (Nov. 2020). DOI: [10.1002/pd.5866](https://doi.org/10.1002/pd.5866).
- [180] Vidya Rajagopalan et al. "Fetuses with Single Ventricle Congenital Heart Disease Manifest Impairment of Regional Brain Growth". In: Prenatal Diagnosis 38 (Oct. 2018). DOI: [10.1002/pd.5374](https://doi.org/10.1002/pd.5374).
- [181] Vidya Rajagopalan et al. "Is fetal MRI ready for neuroimaging prime time? An examination of progress and remaining areas for development". In: Developmental Cognitive Neuroscience 51 (Aug. 2021), p. 100999. DOI: [10.1016/j.dcn.2021.100999](https://doi.org/10.1016/j.dcn.2021.100999).
- [182] Vidya Rajagopalan et al. "Local Tissue Growth Patterns Underlying Normal Fetal Human Brain Gyrfication Quantified In Utero". In: The Journal of neuroscience : the official journal of the Society for Neuroscience 31 (Feb. 2011), pp. 2878–87. DOI: [10.1523/JNEUROSCI.5458-10.2011](https://doi.org/10.1523/JNEUROSCI.5458-10.2011).
- [183] Vidya Rajagopalan et al. "Mapping directionality specific volume changes using tensor based morphometry: An application to the study of gyrogenesis and lateralization of the human fetal brain". In: NeuroImage 63 (Apr. 2012), pp. 947–58. DOI: [10.1016/j.neuroimage.2012.03.092](https://doi.org/10.1016/j.neuroimage.2012.03.092).

- [184] Martin Rajchl et al. “DeepCut: Object Segmentation From Bounding Box Annotations Using Convolutional Neural Networks”. In: IEEE Transactions on Medical Imaging PP (May 2016). DOI: [10.1109/TMI.2016.2621185](https://doi.org/10.1109/TMI.2016.2621185).
- [185] Marta Bianca Maria Ranzini et al. MONAIfbs: MONAI-based fetal brain MRI deep learning segmentation. 2021. arXiv: [2103.13314](https://arxiv.org/abs/2103.13314) [eess.IV].
- [186] Charles Raybaud. “The corpus callosum, the other great forebrain commissures, and the septum pellucidum: Anatomy, development, and malformation”. In: Neuroradiology 52 (June 2010), pp. 447–77. DOI: [10.1007/s00234-010-0696-3](https://doi.org/10.1007/s00234-010-0696-3).
- [187] M. J. REEVES et al. “The current role of fetal magnetic resonance imaging”. In: Fetal and Maternal Medicine Review 19.1 (2008), 33–60. DOI: [10.1017/S096553950800209X](https://doi.org/10.1017/S096553950800209X).
- [188] Marie-Pierre Revel et al. “Magnetic resonance imaging of the fetus: A study of 20 cases performed without curarization”. In: Prenatal Diagnosis 13.9 (1993), pp. 775–799. DOI: <https://doi.org/10.1002/pd.1970130902>. eprint: <https://obgyn.onlinelibrary.wiley.com/doi/pdf/10.1002/pd.1970130902>.
- [189] Bernd Rieger et al. “On curvature estimation of ISO surfaces in 3D gray-value images and the computation of shape descriptors”. In: IEEE transactions on pattern analysis and machine intelligence 26 (Sept. 2004), pp. 1088–94. DOI: [10.1109/TPAMI.2004.50](https://doi.org/10.1109/TPAMI.2004.50).
- [190] Laurent Risser et al. “Simultaneous Multi-scale Registration Using Large Deformation Diffeomorphic Metric Mapping”. In: IEEE transactions on medical imaging 30 (Apr. 2011), pp. 1746–59. DOI: [10.1109/TMI.2011.2146787](https://doi.org/10.1109/TMI.2011.2146787).
- [191] Caitlin Rollins et al. “Regional Brain Growth Trajectories in Fetuses with Congenital Heart Disease”. In: Annals of Neurology 89 (Nov. 2020). DOI: [10.1002/ana.25940](https://doi.org/10.1002/ana.25940).
- [192] François Rousseau et al. “In Vivo Human Fetal Brain Analysis Using MR Imaging”. In: Fetal Development: Research on Brain and Behavior, Environmental Influences, and Emerging Technologies. Ed. by Nadja Reissland and Barbara S. Kisilevsky. Cham: Springer International Publishing, 2016, pp. 407–427. ISBN: 978-3-319-22023-9. DOI: [10.1007/978-3-319-22023-9_20](https://doi.org/10.1007/978-3-319-22023-9_20).
- [193] François Rousseau et al. “BTK: An open-source toolkit for fetal brain MR image processing”. In: Computer methods and programs in biomedicine 109 (Oct. 2012). DOI: [10.1016/j.cmpb.2012.08.007](https://doi.org/10.1016/j.cmpb.2012.08.007).
- [194] François Rousseau et al. “On super-resolution for fetal brain MRI”. In: vol. 13. Sept. 2010, pp. 355–62. ISBN: 978-3-642-15744-8. DOI: [10.1007/978-3-642-15745-5_44](https://doi.org/10.1007/978-3-642-15745-5_44).
- [195] François Rousseau et al. “Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images”. In: Academic radiology 13 (Oct. 2006), pp. 1072–81. DOI: [10.1016/j.acra.2006.05.003](https://doi.org/10.1016/j.acra.2006.05.003).

- [196] Mary Rutherford et al. “MR imaging methods for assessing fetal brain development”. In: Developmental Neurobiology 68.6 (2008), pp. 700–711. DOI: <https://doi.org/10.1002/dneu.20614>.
- [197] Sahar N. Saleem. “Fetal MRI: An approach to practice: A review”. In: Journal of Advanced Research 5.5 (2014), pp. 507–523. ISSN: 2090-1232. DOI: <https://doi.org/10.1016/j.jare.2013.06.001>.
- [198] Seyed Sadegh Salehi et al. “Real-Time Automatic Fetal Brain Extraction in Fetal MRI by Deep Learning”. In: Apr. 2018. DOI: [10.1109/ISBI.2018.8363675](https://doi.org/10.1109/ISBI.2018.8363675).
- [199] Susana Ferreira Santo et al. “Counseling in fetal medicine: agenesis of the corpus callosum”. In: Ultrasound in Obstetrics & Gynecology 40 (2012).
- [200] Jean-Baptiste Schiratti et al. “A Bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations”. In: Journal of Machine Learning Research 18 (Dec. 2017).
- [201] Jean-Baptiste Schiratti et al. “Learning spatiotemporal trajectories from manifold-valued longitudinal data”. In: June 2015.
- [202] Andreas Schuh et al. “Unbiased construction of a temporally consistent morphological atlas of neonatal brain development”. In: bioRxiv (2018). DOI: [10.1101/251512](https://doi.org/10.1101/251512).
- [203] Ernst Schwartz et al. “Modeling Fetal Cortical Expansion Using Graph-Regularized Gompertz Models”. In: Oct. 2016, pp. 247–254. ISBN: 978-3-319-46719-1. DOI: [10.1007/978-3-319-46720-7_29](https://doi.org/10.1007/978-3-319-46720-7_29).
- [204] Ernst Schwartz et al. “The Prenatal Morphomechanic Impact of Agenesis of the Corpus Callosum on Human Brain Structure and Asymmetry”. In: Cerebral Cortex (Apr. 2021). DOI: [10.1093/cercor/bhab066](https://doi.org/10.1093/cercor/bhab066).
- [205] Julia Scott et al. “3D Morphometric Analysis of Human Fetal Cerebellar Development”. In: Cerebellum (London, England) 11 (Dec. 2011), pp. 761–70. DOI: [10.1007/s12311-011-0338-2](https://doi.org/10.1007/s12311-011-0338-2).
- [206] Julia Scott et al. “Growth trajectories of the human fetal brain tissues estimated from 3D reconstructed in utero MRI”. In: International journal of developmental neuroscience 29 (Aug. 2011), pp. 529–36. DOI: [10.1016/j.ijdevneu.2011.04.001](https://doi.org/10.1016/j.ijdevneu.2011.04.001).
- [207] Julia Scott et al. “Volumetric and surface-based 3D MRI analyses of fetal isolated mild ventriculomegaly”. In: Brain structure and function 218 (May 2012). DOI: [10.1007/s00429-012-0418-1](https://doi.org/10.1007/s00429-012-0418-1).
- [208] Ahmed Serag et al. “A Multi-channel 4D Probabilistic Atlas of the Developing Brain: Application to Fetuses and Neonates”. In: Special Issue of the Annals of the British Machine Vision Association (Dec. 2011).
- [209] Justina Seyi-Olajide et al. “Screening Methods for Congenital Anomalies in Low and Lower-Middle Income Countries: A Systematic Review”. In: Journal of Pediatric Surgery 58 (Jan. 2023). DOI: [10.1016/j.jpedsurg.2023.01.038](https://doi.org/10.1016/j.jpedsurg.2023.01.038).

- [210] Ori Shen et al. “Abnormal shape of the cavum septi pellucidi: An indirect sign of partial agenesis of the corpus callosum”. In: Ultrasound in Obstetrics and Gynecology 46 (Jan. 2015). DOI: [10.1002/uog.14776](https://doi.org/10.1002/uog.14776).
- [211] Zhengyang Shen, Francois-Xavier Vialard, and Marc Niethammer. “Region-specific Diffeomorphic Metric Mapping”. In: Advances in Neural Information Processing Systems. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019.
- [212] Filomena Giulia Sileo et al. “Role of prenatal magnetic resonance imaging in fetuses with isolated agenesis of corpus callosum in the era of fetal neurosonography: A systematic review and meta-analysis”. In: Acta Obstetrica et Gynecologica Scandinavica 100.1 (2021), pp. 7–16. DOI: <https://doi.org/10.1111/aogs.13958>.
- [213] Nivaldo da Silva Junior et al. “Magnetic resonance imaging of the fetal brain at 3 Tesla: Preliminary experience from a single series”. In: Medicine 97 (Oct. 2018), e12602. DOI: [10.1097/MD.00000000000012602](https://doi.org/10.1097/MD.00000000000012602).
- [214] Francis W. Smith et al. “NMR imaging in human pregnancy: A preliminary study”. In: Magnetic resonance imaging 2 (Feb. 1984), pp. 57–64. DOI: [10.1016/0730-725X\(84\)90126-7](https://doi.org/10.1016/0730-725X(84)90126-7).
- [215] Stefan Sommer et al. “Sparse Multi-Scale Diffeomorphic Registration: The Kernel Bundle Framework”. In: J. of Mathematical Imaging and Vision 46 (July 2012). DOI: [10.1007/s10851-012-0409-0](https://doi.org/10.1007/s10851-012-0409-0).
- [216] P. Sonigo et al. “MR imaging of fetal cerebral anomalies”. In: Pediatric Radiology 28 (1998), pp. 212–222. URL: <https://api.semanticscholar.org/CorpusID:7371169>.
- [217] Colin Studholme. “Mapping Fetal Brain Development In Utero Using Magnetic Resonance Imaging: The Big Bang of Brain Mapping”. In: Annual review of biomedical engineering 13 (July 2010), pp. 345–68. DOI: [10.1146/annurev-bioeng-071910-124654](https://doi.org/10.1146/annurev-bioeng-071910-124654).
- [218] Colin Studholme and François Rousseau. “Quantifying and Modelling Tissue Maturation in the Living Human Fetal Brain”. In: International journal of developmental neuroscience 32 (July 2013). DOI: [10.1016/j.ijdevneu.2013.06.006](https://doi.org/10.1016/j.ijdevneu.2013.06.006).
- [219] Wei Sun, Wiro J. Niessen, and Stefan Klein. “Wavelet based free-form deformations for non-rigid registration”. In: Medical Imaging 2014: Image Processing. Ed. by Sebastien Ourselin and Martin A. Styner. Vol. 9034. International Society for Optics and Photonics. SPIE, 2014, pp. 934–939. DOI: [10.1117/12.2043846](https://doi.org/10.1117/12.2043846).
- [220] Indrani Suresh et al. “Role of MRI in the Diagnosis of Fetal Anomalies at 18–20 Weeks Gestational Age”. In: Journal of South Asian Federation of Obstetrics and Gynaecology 11 (Oct. 2019), pp. 292–296. DOI: [10.5005/jp-journals-10006-1714](https://doi.org/10.5005/jp-journals-10006-1714).
- [221] Vahid Taimouri et al. “A template-to-slice block matching approach for automatic localization of brain in fetal MRI”. In: 2015 (July 2015), pp. 144–147. DOI: [10.1109/ISBI.2015.7163836](https://doi.org/10.1109/ISBI.2015.7163836).

- [222] Youssef Taleb et al. “Automatic Template-based Brain Extraction in Fetal MR Images”. In: June 2013.
- [223] 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: a publication of the IEEE Signal Processing Society 25 (June 2016). DOI: [10.1109/TIP.2016.2574982](https://doi.org/10.1109/TIP.2016.2574982).
- [224] Tomo Tarui et al. “Comprehensive quantitative analyses of fetal magnetic resonance imaging in isolated cerebral ventriculomegaly”. In: NeuroImage: Clinical 37 (2023), p. 103357. ISSN: 2213-1582. DOI: <https://doi.org/10.1016/j.nicl.2023.103357>.
- [225] Tomo Tarui et al. “Disorganized Patterns of Sulcal Position in Fetal Brains with Agenesis of Corpus Callosum”. In: Cereb Cortex 28.9 (Sept. 2018), pp. 3192–3203. ISSN: 1460-2199. DOI: [10.1093/cercor/bhx191](https://doi.org/10.1093/cercor/bhx191).
- [226] Tomo Tarui et al. “Quantitative MRI Analyses of Regional Brain Growth in Living Fetuses with Down Syndrome”. In: Cerebral cortex (New York, N.Y. : 1991) 30 (July 2019). DOI: [10.1093/cercor/bhz094](https://doi.org/10.1093/cercor/bhz094).
- [227] D’Arcy Wentworth Thompson. On Growth and Form. Canto. Cambridge University Press, 1992. DOI: [10.1017/CBO9781107325852](https://doi.org/10.1017/CBO9781107325852).
- [228] D’Arcy Wentworth Thompson. On Growth and Form. Canto. Cambridge University Press, 1992. DOI: [10.1017/CBO9781107325852](https://doi.org/10.1017/CBO9781107325852).
- [229] Sébastien Tourbier et al. “An Efficient Total Variation Algorithm for Super-Resolution in Fetal Brain MRI with Adaptive Regularization”. In: NeuroImage 118 (June 2015). DOI: [10.1016/j.neuroimage.2015.06.018](https://doi.org/10.1016/j.neuroimage.2015.06.018).
- [230] Sébastien Tourbier et al. “Automated Template-based Brain Localization and Extraction for Fetal Brain MRI Reconstruction”. In: NeuroImage 155 (Apr. 2017). DOI: [10.1016/j.neuroimage.2017.04.004](https://doi.org/10.1016/j.neuroimage.2017.04.004).
- [231] Sébastien Tourbier et al. “Automatic brain extraction in fetal MRI using multi-atlas-based segmentation”. In: vol. 9413. Feb. 2015. DOI: [10.1117/12.2081777](https://doi.org/10.1117/12.2081777).
- [232] F Triulzi, Lucia Manganaro, and Paolo Volpe. “Fetal magnetic resonance imaging: Indications, study protocols and safety”. In: La Radiologia medica 116 (Feb. 2011), pp. 337–50. DOI: [10.1007/s11547-011-0633-5](https://doi.org/10.1007/s11547-011-0633-5).
- [233] Alain Trouvé. “Diffeomorphisms Groups and Pattern Matching in Image Analysis”. In: Int. J. Comput. Vision 28.3 (July 1998), 213–221. ISSN: 0920-5691. DOI: [10.1023/A:1008001603737](https://doi.org/10.1023/A:1008001603737).
- [234] Alain Trouvé and Laurent Younes. “Metamorphoses Through Lie Group Action”. In: Foundations of Computational Mathematics 5 (Apr. 2005), pp. 173–198. DOI: [10.1007/s10208-004-0128-z](https://doi.org/10.1007/s10208-004-0128-z).
- [235] Andrea Urru et al. “An automatic pipeline for atlas-based fetal and neonatal brain segmentation and analysis”. In: (May 2022).

- [236] Michel Vaillant et al. “Statistics on diffeomorphisms via tangent space representations”. In: NeuroImage 23 Suppl 1 (2004), S161–9. ISSN: 1053-8119. DOI: [10 . 1016 / j . neuroimage . 2004 . 07 . 023](https://doi.org/10.1016/j.neuroimage.2004.07.023).
- [237] Sandra Vegar-Zubović et al. “Magnetic Resonance Imaging in the Diagnosis of Fetal Pathology”. In: Acta Informatica Medica 27 (Mar. 2019), p. 50. DOI: [10 . 5455 / aim . 2019 . 27 . 50 - 53](https://doi.org/10.5455/aim.2019.27.50-53).
- [238] Lei Wang et al. “Large Deformation Diffeomorphism and Momentum Based Hippocampal Shape Discrimination in Dementia of the Alzheimer type”. In: IEEE transactions on medical imaging 26 (Apr. 2007), pp. 462–70. DOI: [10 . 1109 / TMI . 2006 . 887380](https://doi.org/10.1109/TMI.2006.887380).
- [239] Zhou Wang et al. “Image Quality Assessment: From Error Visibility to Structural Similarity”. In: Image Processing, IEEE Transactions on 13 (May 2004), pp. 600 –612. DOI: [10 . 1109 / TIP . 2003 . 819861](https://doi.org/10.1109/TIP.2003.819861).
- [240] S. Warfield, Kelly H. Zou, and William M. Wells. “Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation”. In: IEEE Transactions on Medical Imaging 23.7 (2004), pp. 903–921. DOI: [10 . 1109 / TMI . 2004 . 828354](https://doi.org/10.1109/TMI.2004.828354).
- [241] Daniel Warren, Daniel Connolly, and P Griffiths. “Assessment of Sulcation of the Fetal Brain in Cases of Isolated Agenesis of the Corpus Callosum Using In Utero MR Imaging”. In: AJNR. American journal of neuroradiology 31 (June 2010), pp. 1085–90. DOI: [10 . 3174 / ajnr . A1982](https://doi.org/10.3174/ajnr.A1982).
- [242] Eva Weigl et al. “A spatio-temporal latent atlas for semi-supervised learning of fetal brain segmentations and morphological age estimation”. In: Medical image analysis 18 (Aug. 2013), pp. 9–21. DOI: [10 . 1016 / j . media . 2013 . 08 . 004](https://doi.org/10.1016/j.media.2013.08.004).
- [243] Eva Weigl et al. “Learning a spatio-temporal latent atlas for fetal brain segmentation”. In: Sept. 2011.
- [244] Jeffrey C. Weinreb et al. “Magnetic resonance imaging in obstetric diagnosis.” In: Radiology 154.1 (1985). PMID: 3880601, pp. 157–161. DOI: [10 . 1148 / radiology . 154 . 1 . 3880601](https://doi.org/10.1148/radiology.154.1.3880601).
- [245] Christian Weisstanner et al. “MRI of the Fetal Brain”. In: Clinical neuroradiology 25 (June 2015). DOI: [10 . 1007 / s00062 - 015 - 0413 - z](https://doi.org/10.1007/s00062-015-0413-z).
- [246] Janet De Wilde, Angus Rivers, and David L. Price. “A review of the current use of magnetic resonance imaging in pregnancy and safety implications for the fetus”. In: Progress in Biophysics and Molecular Biology 87.2 (2005). Effects of static magnetic fields relevant to human health, pp. 335–353. ISSN: 0079-6107. DOI: [https://doi.org/10 . 1016 / j . pbiomolbio . 2004 . 08 . 010](https://doi.org/10.1016/j.pbiomolbio.2004.08.010).
- [247] Robert Wright et al. “Automatic Quantification of Normal Cortical Folding Patterns from Fetal Brain MRI.” In: NeuroImage 91 (Jan. 2014). DOI: [10 . 1016 / j . neuroimage . 2014 . 01 . 034](https://doi.org/10.1016/j.neuroimage.2014.01.034).

- [248] Robert Wright et al. “Construction of a fetal spatio-temporal cortical surface atlas from in utero MRI: Application of spectral surface matching”. In: *NeuroImage* 120 (June 2015). DOI: [10.1016/j.neuroimage.2015.05.087](https://doi.org/10.1016/j.neuroimage.2015.05.087).
- [249] Jiangjie Wu et al. “Age-specific structural fetal brain atlases construction and cortical development quantification for chinese population”. In: *NeuroImage* 241 (2021), p. 118412. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2021.118412>.
- [250] Jue Wu et al. “Assessment of MRI-Based Automated Fetal Cerebral Cortical Folding Measures in Prediction of Gestational Age in the Third Trimester”. In: *AJNR. American journal of neuroradiology* 36 (June 2015). DOI: [10.3174/ajnr.A4357](https://doi.org/10.3174/ajnr.A4357).
- [251] Yu-Te Wu et al. “Image Registration Using Wavelet-Based Motion Model”. In: *International Journal of Computer Vision* 38 (July 2000), pp. 129–152. DOI: [10.1023/A:1008101718719](https://doi.org/10.1023/A:1008101718719).
- [252] Zhengwang Wu et al. “Construction of spatiotemporal neonatal cortical surface atlases using a large-scale dataset”. In: vol. 2018. Apr. 2018, pp. 1056–1059. DOI: [10.1109/ISBI.2018.8363753](https://doi.org/10.1109/ISBI.2018.8363753).
- [253] Jing Xia et al. “Fetal cortical surface atlas parcellation based on growth patterns”. In: *Human Brain Mapping* 40 (May 2019). DOI: [10.1002/hbm.24637](https://doi.org/10.1002/hbm.24637).
- [254] Xinyi Xu et al. “Spatiotemporal atlas of the fetal brain depicts cortical developmental gradient in Chinese population”. In: (May 2022). DOI: [10.1101/2022.05.09.491258](https://doi.org/10.1101/2022.05.09.491258).
- [255] Laurent Younes. *Shapes and Diffeomorphisms*. Applied Mathematical Sciences. Springer Berlin Heidelberg, 2010. ISBN: 9783642120558. URL: <https://books.google.fr/books?id=SdTBtMGgeAUC>.
- [256] Laurent Younes et al. “Regionally selective atrophy of subcortical structures in prodromal HD as revealed by statistical shape analysis”. In: *Human Brain Mapping* 35.3 (2014), pp. 792–809. DOI: <https://doi.org/10.1002/hbm.22214>.
- [257] Dawood Yousif et al. “Imaging fetal anatomy”. In: *Seminars in Cell and Developmental Biology* 131 (2022). Special issue: Human embryogenesis by Naomi Moris and Marta Shahbazi / Special Issue: Luminogenesis and Hydraulics in Development by Chii Jou Chan, pp. 78–92. ISSN: 1084-9521. DOI: <https://doi.org/10.1016/j.semcd.2022.02.023>.
- [258] Hyuk Jin Yun et al. “Automatic labeling of cortical sulci for the human fetal brain based on spatio-temporal information of gyrification”. In: *NeuroImage* 188 (Dec. 2018). DOI: [10.1016/j.neuroimage.2018.12.023](https://doi.org/10.1016/j.neuroimage.2018.12.023).
- [259] Hyuk Jin Yun et al. “Regional Alterations in Cortical Sulcal Depth in Living Fetuses with Down Syndrome”. In: *Cerebral Cortex* 31.2 (Sept. 2020), pp. 757–767. ISSN: 1047-3211. DOI: [10.1093/cercor/bhaa255](https://doi.org/10.1093/cercor/bhaa255).

- [260] Hyuk Jin Yun et al. “Temporal Patterns of Emergence and Spatial Distribution of Sulcal Pits During Fetal Life”. In: Cerebral cortex (New York, N.Y. : 1991) 30 (Mar. 2020). DOI: [10 . 1093/cercor/bhaa053](https://doi.org/10.1093/cercor/bhaa053).
- [261] Paul Yushkevich, Yang Gao, and Guido Gerig. “ITK-SNAP: An interactive tool for semi-automatic segmentation of multi-modality biomedical images”. In: vol. 2016. Aug. 2016, pp. 3342–3345. DOI: [10 . 1109/EMBC.2016.7591443](https://doi.org/10.1109/EMBC.2016.7591443).
- [262] Jinfeng Zhan et al. “Spatial-Temporal Atlas of Human Fetal Brain Development During the Early Second Trimester.” In: NeuroImage 82 (May 2013). DOI: [10 . 1016 / j . neuroimage.2013.05.063](https://doi.org/10.1016/j.neuroimage.2013.05.063).
- [263] Miaomiao Zhang and P Fletcher. “Finite-Dimensional Lie Algebras for Fast Diffeomorphic Image Registration”. In: vol. 24. Apr. 2015. ISBN: 978-3-319-19991-7. DOI: [10 . 1007/978-3-319-19992-4_19](https://doi.org/10.1007/978-3-319-19992-4_19).
- [264] Li Zhao et al. “Automated 3D Fetal Brain Segmentation Using an Optimized Deep Learning Approach”. In: American Journal of Neuroradiology 43 (Feb. 2022). DOI: [10 . 3174 / ajnr . A7419](https://doi.org/10.3174/ajnr.A7419).
- [265] Reza Zolfaghari et al. “A multiscale LDDMM template algorithm for studying ear shape variations”. In: Jan. 2014. DOI: [10 . 1109/ICSPCS.2014.7021100](https://doi.org/10.1109/ICSPCS.2014.7021100).