Imperial College London
Department of Computing

Computational models of the morphology of the developing neonatal human brain

Andreas Schuh

September 2017

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy of **Imperial College London**

Declarations

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Originality

I declare that the work presented in this thesis is my own, unless specifically acknowledged.

Andreas Schuh

Abstract

Automated medical image analysis has made significant progress over the past decades. With recent advances in acquiring high quality *in vivo* images of the developing human brain, analysing this data for the purpose of understanding brain development is rapidly becoming feasible. Premature birth increases the risk of developing neurocognitive and neurobehavioural disorders. Studying the morphology and function of the brain during maturation, provides us not only with a better understanding of normal development, but may help identify causes for these. A difficulty is to differentiate between neurodevelopmental consequences and normal variation. Reference models are therefore needed.

This thesis presents computational methods used to obtain such models. As a prerequisite, an efficient topology-preserving registration is required. Existing methods have been evaluated mostly on adult brain images, with considerably different shape and appearance. We evaluate approaches for the fast diffeomorphic registration on a publicly available neonatal brain image dataset, and present an improved inverse consistent variant of the stationary velocity free-form deformation algorithm. We employ this algorithm for the construction of a spatio-temporal atlas of the neonatal brain, and compare two different approaches. The first approach is based on the registration of all pairs of images. Residual misalignment thereby still impacts the sharpness of the atlas. More detail is preserved with an iterative refinement of the transformations relating each image to the atlas space. We developed a second approach, which jointly estimates mean shape and longitudinal change iteratively. The final atlas demonstrates increased sharpness and temporal consistency. Finally, we present deformable models for the reconstruction of the neonatal cortex, which correct for common errors observed in state-of-the-art neonatal brain segmentations. Our models were found by experts to be superior to the original segmentation in terms of accurately delineating the cortical anatomy, and form a vital component of image processing pipelines of the Developing Human Connectome Project.

To my parents

Acknowledgements

It has been a long journey leading up to this work, starting with the guidance, opportunities, and encouragement I have received from advisors before starting the Ph.D. program in the Biomedical Image Analysis Group (BioMedIA). Since I joined the group, my supervisor, Prof. Daniel Rueckert, has been a great influence, with an always open door, invaluable insights, and most useful advise, for which I express my sincere gratitude.

I am thankful to all past and present members of BioMedIA, which I had the pleasure to work alongside, for a stimulating environment, discussions over a cup of coffee (or tea), and fruitful collaborations. Special thanks go to (in alphabetical order) Antonis, Christian, Emma, Kanwal, Lisa, Ozan, Ricardo, and Stefan, for their extremely useful help and advise, critical comments, and fun times.

Finally, I am forever grateful for the everlasting support of my family. Without the early encouragement of my mother, Renate Schuh (†1998), and the relentless support of my father, Peter Schuh, I would not have accomplished this.

Contents

Abstract iii

Acknowledgements v

Contents vii

List of tables xi

List of figures xiii

List of abbreviations xvii

I. Preface 1

1. Introduction 3

1.1. Imaging the developing brain 4
1.2. Early brain development 5
1.3. Neonatal brain image analysis 9
1.4. Contributions 12
1.5. Publications 13
1.6. Thesis overview 14

II. Background 15

2. Brain image registration 17

2.1. Image dissimilarity measures 21
2.1.1. Mono-modal image dissimilarities 21
2.1.2. Multi-modal image dissimilarities derived from statistics 22
2.1.3. Multi-modal image dissimilarities from information theory 25
2.2. Transformation models 27
2.2.1. Global transformations 27
2.2.2. Non-parametric small deformation models 29
2.2.3. Non-parametric large deformation models 34

2.2.4. Spline-based small deformation models . . . 42
2.2.5. Locally affine transformations . . . 52
2.2.6. Statistical deformation models . . . 53
2.3. Task-specific regularisation . . . 54
2.4. Symmetric energy formulation . . . 57
2.5. Image registration tools . . . 60

3. Brain atlas construction **63**
3.1. Average brain models . . . 66
3.1.1. Using pairwise registration . . . 67
3.1.2. Iterative template construction . . . 68
3.1.3. Using groupwise registration . . . 69
3.1.4. Using diffeomorphic registration . . . 70
3.1.5. Statistical template deformations . . . 72
3.2. Spatio-temporal models . . . 73
3.3. Volumetric brain atlases . . . 76
3.3.1. Fetal brain atlases . . . 76
3.3.2. Neonatal brain atlases . . . 77
3.3.3. Infant and child brain atlases . . . 79
3.4. Developing cortical surface atlases . . . 80

4. Segmentation of the developing brain **81**
4.1. Challenges of neonatal brain image segmentation . . . 82
4.2. Atlas-based brain image segmentation . . . 83
4.2.1. Using a single atlas . . . 83
4.2.2. Using multiple atlases . . . 84
4.2.3. Using a probabilistic atlas . . . 87
4.3. Unsupervised brain tissue segmentation . . . 97
4.4. Supervised brain tissue segmentation . . . 99
4.5. Deformable models of the cerebral cortex . . . 102
4.5.1. Parametric deformable models . . . 103
4.5.2. Geometric deformable models . . . 111

III. Methods **115**

5. Flexible, fast, and diffeomorphic registration of the neonatal brain **117**
5.1. MIRTK: Medical Image Registration ToolKit . . . 121
5.1.1. Modularity . . . 121
5.1.2. Configurable energy function . . . 122
5.1.3. Fast image registration . . . 122

5.2. Soft invertibility constraints . . . 123
5.3. Unbiased free-form deformation . . . 125
5.3.1. Parametric velocity field . . . 126
5.3.2. Symmetric energy function . . . 128
5.3.3. Energy minimisation . . . 129
5.4. Results . . . 132
5.4.1. Brain image datasets . . . 132
5.4.2. Registration evaluation . . . 134
5.4.3. Evaluation results . . . 136
5.5. Conclusion . . . 146

6. Spatio-temporal atlas of the normal developing brain 147
6.1. Unbiased global normalisation . . . 149
6.2. Pairwise spatio-temporal atlas construction . . . 152
6.2.1. Log-Euclidean mean of SVFFD . . . 155
6.2.2. Age-dependent Log-Euclidean mean . . . 156
6.2.3. Spatio-temporal mean intensity and brain shape . . . 157
6.2.4. Anatomical atlas of mean brain morphology . . . 158
6.2.5. Refinement of cortical correspondences . . . 160
6.3. Groupwise spatio-temporal atlas construction . . . 162
6.3.1. Recursive definition of spatio-temporal atlas . . . 163
6.3.2. Joint estimation of mean shape and growth . . . 164
6.4. Results . . . 167
6.4.1. Image acquisition and reconstruction . . . 167
6.4.2. Preprocessing and whole brain segmentation . . . 167
6.4.3. Brain atlas evaluation criteria . . . 168
6.4.4. Selection of neonatal brain images . . . 171
6.4.5. Selection of model parameters . . . 171
6.4.6. Comparison of atlas construction methods . . . 180
6.4.7. Visual assessment of neonatal atlases . . . 185
6.4.8. Qualitative comparison to previous atlases . . . 191
6.5. Conclusion . . . 192

7. Deformable models of the neonatal cortex 193
7.1. Segmentation of the brain volume . . . 196
7.2. Reconstruction of the cortical surfaces . . . 198
7.2.1. Deformable surface model . . . 198
7.2.2. Initial genus-0 surface model . . . 200
7.2.3. White matter surface model . . . 206
7.2.4. Pial surface model . . . 211

7.3. Results . . . 212
7.3.1. Datasets . . . 212
7.3.2. Manual evaluation . . . 213
7.3.3. Expert opinion . . . 217
7.3.4. Qualitative assessment . . . 219
7.3.5. Quantitative measures . . . 223
7.4. Conclusion . . . 224

IV. Conclusion **225**

8. Summary and future work **227**
8.1. Achievements . . . 228
8.2. Limitations and future work . . . 231

Bibliography **233**

Appendices **288**

A. Non-parametric transformations **289**
A.1. Diffusion model . . . 290
A.2. Curvature model . . . 291
A.3. Elastic and fluid model . . . 292

B. Free-form deformation **295**
B.1. Definition . . . 295
B.2. Derivatives . . . 296
B.3. Bending energy . . . 298

C. MIRTK configuration **301**
C.1. Unbiased diffeomorphic image registration . . . 301

D. Lists of dHCP scan IDs **303**
D.1. Spatio-temporal brain atlas . . . 303
D.2. Cortical surface reconstruction . . . 311

List of tables

3.1. Spatio-temporal atlases of the fetal human brain. 76
3.2. Atlases of the neonatal human brain. 77

5.1. Results of different registration tools for ALBERTs dataset. 144
5.2. Results for OASIS dataset with parameters optimised for ALBERTs. 145

6.1. Sets of adaptive kernel widths for different target kernel widths. 178
6.2. Summary of compared spatio-temporal atlas construction methods. 180
6.3. Overall mean quality measures of constructed atlases. 181

List of figures

1.1. Fetal brain volume reconstructed with and without motion compensation. . 6
1.2. Slices of T2w MR image of fetal brain at 25 weeks GA. 7
1.3. Slices of T1w and T2w MR images of neonatal brain at 40 weeks PMA. . . 8
1.4. Example of an automatic brain image segmentation result. 10

2.1. Illustration of spatial mappings between target and source images. 18
2.2. Comparison of compact RBFs and Gaussian with $\sigma = 0.5$. 47
2.3. Cubic B-spline FFD interpolation weights. 51

3.1. Orthogonal mid-sections of Collins-27 single-subject brain template. 64
3.2. Axial and coronal mid-sections of different adult population brain atlases. . 65
3.3. Illustration of the pairwise atlas construction approach. 67
3.4. Illustration of the groupwise atlas construction approach. 68
3.5. Illustration of non-parametric spatio-temporal atlas construction. 74

4.1. Illustration of single-atlas-based segmentation. 84
4.2. Illustration of multi-atlas-based segmentation. 85
4.3. Illustration of probabilistic atlas-based segmentation. 88

5.1. Penalty functions based on the Jacobian determinant. 125
5.2. Exemplar axial slices of two brain images of the ALBERTs dataset. 132
5.3. Exemplar axial slices of two brain images of the OASIS dataset. 133
5.4. Comparison of mean overlap and folding for LNCC and NMI. 136
5.5. Comparison of parametric gradient approximations for FFD algorithm. . . . 138
5.6. Comparison of parametric gradient approximations for SVFFD algorithm. . 138

5.7. Influence of linear continuation threshold of Jacobian penalties. 139
5.8. Diffeomorphic FFD registration using soft constraints. 140
5.9. Comparison of SVFFD integration schemes and gradient approximations. . 142
5.10. Results of different registration tools for ALBERTs dataset. 144
5.11. Results for OASIS dataset with parameters optimised for ALBERTs. 145

6.1. Illustration of pairwise spatio-temporal atlas construction. 153
6.2. Illustration of groupwise spatio-temporal atlas construction. 162
6.3. Illustration of recursive estimation of mean growth. 164
6.4. Distribution of GA at birth and PMA at scan. 171
6.5. Mean DSC of pairwise SVFFD registrations for different energy weights. . . 172
6.6. Mean DSC for different combinations of SVFFD parameters. 174
6.7. Mean DSC of SVFFD registrations for selected parameters. 174
6.8. Mean DSC corresponding to different combinations of FFD parameters. . . 175
6.9. Number of negative Jacobian determinants of arithmetic mean FFDs. . . . 175
6.10. Number of negative Jacobian determinants for selected FFD parameters. . . 175
6.11. Evaluation criteria for selected FFD parameters. 175
6.12. Evaluation measures for different pairwise atlas constructions. 176
6.13. Temporal consistency of atlas constructed with different kernel widths. . . . 178
6.14. Adaptive regression kernels for different target kernel widths. 178
6.15. Mean quality measures over time for different kernel widths 178
6.16. Convergence of atlas quality measures in case of groupwise construction. . . 179
6.17. Overall mean quality measures of constructed atlases 182
6.18. Overall mean quality measures of constructed atlases over PMA at scan. . . 182
6.19. Atlas quality measures averaged separately for each tissue class. 183
6.20. Atlas quality measures averaged separately for cGM, WM, and dGM. . . . 184
6.21. Mean temporal consistency of atlas brain structures. 184
6.22. Axial mid-sections of globally normalised spatio-temporal atlases. 186
6.23. Comparison of 36 weeks templates constructed with pairwise methods. . . . 187
6.24. Detail view highlighting temporal consistency of our groupwise method. . . 187

6.25. Detail view highlighting temporal consistency of our groupwise method. . . 188
6.26. Detail view highlighting temporal consistency of our groupwise method. . . 188
6.27. Spatio-temporal atlas constructed with the proposed pairwise approach. . . 189
6.28. Spatio-temporal atlas constructed with the proposed groupwise approach. . 190
6.29. Qualitative comparison to previous neonatal atlases. 191

7.1. Illustration of contrast differences in adult and neonatal brain images. . . . 195
7.2. Exemplar brain segmentations used for cortical surface reconstruction. . . . 197
7.3. S-shaped membership function. 199
7.4. Illustration of local adaptive remeshing operations. 200
7.5. Left convex hull deformed towards WM segmentation. 201
7.6. Illustration of initial white matter surface assembly. 202
7.7. Distance of WM segmentation boundary in normal direction. 203
7.8. Small creases with negative Gauss curvature. 205
7.9. Common Draw-EM tissue segmentation errors. 207
7.10. Exemplar T2w neonatal intensity profiles in normal direction. 208
7.11. Initial white matter surface model with non-cortical triangles coloured red. 209
7.12. Flow chart of white matter surface model integration scheme. 210
7.13. Distribution of PMA at scan of used dHCP datasets. 212
7.14. Graphical interface of custom cortical surface evaluation application. 216
7.15. Distribution of expert quality scores. 217
7.16. Comparison of proposed white matter surface model to other models. . . . 218
7.17. Exemplar image patches overlaid with our white matter surface contours. . 220
7.18. Qualitative results of cortical surface reconstruction for one brain scan. . . . 221
7.19. T2w image slices of 12 brain scans with reconstructed surfaces. 222
7.20. Mean cortical surface area, GI, and thickness with increasing age. 223
7.21. Mean cortical measures for bilateral ROIs with increasing age. 223

List of abbreviations

AC anterior commissure 132

AKR adaptive kernel regression 75–77, 80

ANN artificial neural network 88, 89, 101

BCH Baker-Campbell-Hausdorff 40, 41, 130, 137, 141, 143, 165, 228, 231

BE bending energy 122, 172, 173, 176, 298, 299

BS brainstem 196, 197, 201, 209

CA computational anatomy 10–12, 64, 66, 69, 72, 118, 147, 157, 229

CB cerebellum 196, 197, 201, 209

CC cross-correlation 22, 23, 71

cGM cortical grey matter xiv, 83, 86, 92–94, 96, 110, 132, 147, 168, 181–184, 187, 189–191, 195–197, 206–209, 211, 213, 218, 219, 224, 230–232

CNN convolutional neural network 38, 101

CNR contrast-to-noise ratio 83

CR correlation ratio 24, 31

CSF cerobrospinal fluid 3, 9, 79, 82, 83, 86, 90, 93–95, 97, 99, 100, 108–110, 112–114, 181, 183, 194–197, 206–209, 211, 219, 224, 230–232

DBM deformation-based morphometry 11, 12, 62, 65, 66, 147, 157, 229

dGM deep grey matter xiv, 86, 96, 110, 112, 181, 183, 184, 196, 197, 209

DOF degree of freedom 19, 20, 27, 28, 37, 43, 53, 77, 127, 134

DSC Dice similarity coefficient xiv, 101, 135, 170, 172–175, 177, 178

EBS elastic-body spline 47, 48, 61

EM expectation maximisation 86, 90–94, 96, 97, 113, 159, 167, 196

FA fractional anisotropy 7, 8, 10, 77, 78, 100, 101, 114

FCM fuzzy C-means 92, 97, 109

FFD free-form deformation xiii, xiv, 12, 21, 26, 27, 34, 42, 43, 45, 49–53, 55, 56, 58, 60–62, 70, 72, 75–77, 117–119, 122, 123, 125–131, 136–141, 143, 148, 152–154, 168, 173, 175, 176, 180, 183, 185, 187, 191, 192, 229, 295, 297, 298, 302

GA gestational age xiii, xiv, 3, 4, 8, 76, 77, 132, 171, 192, 212, 303–312

GEBS Gaussian elastic-body spline 48

GI gyrification index xv, 223, 224

GM grey matter xvii, 3, 9, 11, 79, 82, 83, 90, 92, 94–100, 108–114, 183, 194, 196, 206

GMM Gaussian Mixture Model 90, 96, 97

GVF gradient vector flow 106, 112

GW gestational weeks 3, 7, 8

ISS inverse scaling and squaring 40

JAC Jacobian-based penalty 172, 173

JE joint entropy 25, 26

KLD Kullback-Leibler divergence 26, 55, 70, 71

kNN k-Nearest Neighbour 84, 86, 87, 94, 95, 100

KR kernel regression 74, 76, 77

LDDKBM large deformation diffeomorphic kernel bundle mapping 39, 42

LDDMM large deformation diffeomorphic metric mapping 11, 36, 38, 41, 42, 59, 61, 62, 66, 73, 77, 78, 118, 120, 125, 128, 130, 146, 160

LNCC local normalised cross-correlation xiii, 21, 24, 35, 42, 59, 62, 71, 128, 130, 136, 143, 146, 172

LS least-squares fit 76

MD mean diffusivity 7, 8, 10, 77, 78

MI mutual information 26, 31, 136

MICE mean inverse consistency error 135, 137

MRF Markov random field 61, 90, 93, 94, 96, 97, 99, 100, 159

MSE mean squared error 22

mWM myelinated white matter 82, 86, 95, 209

NCC normalised cross-correlation 23, 24

NMI normalised mutual information xiii, 21, 26, 42, 62, 122, 130, 134, 136–140, 143, 151, 172, 181, 301

ODE ordinary differential equation 29, 36, 37, 39, 105, 126

PC posterior commissure 132

PCA principle component analysis 53, 72, 196

PD proton density 86

PDE partial differential equation 29, 30, 33–35, 47, 48, 102, 105, 125, 294

PMA post-menstrual age xiii–xv, 8, 77–79, 96, 149, 171, 176–180, 182, 191, 194, 212, 219, 222–224, 303–312

RBF radial basis function xiii, 43–47, 49, 119

ROI region of interest xv, 9, 10, 16, 87, 95, 96, 132–136, 182, 185, 213–215, 217, 218, 223, 227

SAD sum of absolute differences 22

SDM statistical deformation model 53, 72

SNR signal-to-noise ratio 5, 63

SOR successive over-relaxation 31, 35

SS scaling and squaring 40, 42, 126, 127, 130, 141, 301

SSD sum of squared differences 22, 35, 46, 58, 59, 71, 120

SVF stationary velocity field 12, 27, 39, 41, 42, 62, 72, 119, 125–128, 130, 131, 137, 138, 141, 143, 146, 148, 154, 155, 163, 165, 166, 228, 231, 302

SVFFD stationary velocity free-form deformation iii, ix, xiii, xiv, 12–14, 120, 122, 125, 126, 128–131, 137, 138, 140–143, 146, 148, 154–156, 158, 161, 163–166, 168, 172–174, 176, 179, 180, 183, 185, 187, 191, 192, 228, 229, 231, 301, 302

TC temporal consistency 170, 177, 178, 183

TDFFD temporal diffeomorphic free-form deformation 130, 166, 231

TPS thin-plate spline 33, 43, 45, 48

uWM unmyelinated white matter 82, 86, 95, 194, 195, 209

VBM voxel-based morphometry 10–12, 62, 66, 93, 147, 193, 227, 229

WM white matter xiv, xv, xviii, xix, 3, 6–9, 79, 82, 83, 86, 90, 92–100, 108–114, 147, 181, 183, 184, 194–197, 201, 203, 206–209, 211, 213, 217–219, 224, 230–232

Part I.

Preface

Chapter 1.

Introduction

Automated medical image analysis has made significant progress over the past decades. One major application of interest is the automated analysis of magnetic resonance imaging (MRI) data of the human brain, with the aim to better understand normal brain maturation and degeneration, and to help identify causes for diseases affecting normal brain function. In this work, the focus is especially on the development of methods used to obtain computational models of the normal developing human brain shortly after birth.

MRI is a non-invasive and non-ionising imaging technique, which maps the level of water content in brain tissues, e.g., white matter (WM), grey matter (GM), and cerobrospinal fluid (CSF). The density of hydrogen atoms differs between brain tissues and thereby enables the distinction of these based on their MRI parameters (Bihan, 2014). Recent advances in MRI enabled the study of the human brain *in vivo* not only after birth, but as early as 23 weeks after conception *in utero* (Hüppi et al., 1998; Rutherford, 2002; Counsell et al., 2002; Dubois et al., 2008; Clouchoux et al., 2012; Habas et al., 2012; Vinall et al., 2013; Molnár and Rutherford, 2013; Ball et al., 2013; Wright et al., 2014; Im et al., 2017).

An average pregnancy lasts about 40 weeks (9.2 months) and delivery at 38–42 weeks gestational age (GA) is considered normal (Blencowe et al., 2012). Premature birth before 37 weeks of gestation was reported in 11% of all live births recorded in 184 countries in the year 2010. It is the second largest direct cause of deaths in children younger than 5 years (Blencowe et al., 2012). Improvements in perinatal care significantly increased the survival rate of infants born very premature before 30 gestational weeks (GW). However,

up to 10% of these infants will develop spastic motor deficits and at least another 20% will suffer from developmental or behavioural disabilities (Sled and Nossin-Manor, 2013).

Reduced cortical folding was observed in premature-born infants at term-equivalent age (Hüppi et al., 1996). MRI studies of premature newborns, which weigh less than 90 percent of newborns (<10th percentile of birth weights), also revealed that this intrauterine growth restriction is associated with altered cortical folding and reduced cortical volumes, when compared to preterm infants born with a weight at birth within the normal range for this GA, or full-term controls, respectively (Tolsa et al., 2004; Dubois et al., 2008; Molnár and Rutherford, 2013). This implies that the brain develops in a different way *ex utero*, and that the differences when compared to infants born at term may be associated with the increase in neurocognitive and neurobehavioural disorders seen later in premature-born children (Rutherford, 2002).

The aetiology of altered brain development is unknown (Molnár and Rutherford, 2013). A particular difficulty is to differentiate between neurodevelopmental consequences and normal individual variation as the cause for these delays (Rutherford, 2002). This highlights the importance of understanding normal brain development, which may enable the early identification of abnormal development, increase our understanding of their relation to neurological outcomes, and ensure the right care for affected infants and their families.

1.1. Imaging the developing brain

Compared to ultrasound, which is the primary imaging modality employed in prenatal care, MRI provides better brain tissue contrast, higher image resolution, and is unhindered by the skull. Unlike imaging modalities such as computed tomography (CT), MRI is considerably more sensitive to motion due to the longer times required for the acquisition of whole volume data. The resolution and acquisition times in CT would enable the acquisition of high resolved, nearly motion free images. It is unsuitable for studying brain growth *in vivo*, because of poor soft tissue contrast owing to the similar and weak X-ray attenuation of different types of brain matter, and more importantly the potentially harmful exposure of the vulnerable fetus or infant to ionising radiation.

Imaging the adolescent brain using MRI is standard in clinical routine, but only recent advances enable us to also obtain three-dimensional (3D) images of the developing brain, and to study the normal growth of living brains in three dimensions. The major difficulty that had to be overcome was the definition of fast echo sequences for the acquisition of two-dimensional (2D) slices, termed multi-slice snapshot MRI, along with robust volume reconstruction methods, which compensate for the inevitable motion of the fetus or unsedated infant in between single snapshots (Jiang et al., 2007; Kim et al., 2010). To illustrate the effect of motion compensation on image quality, figure 1.1 compares the magnetic resonance (MR) images of a fetal brain imaged *in utero* before and after motion compensation using snapshot MRI with volume reconstruction (SVR) (Jiang et al., 2007). From the stack of axial 2D slices shown in the top row, the misalignment of individual slices in the out-of-plane direction is evident. Without motion compensation, a true volumetric image analysis is not feasible. Imaging the brain in different orthogonal directions provides different views on the same anatomy, where each acquired image stack may be corrupted by motion in a different direction. This redundancy is exploited by volume reconstruction techniques to recover a 3D volume that is consistent with each acquired stack of 2D slices. Moreover, the higher signal-to-noise ratio (SNR) of the motion compensated images can be observed. This increase in SNR is due to the averaging of multiple scans acquired in different orthogonal directions, which reduces the additive Gaussian noise. Malamateniou et al. (2013) recently surveyed current motion compensation techniques for neonatal and fetal MRI. Regarding the general role of MRI in mapping the development of the fetal brain, the review of Studholme (2011) provides a comprehensive summary.

1.2. Early brain development

Fernández et al. (2016) recently summarised the many discoveries that have been made regarding the folding and expansion of the cerebral cortex over the past decades from cellular and molecular mechanisms regulating cortical development and neurogenesis, genetic mechanisms defining the patterns of cortical folds, and the biomechanics involved. For the purposes of characterizing brain development using MRI, the analysis makes use

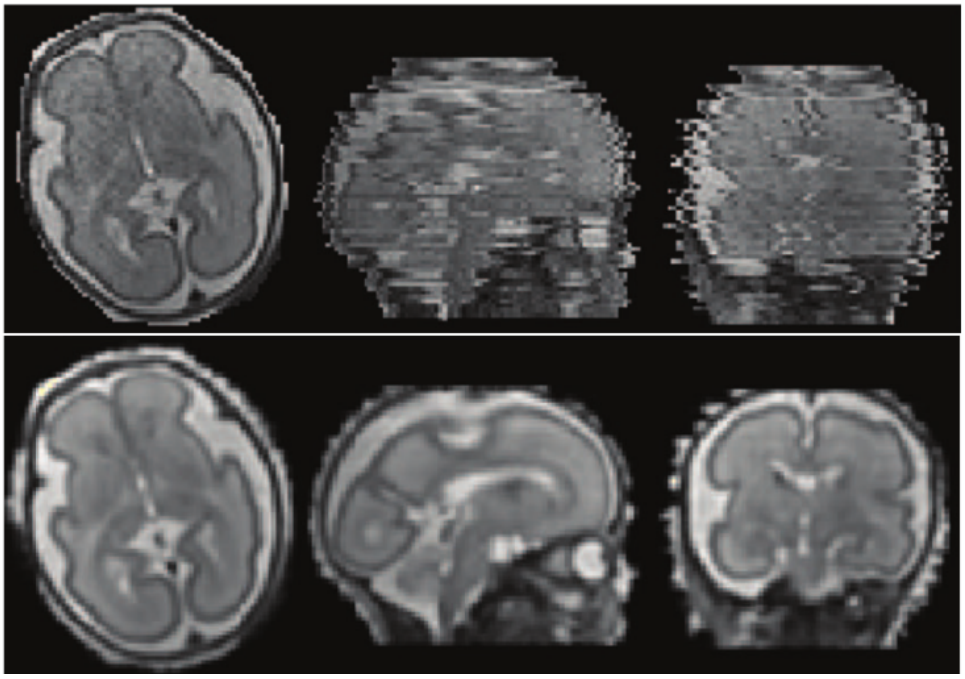
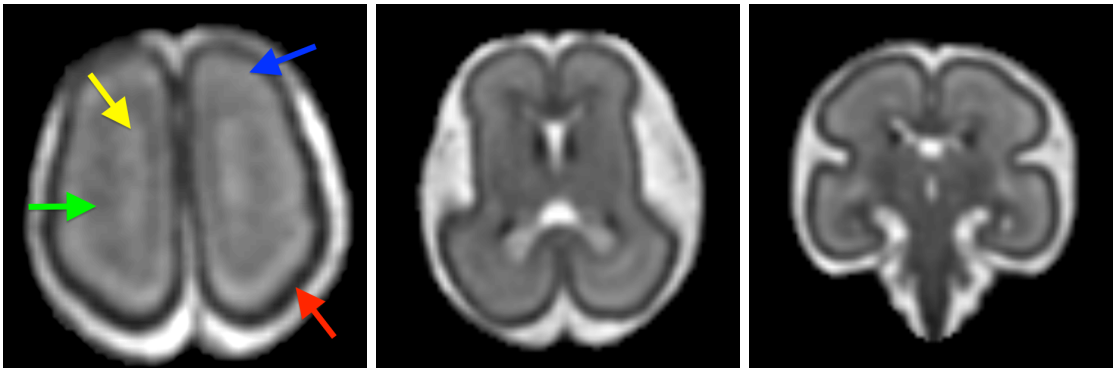


Figure 1.1.: Acquired slices of the brain of a living human fetus *in utero* (top) before, and (bottom) after motion compensation using SVR. Reprinted from Serag (2012).

of morphology, the changing shape of anatomical structures, and dynamic changes of the MRI signal properties to provide insight into the macroscopic and microscopic structural changes during this crucial period (Sled and Nossin-Manor, 2013).

Three key periods of neurogenic development of the human brain can be identified before 26 weeks of gestation, whereby the sequence of events is controlled by extrinsic factors mediated through blood-brain and placental barriers (Stolp et al., 2012; Molnár and Rutherford, 2013): (1) proliferation from 1–7 weeks, (2) migration between 8–15 weeks, and (3) initial differentiation after 16 weeks. At 24–25 weeks gestation, four layers of varying signal intensity (SI) can be identified in the centrum semiovale on T2-weighted (T2w) MR images (Rutherford, 2002): (1) cortex (low SI), (2) subcortical WM (high SI), (3) an intermediate zone of migrating cells (low SI), and (4) a periventricular zone of developing WM (high SI). These layers are highlighted in figure 1.2a.

From about 25 weeks onwards, the cerebral cortex develops from a lissencephalic state, displaying only rudimentary features of central and parieto-occipital cortex (Battin et al., 1998) seen in figure 1.2, to a highly convoluted sulcation pattern until term. This convoluted shape of the cortex can be appreciated in figure 1.3, showing T1-weighted (T1w) and



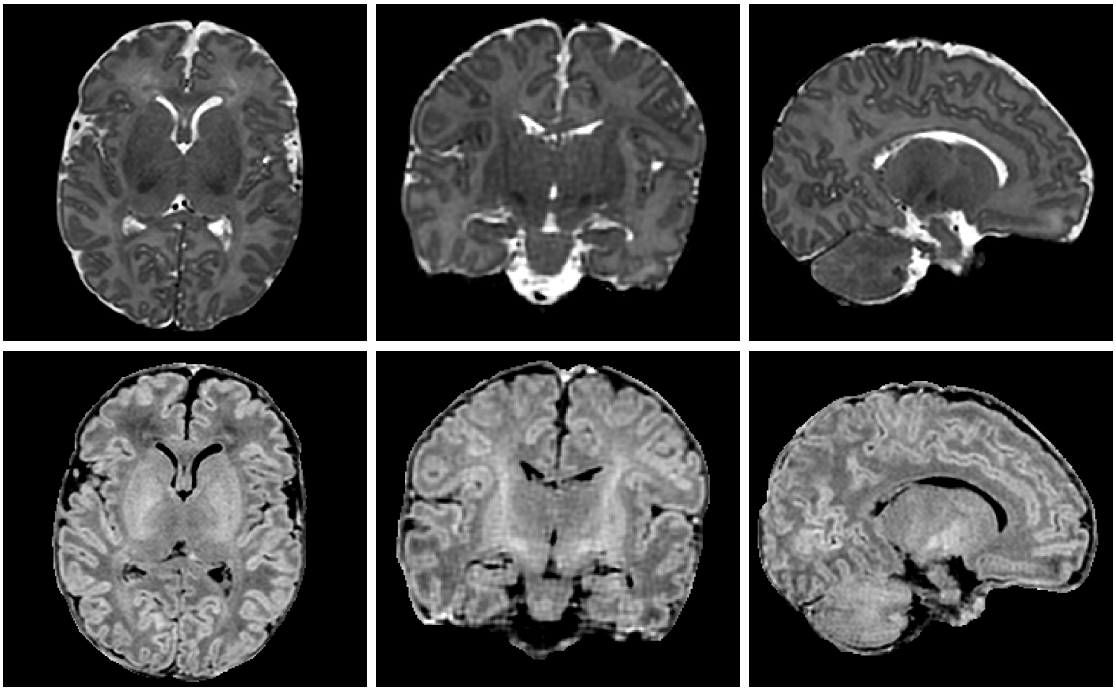
(a) (b) (c)

Figure 1.2.: Slices of T2w MR image of fetal brain at 25 GW. The arrows in (a) highlight the four main layers in the centrum semiovale that can be observed at this age: (red) cortex, (blue) subcortical WM, (green) intermediate zone of migrating cells, and (yellow) a central periventricular band. The smoothness of the cortex is apparent in (a), (b) the axial slice at mid-ventricular level, and (c) a coronal slice at the brain stem level.

T2w brain scans of an infant born and imaged at term (Sled and Nossin-Manor, 2013). All sulci are formed by week 38 of gestation, though they become deeper over the course of the following few weeks (Rutherford, 2002).

A popular explanation for the formation of these sulcal patterns is given by the tension theory of cortical morphogenesis (Van Essen, 1997). It argues that early maturation of long-distance cortical connectivity is the cause for the formation of cortical sulci with gyri developing between these in further consequence (Ball et al., 2013). Results of a recent study by Ball et al. (2013) of cortical micro-structure in premature newborns using diffusion-weighted MRI (DWI), more specifically, mean diffusivity (MD) and fractional anisotropy (FA) derived from diffusion tensor imaging (DTI), confirm this hypothesis.

Briefly, DWI reflects the Brownian motion of water molecules, which is limited by tissue structure particularly in WM tracts, where diffusion along the length of axons is more freely than across them. This anisotropy in water diffusion relative to axon orientation can be seen on DWI images (Rutherford, 2002). MD is a measure reflecting tissue density, whereas FA measures the directionality of water diffusion, reflects axon density and diameter, as well as the extent of WM myelination (Molnár and Rutherford, 2013). The fetal and newborn brain has a higher water content than the adult brain. As a consequence, the diffusion of water molecules is less restricted in all tissue types when compared to the



(a) axial (b) coronal (c) sagittal

Figure 1.3.: Slices of (top) T2w and (bottom) T1w MR images of the brain of a term-born infant scanned shortly after birth at 40 weeks PMA demonstrating the complex folding patterns of the cortex compared to the fetal brain at 25 GW shown in figure 1.2.

diffusion in adult brains (Rutherford, 2002). The MD and FA decrease from 25–40 weeks GA within cortical tissue (Vinall et al., 2013; Ball et al., 2013).

Ball et al. (2013) conclude that the reduction in cortical DTI measures suggests a consistent increase in cellular and synaptic complexity, which starts earlier in sulci than gyri, with rapid elongation of dendrites beginning in sulci and regions of early sulcation, then spreading into gyri and higher-order association cortex. A full understanding of the observed decreases in cortical FA after preterm delivery and how these differ from normal intrauterine cortical development would ideally require a similar DTI-based study of fetal cortical development. Moreover, Molnár and Rutherford (2013) remark that correlating these cortical MRI parameters with functional MRI (fMRI) studies in preterm-born infants could help to ascertain the functional significance of these for neurocognitive development.

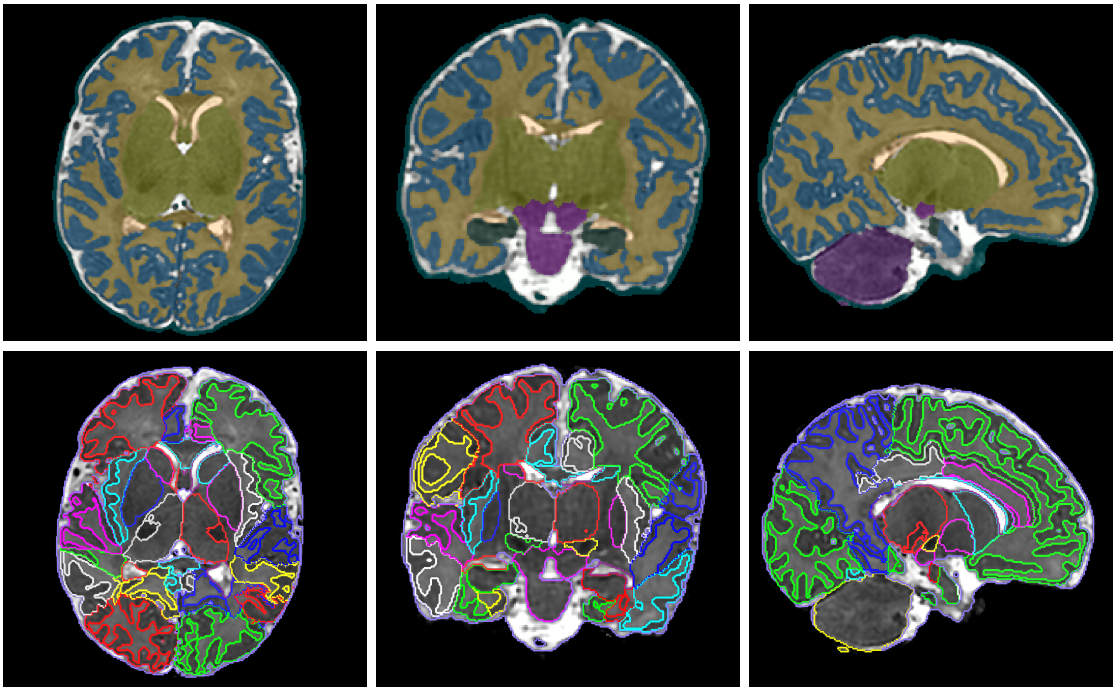
While cortical folding is relatively complete at 38 GW, myelination has at the same time been confined to the brain stem, globus pallidus and ventro-lateral nucleus of the thalamus (Counsell et al., 2002). From about 38 weeks onwards, myelination begins throughout WM

and continues over the first year of life, and to a lesser extent into the second year and adolescence. Myelination is the build up of a sheath of myelin around WM axons to enable the more effective transmission of neural impulses. As a consequence, the MRI contrast of WM changes drastically. First, the T1 relaxation times shorten due to the production of myelin, followed by a shortening of T2 relaxation times occurring at the time of tightening of myelin around the axon (Rutherford, 2002). The absence of myelin in most of the immature brain results in low contrast between WM and GM in T1w images. During the perinatal period, T2w MRI sequences that are more sensitive to water content provide greater tissue contrast (Alexander et al., 2016).

1.3. Neonatal brain image analysis

The main methodological pillars of brain image analysis are: (1) segmentation or brain parcellation, i.e., to identify and delineate individual tissues, structures, or other regions of interest (ROIs) as depicted in figure 1.4, and (2) registration, i.e., to establish a mapping from one brain image to either a different acquisition of the same anatomy (intra-subject registration) or different anatomy (inter-subject registration) for each point of the brain volume. Both tasks are related and may benefit from one another. For example, given a segmentation of a reference image, commonly referred to as brain atlas, the segmentation of a new brain image can be inferred if the mapping between the two imaged anatomies is known. At the same time, knowing the position and boundaries of brain structures may guide the registration process and ensure that corresponding structures are mapped onto each other. Therefore, methods which interleave segmentation with registration have been developed (Ashburner and Friston, 2005; Pohl et al., 2006).

Given the large biological variability of human brain anatomies, and the often subtle micro-structural changes during the early onset of a progressive disease, it is generally difficult to draw a conclusion from one brain image alone. This may only be possible when the disease caused obvious macroscopic changes such as the enlargement of the lateral ventricles and the loss of hippocampal volume (i.e., loss of GM and increase in CSF). These are some of the changes of the ageing brain that may be caused by Alzheimer's



(a) axial (b) coronal (c) sagittal

Figure 1.4.: Automatic brain tissue segmentation (top), and structural whole brain parcellation into 87 ROIs (bottom) of the neonatal T2w brain scan pictured in figure 1.3.

disease (AD) (Merlo Pich et al., 2014). At this point, it may be too late for a successful treatment. A diagnosis without the utilisation of expensive imaging modalities such as MRI may be possible solely based on the symptoms of the disease, which can be observed externally at such late stage. Towards enabling an early diagnosis, it is common to study brain images of entire populations, and to make comparisons between sub-populations, e.g., brain images of healthy control subjects versus pathological brains. The aim is to identify image-based features that are distinct among each group. Such imaging features may be found with: (1) voxel-based morphometry (VBM), (2) deformation-based (or tensor-based) morphometry (DBM), and (3) computational anatomy (CA).

In VBM (Ashburner and Friston, 2000), variations in the MRI parameters of brain tissues within anatomical ROIs are studied on a voxel-by-voxel or ROI-by-ROI basis. This requires a spatial normalisation of all brain images to a stereotactic space using image registration to facilitate comparison of MRI signals across subjects within and across defined ROIs. An application of this methodology is the study by Ball et al. (2013) of the decrease of cortical MD and FA measures in premature born neonates over time.

In DBM (Ashburner et al., 1998), variations of the shape of brain structures are studied based on the deformations needed to normalise the images. It complements VBM, where shape differences are discounted. The deformations encode local changes, which may be correlated to regional atrophy, e.g., loss of cortical GM as has been observed in AD (Merlo Pich et al., 2014) and schizophrenia (Vita et al., 2012). An analogous method was used by Ball et al. (2013) to assess macro-structural growth of the cortex after preterm birth.

DBM is an early approach of CA within the small deformation framework. In CA (Grenander and Miller, 1998; Thompson and Toga, 2002), shapes are studied through the spatial mappings from one anatomy to another. For this, anatomies are represented by one (or more) deformable template(s) delineating the typical structure of the (normal or diseased) anatomy, and a group of topology-preserving transformations, carrying a template to different observed configurations. Shape variability is then explained by probabilities on these diffeomorphisms, i.e., invertible and smooth mappings. Grenander and Miller (1998) identified three principle components of CA:

1. **Computation of large deformation maps** in order to find a transformation, which maps the template onto an observed anatomy. The transformations are the principle method by which anatomical structures are understood in CA, transferring the emphasis from the images to the structural transformations which generate them.

2. **Computation of empirical probability laws** from sub-populations of anatomies. The template and the priors precisely specify the global anatomical relationships between structures as well as how they can vary from one brain to another.

3. **Inference on population and disease testing** using Bayesian methods given this representation in form of prior probabilities on the deformable templates.

Using CA, Bai et al. (2012) studied the differences in brain morphology between 75 Chinese, 73 Malay, and 29 Indian healthy neonates. After creating a cohort specific brain template, the diffeomorphisms relating each neonatal brain image to this template are computed using large deformation diffeomorphic metric mapping (see section 2.2.3). The Jacobian determinant values of these mappings, which are proportional to local volume change, are then utilised for statistical shape comparison across the ethnic groups.

1.4. Contributions

This thesis presents our work on the development of a flexible, fast, and diffeomorphic image registration tool, and its application for the construction of average appearance and shape models of the developing neonatal brain derived from MRI data. Complementary to this volumetric brain model, a method for the construction of geometric models of the neonatal cortex from T2w MR images is developed. Implementations of the proposed methods are freely available, and find application in the Developing Human Connectome Project (dHCP) (Makropoulos et al., 2017b), which aims to create a detailed four-dimensional (4D) connectome of the developing brain. The structural brain models obtained with our methods facilitate a brain image analysis based on VBM, DBM, and CA to help map and understand brain growth. In summary, our contributions are:

1. **A flexible and fast software for neonatal brain image registration**. The Image Registration ToolKit (IRTK), implementing the free-form deformation (FFD) methods of Rueckert et al. (1999) and Schnabel et al. (2001), has been shown to be among the best performing medical image registration tools (Klein et al., 2009a). We developed a related software for neonatal brain images named Medical Image Registration ToolKit (MIRTK) with significantly shorter runtime and improved performance. MIRTK can also be adopted to different applications. It is employed by the atlas-based MALP-EM (Ledig et al., 2015) and Draw-EM (Makropoulos et al., 2014, 2017b) whole brain segmentation tools for adult and neonatal brain images, respectively. Other applications include airway motion estimation (Bates et al., 2017), and inter-slice cardiac motion compensation (Tarroni et al., 2017).

   We further adopt **a diffeomorphic extension of the FFD** based on the one-parameter subgroup of diffeomorphisms generated by stationary velocity fields (SVFs). The SVFFD algorithm enables an efficient registration of images with large anatomical differences, and at the same time ensures invertibility. It facilitates a symmetric optimisation, which is not biased towards one of the anatomies. We investigate several technical details of the algorithm, and evaluate the effect of implementation choices on the registration of neonatal brain images.

2. **Construction of a spatio-temporal atlas of the early developing brain** using the unbiased SVFFD algorithm and the Log-Euclidean mean (Arsigny et al., 2006). We propose and compare two approaches. The first approach improves upon the method of Serag et al. (2012a) in terms of topology-preserving mappings, instead of averaging displacement fields, which assumes small deformations. The second approach follows a groupwise atlas construction (Gholipour et al., 2017). We formulate a new spatio-temporal brain atlas construction, which ensures temporal consistency of the diffeomorphisms used to map each observed anatomy to the 4D average space. Our neonatal atlas depicts significantly greater cortical detail than previous atlases.

3. **Deformable models for the reconstruction of the neonatal cortex** from MR images. These models are inspired by methods for the reconstruction of the adult cortex (Fischl, 2012), but overcome some of the unique challenges encountered in neonatal brain image segmentation. These are in particular the misclassification of tissues near the pial surface due to low tissue contrast and strong partial volume effects. Our deformable models were developed to counter these misclassification by relying on a tissue segmentation only for initialisation. The models are then attracted to local tissue boundaries that are identified in the intensity volumes. Our cerebral cortex models facilitate the study of cortical folding (Wright et al., 2014), and the construction of a neonatal surface atlas (Wright et al., 2015).

1.5. Publications

Parts of our first two contributions, the SVFFD algorithm for neonatal brain atlas construction, were first presented in the workshop paper:

Schuh et al. (2014), **Construction of a 4D brain atlas and growth model using diffeomorphic registration**. *In:* Proceedings of 3rd International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) Workshop on Spatiotemporal Image Analysis for Longitudinal and Time-Series Image Data (STIA), Lecture Notes in Computer Science (LNCS), vol. 8682, pp. 27-37.
doi: 10.1007/978-3-319-14905-9_3

A journal article describing our novel brain atlas construction, including a summary of the SVFFD algorithm detailed in chapter 5, is under review. A preprint is publicly available:

> Schuh et al. (2018), **Unbiased construction of a temporally consistent morphological atlas of neonatal brain development**. *bioRxiv.* **doi:** 10.1101/251512

Chapter 7 describing our third contribution is based on the conference paper:

> Schuh et al. (2017), **A deformable model for the reconstruction of the neonatal cortex**. *In:* IEEE 14th International Symposium on Biomedical Imaging (ISBI), Melbourne, VIC, Australia, 2017, pp. 800–803. **doi:** 10.1109/ISBI.2017.795063

We further co-authored a journal article describing the full structural processing pipeline of the dHCP, of which this cortex reconstruction method constitutes a major part:

> Makropoulos et al. (2018), **The Developing Human Connectome Project: A minimal processing pipeline for neonatal cortical surface reconstruction**. *NeuroImage.* **doi:** 10.1016/j.neuroimage.2018.01.054

1.6. Thesis overview

We start with an introduction to brain image registration in chapter 2, where we highlight different transformation models, energy formulations, and available software tools. Afterwards, methods used to build average shape and intensity models, in particular spatio-temporal models of the developing human brain, are reviewed in chapter 3. Specific brain atlases of the developing brain are summarised in section 3.3. In chapter 4, methods for the segmentation of neonatal brains and the modelling of the cerebral cortex are discussed. A number of these methods utilise the brain atlases presented in the preceding chapter.

Our first contribution, the MIRTK for efficient, unbiased, and diffeomorphic neonatal brain image registration is summarised and evaluated in chapter 5. This is followed by a description of our proposed spatio-temporal brain atlas construction in chapter 6. Our method for the reconstruction of the neonatal cortex is formulated in chapter 7.

We conclude with a summary of our achievements, discuss limitations, and provide directions for future work in chapter 8.

Part II.

Background

In the following chapters, we cover the main concepts and advances in computational medical image analysis, with special attention to techniques derived for the processing and modelling of the developing brain during the neonatal period.

A prerequisite for the construction of brain templates and the analysis of brain images, is the registration to a common anatomical coordinate frame of reference. This problem is formulated in chapter 2, where we give an overview of the methodology and the algorithms involved. The reviewed methods are primarily based on continuous cost functions, which are minimised using techniques from numerical optimisation. Our focus is on methods for the registration of brain images acquired from the anatomy of different individuals. A more comprehensive overview of methods is given in Sotiras et al. (2013).

The construction of volumetric average shape and intensity models of the human brain is discussed in chapter 3. In section 3.3, we give an overview of atlases created with these techniques, including atlases of the immature brain during gestation, the first weeks after birth, and within the first two years of life.

Probabilistic atlases find application in brain image segmentation, where prior information about the distribution of brain tissues and structures is inferred from the atlas. The construction of these atlases in turn requires a segmentation of the brain into different ROIs. Methods for segmenting a neonatal brain image are covered in chapter 4. Deformable models for the segmentation of the cortex are reviewed in section 4.5.

Chapter 2.

Brain image registration

One of the core disciplines in medical image analysis is finding a solution to the problem of establishing anatomical correspondences between two or more images. In case of two images, this is referred to as pairwise registration, while in case of more than two images, it is a groupwise registration problem. We further distinguish between mono-modal and multi-modal image registration. Mono-modal registration refers to the registration of images acquired with the same imaging modality (and acquisition parameters), while multi-modal registration establishes correspondences between images acquired with different imaging modalities (or acquisition parameters). When the anatomy depicted in the images is of the same patient, we speak of intra-subject registration, and inter-subject registration otherwise. If more than one image of each anatomy are used in the registration process, acquired either within the same scanner coordinate system or co-registered beforehand, we refer to it as multi-channel or multi-contrast registration. In this work, we focus on the registration of T2w images of immature brains of different infants, i.e., the single-channel, mono-modal, and inter-subject registration problem.

Formally, given n_I scalar image volumes $I_i : \mathbb{R}^d \supset \Omega_i \to \mathbb{R}$ for $0 \leq i < n_\mathrm{I}$, image registration aims to find spatial mappings, $\mathbf{\Phi}_i : \Omega_0 \to \Omega_i$, of points in one of the d-dimensional images, I_0, to the respective points of the equivalent anatomy depicted in the other images, I_i for $0 < i < n_\mathrm{I}$, where in case of pairwise registration $n_\mathrm{I} = 2$ and $\mathbf{\Phi} = \mathbf{\Phi}_1$. These mappings are illustrated in figure 2.1. The image I_0 is hereby referred to as reference, fixed, or target image, while the remaining images, I_i, are called study, moving, or source

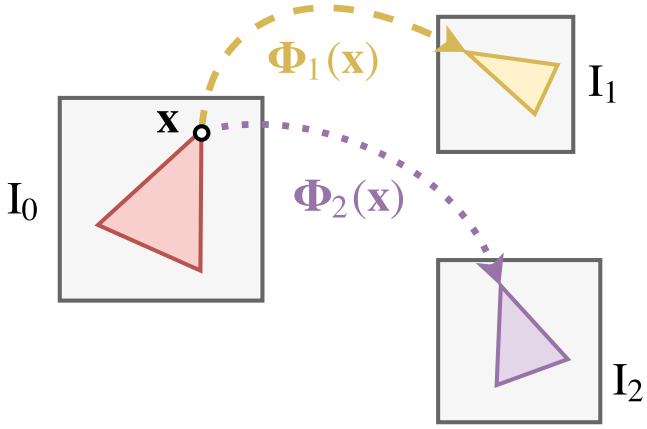


Figure 2.1.: Illustration of spatial mappings between a target image and two source images. The corner of the red triangle at position $\boldsymbol{x} \in \Omega_0$, which is marked by a circle in I_0, is mapped by $\boldsymbol{\Phi}_1$ (yellow dashed arrow) to the corresponding corner at position $\boldsymbol{\Phi}_1(\boldsymbol{x}) \in \Omega_1$ of the triangle in image I_1, and by $\boldsymbol{\Phi}_2$ (purple dotted arrow) to the corner at $\boldsymbol{\Phi}_2(\boldsymbol{x}) \in \Omega_2$ in I_2. The finite image domains $\Omega_{0,1,2}$ are depicted by three grey shaded squares. While a unique mapping between corners with acute angles can be determined, this example demonstrates the particular difficulty of ambiguous correspondences such as for the corners of the edge opposite to $\boldsymbol{x}$. The inter-subject registration problem is further complicated when shapes of anatomical structures are not isomorphic, e.g., due to normal biological variability of anatomies as in case of differing cortical folding observed in human brains.

images. This nomenclature reflects the resampling of source image intensities at the set of discrete voxels of the target anatomy, $\boldsymbol{x}_k \in \Omega_0$, after the spatial mapping. Changing the spatial mapping results in a moving source image, whose anatomical configuration is identical to the fixed reference anatomy, when a perfect registration has been achieved. The existence of such perfect match assumes the existence of an isomorphism between the imaged anatomies, i.e., a bijective mapping between similar shapes, although brain shapes vary considerably. Instead, registration maximises a measure of similarity as to match the anatomies on some macroscopic level of representation with sufficient structural similarity.

The components of a registration algorithm are summarised briefly in the following, in order to provide an overview of the more detailed discussion given in the following sections.

Interpolation strategy

The source image has to be transformed to the target image domain to facilitate a comparison of the images for assessing the quality of the anatomical alignment. This is accomplished by mapping the discrete target points to the source image domain. The intensity

value of the source anatomy at the mapped point, i.e., $I_i(\mathbf{\Phi}_i(\boldsymbol{x}_k))$, is then assigned to this location in the transformed source image. Because mapped points do generally not coincide with points at which the discrete source image is defined, intensity values have to be interpolated. The most frequently used interpolation strategy is linear interpolation, due to its simplicity and efficiency. A discrete image with an interpolation strategy gives a continuous intensity function, denoted in this work by I.

Matching criteria

The quality of the image alignment is assessed by one or more matching criteria. These criteria can be based on the location of either manually or automatically extracted image features, also referred to as landmarks or fiducial markers in medical imaging, a set of curves such as sulcal curves on the cortical surface, surface models such as the inner- or outer-cortical surface, or an intensity-based measure of image similarity. Here, we focus on intensity-based image similarity measures. Because image registration is formulated as cost minimisation problem, the negative of an image similarity measure, which increases when the images become more similar, is used in the cost function formulation. We therefore refer to the intensity-based matching criterion also as image dissimilarity or image distance measure, where we assume the sign of the measure to be such that a smaller value corresponds to a better alignment of the imaged anatomies. Image dissimilarity measures for brain image registration are presented in section 2.1.

Transformation model

A parametrisation of $\mathbf{\Phi}_i(\boldsymbol{x})$ combined with one or more regularisation terms to constrain the space of admissible transformations, is referred to as transformation model. The number of parameters of this model defines the number of degrees of freedom (DOFs) of the optimisation problem. In case of a non-parametric transformation, the registration is either formulated as a variational minimisation of an energy functional (Bajcsy et al., 1983; Amit, 1994; Hermosillo and Faugeras, 2001; Fischer and Modersitzki, 2004), or as maximum *a posteriori* estimation problem using a Bayesian formulation (Christensen, 1994; Dupuis et al., 1998). In a Bayesian formulation, the transformation is given by a discrete random

variable, and the regularisation is given by its prior probability distribution. In case of a parametric transformation model, no additional regularisation may be required when the number of DOFs is low. In general, however, the registration problem is ill-posed, i.e., no unique solution exists, and the function measuring the quality of the image match is highly non-convex with many local minima. Regularisation of the solutions based on properties of the spatial mapping, e.g., smoothness, is thus needed. Different types of transformation models are summarised in section 2.2, followed by a brief discussion of additional task-specific regularisations in section 2.3.

When the transformation can vary locally, the registration is called non-rigid, non-linear, or deformable, where the terms non-rigid and non-linear do not state clearly that the registration is not restricted to an affine transformation, which is neither rigid nor linear in the original image space. We thus prefer the term "deformable registration".

Energy function

The brain image registration problem is stated as a cost minimisation problem, where the sum of matching criteria, whose values quantify the mismatch of the anatomies, and a regularisation term is minimised. In the context of image registration, this cost function is called registration energy. When the transformation is non-parametric, it is a functional that assigns a cost or energy to any continuous spatial mapping $\mathbf{\Phi}$.

For asymmetric, intensity-based, and pairwise image registration, it is given by

$$\mathcal{E}\left(\mathbf{\Phi}\right) = \mathcal{D}_{\star}\left(I_0, I_1 \circ \mathbf{\Phi}\right) + \alpha \mathcal{P}_{\star}\left(\mathbf{\Phi}\right) \tag{2.1}$$

where $\mathcal{D}_{\star}$ is the image dissimilarity measure, $\mathcal{P}_{\star}$ is the (sum of) regularisation term(s), which penalise irregular or undesired transformations, and $\alpha \in \mathbb{R}_{\geq 0}$ is a weighting factor that balances image matching and transformation regularity. The transformation found by minimising a non-linear energy function given by equation (2.1) depends on the order of the images. When the target and source images are exchanged, the result will not necessarily be the inverse of this transformation. Symmetric formulations of the image registration problem, which enforce this inverse consistency, are discussed in section 2.4.

Optimisation method

The transformation, which maximises the image similarity while conforming to the required regularity of the solution, is found by minimising the registration energy. In general, no closed form solution exists for the optimal transformation. Hence, a numerical method has to be employed. Common approaches include steepest gradient descent (Rueckert et al., 1999; Hermosillo et al., 2001), conjugate gradient descent (Hestenes and Stiefel, 1952; Press et al., 2007), stochastic gradient descent (Klein et al., 2009b), and Gauss-Newton methods such as the Levenberg-Marquardt algorithm (Press et al., 2007).

A comparison of optimisation methods for the free-form deformation model, which is closely related to this work, can be found in Kybic and Unser (2003) and Klein et al. (2007). For a comprehensive overview of optimisation techniques used for the deformable registration of medical images in general, including discrete methods for graph-based models not discussed here, see Sotiras et al. (2013).

2.1. Image dissimilarity measures

In this section, we define common image dissimilarity measures used in pairwise mono- and multi-modal brain image registration. The target and source images are hereby assumed to be defined on the same image domain Ω. For this, the source image is transformed by the current spatial mapping $\mathbf{\Phi} : \Omega_0 \rightarrow \Omega_1$ to the domain of the target image, i.e., $\Omega = \Omega_0$. The image dissimilarity is thus given by $\mathcal{D}_\star \left(I_0, I_1 \circ \mathbf{\Phi}(\boldsymbol{x})\right)$. In case of a symmetric approach based on a velocity-based transformation model, both images may be transformed.

Of the reviewed dissimilarity measures, local normalised cross-correlation and normalised mutual information are employed in our work.

2.1.1. Mono-modal image dissimilarities

The following image dissimilarity measures directly compare the image intensities, and therefore rely on similar intensities in both images for the same structures.

Sum of squared differences

Undoubtedly the most common distance measure in mono-modal image registration is the sum of squared differences (SSD) due to its simplicity, which makes it easy to compute in parallel, and its differentiability. It is also known as mean squared error (MSE), i.e.,

$$\mathcal{D}_{\text{SSD}}(I_0, I_1) = \frac{1}{2|\Omega|} \int_\Omega |I_0(\boldsymbol{x}) - I_1(\boldsymbol{x})|^2 d\boldsymbol{x} \tag{2.2}$$

Sum of absolute differences

A voxel-wise image dissimilarity measure, which does not emphasis larger intensity errors, but penalises errors linearly proportional to the intensity difference, is given by the sum of absolute differences (SAD), i.e.,

$$\mathcal{D}_{\text{SAD}}(I_0, I_1) = \frac{1}{2|\Omega|} \int_\Omega |I_0(\boldsymbol{x}) - I_1(\boldsymbol{x})| d\boldsymbol{x} \tag{2.3}$$

This measure is more robust to outliers, but is not differentiable for $I_0(\boldsymbol{x}) = I_1(\boldsymbol{x})$.

2.1.2. Multi-modal image dissimilarities derived from statistics

Image dissimilarities that are robust to differences in the MRI signal evaluate the similarity of two images based on the statistical distributions of the intensities as opposed to directly comparing their absolute intensity values. A linear dependence of intensities is assumed for image regions within which statistics are computed (Sotiras et al., 2013).

Cross-correlation

The cross-correlation (CC) measure was originally used for 2D template matching, where the most likely position of a template, T, is searched for within an image, I, e.g., images of an object within a photograph (Lewis, 1995). It is motivated by the MSE of the images within the template region Ω for different translations $\boldsymbol{u} \in \mathbb{R}^d$ of the image, i.e.,

$$MSE(T, I; \boldsymbol{u}) = \frac{1}{2|\Omega|} \int_\Omega |T(\boldsymbol{x}) - I(\boldsymbol{x} + \boldsymbol{u})|^2 d\boldsymbol{x} \tag{2.4}$$

which can be expanded to

$$MSE(T, I; \boldsymbol{u}) = \frac{1}{2|\Omega|} \int_{\Omega} |T(\boldsymbol{x})|^2 - 2T(\boldsymbol{x})I(\boldsymbol{x}+\boldsymbol{u}) + |I(\boldsymbol{x}+\boldsymbol{u})|^2 d\boldsymbol{x} \tag{2.5}$$

where $\int_{\Omega} |T(\boldsymbol{x})|^2 d\boldsymbol{x}$ is constant. It is now assumed that also the term $\int_{\Omega} |I(\boldsymbol{x}+\boldsymbol{u})|^2 d\boldsymbol{x}$ is approximately constant for all translations. The remaining term is the cross-correlation of the image and the translated template. When the template image corresponds to a sub-region of the image I, this value takes on its maximum for a translation $\boldsymbol{u}$, which perfectly aligns the template with this sub-region. The cross-correlation for all possible translations can be computed by convolving the image with the template.

The CC image dissimilarity measure of a target image, I_0, and the co-registered source image, I_1, is given by the negative cross-correlation of these images for $\boldsymbol{u} = 0$, i.e.,

$$\mathcal{D}_{\mathrm{CC}}(I_0, I_1) = -\frac{1}{|\Omega|} \int_{\Omega} I_0(\boldsymbol{x}) I_1(\boldsymbol{x}) d\boldsymbol{x} \tag{2.6}$$

Normalised cross-correlation

With changing lighting conditions, the assumption that $\int_{\Omega} |I(\boldsymbol{x}+\boldsymbol{u})|^2 d\boldsymbol{x}$ in equation (2.4) is constant for all translations $\boldsymbol{u}$ is no longer valid (Lewis, 1995; Di Stefano et al., 2003). A measure, which is robust to additive and multiplicative intensity dependence, is obtained by subtracting the mean intensity from the images, and by dividing the CC of these images by their auto-correlation value.

The normalised cross-correlation (NCC) image dissimilarity is given by

$$\mathcal{D}_{\mathrm{NCC}}(I_0, I_1) = -\frac{\left[\int_{\Omega} \tilde{I}_0(\boldsymbol{x}) \tilde{I}_1(\boldsymbol{x}) d\boldsymbol{x}\right]^2}{2\left[\int_{\Omega} \tilde{I}_0(\boldsymbol{x})^2 d\boldsymbol{x}\right]\left[\int_{\Omega} \tilde{I}_1(\boldsymbol{x})^2 d\boldsymbol{x}\right]} \tag{2.7}$$

where $\tilde{I}_i(\boldsymbol{x}) = I_i(\boldsymbol{x}) - \mu_i$ are the zero mean images, and $\mu_i = \frac{1}{|\Omega|} \int_{\Omega} I_i(\boldsymbol{x}) d\boldsymbol{x}$ is the mean intensity of image I_i. The NCC is squared to avoid the square root in the denominator. The nominator corresponds to the covariance of the zero mean intensity distributions, whereas the denominator is the product of the variance of image intensities in each image. This ratio is also known as the correlation coefficient (Roche et al., 1998).

Local normalised cross-correlation

A measure, which is robust to locally varying intensities, is obtained by integrating the NCC for small image regions $\omega(\boldsymbol{x}) \subset \Omega$ centred at $\boldsymbol{x} \in \Omega$ over the whole image domain (Bajcsy and Kovacic, 1989; Avants et al., 2008; Lorenzi et al., 2013). The local statistics can be efficiently computed by convolution with a window function (Cachier et al., 2003).

The local normalised cross-correlation (LNCC) image dissimilarity measure is given by

$$\mathcal{D}_{\mathrm{LNCC}}(I_0, I_1) = -\frac{1}{2}\int_{\Omega} \frac{\left[\int_{\omega(\boldsymbol{x})} \tilde{I}_0(\boldsymbol{y})\tilde{I}_1(\boldsymbol{y})d\boldsymbol{y}\right]^2}{\left[\int_{\omega(\boldsymbol{x})} \tilde{I}_0(\boldsymbol{y})^2 d\boldsymbol{y}\right]\left[\int_{\omega(\boldsymbol{x})} \tilde{I}_1(\boldsymbol{y})^2 d\boldsymbol{y}\right]} d\boldsymbol{x} \tag{2.8}$$

where $\tilde{I}_i(\boldsymbol{x}) = I_i(\boldsymbol{x}) - \mu_i(\boldsymbol{x})$, and the local mean intensity is given by

$$\mu_i(\boldsymbol{x}) = \frac{1}{|\omega(\boldsymbol{x})|}\int_{\omega(\boldsymbol{x})} I_i(\boldsymbol{y})d\boldsymbol{y} \tag{2.9}$$

Correlation ratio

The correlation ratio (CR) image dissimilarity measure was proposed for multi-modal image registration by Roche et al. (1998). Given the sets $\omega(a) = \{\boldsymbol{x}|I_0(\boldsymbol{x}) = a\}$ of points for which the target image has the discrete intensity value a, the ratio of the mean variance of source image intensities within each image region $\omega(a)$, and the variance of source image intensities over the whole image domain is given by

$$\mathcal{D}_{\mathrm{CR}}(I_0, I_1) = \frac{\sum_a n_a \sigma_a^2}{n\sigma^2} \tag{2.10}$$

where $n_a = |\omega(a)|$, $n = |\Omega|$, and the source intensity variance values are given by

$$\sigma = \frac{1}{n}\int_{\Omega} (I_1(\boldsymbol{x}) - \mu)^2 \, d\boldsymbol{x} \quad \text{where} \quad \mu = \frac{1}{n}\int_{\Omega} I_1(\boldsymbol{x}) \, d\boldsymbol{x} \tag{2.11}$$

and

$$\sigma_a = \frac{1}{n_a}\int_{\omega(a)} (I_1(\boldsymbol{x}) - \mu_a)^2 \, d\boldsymbol{x} \quad \text{where} \quad \mu_a = \frac{1}{n_a}\int_{\omega(a)} I_1(\boldsymbol{x}) \, d\boldsymbol{x} \tag{2.12}$$

Contrary to most image dissimilarity measures, $\mathcal{D}_{\mathrm{CR}}(I_0, I_1) \neq \mathcal{D}_{\mathrm{CR}}(I_1, I_0)$.

2.1.3. Multi-modal image dissimilarities from information theory

In the following, we consider the random variables A and B of image intensities. The joint probability distribution, $p_{AB}(a, b)$, of intensities $a \in I_0$ and $b \in I_1$ can be estimated from the discrete images using Parzen window interpolation (Thévenaz and Unser, 2000) or generalised partial volume estimation (Hua-mei Chen and Varshney, 2003). A comparison of these estimation techniques is given in Loeckx et al. (2006).

The marginal probability distributions are given by

$$\begin{aligned} p_A(a) &= \int_B p_{AB}(a, b)\, db \\ p_B(b) &= \int_A p_{AB}(a, b)\, da \end{aligned} \tag{2.13}$$

The joint and marginal probability distributions are used to define the following quantities. A survey of these measures and their properties is given in Pluim et al. (2003).

Joint entropy

Minimising Shannon's joint entropy (JE) given by

$$\begin{aligned} \mathcal{D}_{\mathrm{JE}}(I_0, I_1) &= H(A, B) \\ &= -\int_A \int_B p_{AB}(a, b) \log\left(p_{AB}(a, b)\right)\, db\, da \end{aligned} \tag{2.14}$$

reduces the dispersion of the joint probability distribution. When the images are in alignment, it has a few sharply defined, dominant peaks (Pluim et al., 2003). Both Studholme et al. (1995) and Collignon et al. (1995) proposed it for multi-modal image registration. This measure is highly sensitive to the image overlap of the anatomies. Low values can be achieved with a complete misalignment of the images, where only the background regions of the two images overlap (Pluim et al., 2003).

Mutual information

Viola and Wells (1995) and Maes et al. (1997) independently introduced an information theoretic quantity, which measures the balance between marginal and joint entropies. This quantity is the Kullback-Leibler divergence (KLD) of $p_{AB}(a, b)$ from the joint distribution of independently distributed random variables given by $p_A(a)p_B(b)$ (Pluim et al., 2003). In terms of entropy, the negative mutual information (MI) is

$$\mathcal{D}_{\mathrm{MI}}(I_0, I_1) = -MI(A, B) = -\left(H(A) + H(B) - H(A, B)\right) \tag{2.15}$$

with joint entropy given by equation (2.14), and marginal entropies

$$H(X) = -\int_X p_X(a) \log p_X(a)\, da \tag{2.16}$$

where $X \in \{A, B\}$, and the marginal probability distributions are given by equation (2.13).

Similar to the JE measure, MI is sensitive to the image overlap. However, as the marginal entropies have to be maximised, while the joint entropy is to be minimised, the MI penalises complete misalignments where only background overlaps (Pluim et al., 2003).

Normalised mutual information

A normalised variant of MI was proposed by Studholme et al. (1999). This normalised mutual information (NMI) is invariant to the amount of overlap between the images. It is thus more robust than MI, and therefore more commonly used for the global alignment of multi-modal images. The NMI is the ratio of the sum of marginal entropies and the joint entropy, i.e.,

$$\mathcal{D}_{\mathrm{NMI}}(I_0, I_1) = -NMI(A, B) = -\frac{H(A) + H(B)}{H(A, B)} \tag{2.17}$$

An algorithm for the parallel computation of the derivatives with respect to the coefficients of the cubic B-spline free-form deformation model was developed by Modat et al. (2009). The robustness of the NMI measure to both anatomical overlap and intensity variations, makes it suitable for the registration of brain images of the early developing brain, where the MR appearance of brains at different ages varies greatly, as seen in chapter 1.

2.2. Transformation models

In this section, we review common transformation models used in brain image registration, including global mappings, non-parametric small and large deformations, and parametric spline models. Of these, the global transformations, cubic B-spline free-form deformation, and exponential map of a stationary velocity field are most relevant to our work.

2.2.1. Global transformations

A global transformation of $\boldsymbol{x} \in \Omega_0 \subset \mathbb{R}^d$ is given by

$$\boldsymbol{\Phi}_{\text{global}}(\boldsymbol{x}) = \boldsymbol{M}_\star \boldsymbol{x} + \boldsymbol{u} \tag{2.18}$$

where $\boldsymbol{u} \in \mathbb{R}^d$ is a translation vector, and $\boldsymbol{M}_\star \in \mathbb{R}^{d \times d}$ is a linear mapping. This mapping can be formulated as a single matrix product using homogeneous coordinates such that the composition of global transformations is simply a matrix product.

The homogeneous coordinate transformation is

$$\begin{pmatrix} \boldsymbol{\Phi}_{\text{global}}(\boldsymbol{x}) \\ 1 \end{pmatrix} = \underbrace{\begin{pmatrix} & \boldsymbol{M}_\star & & \boldsymbol{u} \\ 0 & 0 & 0 & 1 \end{pmatrix}}_{\boldsymbol{A}} \begin{pmatrix} \boldsymbol{x} \\ 1 \end{pmatrix} \tag{2.19}$$

Global transformations are therefore also referred to as linear transformations, although the translation of $\boldsymbol{x}$ is not a linear mapping. In the following, we define different mappings $\boldsymbol{M}_\star$ of common global transformations for $d = 3$ with increasing number of DOFs as the matrix product of elementary coordinate transformations, i.e., rotation, scaling, and shearing. Note that the order in which these are composed may differ for different registration tools, in which case the optimal parameter values for a given pair of images may be different. This is because matrix multiplication is not commutative.

Rigid transformation

This transformation consists of a sequence of translations and rotations, e.g.,

$$\boldsymbol{M}_{\text{rigid}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_{\text{x}} & -\sin\theta_{\text{x}} \\ 0 & \sin\theta_{\text{x}} & \cos\theta_{\text{x}} \end{pmatrix} \begin{pmatrix} \cos\theta_{\text{y}} & 0 & \sin\theta_{\text{y}} \\ 0 & 1 & 0 \\ -\sin\theta_{\text{y}} & 0 & \cos\theta_{\text{y}} \end{pmatrix} \begin{pmatrix} \cos\theta_{\text{z}} & -\sin\theta_{\text{z}} & 0 \\ \sin\theta_{\text{z}} & \cos\theta_{\text{z}} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.20}$$

is the product of orthonormal matrices of a rotation around each axis, with angles $\theta_{\text{x,y,z}}$. Alternatively, Euler rotations can be parametrised by a single quaternion (Faugeras and Hebert, 1986; Flifla et al., 1992). Rigid transformations are used in intra-subject registration, and for compensation of differences in scanner coordinates and patient positioning.

Similarity transformation

A rigid transformation with isotropic scaling is given by

$$\boldsymbol{M}_{\text{sim}} = \boldsymbol{M}_{\text{rigid}}\, \boldsymbol{I}s \tag{2.21}$$

with scaling factor s, and $\boldsymbol{I}$ denoting the identity matrix. This transformation is useful for initial "rigid" alignment of neonatal brain images, where brain size can vary clearly.

Affine transformation

The most general of global transformations consists of a rigid transformation, anisotropic scaling, and a shearing of the coordinate axes, i.e.,

$$\boldsymbol{M}_{\text{affine}} = \boldsymbol{M}_{\text{rigid}} \begin{pmatrix} s_{\text{x}} & 0 & 0 \\ 0 & s_{\text{y}} & 0 \\ 0 & 0 & s_{\text{z}} \end{pmatrix} \begin{pmatrix} 1 & \tan\theta_{\text{xy}} & \tan\theta_{\text{xz}} \\ 0 & 1 & \tan\theta_{\text{yz}} \\ 0 & 0 & 1 \end{pmatrix} \tag{2.22}$$

where $s_{\text{x,y,z}}$ are the anisotropic scaling factors in each dimension, and $\frac{\pi}{2} - \theta_{\text{xy}}$, $\frac{\pi}{2} - \theta_{\text{xz}}$, and $\frac{\pi}{2} - \theta_{\text{yz}}$ are the angles between the respective pair of coordinate axes after the shear. When all shearing angles are zero, the affine transformation has 9 DOFs, otherwise 12.

2.2.2. Non-parametric small deformation models

Since the early deformable image registration techniques in medical image analysis, non-parametric transformation models have been employed (Broit, 1981; Bajcsy et al., 1983). Following the notation of Fischer and Modersitzki (2004) and Cahill et al. (2009), the registration problem is formulated in terms of a continuous energy functional, i.e.,

$$\mathcal{E}(\boldsymbol{u}) = \mathcal{D}_\star\left(I_0, I_1 \circ (\mathrm{Id} + \boldsymbol{u})\right) + \alpha \mathcal{S}_\star(\boldsymbol{u}) \tag{2.23}$$

where $\boldsymbol{u}$ is the continuous displacement field of the spatial mapping $\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{x} + \boldsymbol{u}(\boldsymbol{x})$, and $\alpha \in \mathbb{R}_{\geq 0}$ is a weighting factor, which compromises between data fidelity and regularity. The registration task is to find the function $\hat{\boldsymbol{u}}$, which minimises this energy, i.e., $\hat{\boldsymbol{u}} = \arg\min_{\boldsymbol{u}} \mathcal{E}(\boldsymbol{u})$. The regularisation term, $\mathcal{S}_\star(\boldsymbol{u})$, models the desired properties of the deformation. It therefore constitutes the non-parametric transformation model.

A necessary condition for a displacement field $\boldsymbol{u}$ to minimise the energy functional $\mathcal{E}(\boldsymbol{u})$ (under suitable boundary conditions) is given by an Euler-Lagrange equation of the form

$$\alpha \mathcal{A}_\star[\hat{\boldsymbol{u}}](\boldsymbol{x}) = \boldsymbol{f}_\star(\boldsymbol{x}) \tag{2.24}$$

where $\mathcal{A}_\star$ is a partial differential operator arising from the Gâteaux derivative of $\mathcal{S}_\star$ with respect to a small variation of $\boldsymbol{u}$, and the force vector $\boldsymbol{f}_\star$ is the negative gradient field of the image dissimilarity measure. A numerical method for solving this partial differential equation (PDE) finds the deformation for which the internal forces of the transformation model on the left-hand side of equation (2.24) are in equilibrium with the external forces arising from the image dissimilarity measure on the right. Solving the Euler-Lagrange equation directly may be difficult or even impossible. A suboptimal solution is usually obtained instead, by numerical integration of the ordinary differential equation (ODE) $\frac{\partial \boldsymbol{u}}{\partial t} = -\nabla \mathcal{E}(\boldsymbol{u})$ with initial value $\boldsymbol{u} = \mathrm{Id}$ (Hermosillo and Faugeras, 2001).

Because registration of medical images involves deforming biological structures such as organs, researchers have derived transformation models based on physical models, which are plausible for deforming these entities, and reflect how the soft tissues that make up

our body would respond to such deformation. The majority of non-parametric small deformation models are therefore derived from the continuum mechanics of elastic bodies. To simplify the mathematical derivations, and to avoid costly and complex formulations, regularisation terms are usually based on linearised physical models.

Physics-based small deformation models are given by a generalised differential operator

$$\mathcal{A}[\boldsymbol{u}](\boldsymbol{x}) = a\Delta\boldsymbol{u}(\boldsymbol{x}) + b\nabla\left(\nabla\cdot\boldsymbol{u}(\boldsymbol{x})\right) - c\boldsymbol{u}(\boldsymbol{x}) \tag{2.25}$$

where $\nabla = \left(\frac{\partial}{\partial x_1}, \ldots, \frac{\partial}{\partial x_d}\right)$ denotes the gradient, $\Delta = \nabla^2$ is the Laplace operator, and $\nabla\cdot\boldsymbol{u}$ is the divergence field of the displacement field $\boldsymbol{u}(\boldsymbol{x})$. Different combinations of non-negative weights a, b, and c correspond to different variants of this small deformation model (Christensen et al., 1997). It can be noted that the restoring internal forces are proportional to the displacement. This prevents large deformations and limits the application of these models for the deformable registration of fetal, neonatal, and infant brain images. The smoothness imposed by such model does not ensure that the deformation maintains the topology of the anatomy for locally large deformations (Christensen et al., 1997). For this, an additional regularisation term is needed, as discussed in section 2.3.

Differential operators proposed for non-parametric registration under the assumption of small deformations are summarised in the following, along with their energy functionals whose Gâteaux derivative gives rise to the linear differential operator (Fischer and Modersitzki, 2004). Detailed derivations are given in appendix A.

Linear elastic deformation model

In his seminal work, Broit (1981) derives a transformation model from the continuum mechanics of elastic bodies. The differential operator corresponding to the Navier-Cauchy PDE describing the equilibrium condition of an elastic body under small deformations $\boldsymbol{u}(\boldsymbol{x})$ is given by (Gee and Bajcsy, 1999)

$$\mathcal{A}_{\text{elastic}}[\boldsymbol{u}](\boldsymbol{x}) = \mu\Delta\boldsymbol{u}(\boldsymbol{x}) + (\lambda+\mu)\nabla\big(\nabla\cdot\boldsymbol{u}(\boldsymbol{x})\big) \tag{2.26}$$

where the properties of the elastic body are defined by the Lamé coefficients $\mu \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}$, which are related to Young's modulus of elasticity and Poisson's ratio $\nu = \frac{\lambda}{2(\mu+\lambda)}$.

The elastic potential is given by (Fischer and Modersitzki, 2004)

$$\mathcal{S}_{\text{elastic}}(\boldsymbol{u}) = \int_{\Omega} \frac{\mu}{4} \sum_{j,k=1}^{d} \left(\frac{\partial u_k}{\partial x_j} + \frac{\partial u_j}{\partial x_k} \right)^2 + \frac{\lambda}{2} (\nabla \cdot \boldsymbol{u})^2 \, d\boldsymbol{x} \tag{2.27}$$

A multi-resolution approach for 3D images based on this elastic deformation model was first developed by Bajcsy and Kovacic (1989). Christensen et al. (1994a) presented a Bayesian formulation of the elastic registration, where the displacement vectors $\boldsymbol{u}(\boldsymbol{x}_k)$ for voxels $\boldsymbol{x}_k \in \Omega$ are determined by a stochastic differential equation based on a stochastic process known as the Wiener process. This process is related to the Brownian motion of particles suspended in a fluid. Christensen and Johnson (2001) developed a topology preserving and inverse consistent method, whereas Hermosillo et al. (2001) utilised the linear elastic model in a multi-modal registration based on the CR and MI image similarity measures. Another early application of the linear elastic model is given by the surface-based brain image registration method of Davatzikos (1996). In their method, a displacement is computed for each point in the target image, given point correspondences between cortical surfaces that were extracted with a deformable surface model (Davatzikos and Bryan, 1996). A more recent elastic brain registration method, based on an initial cortical surface registration in the spherical domain, was presented by Postelnicu et al. (2009). A filter approximating the linear elasticity constraint for the convolution of the dense displacement field was derived by Bro-Nielsen and Gramkow (1996) in case of the fluid registration, where the same filter is applied to the instantaneous velocity field by multiplication in the Fourier domain (see large deformation models below). This is more efficient than previous approaches based on numerical solvers such as successive over-relaxation (SOR).

Diffusion model

The diffusion model was motivated as an analogy to Maxwell's demons and the thermodynamics of heat diffusion by Thirion (1998). In his seminal approach, an external force derived from the optical flow equations (Horn and Schunck, 1981) is used. The composi-

tion of force field and current deformation is smoothed with a Gaussian filter to obtain a new deformation estimate. These iterative steps led to a fast algorithm, which was lacking sound mathematical derivation at the time. Several researchers have studied Thirion's Demons algorithm since (Cachier et al., 1999; Pennec et al., 1999; Cachier et al., 2003). It is formalised as an alternating optimisation problem by Cachier et al. (2003). The authors show that elastic and fluid registration correspond to two coupled diffusion equations. Stefanescu et al. (2004) proposed an adaptive filtering to model spatially varying elastic properties of different tissues, and re-introduced the compositive update scheme of Thirion's approach, which had been replaced by an additive scheme in Cachier et al. (1999). While Thirion's method has been shown to be an approximation of the linear elastic model, which inspired the fast elastic registration using convolution with a linear filter (Bro-Nielsen and Gramkow, 1996)[1], later approaches derived from this algorithm correspond to a visco-elastic transformation model, which combines the elastic model with the viscous-fluid model of Christensen et al. (1996). Because these methods can recover large deformations, they are discussed in the context of large deformation models.

A well-defined diffusion model for medical image registration, which has already been used in optical flow approaches (Horn and Schunck, 1981), was introduced by Fischer and Modersitzki (2002). The regularisation term of their diffusion model is given by

$$\mathcal{S}_{\text{diff}}(\boldsymbol{u}) = \frac{1}{2}\int_{\Omega} \sum_{k=1}^{d} \|\nabla u_k(\boldsymbol{x})\|^2 d\boldsymbol{x} \tag{2.28}$$

with corresponding linear differential operator of the Euler-Lagrange equation

$$\mathcal{A}_{\text{diff}}[\boldsymbol{u}](\boldsymbol{x}) = \Delta \boldsymbol{u}(\boldsymbol{x}) \tag{2.29}$$

Fischer and Modersitzki (2002) compare their diffusion model to Thirion's algorithm and show that the Gaussian filtering of the dense displacement field is a low order approximation. To solve the linear system of equations of a finite difference approximation

[1]The actual method proposed by Bro-Nielsen and Gramkow (1996) is based on the fluid deformation model of Christensen et al. (1996), where the linear elastic filter is used to smooth the instantaneous velocity field. The same filter can, however, be applied in an efficient elastic registration.

to the diffusion PDE, Fischer and Modersitzki (2002) developed a solver which has linear complexity in the number of points. Their diffusion registration method is thus very fast.

Thin-plate model

Unlike previous regularisations that penalise affine transformations, a term based on an approximation to the curvature of the deformation was proposed by Fischer and Modersitzki (2003). This curvature registration does not require a prior rigid registration, when the anatomies have a sufficient overlap[2]. In most cases, a pre-alignment is still beneficial to avoid suboptimal solutions due to noise and local minima of the image dissimilarity.

The curvature registration model is given by the energy functional

$$\mathcal{S}_{\text{curv}}(\boldsymbol{u}) = \frac{1}{2}\int_{\Omega} \sum_{k=1}^{d} \left(\Delta u_k(\boldsymbol{x})\right)^2 d\boldsymbol{x} \tag{2.30}$$

with linear differential operator corresponding to the biharmonic equation, $\Delta^2 \boldsymbol{u} = \boldsymbol{0}$, i.e.,

$$\mathcal{A}_{\text{curv}}[\boldsymbol{u}](\boldsymbol{x}) = \Delta^2 \boldsymbol{u}(\boldsymbol{x}) \tag{2.31}$$

A registration based on this non-parametric model is closely related to the parametric thin-plate spline (TPS) model of Bookstein (1989) discussed in the following section, which minimises the bending energy given by

$$\begin{aligned} \mathcal{S}_{\text{BE}}(\boldsymbol{\Phi}) &= \frac{1}{2}\int_{\Omega} \sum_{k=1}^{d} \|D^2 u_k(\boldsymbol{x})\|_F^2 \, d\boldsymbol{x} \\ &= \frac{1}{2}\int_{\Omega} \sum_{i,j,k=1}^{d} \left(\frac{\partial^2 u_k(\boldsymbol{x})}{\partial x_i \partial x_j}\right)^2 d\boldsymbol{x} \end{aligned} \tag{2.32}$$

where $D^2 u_k(\boldsymbol{x})$ denotes the Hessian matrix, and $\|\cdot\|_F$ is the Frobenius matrix norm.

Henn (2006) demonstrated that while the regularisation term given by equation (2.30) does not penalise affine transformations, the Dirichlet or Neumann boundary conditions used to solve the associated PDE do penalise these. To overcome this problem, the authors

[2] An overlap of brain anatomies can be ensured after a translation which aligns their centres of mass.

propose to directly minimise the bending energy as opposed to the approximation proposed by Fischer and Modersitzki (2003). A regularisation based on the bending energy was also proposed by Rueckert et al. (1999) for the parametric free-form deformation model.

2.2.3. Non-parametric large deformation models

Large deformation models were pioneered by Christensen et al. (1996), after realising the limitations of elastic deformation models caused by the direct regularisation of the total displacements. The physical models of the restoring internal forces resulting from the strain of an elastic body under external load, i.e., the forces of the image dissimilarity term, further assume small deformations. When the assumption of small deformations is violated, the modelled internal force field no longer corresponds to a spatial mapping that preserves the topology of the anatomy. To obtain a diffeomorphic mapping that is differentiable with differentiable inverse, so that connected sets in the image remain connected, surfaces are mapped as surfaces, and the global relationships between structures are maintained (Grenander and Miller, 1998), strong elasticity must be enforced by increasing the shear modulus, μ, defining the local rigidity.

Viscous fluid model

The composition of a sequence of small diffeomorphic deformations is able to model large diffeomorphic mappings. By introducing an auxiliary vector field representing the instantaneous velocity of a viscous fluid at a time point t, the regularisation based on linear elasticity is applied to each time step individually. The restoring internal forces therefore relax over time, and do not penalise large total deformations.

The model of Christensen et al. (1996) corresponds to the Navier-Stoke PDE of a viscous fluid, which is just the linear elastic PDE of Navier-Cauchy, but with the displacement field $\boldsymbol{u}(\boldsymbol{x})$ substituted by the instantaneous velocity field $\boldsymbol{v}(\boldsymbol{x}, t)$, i.e.,

$$\mathcal{A}_{\text{fluid}}[\boldsymbol{u}](\boldsymbol{x}, t) = \mathcal{A}_{\text{elastic}}[\boldsymbol{v}](\boldsymbol{x}, t) \tag{2.33}$$

where the velocity and displacement vector fields are related by the material derivative

$$\boldsymbol{v}(\boldsymbol{x},t) = \frac{d\boldsymbol{u}(\boldsymbol{x},t)}{dt} = \frac{\partial \boldsymbol{u}(\boldsymbol{x},t)}{\partial t} + \left[\nabla \boldsymbol{u}(\boldsymbol{x},t)\right]^{\mathrm{T}} \boldsymbol{v}(\boldsymbol{x},t) \tag{2.34}$$

At each time point, the instantaneous velocity field is obtained by solving the elastic PDE using SOR, and composed with the current transformation in a forward Euler scheme.

To avoid negative Jacobian determinants due to the discretisation of the deformation, the deformed target image domain, which is no longer sampled on a rectilinear grid after the deformation, has to be "regridded" each time the Jacobian determinant for any point falls below a predefined threshold (Christensen et al., 1996).

The greedy viscous fluid registration proposed by Christensen et al. (1996) results in a diffeomorphic registration approach that is not limited to small deformations. Bro-Nielsen and Gramkow (1996) reduced the runtime of this method significantly by obtaining the instantaneous velocity field by convolving the external force field with a linear filter derived from the elastic PDE, instead of using SOR. Further, Chefd'Hotel et al. (2002) proposed a fluid flow registration using LNCC instead of SSD for multi-modal image registration. A fluid registration approach with additional shape boundary constraints derived from a statistical shape model was developed by Wang and Staib (2000).

Visco-elastic model

Registration methods derived from Thirion's algorithm (Thirion, 1998), employ both fluid and elastic type regularisation. In a visco-elastic model, the image dissimilarity forces are smoothed with a Gaussian filter. This update field is then either added (Pennec et al., 1999) or composed with (Stefanescu et al., 2004) the previous transformation estimate. This is followed by a diffusion-like Gaussian smoothing of the total displacement field. This last step prevents folding of the total deformation such that no "regridding" has to be performed (Stefanescu et al., 2004). Both regularisation steps correspond to an alternating optimisation of an energy function based on an auxiliary vector field of "raw" correspondences (Cachier et al., 2003). Similarly, Cahill et al. (2009) present Demons variants based on Euler-Lagrange equations that are given by two coupled diffusion equations.

These algorithms realise either diffusion, elastic, or fluid registration.

For diffusion-like regularisation, the Demons energy is (Vercauteren et al., 2007)

$$\mathcal{E}(\boldsymbol{u}) = \frac{1}{\sigma_{\mathrm{I}}}\mathcal{D}_{\star}\left(I_0, I_1 \circ (\mathrm{Id} + \boldsymbol{c})\right) + \frac{1}{\sigma_{\mathrm{x}}}\|\boldsymbol{u} - \boldsymbol{c}\|^2 + \frac{1}{\sigma_{\Phi}}\|\nabla \boldsymbol{u}\|^2 \tag{2.35}$$

where σ_{I} accounts for the noise in the images, σ_{x} reflects spatial uncertainty of the correspondence field $\boldsymbol{c}$, and σ_{Φ} controls the amount of regularisation of the spatial mapping.

Large deformation diffeomorphic metric mapping

Using Lie group theory, Trouvé (1995, 1998) studied diffeomorphic mappings generated by velocity fields, i.e., elements of an infinite dimensional Lie algebra. In parallel work, Dupuis et al. (1998) extended the viscous fluid model by considering it as a problem in optimal control theory. Both works define a metric, which measures the distance between diffeomorphic mappings, and derive necessary conditions for the velocity field to generate a flow of diffeomorphisms, i.e., that the solution at $t = 1$ of the ODE

$$\frac{\partial \boldsymbol{\varphi}(\boldsymbol{x}, t)}{\partial t} = \boldsymbol{v}(\boldsymbol{\varphi}(\boldsymbol{x}, t), t) \tag{2.36}$$

with initial condition $\boldsymbol{\varphi}(\boldsymbol{x}, 0) = \boldsymbol{x}$ exists, and where $\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{\varphi}(\boldsymbol{x}, 1)$ is diffeomorphic[3]. The length of the path on the manifold is given by (Trouvé, 1998; Dupuis et al., 1998)

$$d_V\left(\mathrm{Id}, \boldsymbol{\Phi}\right) = \int_0^1 \|\boldsymbol{v}_t\|_V \, dt \tag{2.37}$$

where $\boldsymbol{v}_t(\boldsymbol{x}) = \boldsymbol{v}(\boldsymbol{x}, t)$. The norm $\|\boldsymbol{v}_t\|_V = \|\mathcal{L}\boldsymbol{v}_t\|$ on the tangent space $V \subset \mathbb{R}^d$ enforces the required smoothness of $\boldsymbol{v}_t \in V$ through a linear differential operator $\mathcal{L}$ such that $d_V\left(\mathrm{Id}, \boldsymbol{\Phi}\right) < \infty$. The inner product of the vector space V inducing this norm is given by $\langle \boldsymbol{v}, \boldsymbol{w} \rangle_V = \langle \mathcal{L}\boldsymbol{v}, \mathcal{L}\boldsymbol{w} \rangle_{L^2} = \langle \mathcal{L}^{\dagger}\mathcal{L}\boldsymbol{v}, \boldsymbol{w} \rangle_{L^2}$. The metric d_V depends on the choice of $\mathcal{L}$, and a common choice is the diffusion model given by $\mathcal{L}\boldsymbol{v} = \gamma\boldsymbol{v} - \alpha\Delta\boldsymbol{v}$, with positive γ to ensure non-singularity, and periodic boundary conditions such that $\mathcal{L}^{\dagger} = \mathcal{L}$ (Beg et al., 2005).

Following this theory, the large deformation diffeomorphic metric mapping (LDDMM)

[3]This ODE uses a Lagrangian frame of reference, whereas equation (2.34) is the Eulerian transport.

(Beg et al., 2005) regularisation term is given by

$$\mathcal{S}_{\text{LDDMM}}(\boldsymbol{u}) = \frac{1}{2}\left(d_V\left(\text{Id}, \boldsymbol{\Phi}\right)\right)^2 = \frac{1}{2}\int_0^1 \|\boldsymbol{v}_t\|_V^2 \, dt \tag{2.38}$$

such that the optimal flow is a geodesic from the identity to the diffeomorphism that minimises the image dissimilarity. Finding this flow requires a full space-time optimisation, where not only the velocity at the latest time point is optimised, but the update is propagated through time. In comparison, the greedy viscous-fluid model yields a flow of diffeomorphisms, which is not necessarily the shortest path (Beg et al., 2005).

Geodesic shooting

Miller et al. (2006) introduced the conservation of momentum as a fundamental property of geodesic flows. The velocity and displacement fields at any time point t are hereby prescribed by a set of ODEs known as Hamiltonian's equations of geodesics. In the geodesic shooting literature, these ODEs are referred to as EPDiff equations (Younes et al., 2009). The initial value is given by the momentum $\boldsymbol{m}_0 = \mathcal{L}^\dagger\mathcal{L}\boldsymbol{v}_0$, and convolution of the momentum with the kernel K of the inverse operator yields the velocity field, i.e., $\boldsymbol{v}_t = K \star \boldsymbol{m}_t$.

Due to the conservation of momentum, the regularisation is now simply given by

$$\mathcal{S}_{\text{GS}}(\boldsymbol{u}) = \frac{1}{2}\|\boldsymbol{m}_0\|^2 \tag{2.39}$$

Miller et al. (2006) further show that the optimal initial momentum at any target point is given by a scalar value that is multiplied with the gradient of the target image, i.e., $\boldsymbol{m}_0(\boldsymbol{x}) = a(\boldsymbol{x})\nabla I_0(\boldsymbol{x})$. Instead of having to optimise a velocity vector for every point in time, a single scalar value for each discrete point in the target image is optimised. This greatly reduces the number of DOFs, i.e., the dimension of the parameter space, which in turn results in a more tractable optimisation problem. In more recent work, Singh et al. (2013) suggested to optimise a vector-valued initial momentum field to avoid the propagation of noise in the target image, which is further amplified by finite difference approximations of the gradient. The compact representation of the diffeomorphic mappings

from a template to different anatomies by initial momenta, enables the computation of statistics in the Euclidean tangent space at the identity (Vaillant et al., 2004).

Methods to reduce the runtime of LDDMM have been proposed in recent years. Instead of propagating the velocity update through time with a slow gradient descent, Ashburner and Friston (2011) use a faster Gauss-Newton method to optimise the initial velocity. Inspired by this, Hernandez (2014) derive a preconditioned optimisation. More recently, Zhang and Fletcher (2015) used a low-dimensional discretisation of $\boldsymbol{v}_t$ as band-limited signals in the Fourier domain. Most recently, Yang et al. (2016) trained a convolutional neural network (CNN) on pairs of image patches to predict the initial momentum for each target point in only a fraction of the time required by optimisation-based techniques.

Kernel bundle framework

Earlier we noted briefly, that the norm $\|\cdot\|_V$ is associated with a spatially dependent matrix-valued kernel $\boldsymbol{K} : \Omega \times \Omega \to \mathbb{R}^{d\times d}$. This kernel connects the momentum with the velocity field through the mapping $\boldsymbol{m} \mapsto \boldsymbol{v}(\boldsymbol{x}) = \int_\Omega \boldsymbol{K}(\boldsymbol{x}, \boldsymbol{y})\boldsymbol{m}(\boldsymbol{y})\, d\boldsymbol{y}$. When the kernel is translation and rotation invariant, an efficient convolution of the momentum field with this kernel in the Fourier domain can be used to compute the velocity field. A common choice is therefore an isotropic Gaussian kernel (Younes, 2010; Sommer et al., 2012), i.e.,

$$\boldsymbol{K}_r(\boldsymbol{x}, \boldsymbol{y}) = \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{y}\|^2}{2r^2}\right)\boldsymbol{I} \tag{2.40}$$

where $\boldsymbol{I} \in \mathbb{R}^{d\times d}$ is the identity matrix, and r controls the spatial extent of the smoothing, i.e., the amount of regularisation. The scale at which structures in the image can be aligned thus depends on this parameter.

To enable the simultaneous registration at multiple scales, Risser et al. (2011) used a mixture of n_K Gaussian kernels with parameters r_i and weights w_i, i.e., $\mathbf{MK} = \sum_{i=1}^{n_\mathrm{K}} w_i \boldsymbol{K}_{r_i}$. This, however, only results in a differently shaped kernel, but does not decouple the momenta corresponding to the deformation at different scales. Such decoupling would on one hand allow the statistical analysis of the deformation separately for each scale, and on the other hand enable a sparse representation, where the deformation is represented

compactly at the relevant scales only. To obtain such sparse multi-scale representation, Sommer et al. (2011a,b, 2012) introduced the large deformation diffeomorphic kernel bundle mapping (LDDKBM). In this framework, the vector field, which defines the flow of diffeomorphisms, is given by $\boldsymbol{v} = \int_{I_W} \boldsymbol{v}_r \, dr$, where I_W is the interval of scales, and $\boldsymbol{v}_r \in V_r$ is an element of the tangent space associated with the kernel of scale r. The inner product and norm of the vector space W with elements $\boldsymbol{w} = \{\boldsymbol{v}_r\}$, as well as the energy of the kernel bundle flow, are defined with similar integrals over I_W in Sommer et al. (2012).

In recent work, Pai et al. (2015) used the Wendland kernels (Wendland, 1995) with compact support instead of the Gaussian kernels within this kernel bundle framework. Moreover, they use a stationary velocity field as opposed to a time-varying velocity field. Large deformation models parametrised by a stationary velocity field are discussed next.

Exponential map of stationary velocity field

At the time where geodesic shooting with initial momenta had only been developed for landmark-based registration, Arsigny et al. (2006) generalised the principle logarithm to non-linear deformations, and established the theory of flows of diffeomorphisms generated by stationary velocity fields (SVFs). This flow defines the group exponential map, $\exp : \mathfrak{g} \to G$, of an element of the infinite-dimensional Lie algebra $\mathfrak{g}$ to the Lie group of diffeomorphic mappings G. Specifically, the group exponential map, $\boldsymbol{\Phi} = \exp(\boldsymbol{v})$, is the solution at time 1 of the stationary ODE

$$\frac{\partial \boldsymbol{\varphi}(\boldsymbol{x}, t)}{\partial t} = \boldsymbol{v}(\boldsymbol{\varphi}(\boldsymbol{x}, t)) \tag{2.41}$$

with initial condition $\boldsymbol{\varphi}(\boldsymbol{x}, 0) = \boldsymbol{x}$, i.e., the transformation is given by $\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{\varphi}(\boldsymbol{x}, 1)$. Because the velocity at a point $\boldsymbol{x}$ is constant over time, the SVF $\boldsymbol{v} \in V$ is the infinitesimal generator of a one-parameter subgroup of diffeomorphisms $\boldsymbol{\varphi}_t = \exp(t \cdot \boldsymbol{v})$ (Arsigny et al., 2006), where V is the tangent space at the identity. The inverse map is the group logarithm, where $\boldsymbol{v} = \log(\boldsymbol{\varphi})$ is the element of the Lie algebra whose exponential is equal to $\boldsymbol{\varphi}$. Hence, the group exponential and logarithm map elements from the tangent space to the manifold of diffeomorphisms and vice versa (Bossa and Olmos, 2008).

The flow $\boldsymbol{\varphi}_t$, and therefore the group exponential, can be computed using a numerical integration scheme. For example, Pai et al. (2015) use a forward Euler integration. The advantage of a stationary parametrisation over a time-varying velocity field is, however, due to a method known as scaling and squaring (SS) (Arsigny et al., 2006), i.e.,

$$\exp(\boldsymbol{v}) = \exp\left(2^{-n}\boldsymbol{v}\right)^{2^n} \tag{2.42}$$

where

$$\exp(\boldsymbol{v})^2 = \exp(2\boldsymbol{v}) = \exp(\boldsymbol{v}) \circ \exp(\boldsymbol{v}) \tag{2.43}$$

and $\exp(2^{-n}\boldsymbol{v}) \approx \mathrm{Id} + 2^{-n}\boldsymbol{v}$. Instead of n Euler steps, only $\log_2 n$ squaring steps are needed. A disadvantage is that the exponential cannot be evaluated for a point without computing the map for the whole domain. For this, a Runge-Kutta method is better suited. Bossa et al. (2008) compared different methods for computing the group exponential.

The inverse scaling and squaring (ISS) for the group logarithm is (Arsigny et al., 2006)

$$\log(\boldsymbol{\varphi}) = 2^n \log\left(\boldsymbol{\varphi}^{2^{-n}}\right) \approx 2^n\left(\boldsymbol{\varphi}^{2^{-n}} - \mathrm{Id}\right) \tag{2.44}$$

with a recursive computation of the square root $\boldsymbol{T} = \boldsymbol{\varphi}^{\frac{1}{2}}$, which is found by minimising $\mathcal{E}_{\mathrm{sqrt}}(\boldsymbol{T}) = \frac{1}{2}\int_\Omega \|\boldsymbol{T} \circ \boldsymbol{T}(\boldsymbol{x}) - \boldsymbol{\varphi}(\boldsymbol{x})\|^2 \, d\boldsymbol{x}$ with initial value $\boldsymbol{T} = \frac{1}{2}\boldsymbol{\varphi}$ (Arsigny et al., 2006).

The ISS method requires the relatively costly computation of about $n \approx 10$ square roots (Bossa and Olmos, 2008). Bossa et al. (2007) and Bossa and Olmos (2008) proposed a faster method for computing the group logarithm using the Baker-Campbell-Hausdorff (BCH) formula, which gives a solution for $\boldsymbol{u} = \log(\exp(\boldsymbol{v}) \circ \exp(\boldsymbol{w}))$ for $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{w} \in \mathfrak{g}$, i.e.,

$$\begin{aligned}\boldsymbol{u} = \boldsymbol{v} + \boldsymbol{w} &+ \frac{1}{2}[\boldsymbol{v}, \boldsymbol{w}] + \frac{1}{12}[\boldsymbol{v}, [\boldsymbol{v}, \boldsymbol{w}]] + \frac{1}{12}[[\boldsymbol{v}, \boldsymbol{w}], \boldsymbol{w}] \\ &+ \frac{1}{48}[[\boldsymbol{v}, [\boldsymbol{v}, \boldsymbol{w}]], \boldsymbol{w}] + \frac{1}{48}[\boldsymbol{v}, [[\boldsymbol{v}, \boldsymbol{w}], \boldsymbol{w}]] + O\left((\|\boldsymbol{v}\| + \|\boldsymbol{w}\|)^5\right)\end{aligned} \tag{2.45}$$

where $[\cdot, \cdot]$ is the Lie bracket, which is a bilinear operation defined in the tangent space such that $[\boldsymbol{v}, \boldsymbol{w}] = -[\boldsymbol{w}, \boldsymbol{v}]$. It fulfils the Jacobi identity $[\boldsymbol{u}, [\boldsymbol{v}, \boldsymbol{w}]] + [\boldsymbol{v}, [\boldsymbol{w}, \boldsymbol{u}]] + [\boldsymbol{w}, [\boldsymbol{v}, \boldsymbol{u}]] = 0$.

In the group of diffeomorphisms, this Lie bracket is defined as the Lie derivative, i.e.,

$$[\boldsymbol{v}, \boldsymbol{w}] = D\boldsymbol{w}\, \boldsymbol{v} - D\boldsymbol{v}\, \boldsymbol{w} \tag{2.46}$$

where $D\boldsymbol{v}$ and $D\boldsymbol{w}$ are the Jacobi matrices of first order derivatives of the vector fields $\boldsymbol{v}$ and $\boldsymbol{w}$, respectively, with entries $\frac{\partial v_i}{\partial x_j}$ for row index $1 \leq i \leq d$, and column index $1 \leq j \leq d$.

Statistics such as the mean deformation can now be computed in the tangent space. Arsigny et al. (2006) defined an inversion-invariant distance between two diffeomorphisms generated by SVFs as $d_{\mathrm{A}}(\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2) = \| \log(\boldsymbol{\Phi}_1) - \log(\boldsymbol{\Phi}_2) \|$. The mean deformation which minimises a weighted sum of mean distances to n diffeomorphisms $\boldsymbol{\Phi}_i$ is then given by the Log-Euclidean mean, i.e., $\bar{\boldsymbol{\Phi}} = \exp\left(\sum_{i=1}^{n} w_i \log \boldsymbol{\Phi}_i\right)$ where $\sum_{i=1}^{n} w_i = 1$.

Bossa et al. (2007) proposed a different distance $d_{\mathrm{B}}(\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2) = \| \log\left(\boldsymbol{\Phi}_2 \circ \boldsymbol{\Phi}_1^{-1}\right) \|_V$ that is also invariant under the composition of diffeomorphisms, i.e., translation invariant. Given this distance and the deformations $\boldsymbol{\Phi}_i = \exp(\boldsymbol{v}_i)$, Bossa et al. (2007) compute the mean deformation $\bar{\boldsymbol{\Phi}} = \exp(\bar{\boldsymbol{v}})$ using a recursive definition, i.e.,

$$\exp(\bar{\boldsymbol{v}}_k) = \exp\left(\sum_{i=1}^{n} w_i \log\left(\exp(\boldsymbol{v}_i) \circ \exp\left(-\bar{\boldsymbol{v}}_{k-1}\right)\right)\right) \circ \exp\left(\bar{\boldsymbol{v}}_{k-1}\right) \tag{2.47}$$

At each iteration, the BCH formula is applied $n+1$ times to approximate the compositions.

Hernandez et al. (2007, 2009) and Yang et al. (2015) specialised LDDMM for the SVF case. This limits the solutions to a subset of the group of diffeomorphisms (Ashburner, 2007). A comparison with non-stationary LDDMM for synthetic and real brain images has, however, shown that this does not appear to be a practically relevant limitation. The regularisation of the displacement field $\boldsymbol{u} = \exp(\boldsymbol{v}) - \mathrm{Id}$ is given by

$$\mathcal{S}_{\mathrm{SVF}}(\boldsymbol{u}) = \frac{1}{2}\|\boldsymbol{v}\|_V^2 \tag{2.48}$$

Independently, Ashburner (2007) developed an algorithm for diffeomorphic registration of images using a SVF parametrised by a linear combination of first order B-spline basis functions with periodic boundary condition. His method is the popular "Diffeomorphic Anatomical Registration using Exponentiated Lie algebra" (DARTEL) algorithm imple-

mented in the Statistical Parameter Mapping (SPM) toolbox. The optimisation of the spline coefficients is performed using the Levenberg-Marquardt algorithm, where either diffusion, thin-plate, or linear elastic energy are used for regularisation.

Based on the Demons formulation of Cachier et al. (2003), Vercauteren et al. (2007) developed a diffeomorphic Demons method using composition with the exponential of the update field. A comparison to stationary LDDMM was done by Hernandez et al. (2008). A symmetric approach, where also the total deformation is the exponential of a SVF, was presented in Vercauteren et al. (2008, 2009). Further extensions include the use of LNCC (Lorenzi et al., 2013) or NMI (Modat et al., 2012b) for multi-modal image registration, and the Spectral Log-Demons method of Lombaert et al. (2013a) for very large deformations.

A multi-scale SVF using Wendland kernels and the LDDKBM framework was proposed by Pai et al. (2015). Modat et al. (2012b) used a uniform lattice of cubic B-spline kernels to represent the SVF, and performed a fast SS on the coarser control point lattice.

2.2.4. Spline-based small deformation models

Interpolating spline-based deformation models are given by a superposition of translated basis functions, which are centred at a sparse set of control points in the target image. The majority of these models are used for landmark-based image registration. They can, however, also be used for intensity-based registration, where the control point displacement is iteratively updated in order to minimise the image dissimilarity. Most common spline-based models in intensity-based registration are free-form deformation models with control points placed at the vertices of a regular lattice. These are discussed in more detail after the review of spline-based interpolation techniques for landmark-based registration.

In general, a spline-based deformation model is given by

$$\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{\Phi}_m(\boldsymbol{x}) + \underbrace{\sum_{l=1}^{n_\mathrm{c}} \boldsymbol{G}_\star(\boldsymbol{x} - \boldsymbol{p}_l)\boldsymbol{c}_l}_{\boldsymbol{u}(\boldsymbol{x})} \tag{2.49}$$

where $\boldsymbol{x} \in \Omega_0 \subset \mathbb{R}^d$. The global component $\boldsymbol{\Phi}_m(\boldsymbol{x})$ is a vector of three scalar polynomials $\phi_m(\boldsymbol{x}; \boldsymbol{a}_i) = \boldsymbol{a}_i \cdot \boldsymbol{b}_m(\boldsymbol{x})$ of degree m with polynomial basis $\boldsymbol{b}_m(\boldsymbol{x}) \in \mathbb{R}^{n_\mathrm{b}}$, e.g., for $m = 1$,

$\boldsymbol{b}_1(\boldsymbol{x}) = (1, x_1, \ldots, x_d)^{\mathrm{T}}$, and coefficients $\boldsymbol{a}_i \in \mathbb{R}^{n_\mathrm{b}}$ for $i \in \{1, \ldots, d\}$. The number of control points (landmarks) of the local deformation, $\boldsymbol{u}(\boldsymbol{x})$, is n_c, and $\boldsymbol{G}_\star(\boldsymbol{x} - \boldsymbol{p}_l) \in \mathbb{R}^{d\times d}$ is a matrix-valued basis function centred at the l-th control point with spatial coordinates $\boldsymbol{p}_l \in \Omega_0$ and spline coefficients $\boldsymbol{c}_l \in \mathbb{R}^d$. This spline-based model has $d(n_\mathrm{b} + n_\mathrm{c})$ DOFs.

In landmark-based registration, the coefficients $\boldsymbol{c}_l$ can be determined by solving a system of linear equations based on the interpolation criterion $\boldsymbol{\Phi}(\boldsymbol{p}_l) = \boldsymbol{q}_l$, where $\boldsymbol{q}_l$ is the landmark location in the source image. The polynomial part ensures a certain "polynomial precision", which is required by some models to ensure solvability of the interpolation problem (Fornefett et al., 2001). For $m = 1$, the polynomial part is an affine transformation with 12 DOFs. Additional constraints to the interpolation criterion are therefore required to obtain a fully determined system. These not only ensure solvability, but also that $\boldsymbol{u}(\boldsymbol{x})$ tends to zero away from the landmarks. If $\boldsymbol{G}_\star(\boldsymbol{x})$ is positive definite, i.e., non-zero at the origin and monotonically decreasing for increasing distance, the parameters $\boldsymbol{a}_i$ can be determined by a separate affine registration, or $\boldsymbol{a}_i^{\mathrm{T}} = \left(0, \boldsymbol{e}_i^{\mathrm{T}}\right)$, i.e., $\boldsymbol{\Phi}_m(\boldsymbol{x}) = \boldsymbol{x}$.

Three types of basis functions, $\boldsymbol{G}_\star(\boldsymbol{x})$, are used in spline-based deformation models. The first type of basis functions is given by a single univariate radial basis function (RBF), i.e., $\boldsymbol{G}_\star(\boldsymbol{x}) = \boldsymbol{I}\, U_\star(\|\boldsymbol{x}\|)$. An example are the well-known thin-plate splines (Bookstein, 1989). Another class of basis functions are given by a tensor product of a one-dimensional (1D) basis function, e.g., $\boldsymbol{G}_\star(\boldsymbol{x}) = \boldsymbol{I} \prod_{i=1}^{d} U_\star^{(i)}(x_i)$. A prominent example is the free-form deformation model (Sederberg and Parry, 1986; Rueckert et al., 1999). The third family of functions are the non-diagonal basis functions of elastic-body splines (Davis et al., 1997). Specific examples of each type of basis functions are discussed in more detail below.

Given the vector of source landmark positions $\boldsymbol{q}^{\mathrm{T}} = \left(\boldsymbol{q}_1^{\mathrm{T}}, \ldots, \boldsymbol{q}_{n_\mathrm{c}}^{\mathrm{T}}\right)$, the affine parameters $\boldsymbol{a}^{\mathrm{T}} = \left(\boldsymbol{a}_0^{\mathrm{T}}, \ldots, \boldsymbol{a}_d^{\mathrm{T}}\right)$ and local spline coefficients $\boldsymbol{c}^{\mathrm{T}} = \left(\boldsymbol{c}_1^{\mathrm{T}}, \ldots, \boldsymbol{c}_{n_\mathrm{c}}^{\mathrm{T}}\right)$, which interpolate these landmark positions for a given basis function $\boldsymbol{G}_\star(\boldsymbol{x})$, are the solution of the linear system

$$\begin{pmatrix} \boldsymbol{W} & \boldsymbol{B} \\ \boldsymbol{B}^{\mathrm{T}} & \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{c} \\ \boldsymbol{a} \end{pmatrix} = \begin{pmatrix} \boldsymbol{q} \\ \boldsymbol{0} \end{pmatrix} \tag{2.50}$$

where

$$\boldsymbol{W} = \begin{pmatrix} \boldsymbol{G}_\star(\boldsymbol{p}_1 - \boldsymbol{p}_1) & \dots & \boldsymbol{G}_\star(\boldsymbol{p}_1 - \boldsymbol{p}_{n_c}) \\ \vdots & \ddots & \vdots \\ \boldsymbol{G}_\star(\boldsymbol{p}_{n_c} - \boldsymbol{p}_1) & \dots & \boldsymbol{G}_\star(\boldsymbol{p}_{n_c} - \boldsymbol{p}_{n_c}) \end{pmatrix} \tag{2.51}$$

and polynomial bases matrices

$$\boldsymbol{B} = \begin{pmatrix} \boldsymbol{B}_m(\boldsymbol{p}_1) \\ \vdots \\ \boldsymbol{B}_m(\boldsymbol{p}_{n_c}) \end{pmatrix} \tag{2.52}$$

$$\boldsymbol{B}_m(\boldsymbol{x}) = \begin{pmatrix} \boldsymbol{b}_m(\boldsymbol{x}) & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{b}_m(\boldsymbol{x}) & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{b}_m(\boldsymbol{x}) \end{pmatrix}^{\mathrm{T}} \tag{2.53}$$

of the polynomial part given by

$$\boldsymbol{\Phi}_m(\boldsymbol{x}) = \boldsymbol{B}_m(\boldsymbol{x})\boldsymbol{a} \tag{2.54}$$

Uncertainties in landmark localisation can be incorporated by relaxing the interpolation criterion by weighting the influence of each landmark such that displacements of landmarks with greater uncertainty (i.e., lower weight) have less influence on the deformation. This results in a different system of linear equations depending on the particular spline model. When $\boldsymbol{G}_\star(\boldsymbol{x})$ is based on a RBF, the linear system is given by (Fornefett et al., 2001)

$$\begin{pmatrix} \boldsymbol{W} + \lambda\boldsymbol{\Sigma}^{-1} & \boldsymbol{B} \\ \boldsymbol{B}^{\mathrm{T}} & \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{c} \\ \boldsymbol{a} \end{pmatrix} = \begin{pmatrix} \boldsymbol{q} \\ \boldsymbol{0} \end{pmatrix} \tag{2.55}$$

with

$$\boldsymbol{\Sigma} = \mathrm{diag}\left(\frac{1}{\sigma_1^2}\boldsymbol{1}_d^{\mathrm{T}}, \dots, \frac{1}{\sigma_{n_c}^2}\boldsymbol{1}_d^{\mathrm{T}} \right) \tag{2.56}$$

where σ_l is proportional to the localisation error of the l-th landmark, and the d-dimensional vector $\boldsymbol{1}_d = (1, \dots, 1)^{\mathrm{T}} \in \mathbb{R}^d$. In case of the elastic-body splines, the modifications of the

linear system given by equation (2.50) used for interpolation, in order to incorporate landmark approximation weights, is considerably more involved (Wörz and Rohr, 2008).

We first briefly discuss basis functions that are commonly applied in landmark-based image registration. Afterwards, we focus on free-form deformation models, which are more commonly used in intensity-based registration of medical images.

Radial basis functions with global support

The first spline-based models which have been proposed for the smooth deformation of images are based on RBFs with global support, i.e., functions whose value is non-zero everywhere. Widely used RBFs include (see Fornefett et al. (2001), Siddiqui et al. (2009), and references therein):

- Gaussian with $\sigma > 0$:

$$U_{\text{Gauss}}(r) = \exp\left(-\frac{r^2}{2\sigma^2}\right) \tag{2.57}$$

- Multiquadrics ($\mu > 0$) and inverse multiquadrics ($\mu < 0$) with $R > 0$ and $\mu \neq 0$:

$$U_{\text{MQ}}(r) = (r^2 + R^2)^{\mu} \tag{2.58}$$

- Thin-plate splines (TPS):

$$U_{\text{TPS}}(r) = \begin{cases} r^{4-d} \ln r & \text{if } 4 - d \in 2\mathbb{N} \\ r^{4-d} & \text{otherwise} \end{cases} \tag{2.59}$$

Of these, the TPS model was originally introduced by Bookstein (1989) for landmark-based registration of medical images in 2D, and later applied in 3D by Evans et al. (1991). In this model, the RBF is a fundamental solution of the biharmonic equation $\Delta^2 \boldsymbol{u}(\boldsymbol{x}) = \boldsymbol{0}$. Hence, the deformation minimises the bending energy given by equation (2.32), which we discussed earlier in the context of the non-parametric curvature (thin-plate) model.

The drawback of deformation models based on basis functions with global support is, that moving a single control point changes the deformation globally, in particular also

for points distant from this point. Similarly, adding and removing landmarks, e.g., to interactively adjust the deformation, may have unexpected and undesired side effects. It also requires solving a dense linear system, which is more time and memory consuming than solving a system of linear equations with a sparse system matrix $\boldsymbol{W}$.

Radial basis functions with compact support

To overcome the drawbacks of deformation models based on global RBFs, a number of RBFs with compact support in $[0, 1)$ have been proposed. The RBF in equation (2.49) is hereby defined as $U_{\mathrm{c}}(r) = U_{\star}(r/R)$, where R is the radius of the support region, and $U_{\star}(r)$ is one of the RBFs defined below.

Fornefett et al. (2001) proposed the use of the Ψ-functions by Wendland (1995) for elastic registration of medical images. In particular, they suggest the use of $\Psi_{3,1}$, which has a shape similar to the Gaussian function, but compact support, i.e.,

$$U_{\Psi_{3,1}}(r) = \begin{cases} (1-r)^4(4r+1) & \text{if } 0 \leq r < 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.60}$$

Wen et al. (2006) used a RBF referred to as locally constrained deformation (LCD) for the elastic registration of medical 2D images based on manually selected landmarks, followed by a minimisation of the intensity-based SSD criterion. It is given by

$$U_{\mathrm{LCD}}(r) = \begin{cases} (1-r^2)^3 & \text{if } 0 \leq r < 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.61}$$

Siddiqui et al. (2009) proposed an approximately linear RBF with nearly equal distribution of the deformation within the full support range $[0, 1)$. It has the form

$$U_{\cos}(r) = \begin{cases} 0.5\,(1+\cos(r\pi)) & \text{if } 0 \leq r < 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.62}$$

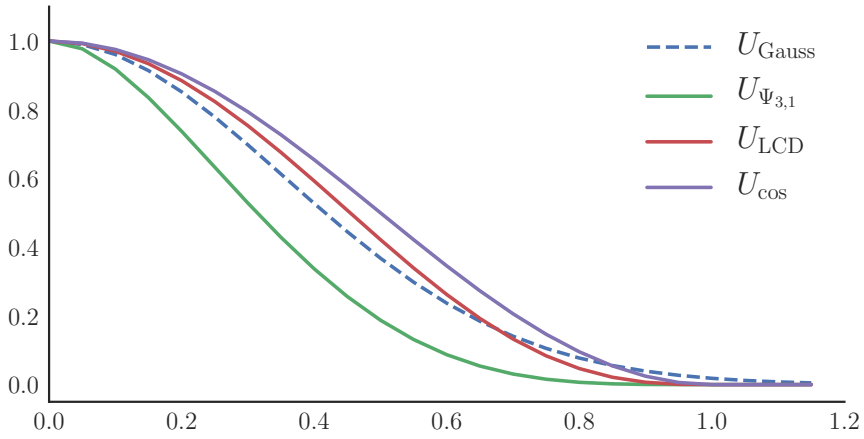


Figure 2.2.: Comparison of compact RBFs and Gaussian with $\sigma = 0.5$.

This RBF is compared by Siddiqui et al. (2009) to Wendland's RBF used by Fornefett et al. (2001), and the LCD basis function used by Wen et al. (2006). For a single displaced landmark, the deformation is more smoothly distributed within the support region with this new RBF. The evaluation of the cosine function is, however, also more costly. A comparison of the shape of these compactly supported RBFs is shown in figure 2.2.

Elastic body splines

This class of interpolating splines was introduced by Davis et al. (1997), and is based on matrix-valued basis functions, $\boldsymbol{G}_\star(\boldsymbol{x})$, which are solutions to the Navier-Cauchy PDE of linear elasticity, $\mathcal{A}_{\text{elastic}}[\boldsymbol{u}](\boldsymbol{x}) = \boldsymbol{f}_\star(\boldsymbol{x})$, with the linear differential operator given by equation (2.26). These deformation models are called elastic-body splines, whereby the basis function depends on the choice of force field $\boldsymbol{f}_\star(\boldsymbol{x})$. The difference to RBF-based spline models is, that the deformation is not independent for each dimension, which is more plausible than independent deformation in each dimension.

Davis et al. (1997) derived the basis functions for polynomial forces given by

$$\boldsymbol{f}_{\text{EBS}}(\boldsymbol{x}) = \boldsymbol{c}_l r^{2k+1} \tag{2.63}$$

where $r = \|\boldsymbol{x}\|$, $k \in \mathbb{Z}$, and $\boldsymbol{c}_l \in \mathbb{R}^d$ is the direction and magnitude of the force in unit distance, with increasing force magnitude away from the origin for $k \geq 0$. The basis

function for rational forces, i.e., $\boldsymbol{f}_{\mathrm{EBS}}(\boldsymbol{x})$ for $k = -1$, has similarity with the basis function of the TPS model. The deformation field that satisfies the linear elasticity PDE is found to be $\boldsymbol{u}(\boldsymbol{x}) = \boldsymbol{G}_{\mathrm{EBS}}(\boldsymbol{x})\boldsymbol{c}_l$ using the Galerkin vector method, where

$$\boldsymbol{G}_{\mathrm{EBS}}(\boldsymbol{x}) = \begin{cases} \alpha r^3 \boldsymbol{I} - 3r\boldsymbol{x}\boldsymbol{x}^{\mathrm{T}} & \text{if } k = 0 \\ \beta r \boldsymbol{I} - r^{-1}\boldsymbol{x}\boldsymbol{x}^{\mathrm{T}} & \text{if } k = -1 \end{cases} \tag{2.64}$$

with $\alpha = 12(1-\nu)-1$, $\beta = 8(1-\nu)-1$, and ν denoting Poisson's ratio. For the landmark-based breast image registration results reported by Davis et al. (1997), the elastic-body spline (EBS) for $k = 0$ gives a better image similarity than both the EBS for $k = -1$ and the TPS. A problem of the polynomial force ($k \geq 0$) is, however, that it is physically not plausible, because it increases with the distance from the point at which the force is exerted, whereas the rational force ($k < 0$) diverges at this point.

Kohlrausch et al. (2005) derived the basis function for physically more plausible Gaussian forces, whose local influence can be controlled by the standard deviation σ_{f}, i.e.,

$$\boldsymbol{f}_{\mathrm{GEBS}}(\boldsymbol{x}) = \frac{1}{\left(\sqrt{2\pi}\sigma_{\mathrm{f}}\right)^3} \exp\left(-\frac{r^2}{2\sigma_{\mathrm{f}}^2}\right) \tag{2.65}$$

The basis function of Gaussian elastic-body splines (GEBS) is (Wörz and Rohr, 2008)

$$\begin{aligned} \boldsymbol{G}_{\mathrm{GEBS}}(\boldsymbol{x}) = \frac{1}{16\pi\mu}\frac{1}{1-\nu}\Bigg[\left((3-4\nu)\frac{\operatorname{erf}(\hat{r})}{r} - \sigma_{\mathrm{f}}\sqrt{\frac{2}{\pi}}\frac{\exp\left(\hat{r}^2\right)}{r^2} + \sigma_{\mathrm{f}}^2\frac{\operatorname{erf}(\hat{r})}{r^3}\right)\boldsymbol{I} \\ + \left(\frac{\operatorname{erf}(\hat{r})}{r^3} + 3\sigma_{\mathrm{f}}\sqrt{\frac{2}{\pi}}\frac{\exp\left(\hat{r}^2\right)}{r^4} - 3\sigma_{\mathrm{f}}^2\frac{\operatorname{erf}(\hat{r})}{r^5}\right)\boldsymbol{x}\boldsymbol{x}^{\mathrm{T}}\Bigg] \end{aligned} \tag{2.66}$$

where $\hat{r} = \frac{r}{\sqrt{2}\sigma_{\mathrm{f}}}$, and the error function $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x \exp\left(-t^2\right) dt$. It is apparent from equation (2.66) that a physically sound model comes at a considerable cost, in terms of the complexity of the basis functions involved. The linear system of equations corresponding to approximating as opposed to interpolating GEBS is given in Wörz and Rohr (2008). To improve the numerical stability for $r \to 0$, and to reduce the computations for $r \gg \sigma_{\mathrm{f}}$, Wörz and Rohr (2008) approximate $\boldsymbol{G}_{\mathrm{GEBS}}(\boldsymbol{x})$ by simplified formulae for these limits.

Free-form deformation

The spline-based model that is more commonly used in intensity-based, as opposed to landmark-based, image registration is the free-form deformation (FFD) model. It has originally been proposed by Sederberg and Parry (1986) for the deformation of geometric solid objects. Different shapes of geometric objects are hereby modelled by deforming the space in which these objects are embedded. In contrast to other spline-based models for the interpolation or approximation of landmark displacements, the FFD is defined by a lattice of uniformly spaced control points within this space. An object is deformed by manipulating the underlying mesh of control points. The resulting deformation produces a smooth and continuous spatial mapping. Sederberg and Parry (1986) used the Bernstein polynomial as basis. Szeliski and Coughlan (1994) proposed B-spline FFD models of various orders for motion estimation from image sequences of real scenes.

The basis function of the cubic B-spline FFD model is given by

$$\boldsymbol{G}_{\mathrm{FFD}}(\boldsymbol{x}) = \boldsymbol{I} \prod_{i=1}^{d} U_{\mathrm{FFD}}\left(\hat{x}_i\right) \tag{2.67}$$

where $\hat{\boldsymbol{x}} = \mathrm{diag}(s_1, \ldots, s_d)^{-1}\boldsymbol{R}\boldsymbol{x} + \boldsymbol{t}$ are the local coordinates of the target point $\boldsymbol{x} \in \Omega_0$ with respect to the control point lattice of size $n_{\mathrm{c}} = \prod_{i=1}^{d} n_i^c$ given by a rotation $\boldsymbol{R}$, translation $\boldsymbol{t}$, and anisotropic scaling with inverse scaling factors corresponding to the uniform spacing s_i of control points in the i-th dimension of the rotated lattice domain. The RBF $U_{\mathrm{FFD}}(r)$ is a B-spline of order 4, i.e., the piecewise function shown left in figure 2.3, where each piece is a polynomial of degree 3 with knots at -2, -1, 0, 1, and 2. This cubic B-spline basis function has compact support in $|r| < 2$, and is given by

$$U_{\mathrm{FFD}}(r) = \begin{cases} B_0(r) & \text{if } -2 \leq r < -1 \\ B_1(r) & \text{if } -1 \leq r < 0 \\ B_2(r) & \text{if } 0 \leq r < 1 \\ B_3(r) & \text{if } 1 \leq r < 2 \\ 0 & \text{otherwise} \end{cases} \tag{2.68}$$

The pieces $B_i(r)$ are derived by Cox-de Boor's recursion formula (de Boor, 2001), i.e.,

$$\begin{aligned} N_i^{(0)}(r) &= 1 \quad \text{if} \quad \xi_i \leq r < \xi_{i+1} \quad \text{and} \quad 0 \quad \text{otherwise} \\ N_i^{(k)}(r) &= \frac{r - \xi_i}{\xi_{i+k} - \xi_i} N_i^{(k-1)}(r) + \frac{\xi_{i+k+1} - r}{\xi_{i+k+1} - \xi_{i+1}} N_{i+1}^{(k-1)}(r) \end{aligned} \tag{2.69}$$

with knot vector $\boldsymbol{\xi} \in \mathbb{R}^{n_\mathrm{k}}$. The B-spline $N_0^{(n_\mathrm{k}-2)}(r)$ of degree $n_\mathrm{k} - 1$ is a non-overlapping linear combination of rectangle functions $N_i^{(0)}(r)$, where the weights are its polynomial pieces. For knot vector $\boldsymbol{\xi} = (-2, -1, 0, 1, 2)$, the cubic B-spline polynomial pieces are

$$\begin{aligned} B_0(r) &= (r+2)^3 \,/\, 6 & B_1(r) &= \left(-3r^3 - 6r^2 + 4\right) / \, 6 \\ B_2(r) &= \left(3r^3 - 6r^2 + 4\right) / \, 6 & B_3(r) &= (2-r)^3 \,/\, 6 \end{aligned} \tag{2.70}$$

While the formulation as superposition of n_c translated cubic B-spline basis functions is mathematically convenient, and emphasises the similarities with other spline-based deformation models, the FFD is more commonly expressed in 3D as (Rueckert et al., 1999)

$$\boldsymbol{\Phi}_{\mathrm{FFD}}(\boldsymbol{x}) = \boldsymbol{\Phi}_{\mathrm{global}}(\boldsymbol{x}) + \underbrace{\sum_{l,m,n=0}^{3} w_l(\tilde{x}_1) w_m(\tilde{x}_2) w_n(\tilde{x}_3) \boldsymbol{c}_{\{(i_1+l)+n_1^\mathrm{c}((i_2+m)+n_2^\mathrm{c}\,(i_3+n))\}}}_{\boldsymbol{u}(\boldsymbol{x}) = \boldsymbol{u}_{\mathrm{FFD}}(\boldsymbol{x};\boldsymbol{c})} \tag{2.71}$$

where $\tilde{\boldsymbol{x}} = \hat{\boldsymbol{x}} - \lfloor \hat{\boldsymbol{x}} \rfloor$ are the coordinates of $\boldsymbol{x}$ relative to the edges of the cell of the lattice that it is contained in, and $\boldsymbol{i} = \lfloor \hat{\boldsymbol{x}} \rfloor - 1$ is the index of the "leftmost" control point, whose B-spline value at $\hat{\boldsymbol{x}}$ is non-zero. Figure 2.3 illustrates the relation between B-spline polynomials $B_i(r)$ and weights $w_l(\tilde{x}_i)$ associated with the control points, i.e.,

$$\begin{aligned} w_0(\tilde{x}_i) &= U_{\mathrm{FFD}}(\tilde{x}_i + 1) = B_3(\tilde{x}_i + 1) \\ w_1(\tilde{x}_i) &= U_{\mathrm{FFD}}(\tilde{x}_i + 0) = B_2(\tilde{x}_i + 0) \\ w_2(\tilde{x}_i) &= U_{\mathrm{FFD}}(\tilde{x}_i - 1) = B_1(\tilde{x}_i - 1) \\ w_3(\tilde{x}_i) &= U_{\mathrm{FFD}}(\tilde{x}_i - 2) = B_0(\tilde{x}_i - 2) \end{aligned} \tag{2.72}$$

This formulation considers the regular placement of control points, and highlights that the deformation at point $\boldsymbol{x}$ depends only on the coefficients of 4^d control points.

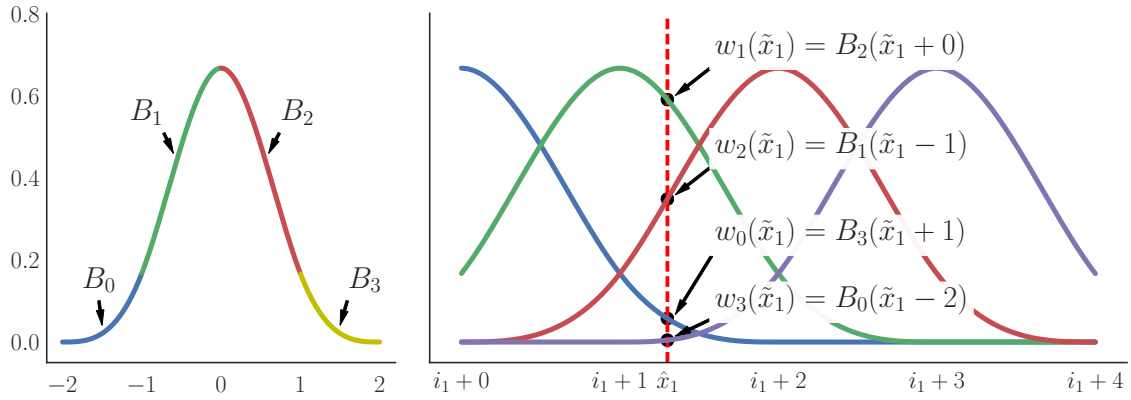


Figure 2.3.: Cubic B-spline FFD weights: (left) piecewise polynomials of B-spline function, (right) weights of control points for evaluation of $\mathbf{\Phi}_{\text{FFD}}(\boldsymbol{x})$ in the first lattice dimension.

An efficient algorithm for the approximation of a FFD based on a cubic B-spline basis function from a set of landmarks in 2D was developed by Lee et al. (1997), which can be easily extended to higher dimensions. Because the initial position of control points does generally not coincide with the target landmark locations, a FFD with pre-specified control point spacing can in general not interpolate the source landmarks exactly. Moreover, in contrast to other approximating spline models, a landmark localisation error is not considered by the FFD approximation. This model is better suited for object shape modelling and intensity-based image registration. It should be noted, though, that Lee et al. (1997) also developed a hierarchical algorithm for interpolation of scattered data.

The advantage of the B-spline basis over the Bernstein polynomial is that it can be subdivided, i.e., the number of control points in each dimension doubled, without a change in the deformation. This property results in an efficient multi-level approach for intensity-based image registration, where at the lower levels a coarse deformation is estimated first, and successively refined on higher levels to capture more localised deformations.

Kybic and Unser (2003) extended their previous FFD-based motion estimation, and applied it to the registration of brain images. The cubic B-spline FFD was introduced to the medical image registration community by Rueckert et al. (1999), who applied it to breast images. This FFD algorithm was extended by Schnabel et al. (2001) to model non-uniform control point spacings by marking points of the uniform lattice as either “active” or

"passive", where the coefficients of "passive" control points remain fixed. The smoothness of the FFD is implicitly controlled by the spacing of the control point lattice, and explicitly by using the bending energy, $\mathcal{S}_{\text{BE}}(\boldsymbol{\Phi})$, given by equation (2.32) as regularisation term. The derivative of this penalty with respect to the FFD coefficients is given in appendix B.3.

2.2.5. Locally affine transformations

The spline-based models offer a rich model with fewer parameters than non-parametric models for both intra- and inter-subject registration. Neither of these deformation models guarantees invertibility, although topology preservation may be encouraged by an additional penalty. Moreover, because these models are quite flexible regarding the deformations of the anatomy that can be accomplished, it is difficult to model the behaviour of rigid structures such as bones. Approaches to preserve the topology, volume, and local rigidity based on additional regularisation terms are considered in section 2.3.

A thin-plate spline model with a spatially adaptive basis function for the deformation of segmented rigid structures was first proposed by Little et al. (1997). A transformation model which interpolates rigid or affine transformations, defined at a number of anchor points, was proposed by Arsigny et al. (2005, 2009). These transformations are guaranteed invertible, and therefore preserve topology. The poly-rigid model is moreover locally rigid. As an extension of equation (2.49), Gu et al. (2014) proposed a FFD with a different affine transformation at each control point for point-based registration of images. The authors note that their B-spline affine transformation (BSAT) may be regarded as a special case of the poly-affine model (Arsigny et al., 2009). It does, however, not ensure invertibility.

2.2.6. Statistical deformation models

Spline-based deformation models can considerably reduce the number of DOFs compared to non-parametric models, thereby reducing the computational cost. But the number of DOFs of a FFD may still be significant. A statistical deformation model (SDM), on the other hand, reduces the number of DOFs by orders of magnitude and restricts the space of admissible solutions to statistically plausible deformations, which are based on the statistical analysis of mappings from a template image to different observed anatomies.

A number of SDMs based on the principle component analysis (PCA) of the FFD coefficients have been proposed (Rueckert et al., 2001, 2003; Loeckx et al., 2003; Pszczolkowski et al., 2012; Onofrey et al., 2013). The PCA hereby learns correlations between the coefficients of the control points. When registering an unseen image to the template, control points are no longer moved independently, but based on the principle components of the previously observed deformations. Wouters et al. (2006) used PCA to learn a SDM from non-parametric viscous-fluid type deformations. Xue et al. (2006) performed a PCA of the wavelet decomposition of non-parametric displacement fields. Kim et al. (2008) used the principle components learned from non-parametric deformations to generate an intermediate template image, which is more similar to the unseen image.

As an example of a SDM, the transformation model of Rueckert et al. (2003), which is based on the principle components of n_Φ FFD coefficient vectors $\boldsymbol{c}^{(i)} \in \mathbb{R}^{dn_c}$, is given by

$$\boldsymbol{\Phi}_{\text{SDM}}(\boldsymbol{x}) = \boldsymbol{\Phi}_{\text{global}}(\boldsymbol{x}) + \boldsymbol{u}_{\text{FFD}}\left(\boldsymbol{x}; \bar{\boldsymbol{c}} + \boldsymbol{C}\boldsymbol{w}\right) \tag{2.73}$$

where $\bar{\boldsymbol{c}} = \frac{1}{n_\Phi}\sum_{i=1}^{n_\Phi} \boldsymbol{c}^{(i)}$ is the average coefficient vector, $\boldsymbol{C}$ is a matrix whose columns are the principle components of the covariance matrix $\frac{1}{n_\Phi - 1}\sum_{i=1}^{n_\Phi}(\boldsymbol{c}^{(i)} - \bar{\boldsymbol{c}})(\boldsymbol{c}^{(i)} - \bar{\boldsymbol{c}})^{\mathrm{T}}$, $\boldsymbol{w}$ is the vector of parameters of the statistical model, i.e., the weights of the principle components, and $\boldsymbol{u}_{\text{FFD}}(\boldsymbol{x}; \boldsymbol{c})$ is the local deformation of the FFD model as defined in equation (2.71). Non-parametric PCA-based SDMs are defined in a similar fashion. The SDM may further be composed with another deformation, e.g., $\boldsymbol{\Phi}_{\text{FFD}} \circ \boldsymbol{\Phi}_{\text{SDM}}(\boldsymbol{x})$ (Onofrey et al., 2013), in order to increase the space of possible deformations and allow the model to adapt to anatomies that are not (sufficiently) represented in the set of training images.

2.3. Task-specific regularisation

The deformation models presented in section 2.2 may be augmented with penalty terms of the energy function, in order to adapt the model to a specific application.

A number of constraints proposed for medical image registration, such as volume preservation and rigidity of bony structures, do not apply to the inter-subject brain image registration task, because skull is removed prior to the registration of soft brain tissues, and anatomies of different subjects vary significantly. These constraints are more common in intra-subject registration. For example, for the registration of breast images with contrast-enhancing lesions (Tanner et al., 2000; Rohlfing et al., 2003), abdominal images containing bony structures (Little et al., 1997; Loeckx et al., 2004; Staring et al., 2007), for strain estimation from a sequence of cardiac images (Mansi et al., 2010), and intra-operative registration of brain images to live images during the procedure (Edwards et al., 1998).

Local rigidity, if not inherent to the model (Little et al., 1997; Arsigny et al., 2009), can be accomplished for spline-based models by coupling the coefficients within a rigid object (Tanner et al., 2000). In general, volume preservation and local rigidity can be enforced based on the Jacobian of the transformation (Rohlfing et al., 2003; Loeckx et al., 2004). The Jacobian determinant is proportional to the local change in volume, where a value of one indicates no change of volume, a value below one indicates volume loss or contraction, and a value greater than one reflects volume increase or expansion. A penalty based on this determinant was combined with further constraints derived from curvature and orthonormality requirements of rigid transformations by Staring et al. (2007).

Of main interest for the registration of images of the developing brain are topology preservation and inverse consistency. Two transformations $\mathbf{\Phi}_0 : \Omega_1 \to \Omega_0$ and $\mathbf{\Phi}_1 : \Omega_0 \to \Omega_1$ are inverse consistent when $\mathbf{\Phi}_0 \circ \mathbf{\Phi}_1 = \mathbf{\Phi}_1 \circ \mathbf{\Phi}_0 = \mathrm{Id}$. That a change in the registration direction results in the inverse mapping is an important property for the construction of unbiased brain atlases, and is ensured by approaches with a symmetric energy formulation. Regularisation terms used to encourage inverse consistency are described in section 2.4, where we discuss symmetric methods. Inverse consistency implies invertible transformations because $\mathbf{\Phi}_0 \circ \mathbf{\Phi}_1 = \mathrm{Id} \iff \mathbf{\Phi}_0 = \mathbf{\Phi}_1^{-1}$, and thus must be either combined with a

model that guarantees invertibility or a regularisation that enforces non-negative Jacobian determinants. Smoothness alone is not sufficient (Christensen et al., 1997).

Penalty terms enforcing invertibility are similar to penalties used for volume preservation. This is because these penalties constrain the range of Jacobian determinants, denoted by $J_{\boldsymbol{\Phi}}(\boldsymbol{x}) = |D\boldsymbol{\Phi}(\boldsymbol{x})|$. The constraints used for topology preservation discourage negative and extreme values, whereas the range of values for volume preservation is narrower.

Kybic and Unser (2000) penalise negative and small determinants using the penalty

$$\mathcal{P}_{\text{TP}}(\boldsymbol{\Phi}) = \int_{\Omega} \exp\left(-\alpha\, J_{\boldsymbol{\Phi}}(\boldsymbol{x})\right)\, d\boldsymbol{x} \tag{2.74}$$

where α defines how quickly the penalty increases with increasing negative value. The penalty is non-zero also for $J_{\boldsymbol{\Phi}}(\boldsymbol{x}) = 1$, which corresponds to no local deformation.

Rohlfing et al. (2003) proposed an incompressibility constraint for the FFD model, which is given by the absolute value of the log-transformed Jacobian determinant. Modat et al. (2012b) used a similar but differentiable term, i.e.,

$$\mathcal{P}_{\text{TP}}(\boldsymbol{\Phi}) = \int_{\Omega} \left(\log\left(J_{\boldsymbol{\Phi}}(\boldsymbol{x})\right)\right)^2\, d\boldsymbol{x} \tag{2.75}$$

This penalty of the local volume change was also derived by Nielsen et al. (2002) in a Bayesian formulation of non-parametric deformations based on the principles of Brownian motion. Their prior model of Brownian warps includes further Jacobian-based penalties for local non-rigidity (skewness), and local rotation.

Yanovsky et al. (2007) derived a regularisation term from the symmetric KLD of the distribution of Jacobian determinant values from their respective distribution in case of the identity transformation, i.e., the Dirac delta function $\delta(J-1)$. Hence, it measures the amount of deformation based on the local volume change, and is unbiased with respect to the log-transformed values. This regularisation has been applied to both elastic (Baluwala et al., 2010) and fluid (Yanovsky et al., 2007; Baluwala et al., 2011) models. It is given by

$$\mathcal{P}_{\text{TP}}(\boldsymbol{\Phi}) = \int_{\Omega} \left(J_{\boldsymbol{\Phi}}(\boldsymbol{x}) - 1\right) \log J_{\boldsymbol{\Phi}}(\boldsymbol{x})\, d\boldsymbol{x} \tag{2.76}$$

Christensen and Johnson (2001) use a quadratic penalty of small and large values, i.e.,

$$\mathcal{P}_{\text{TP}}(\boldsymbol{\Phi}) = \int_{\Omega} J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2 + \frac{1}{J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2}\, d\boldsymbol{x} \tag{2.77}$$

To avoid folding of the FFD model, Rueckert et al. (2006) use a similar penalty term for Jacobian determinant values below a small positive threshold γ that was first proposed by Edwards et al. (1998) for a finite elements based approach, i.e.,

$$\mathcal{P}_{TP}(\boldsymbol{\Phi}) = \int_{\Omega} \mathcal{P}_{TP}(\boldsymbol{x}; \boldsymbol{\Phi})\, d\boldsymbol{x} \tag{2.78}$$

where

$$\mathcal{P}_{TP}(\boldsymbol{x}; \boldsymbol{\Phi}) = \begin{cases} \dfrac{J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2}{\gamma^2} + \dfrac{\gamma^2}{J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2} - 2 & \text{if } J_{\boldsymbol{\Phi}}(\boldsymbol{x}) < \gamma \\ 0 & \text{otherwise} \end{cases} \tag{2.79}$$

If a diffeomorphic mapping must be guaranteed, Rueckert et al. (2006) proposed to instead constrain the maximum control point coefficient using the bounds derived earlier by Choi and Lee (2000). Under the constraint $c_d^{(i)} < 0.4 s_d \; \forall i \in \{1, \ldots, n_{\text{c}}\}$, $\boldsymbol{\Phi}_{\text{FFD}}(\boldsymbol{x})$ is guaranteed to be diffeomorphic. To be able to still capture larger deformations, the total transformation is given by a composition of small diffeomorphic FFDs, where a new FFD is added, when the displacement of a control point of the last added FFD would otherwise exceed the maximum allowed control point displacement. This approach is similar to the greedy non-parametric fluid flow method of Christensen et al. (1996), where the addition of a new FFD with initially zero control point displacements can be compared to the automatic "regridding" step of Christensen et al. (1996).

The downside of the representation of the transformation by the composition of small FFDs is, that it increases the memory requirements and time needed to evaluate the composite FFD. This can be avoided when the composition of FFDs with already fixed parameters is stored as a dense displacement field, $\boldsymbol{u}(\boldsymbol{x})$, and the transformation $\boldsymbol{\Phi}_{\text{FFD}}\left(\boldsymbol{x} + \boldsymbol{u}(\boldsymbol{x})\right)$ is optimised until the constraint on the coefficient values is violated. Approximating the composition by an interpolated dense displacement field may, however, introduce some folding for very small Jacobian determinants, because of the interpolation error.

2.4. Symmetric energy formulation

In previous sections, we considered image registration as minimisation problem, where an energy function of the form $\mathcal{E}(\mathbf{\Phi}) = \mathcal{D}_\star(I_0, I_1 \circ \mathbf{\Phi}) + \alpha\mathcal{P}_\star(\mathbf{\Phi})$ is given. It can be seen that in this formulation, the image dissimilarity measure is asymmetric, because one image is considered fixed, and the other being deformed. Besides the smoothing effect caused by interpolating the moving image on the discrete lattice of the fixed image, this asymmetry causes more fundamental bias. Cachier and Rey (2000) note that such asymmetric formulation favours contracting regions over expanding structures, such as evolving lesions. This is because a shrinking region occupies a relatively larger space in the target domain than an expanding region, which is larger in the source image. Because of the integration of the point-wise image dissimilarity measures over the target image domain, regions of greater extent in the target image contribute more to the mismatch criterion than smaller regions, which have yet to expand through the deformation. Similarly, the smoothing kernel of an elastic deformation model stronger regularises the expansion of smaller regions than the contraction of larger regions (Cachier and Rey, 2000). As a result, the spatial mapping after exchanging the images will not be the inverse of the otherwise obtained mapping.

An inverse consistent registration based on the elastic deformation model was first proposed by Christensen and Johnson (2001). Two transformations are hereby estimated simultaneously, where the first transformation maps points from the target to the source image space, and the second transformation maps points in the reverse direction, i.e., $\mathbf{\Phi}_0 : \Omega_1 \to \Omega_0$ (forward) and $\mathbf{\Phi}_1 : \Omega_0 \to \Omega_1$ (backward). Given these transformations, the image dissimilarity is evaluated twice, i.e.,

$$\mathcal{D}_\star^{\mathrm{FB}}(I_0, I_1) = \mathcal{D}_\star(I_0 \circ \mathbf{\Phi}_0, I_1) + \mathcal{D}_\star(I_0, I_1 \circ \mathbf{\Phi}_1) \tag{2.80}$$

and the transformations are coupled using an inverse consistency term encouraging the two spatial mappings to be inverses of each other, i.e.,

$$\mathcal{P}_{\mathrm{IC}}(\mathbf{\Phi}_0, \mathbf{\Phi}_1) = \int_{\Omega_0} \|\mathbf{\Phi}_0(\boldsymbol{x}) - \mathbf{\Phi}_1^{-1}(\boldsymbol{x})\|^2 \, d\boldsymbol{x} + \int_{\Omega_1} \|\mathbf{\Phi}_1(\boldsymbol{x}) - \mathbf{\Phi}_0^{-1}(\boldsymbol{x})\|^2 \, d\boldsymbol{x} \tag{2.81}$$

This inverse consistency criterion requires the numerical inversion of the transformations. It can be avoided, by using the constraint proposed by Feng et al. (2009), which penalises the deviation of the composition from the identity instead, i.e.,

$$\mathcal{P}_{\text{IC}}(\boldsymbol{\Phi}_0, \boldsymbol{\Phi}_1) = \int_{\Omega_0} \|\boldsymbol{\Phi}_1 \circ \boldsymbol{\Phi}_0(\boldsymbol{x}) - \boldsymbol{x}\|^2 \, d\boldsymbol{x} + \int_{\Omega_1} \|\boldsymbol{\Phi}_0 \circ \boldsymbol{\Phi}_1(\boldsymbol{x}) - \boldsymbol{x}\|^2 \, d\boldsymbol{x} \tag{2.82}$$

An analytic derivation is possible with this formulation, when a spline-based model, e.g., a B-spline FFD, is used (Feng et al., 2009). The resulting inverse consistent method is only asymptotically symmetric, where the inverse-consistency can be violated if another term of the energy function is weighted more importantly (Cachier and Rey, 2000).

Avoiding the separate forward mapping, Cachier and Rey (2000) derived symmetric image dissimilarity and elastic regularisation terms. In case of SSD, the symmetric image dissimilarity measure given by equation (2.80) can be written with a change of variable as

$$\mathcal{D}_{\text{SSD}}^{\text{FB}}(I_0, I_1) = \frac{1}{2} \int_{\Omega_0} \left(1 + J_{\boldsymbol{\Phi}}(\boldsymbol{x})\right) \left(I_0(\boldsymbol{x}) - I_1 \circ \boldsymbol{\Phi}(\boldsymbol{x})\right)^2 \, d\boldsymbol{x} \tag{2.83}$$

for $\boldsymbol{\Phi}_0(\boldsymbol{x}) = \boldsymbol{\Phi}^{-1}(\boldsymbol{x})$, and $\boldsymbol{\Phi}_1(\boldsymbol{x}) = \boldsymbol{\Phi}(\boldsymbol{x})$. A similar derivation for the elastic regularisation term yields an inversion-invariant energy functional. Because of the necessary discretisation, the registration algorithm retains some asymmetric bias, however. Using a direct derivation of the gradient of equation (2.80) based on a first order approximation of the Taylor expansion of the composition of backward and forward mappings, which is required to be equal the identity mapping, Leow et al. (2005) obtained a gradient of the symmetrised image dissimilarity, $\mathcal{D}_{\star}^{\text{FB}}$, with respect to the backward mapping $\boldsymbol{\Phi}_1$, for energy terms depending on both backward and inverse transformation.

He and Christensen (2003) extended the inverse consistent method proposed by Christensen and Johnson (2001) to recover large deformations. Their energy function is based on a periodic sequence of images, where the first and last images correspond to the target, and the middle image corresponds to the source image. These images are deformed by a sequence of inverse consistent small elastic deformations. The composition of individual mappings yields the total forward and backward transformations of the image pair.

Noblet et al. (2008) transform both images "halfway" by requiring that the sum of the displacement fields of half forward and backward transformations is zero. This idea was generalised to a groupwise registration, where the sum of displacements from the common mean image space to each image is required to be zero. This constraint has previously been proposed by Bhatia et al. (2004) for groupwise image registration. It was also adopted by Sotiras and Paragios (2012) for a discrete symmetric registration of image pairs.

Given a velocity-based model, the mapping between any two time points $t_{1,2} \in [0, 1]$ is

$$\boldsymbol{\varphi}_{t_1}(\boldsymbol{x}, t_2) = \boldsymbol{x} + \int_{t_1}^{t_2} \boldsymbol{v}\left(\boldsymbol{\varphi}_{t_1}(\boldsymbol{x}, \tau), \tau\right) \, d\tau \tag{2.84}$$

Two variants of symmetric LDDMM were proposed by Beg and Khan (2007). The first variant uses an inverse consistent image dissimilarity term named consistent-integral-cost (CIC) that compares the images at all time points, whereas the second variant is based on a term named consistent-midpoint-cost (CMC), which evaluates the match at $t = \frac{1}{2}$, i.e.,

$$\mathcal{D}_\star^{\mathrm{CIC}}(I_0, I_1; \boldsymbol{v}) = \int_0^1 \mathcal{D}_\star\left(I_0 \circ \boldsymbol{\varphi}_0(\cdot, t), I_1 \circ \boldsymbol{\varphi}_1(\cdot, t)\right) \, dt \tag{2.85}$$

$$\mathcal{D}_\star^{\mathrm{CMC}}(I_0, I_1; \boldsymbol{v}) = \mathcal{D}_\star\left(I_0 \circ \boldsymbol{\varphi}_0(\cdot, 0.5), I_1 \circ \boldsymbol{\varphi}_1(\cdot, 0.5)\right) \tag{2.86}$$

Both symmetric variants achieve a better anatomical overlap than asymmetric LDDMM. Comparing the images at the midpoint of the flow is only slightly more expensive, whereas comparisons at many instants in time increases the computational burden noticeably. An optimisation approach similar to the fluid registration of Christensen et al. (1996) for the symmetric LDDMM with CMC using LNCC was proposed by (Avants et al., 2008). This method is the well-known greedy symmetric normalisation method of ANTs (SyN).

Hernandez et al. (2009) proposed a symmetric stationary LDDMM with two image dissimilarity terms as in equation (2.80), where $\boldsymbol{\Phi}_0 = \exp(-\boldsymbol{v})$, and $\boldsymbol{\Phi}_1 = \exp(\boldsymbol{v})$. A similar approach was used by Vercauteren et al. (2008) in a symmetric Log-Demons method. Lorenzi et al. (2013) employed a single symmetric LNCC term similar to the SSD-based CMC of Beg and Khan (2007), i.e., the image dissimilarity term is given by equation (2.86), where $\boldsymbol{\varphi}_{t_1}(\cdot, t_2) = \exp\left((t_2 - t_1)\boldsymbol{v}\right)$, and $\mathcal{D}_\star(I_0, I_1) = \mathcal{D}_{\mathrm{LNCC}}(I_0, I_1)$.

2.5. Image registration tools

Many researchers made implementations of their image registration methods publicly available. Here, we briefly summarise the most popular among these tools, which have further been included in a survey of software tools for deformable registration of medical images by Keszei et al. (2017). Please refer to this survey for a more comprehensive summary. A number of these methods were evaluate by Klein et al. (2009a) in a comparative study. Best mean performance was achieved by ANTs (SyN), ART, IRTK, and DARTEL (SPM). More recently, Ou et al. (2014) compared DRAMMS on multiple brain image datasets among others to AIR, ANTs, ART, Demons, Drop, FNIRT, and IRTK.

AIR[4]. The Automated Image Registration (AIR) tools were developed by Woods et al. (1998). The package includes tools for rigid, affine, and deformable registration. The non-rigid transformation model is based on polynomial basis functions up to fifth order.

ANTs[5]. The Advanced Normalization Tools (ANTs) are based on the Insight Segmentation and Registration Toolkit (ITK)[6]. Many of the algorithms are now part of ITK 4. The ANTs package implements the SyN approach of Avants et al. (2008), and a number of other large deformations models (Tustison and Avants, 2013), including a Demons type visco-elastic model, B-spline FFD, and time-varying and stationary velocity field models.

ART[7]. The Automatic Registration Toolbox (ART) of Ardekani et al. (2005) implements a deformable registration using an asymmetric patch similarity measure. Matching source patches are searched within a predefined radius of a target patch centre. The deformation is modelled as linear combination of the tensor product of Legendre polynomials.

DRAMMS[8]. The Deformable Registration via Attribute Matching and Mutual-Saliency weighting (DRAMMS) was developed by Ou et al. (2011). It is a discrete registration approach based on multi-scale Gabor features, which was inspired by HAMMER and Drop.

[4] https://www.nitrc.org/projects/air/ (last accessed: Jun 19, 2017)
[5] http://stnava.github.io/ANTs/ (last accessed: Jun 19, 2017)
[6] https://itk.org (last accessed: Feb 19, 2018)
[7] https://www.nitrc.org/projects/art (last accessed: Jun 19, 2017)
[8] https://www.cbica.upenn.edu/sbia/software/dramms/index.html (last accessed: Jun 19, 2017)

Drop[9]**.** The Drop registration tool is based on the Markov random field (MRF) model of Glocker et al. (2011). It implements the pioneering discrete registration based on a B-spline FFD model. For further details on discrete models in medical image registration, please refer to Glocker et al. (2011) and Sotiras et al. (2013).

Elastix[10]**.** The `elastix` toolbox developed by Klein et al. (2010) is based on ITK. Recently, an extension for SimpleITK named SimpleElastix has been released (Marstal et al., 2016). These libraries include a variety of similarity measures and transformation models, e.g., B-spline FFD, thin-plate splines, and EBS. This software implements a FFD registration using stochastic gradient descent (Klein et al., 2009b).

FAIR[11]**.** The Flexible Algorithms for Image Registration (FAIR) tools are based on the variational formulation of elastic and fluid registration of Fischer and Modersitzki (2004). It implements their fast diffusion (Fischer and Modersitzki, 2002) and curvature (Fischer and Modersitzki, 2003) registration methods. Very recently, a new add-on for stationary and non-stationary LDDMM has been published (Mang and Ruthotto, 2017).

FNIRT[12]**.** FMRIB's Non-Linear Image Registration Tool (FNIRT) is part of the FMRIB Software Library (FSL) (Jenkinson et al., 2012). It is written by Andersson et al. (2007), and uses a diffusion model. The transformation is given by a linear combination of basis functions, i.e., the discrete cosine transform or cubic B-spline kernels.

HAMMER[13]**.** The Hierarchical Attribute Matching Mechanism for Elastic Registration (HAMMER) proposed by Shen and Davatzikos (2002) is based on attribute vectors of intensities and segmentation-based geometric features. To improve the matching performance and avoid local minima, a hierarchical selection of a subset of driving features that have distinct attribute vectors is employed. The curvature model is used for regularisation. A recent publicly available implementation is based on ITK.

[9] http://www.mrf-registration.net/ (last accessed: Jun 19, 2017)
[10] http://elastix.isi.uu.nl/ (last accessed: Jun 19, 2017)
[11] http://www.siam.org/books/fa06/ (last accessed: Jun 19, 2017)
[12] https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FNIRT (last accessed: Jun 19, 2017)
[13] https://www.med.unc.edu/bric/ideagroup/free-softwares/fast-hammer (last accessed: Jun 19, 2017)

IRTK[14]. The Image Registration ToolKit (IRTK) was developed by Rueckert et al. (1999) and Schnabel et al. (2001). It implements the cubic B-spline FFD algorithm proposed by Rueckert et al. (1999). A number of mono-modal and multi-modal image dissimilarity measures are available. The optimisation is based on finite difference approximations of the energy gradient. Therefore, the runtime is relatively long compared to other, more recent implementations of the FFD algorithm, e.g., those included in NiftyReg or `elastix`. In this dissertation, we present the new Medical Image Registration ToolKit (MIRTK), which is a complete rewrite of the original IRTK, including a new modular registration framework based on analytic gradient calculations, and fast multi-threaded execution.

NiftyReg[15] . NiftyReg implements the FFD algorithm of Rueckert et al. (1999), but with a fast parallel analytic NMI gradient computation as proposed by Modat et al. (2009) instead of finite difference approximations. It also implements symmetric (Modat et al., 2012a) and SVF-based (Modat et al., 2012b) FFD methods. It provides a fast multi-threaded implementation of the algorithms, with an optional highly parallel execution on the general purpose graphics processing unit (GPGPU) (Modat et al., 2010).

SPM[16]. The Statistical Parametric Mapping (SPM) software (Ashburner, 2012) is a MATLAB® package for spatial normalisation, brain segmentation (Ashburner and Friston, 2005), voxel-based morphometry (VBM), and deformation-based morphometry (DBM). It includes algorithms for elastic registration (Ashburner and Friston, 1999), DARTEL (Ashburner, 2007), and LDDMM (Ashburner and Friston, 2011).

Symmetric Log-Domain Demons[17]. An ITK implementation of the symmetric diffeomorphic Log-Demons method proposed by Vercauteren et al. (2008) is publicly available. A variant known as LCC LogDemons (Lorenzi et al., 2013), which uses LNCC as similarity measure, is also publicly available[18].

[14] https://biomedia.doc.ic.ac.uk/software/irtk/ (last accessed: Jun 19, 2017)
[15] https://sourceforge.net/projects/niftyreg/ (last accessed: Jun 19, 2017)
[16] http://www.fil.ion.ucl.ac.uk/spm/ (last accessed: Jun 19, 2017)
[17] https://github.com/BRAINSia/LogSymmetricDemons (last accessed: Jun 19, 2017)
[18] https://github.com/Inria-Asclepios/LCC-LogDemons (last accessed: Jun 19, 2017)

Chapter 3.

Brain atlas construction

Among the first applications of a reference brain image, and a use case of image registration, is to employ it as a map of brain structures for the automatic segmentation of these structures in another brain scan (Collins et al., 1995; Gee and Bajcsy, 1999). The brain atlas utilised by Collins et al. (1995) consisted of a manually annotated brain scan of a single subject, namely Collins himself. Averaging multiple brain scans of the same individual's brain, seen in figure 3.1, has further been shown to increase SNR and markedness of tissue boundaries (Holmes et al., 1998). To that end, 27 brain scans of the same brain anatomy were affinely registered to a different brain atlas, which was previously constructed at the Montreal Neurological Institute (MNI) by averaging the T1w brain MR images of 305 adults (Evans et al., 1993). This MNI-305 template is affinely aligned with the Talairach & Tournoux atlas, a traditional stereotaxic space, which has been defined based on the post-mortem histology of a single human brain (Talairach and Tournoux, 1988). The MNI-305 atlas is one of the first adult brain atlases created from *in vivo* MRI scans of an entire population, as opposed to a single subject. Such population average not only provides a stereotaxic space similar to the Talairach & Tournoux atlas for spatial normalisation of brain images, but also represents the large biological variability of brain anatomies. Expressing this uncertainty of the prior information about the spatial distribution of tissues is important for applications such as brain tissue segmentation. The use of a population average reduces the bias towards a specific anatomy in the automatic brain image analysis, which is otherwise introduced when a single subject brain is used

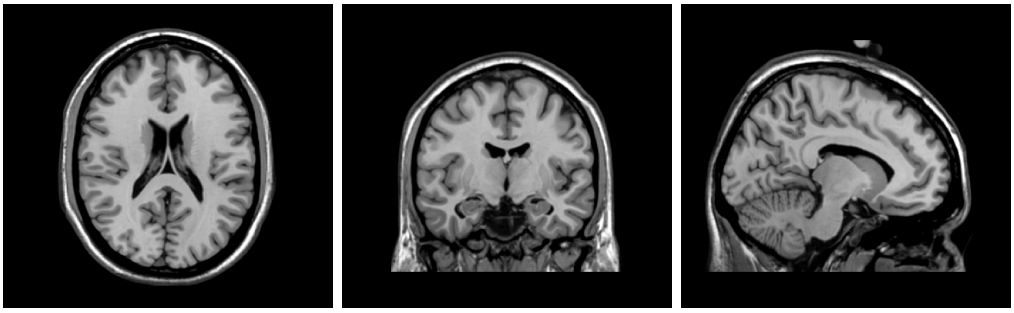
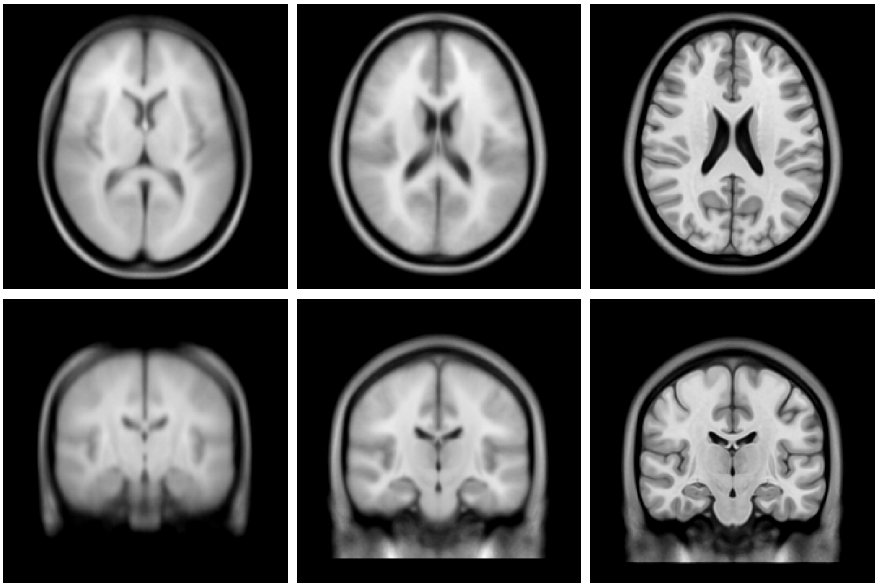


Figure 3.1.: Orthogonal mid-sections of Collins-27 single-subject brain template.

as template. An improved version of the MNI-305 template was created as part of the International Consortium for Brain Mapping (ICBM) project from 152 T1w brain scans of a subset of the 305 subjects that were acquired with higher resolution and better contrast than those used to build the MNI-305 template (Mazziotta et al., 2001). However, because of the residual misalignment of brain structures after affine registration, owing to the anatomical variations of human brains, average brain models such as the MNI-305 and affine ICBM-152 templates retain few cortical features. This degrades the accuracy of the registration of the template to a given brain image, limits its application in whole brain as opposed to tissue segmentation, and makes it unsuitable for the study of geometric differences of brain anatomies based on shape deformations in the field of CA briefly outlined in section 1.3. Adult brain templates with markedly greater detail were created from the same 152 images using deformable registration by Grabner et al. (2006) and more recently Fonov et al. (2011). The different adult brain templates created at the MNI are publicly available[1]. The higher detail of the ICBM-152 population average created by Fonov et al. (2011) using non-rigid registration, as opposed to affine transformations only, is apparent in the images shown in figure 3.2. The ICBM-152 template is thus well suited for spatial normalisation of adult brain images. A symmetric version with shape and appearance of both hemispheres made equal is also available for the study of inter-hemispheric differences (Fonov et al., 2011).

But not only the sharpness of the template brain images used in an analysis is crucial. The interest in average brain models reflecting the mean shape of a whole population of anatomies is driven by the fact that registration between images of similar anatomies

[1] http://nist.mni.mcgill.ca/?page_id=714 (last accessed: May 20, 2017)



(a) MNI-305 (b) affine ICBM-152 (c) ICBM-152

Figure 3.2.: Axial and coronal mid-sections of different adult population brain atlases.

is more accurate. Moreover, DBM generally assumes diffeomorphic mappings between each observed anatomy and the template brain, at least at a macroscopic rather than cellular scale. This assumption is not fulfilled when registering an immature brain to an adult template. The spatial normalisation of a child brain scan using an adult brain template, created from a population imaged at a significantly older age than the brain under investigation, has been shown to be negatively influenced by the differences in tissue distribution (Wilke et al., 2002). These inaccuracies in the spatial normalisation might introduce a strong bias in quantitative measures. This is even more significant when an adult brain template is used as reference for the co-registration of a fetal, neonatal, or infant brain scan (Shi et al., 2011b). At this early period of life, the brain undergoes rapid shape and appearance changes, such as changes in cortical folding, involution of the germinal layer, premyelination changes within white matter, myelination, iron deposition, and the growth of different regions of the brain (Rutherford, 2002). These large anatomical differences between developing and mature brains demand the construction of brain atlases for different developmental phases. Spatio-temporal models of early brain development during gestation and shortly after birth are needed to capture these dynamic structural brain changes. A 4D model not only increases the accuracy of spatial normalisation, and

atlas-based segmentation, but also facilitates the study of normal brain development, and gives valuable insights into structural brain maturation.

Different techniques for constructing average brain models from a set of brain images have been proposed over the past few decades. The most influential of these are summarised in the following section, followed by a discussion of more recent approaches to create spatio-temporal atlases of the developing brain imaged during the intra-uterine and neonatal period. Brain atlases built for the study of fetal (gestation), neonatal (0-1 month old), infant (1-12 months old), and child (1-21 years old) brains are presented thereafter.

3.1. Average brain models

The first population atlases of the adult brain were models of the average intensity at a reference point corresponding to the expected response of brain tissues to the acquisition parameters observed at this location (Guimond et al., 2000). For this, brain images are aligned to a reference brain using affine or deformable registration. The reference may be a single-subject (e.g., Collins' brain) or previous population atlas such as MNI-305 (Evans et al., 1993; Holmes et al., 1998; Mazziotta et al., 2001; Grabner et al., 2006). This approach is related to the study of intensity variations within anatomical regions termed VBM (Ashburner and Friston, 2000). The use of a single reference results in a brain model, whose topology is similar to the reference brain (Rueckert et al., 2003). Brain studies using this atlas are thus biased towards this specific anatomy.

The emerging discipline of CA introduced the idea of average geometry instead of average intensity models (Grenander and Miller, 1998). In CA, each anatomy is represented as a diffeomorphism, which maps a coordinate of the deformable atlas to its corresponding anatomical point in the observation (Thompson and Toga, 2002). Anatomical differences are studied based on these shape deformations, as opposed to the intensity variations, and is known as DBM. LDDMM (Trouvé, 1998; Dupuis et al., 1998; Beg et al., 2005) equips the field of CA with a distance metric between shapes, which is based on the deformations that capture their geometric differences. A template brain anatomy used in CA is defined as the brain model, which is on average geometrically closest to all observations.

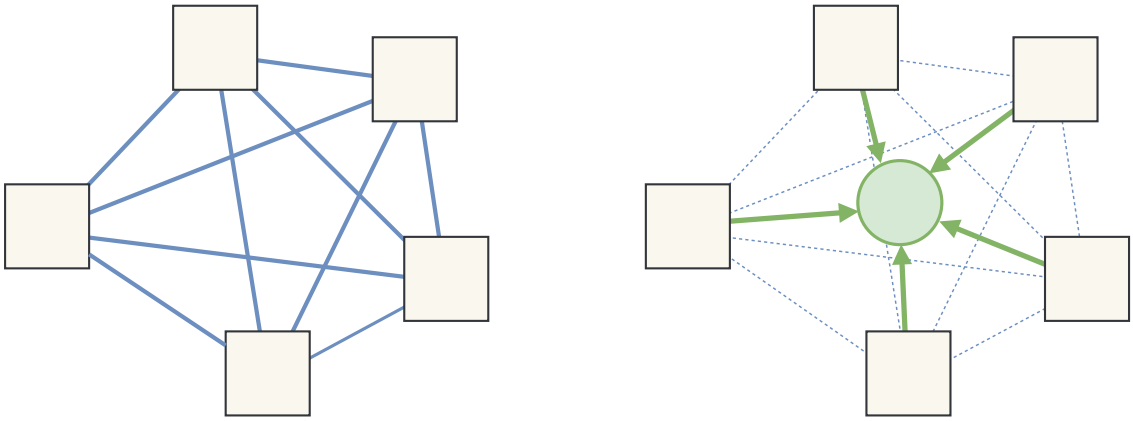


Figure 3.3.: Illustration of pairwise atlas construction. First, the transformations between all pairs of images (squares) are computed (blue lines). A mean transformation (green arrows) is then derived for each image which maps it to the average space (circle).

3.1.1. Using pairwise registration

Woods et al. (1998) were among the first to create a brain model of both mean shape and intensity from 22 adult brain MR images using affine registration between all possible pairs of images. For two images, the affine transformation corresponding to the square root of the respective homogeneous coordinate transformation matrix, and its inverse, respectively, maps each image midway to their common average space. This reconciliation process is repeated considering all transformations between a given image and the others. The result is a mean affine transformation, which maps this image to the natural coordinate system, that is common to all images. The obtained brain atlas has the average orientation, size, and shape of all imaged brains. Moreover, by averaging the images after mapping each to this average space by its respective mean transformation, the resulting average shape and intensity model is not biased towards any of the brain images from which the atlas is created. A similar approach using deformable instead of affine registration was proposed by Seghers et al. (2004). An approach based on direct registration of all (unique) pairs of images is herein referred to as pairwise atlas construction, as illustrated in figure 3.3.

While the pairwise approach is completely unbiased, it suffers from a few drawbacks. The major drawback is its quadratic computational complexity, which especially in case of deformable registration does not scale well. Another difficulty is how to derive a mean

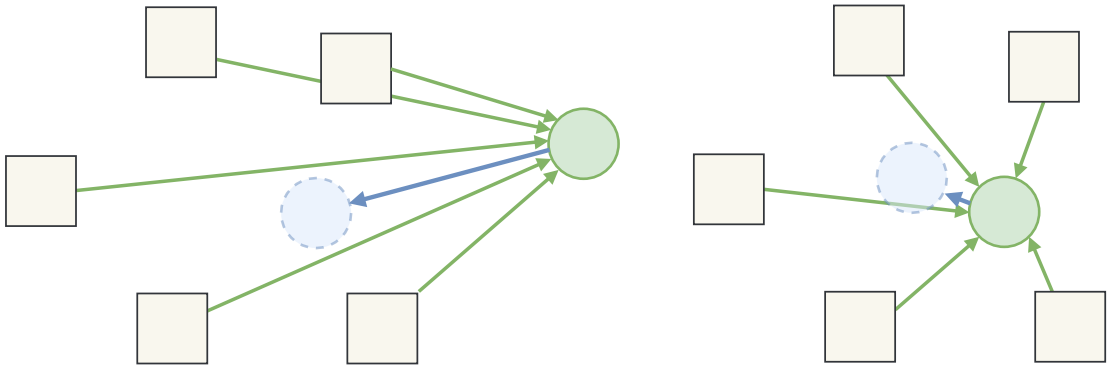


Figure 3.4.: Illustration of two steps of the groupwise atlas construction. At each iteration, the individual anatomies (squares) are registered to the reference anatomy (green circle). The average intensity image is subsequently deformed by the average of the deformations (blue arrow) to form a new mean brain shape and intensity model (dashed blue circle) which is the reference for the next iteration (right). This iteratively minimises the distance of the reference to each anatomy until the residual deformation approaches zero.

transformation for each image, which indeed maps the individual anatomy to the common anatomical coordinates of all images, i.e., the mean must be consistent across all images. In case of deformable atlas construction, the estimation of an average deformation from many individual deformations moreover might smooth high spatial frequency components in the deformation fields, and thus result in a sub-optimal alignment of anatomical details, such as the cortical folding patterns of human brains, in the common reference space. An improved anatomical alignment of the images in the average space can be achieved by iterating the pairwise construction steps. This reduces residual anatomical misalignment of the images, which is caused by registration errors and possibly inconsistent mean transformations. The high computational complexity limits the practicality of such iterative refinement.

3.1.2. Iterative template construction

Instead of selecting each image in turn as target for the pairwise registrations, and to derive a separate mean transformation for each image, Guimond et al. (1998) proposed to choose a single image as reference. The authors further built an average shape and intensity brain model using the Demons algorithm (Thirion, 1998), i.e., an elastic image registration method instead of affine registration. The bias towards the selected reference brain shape is subsequently reduced by applying the mean deformation to the intensity

average. This process is repeated, as illustrated in figure 3.4, where the average brain image of the previous iteration is the new reference image for the deformable registration at the next iteration. The convergence of the template for different choices of initial reference has been evaluated by Guimond et al. (2000).

For CA, where a single brain template with greater cortical detail is a better target for the image registration than an average intensity image, Kochunov et al. (2001) developed a method where a so-called best individual target (BIT) is selected among all brain scans. The selection of the BIT brain is based on a measure of deformation needed to align two images in the data set. This scalar quantity is evaluated for each pair of images. It has therefore the same computational complexity as the pairwise atlas construction. To reduce this computational cost, Kochunov et al. (2001) proposed to approximate the scalar deformation measure for an image pair A and C using the individual measures for the deformation of A to B and B to C, respectively. The average displacements of the deformations to the BIT brain are subsequently used to deform this image, and to thereby obtain the minimal deformation target (MDT). The MDT brain is the template brain used for further analysis of the brain images. Although both of the techniques proposed by Guimond et al. (1998) and Kochunov et al. (2001) reduce the bias towards the initial reference anatomy, these brain models will always depend on the choice of reference.

3.1.3. Using groupwise registration

Methods developed for the simultaneous registration of multiple images, referred to as groupwise as opposed to pairwise image registration, can be used to spatially normalise brain images, without any bias towards a reference (Studholme, 2003; Studholme and Cardenas, 2004; Bhatia et al., 2004, 2007; Zöllei, 2006; Noblet et al., 2008; Wu et al., 2012). These techniques use an image similarity term, which quantifies the joint similarity of multiple aligned images instead of only one pair of images, e.g., based on a high dimensional joint histogram (Studholme, 2003), the comparison to the current average (Bhatia et al., 2007), or the sum of all possible pairwise similarities (Wu et al., 2012). Moreover, an either soft (Studholme, 2003) or hard (Bhatia et al., 2004) constraint requires that the sum of displacement vectors to the different images at each spatial point is zero. The

anatomical reference coordinate system maximising the joint similarity of the images, and at the same time fulfilling this constraint, thereby provides a natural space for spatial normalisation and unbiased average brain model construction. The main difference between groupwise registration and the brain atlas construction of Guimond et al. (1998) is, that instead of comparing every image to the same reference image, a single similarity is evaluated jointly for the whole image set without preference for any of the images. The iterative maximisation of this similarity, on the other hand, is equivalent to the iterative atlas construction, except for the smaller template update steps. Scaling the mean deformation field by a small step size can be used, however, also in the approach of Guimond et al. (1998) to better control each update of the template shape. Bhatia et al. (2007) later extended their groupwise FFD algorithm to a registration interleaved with tissue segmentation, where the Kullback-Leibler divergence (KLD) of the tissue distributions is used to bring the images into alignment. The method was developed with the aim to create an unbiased probabilistic atlas of brain tissue distributions.

3.1.4. Using diffeomorphic registration

A shortcoming of the methods discussed so far is, that these are based on the small deformation assumption, where deformations are close to the identity, and the addition of two displacement vector fields is a valid approximation of the composite deformation. Given the large geometric variability of human brains, this assumption does not hold, which in turn prohibits the parametrisation of brains by a flat Euclidean space, where simple averaging can be applied to generate a mean representation (Joshi et al., 2004). Joshi et al. (2004) therefore proposed an unbiased large deformation diffeomorphic template estimation, where transformations are elements of the infinite dimensional group of diffeomorphisms (Miller et al., 2002) parametrised by velocity vector fields. Because simple averaging of diffeomorphisms does not generally yield an element of this group, the requirement that deformations sum to zero is substituted by a metric of the distance of the diffeomorphic mappings from the identity, whereby this distance metric is based on velocities as opposed to directly on the corresponding deformations. The initial greedy minimisation of the joint energy functional of image to average maps and average intensity

model was later replaced by a full, more costly space-time optimisation (Lorenzen et al., 2005), following the work of Beg et al. (2005). A parallel GPGPU implementation of this atlas construction, with speed ups of 20-60 times compared to a single CPU computation, was developed more recently by Ha et al. (2009). Because the brain template estimation as formalised by Joshi et al. (2004) is based on the SSD image similarity, it is unsuitable for constructing an average brain model from multiple modalities, or images with wide appearance changes as observed during early brain development. An extension for multi-modal atlas formation based on the KLD is presented in Lorenzen et al. (2006).

In a related work by Avants and Gee (2004), an alternating brain template optimisation similar to the approach of Guimond et al. (1998) is presented, that is based on a pairwise large deformation registration method and mean diffeomorphic deformations obtained by averaging velocity fields. In their initial work, the authors also choose a single brain image as reference, and iteratively deform it by the sequence of mean deformations to obtain an average brain shape template. They extended this work in Avants et al. (2010) to an unbiased approach referred to as symmetric groupwise normalisation (SyGN) using their pairwise SyN registration method (Avants et al., 2008) with LNCC as similarity metric. Besides the update of the template geometry by deforming it with the mean deformation, the appearance of the model is optimised at each iteration by minimising the CC of the average intensity image with respect to each co-registered brain image. This is done instead of a simple averaging of intensities to model local tissue inhomogeneities and slowly varying distortions in tissue appearance.

Another approach to counter the fuzziness of population averages obtained by simple averaging of intensities was proposed by Zhang et al. (2016b). For this, a wavelet decomposition of the co-registered images is computed. A sparse dictionary of representative image patches is then learned for each frequency sub-band. The dictionaries are learned such that these minimise the reconstruction error with respect to the corresponding frequency sub-band of the decomposed mean intensity image. Composing the individually reconstructed frequency sub-bands results in a sharper brain template than the mean intensity image. This patch-based method was extended to also include a temporal consistency term for the construction of a longitudinal atlas in Zhang et al. (2016a).

A template construction, which simultaneously estimates for each image the bias of intensity inhomogeneities, a diffeomorphic mapping relating it to the atlas space, and the intensity variance across the images, was recently proposed by Pai et al. (2017). Instead of relying on an image similarity measure that is robust to intensity variations, this method explicitly incorporates statistical models of these variations. In this approach, the template deformations are parametrised by SVFs at multiple scales using the kernel bundle framework (Sommer et al., 2012).

3.1.5. Statistical template deformations

The study of biological variability based on the statistical analysis of the deformations relating each observed anatomy to a template anatomy has also inspired brain models, which not only depict average morphology, but also the modes of variation encoded by the deformations (Joshi, 1998; Gee and Bajcsy, 1999). A SDM based on the FFD model was proposed by Rueckert et al. (2003). It captures the statistical distribution of anatomical template coordinates in a population. Whereas average intensity or probabilistic atlases depict this distribution only partially by means of a blurred average intensity image and tissue probability maps, a SDM preserves this information of shape differences encoded by the template deformations, and parametrises these by their principle modes. The PCA of the FFD coefficients results in a compact representation, and focuses the attention on anatomical differences that are most prevalent across all observed anatomies of the studied population. The resulting generative brain model can be used to create different biologically plausible examples of the reference anatomy. Applications of SDMs include the regularisation of the spatial normalisation of brain images (Pszczolkowski et al., 2012), and the creation of intermediate target images that are closer to the brain image to be normalised (Kim et al., 2008; Tang et al., 2009).

Most techniques discussed in this section were developed for the construction of an average brain model of the mature brain for CA, or the creation of a probabilistic atlas for use in image segmentation. In early life, the brain undergoes rapid changes (Rutherford, 2002) that cannot be captured by a 3D model. For this reason, spatio-temporal models of the developing human brain, which capture these dynamic changes, have been developed.

3.2. Spatio-temporal models

Purely shape-driven generative models have been proposed by Niethammer et al. (2011) and Singh et al. (2013) for the study of time series data and spatio-temporal atlas building based on an extension of LDDMM for geodesic regression. These models parametrise a time series of generally adult brain images by an initial image and initial momentum of the longitudinal mean shape deformation. Due to the significant age-related intensity variations, representing an atlas time series of the early developing brain by a single deformable template without a model of the growth-related intensity changes is not suitable, though.

Other spatio-temporal models that explicitly make use of the longitudinal information of time series data, e.g., brain scans of the same anatomy shortly after birth and at one and two years of age, were proposed by Durrleman et al. (2009) and Liao et al. (2012). Also in the work by Liao et al. (2012), the brain degeneration of an ageing population has been modelled. Longitudinal brain scans of normal brain growth of the same developing brain are not usually available, because imaging an infant at this early age for research-only purposes is hard to justify if there are no concerns regarding pathologies. This is more common for babies who were born very premature in order to assess the extra-uterine development of their brain at term-equivalent age and after.

A comprehensive review of atlases of the developing brain, with focus on spatio-temporal models of the fetal and neonatal brain, was recently conducted by Benkarim et al. (2017). The first spatio-temporal model of the fetal brain was constructed by Habas et al. (2010a) for atlas-based tissue segmentation. To cope with the morphological differences of fetal brains at different ages, and the low tissue contrast of structural MR images, the deformable spatial alignment of the brain images was driven by the segmented tissue boundaries found to be common across the age range, using the template-free approach of Studholme (2003). Temporal polynomial models of voxel displacement, intensity, and tissue probabilities are for each voxel in the group-average space individually subsequently fit to the displacements relating it to each anatomy, the co-registered image intensities, and the tissue segmentations, respectively. Given these voxel-wise parametric models, a brain image and tissue priors can be regressed for every age within the captured age range.

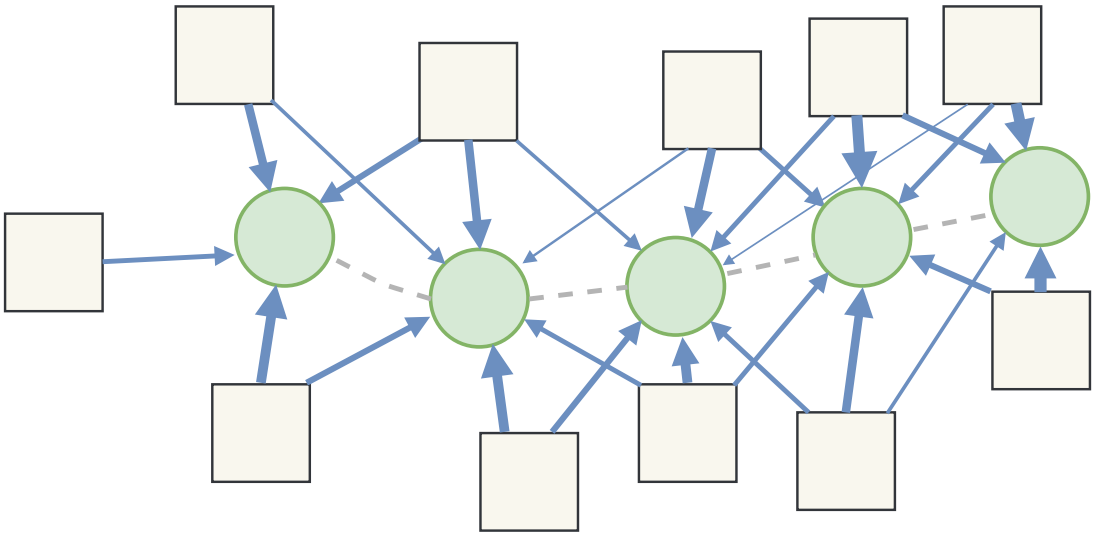


Figure 3.5.: Illustration of non-parametric spatio-temporal atlas construction from random design data. The influence of each brain image (squares) on the spatio-temporal average (circles) is weighted by a temporal kernel (blue arrows; thickness visualises weight) according to the distance in time (horizontal axis).

A non-parametric temporal regression for longitudinal atlas construction from brain scans of different individuals was proposed by Davis et al. (2007). This approach extends the groupwise diffeomorphic atlas construction of Joshi et al. (2004) to a spatio-temporal model using Nadaraya-Watson kernel regression (KR). The temporal regression kernels hereby weight the deformation-based distance of each anatomical sample relative to the population average, such that at any given time, this average is closer to those anatomical samples that are also closer in time. Brain anatomies from an earlier or later developmental phase, respectively, have little to no influence on the regressed mean brain shape at a specific age. This is illustrated schematically in figure 3.5. It makes the approach well suited for spatio-temporal atlas construction at a very early time of life, such as the intra-uterine and neonatal period, where anatomies that are only few weeks apart have a very distinct shape and appearance.

Kuklisova-Murgasova et al. (2011) constructed a spatio-temporal atlas of the preterm-born neonatal brain using a similar non-parametric kernel regression. The atlas was constructed by affine registration of the brain images to a template brain. The bias towards this reference anatomy is then removed by composing each affine subject-to-reference transformation with the regressed age-dependent Log-Euclidean mean transformation (Arsigny et al., 2007). This mean transformation adjusts the global shape and size of the reference to reflect the average shape and brain size of the population. The resulting probabilistic atlas is applied in neonatal brain tissue segmentation (Cardoso et al., 2011a; Wang et al., 2012). Because of the affine-only alignment, cortical details are not retained.

A spatio-temporal neonatal brain atlas with markedly greater cortical definition was created shortly after by Serag et al. (2012a). Their approach extends the pairwise atlas construction of Seghers et al. (2004) from a 3D to 4D model using kernel regression in time. Using pairwise deformable registration, the temporal kernel regression not only shifts the mean anatomy towards age-related anatomical samples, but further reduces the influence of registration errors. These errors are more significant when registering distinctively dissimilar images of brains at very different ages (Serag et al., 2012b). Serag et al. (2012a) also proposed the use of an adaptive kernel width to balance the often uneven distribution of available neonatal brain scans across the age range, where most babies are imaged around term. Doing so, an equal number of images contributes to each average, and a similar level of detail is achieved for all time points. This adaptive kernel regression (AKR) was adopted by Gholipour et al. (2014, 2017) to create a 4D atlas of the fetal brain using the diffeomorphic SyGN approach (Avants et al., 2010) instead of the FFD algorithm. These spatio-temporal atlas construction methods based on non-parametric kernel regression are revisited in more detail in the context of our related methods in chapter 6.

3.3. Volumetric brain atlases

In the following, we summarise published work in which a model of the developing brain has been constructed. A summary of brain atlases of the fetal and neonatal brain is given in tables 3.1 and 3.2, respectively. When available, a link to the online resource, where the atlas can be obtained, is provided as footnote next to the column entry of the citation.

3.3.1. Fetal brain atlases

	mod.	no	age	temp.	spatial normalisation
Habas et al. (2010a)[2]	T2w	20	21–25	LS	groupwise, elastic (Studholme and Cardenas, 2004)
Dittrich et al. (2014)	T2w	32	20–30	KR	FFD, elastix toolbox (Klein et al., 2010)
Serag et al. (2012c)[3]	T1w T2w	80	23–37	AKR	FFD, IRTK (Rueckert et al., 1999)
Gholipour et al. (2017)[4]	T2w	81	21–37	AKR	SyGN, ANTs (Avants et al., 2010)

Table 3.1.: Spatio-temporal atlases of the fetal human brain. Columns from left to right: Citation, imaging modalities (mod.), number of subjects (no), gestational age in weeks (age), method used for temporal regression (temp.), and spatial normalisation approach.

Following the development of faster and high quality imaging sequences as well as improved motion correction techniques (Studholme, 2011), atlases of the normal fetal brain became available more recently (Benkarim et al., 2017). The rapid development of the brain during gestation hereby requires a model, which captures these dynamic changes. The first generative spatio-temporal model was created by Habas et al. (2010a) from only a limited set of 20 manually annotated fetal brain images for GAs ranging from 21 to 25 weeks. Dittrich et al. (2014), on the other hand, focused on the construction of a dynamic probabilistic tissue atlas from incomplete annotations. Their method relies on a previous spatial normalisation of the fetal brain images to age-specific templates corresponding to the weighted intensity average after global normalisation, where weights are given by a temporal regression kernel. Serag et al. (2012c) created a spatio-temporal atlas from a set

[2] http://depts.washington.edu/bicg/research/fba.php (last accessed: May 20, 2017)
[3] http://brain-development.org/brain-atlases/ (last accessed: May 20, 2017)
[4] http://crl.med.harvard.edu/research/fetal_brain_atlas/ (last accessed: May 20, 2017)

of 80 imaged healthy fetal brains for a wider GA range from 23 to 37 weeks, using the pairwise atlas construction approach with adaptive kernel regression. A spatio-temporal atlas from a similarly sized, but different set of 81 fetal brain images, also using adaptive kernel regression within the groupwise diffeomorphic approach, was built by Gholipour et al. (2017) for the GA range from 21 to 37 weeks.

3.3.2. Neonatal brain atlases

	mod.	no	age	temp.	spatial normalisation
Kazemi et al. (2007)[5]	T1w	7	39–42	–	non-rigid, SPM2 (Ashburner and Friston, 1999)
Noorizadeh et al. (2013)[5]	T1w	14	39–42	–	non-rigid, SPM8 (Ashburner and Friston, 2005)
Ghadimi et al. (2017)[5]	T1w CT	7	39–42	–	non-rigid, SPM8 (Ashburner and Friston, 2005)
Oishi et al. (2011)[6]	T1w T2w DTI	14 20	38-41	–	dual-channel LDDMM of FA and MD images (Ceritoglu et al., 2009)
Alexander et al. (2016)[7]	T1w T2w	10	40–43	–	SyGN, ANTs (Avants et al., 2010)
Gousias et al. (2012)[3]	T2w	15	26–35	–	–
Gousias et al. (2012)[3]	T2w	5	39–45	–	–
Shi et al. (2011b)[8]	T1w T2w	95	39–46	–	groupwise, HAMMER (Wu et al., 2012)
Kuklisova-Murgasova et al. (2011)[3]	T2w	142	29–44	KR	affine (12 DOFs), IRTK
Serag et al. (2012c)[3]	T1w T2w	204	28–44	AKR	pairwise, FFD, IRTK (Rueckert et al., 1999)
Makropoulos et al. (2016)[3]	T2w	420	28–44	AKR	FFD, IRTK (Rueckert et al., 1999)

Table 3.2.: Atlases of the neonatal human brain. Columns from left to right: Citation, imaging modalities (mod.), number of subjects (no), post-menstrual age in weeks (age), method used for temporal regression (temp.), and spatial normalisation approach.

A first population average for spatial normalisation of neonatal brain scans was created by Kazemi et al. (2007) using the SPM toolbox (Ashburner and Friston, 1999, 2005). This template brain was constructed from seven neonatal brains scanned at 39 to 42 weeks post-

[5] http://www.u-picardie.fr/labo/GRAMFC (last accessed: May 20, 2017)
[6] http://www.mri.kennedykrieger.org/softwaredata.html (last accessed: May 20, 2017)
[7] https://github.com/DevelopmentalImagingMCRI/M-CRIB_atlas (last accessed: May 20, 2017)
[8] https://www.med.unc.edu/bric/ideagroup/free-softwares/unc-infant-0-1-2-atlases (last accessed: May 20, 2017)

menstrual age (PMA). Validation of spatial normalisation with SPM using this template instead of the adult MNI template on another seven neonatal brain scans showed, that a neonatal template results in better co-registration. The same research group also provided a symmetric template built from all 14 term-born neonates for asymmetry studies of language development (Noorizadeh et al., 2013). More recently, a bimodal neonatal MR-CT brain template was added to the collection (Ghadimi et al., 2017).

A set of 14 T1w, 20 T2w, and 20 DTI scans of neonatal brains acquired within 38 to 41 weeks PMA were used by Oishi et al. (2011) to create different neonatal brain templates after affine and non-rigid alignment. In contrast to previous approaches, the non-rigid registration was done using a dual-channel LDDMM method (Ceritoglu et al., 2009) to match the MD and FA maps derived from DTI scans.

A parcellation of the cortex of 10 healthy neonates born at term, and scanned between 40 and 43 weeks PMA, which is consistent with the cortical and subcortical adult brain atlas used in FreeSurfer (Desikan et al., 2006), was done by Alexander et al. (2016). Structural T1w and T2w brain templates for spatial normalisation of a neonatal brain scan to this cortical parcellation atlas were created using SyGN.

Kuklisova-Murgasova et al. (2011) created an affine spatio-temporal atlas consisting of T2w average images of 142 preterm-born neonates and corresponding tissue probability maps for 29 to 44 weeks PMA. Using the non-rigid pairwise approach, Serag et al. (2012c) built a spatio-temporal model of average T1w and T2w intensities from 204 preterm-born neonates, also with accompanying tissue probability maps. Makropoulos et al. (2016) furthermore non-rigidly registered automatic whole brain segmentations of 87 brain structures of 420 mostly preterm-born neonates, of which 40 were born at term, to the age-matched T2w templates created by Serag et al. (2012c), in order to built a fine-granular probabilistic atlas. These automatic segmentations are based on 20 single-subject atlases of 15 preterm-born neonates scanned at term-equivalent age, and 5 term-born neonates. Each image was manually segmented by Gousias et al. (2012) to create an atlas of 50 anatomical regions, referred to as "a label-based encephalic ROI template" (ALBERT). These were subdivided into cortical and subcortical parts by Makropoulos et al. (2016).

3.3.3. Infant and child brain atlases

Wilke et al. (2002, 2003) created the first brain template for spatial normalisation and segmentation of paediatric brain images from 148 T1w brain scans of children at the age of 5 to 19 years. The researchers at the Cincinnati Children's Hospital Medical Center (CCHMC) noticed considerable differences between the distribution of tissues, when compared to the adult MNI template used by SPM. This first CCHMC template was constructed using affine registration. A toolbox for the construction of paediatric brain templates, using the non-rigid unified segmentation (Ashburner and Friston, 2005) approach of SPM without the adult tissue priors, was later released under the name "Template-O-Matic" (Wilke et al., 2008). It has previously been used in Altaye et al. (2008) to create a brain template from 76 T1w images of 9 to 15 month old infants.

Gousias et al. (2008) automatically segmented brain images of 33 preterm-born 2-year-old children into 83 anatomical regions using multiple annotated adult brain images within a multi-atlas segmentation approach. The atlases were used for automatic multi-atlas-based whole brain segmentation of new paediatric brain images (Gousias et al., 2010).

Shi et al. (2011b) constructed an infant brain atlas from 95 time series data of brains acquired after birth (39–46 weeks PMA), at the age of one year (88-109 weeks), and at two years (131-163 weeks). A brain template and tissue maps were created for each group individually. After affine alignment of all images within an age group, a groupwise extension of HAMMER (Shen and Davatzikos, 2002; Wu et al., 2012) was employed to obtain an unbiased average space, where registration was based on a sparse set of features consisting of image intensity, edge type, and geometric moment invariants of WM, GM, and CSF maps (Wu et al., 2012). Along with brain templates and tissue probability maps, Shi et al. (2011b) derived a longitudinal mean deformation from the individual growth trajectories obtained by longitudinal registration of the time series of each subject.

Numerous paediatric child brain atlases for various ages were created at the MNI by Fonov et al. (2011) from a collection of brain scans obtained from 433 children aged between 4.5 and 18.5 years. The atlas construction follows the iterative approach of Guimond et al. (1998, 2000) with the adult ICBM-152 template as reference.

3.4. Developing cortical surface atlases

The brain models discussed in the previous sections are implicit models of brain shape, by modelling the spatial distribution of the average MR intensity and brain tissues. For the study of cortical folding patterns, few studies have constructed a cortical surface atlas to enable on one hand the comparison of mean measures of curvature calculated directly from the explicit geometric models of individual cortical surfaces, and on the other hand provide a common parcellation of the cortical sheet into inter-related functional units.

A geometric 4D model of the fetal cortex was constructed by Wright et al. (2015) from 80 healthy fetuses imaged between 22 and 39 weeks of gestation. These scans were previously used by Serag et al. (2012c) to built a volumetric atlas. A 4D cortex model is obtained using the same adaptive kernel regression approach. Point correspondences across all cortical surfaces were hereby established using a groupwise extension of a diffeomorphic spectral surface matching technique (Lombaert et al., 2013b; Wright et al., 2015).

Hill et al. (2010) created a surface atlas of the neonatal cortex for the study of hemispheric asymmetries in cortical folding at term. Correspondences between the cortical surfaces of 12 neonates were established with a spherical registration based on manually placed sulcal and gyral landmarks, following a procedure previously used to create a surface atlas of the adult cortex (Van Essen, 2005).

Li et al. (2014) built a spatio-temporal spherical surface atlas of cortical folding measures from surfaces extracted from 202 brain MR images of 35 healthy subjects. Spherical maps of cortical folding were constructed from this longitudinal data set for 7 time points between 1 and 24 months after birth. To exploit the longitudinal information and ensure temporal consistency, correspondences between individual cortices were established in multiple steps. First, intra-subject correspondences for all time points were computed using groupwise Spherical Demons (Yeo et al., 2010). The resulting mean cortical folding maps, one for each subject, are then registered with a second groupwise spherical registration to establish unbiased and age-independent inter-subject correspondences. Finally, the spherical maps corresponding to a specific time point are registered to this unbiased population average to create the respective atlas for this age.

Chapter 4.

Segmentation of the developing brain

In the following, we first discuss challenges of methods trying to automatically identify the cortex in structural MR images of the developing brain. These challenges of neonatal brain image analysis impede the direct application of methods developed for the adult brain. We then draw our attention towards brain segmentation methods in the literature, with focus on methods devised for segmenting the neonatal brain, while highlighting important related work on adult brain segmentation preceding these approaches. A number of authors have participated in the MICCAI Grand Challenge on Neonatal Brain Segmentation 2012 (NeoBrainS12), with results of the on-site challenge summarised in Išgum et al. (2015). Makropoulos et al. (2017a) also reviewed techniques specific to the developing brain.

Discussed approaches are broadly categorised into atlas-based and atlas-free methods. Atlas-free methods are divided into unsupervised methods using pattern recognition techniques or algorithms from mathematical morphology, methods applying supervised classification given a training set, and deformable models. The latter are considered in section 4.5. Note that (non-parametric) supervised classification may be considered atlas-based when the training set, consisting of pre-segmented brain images, is needed to segment an unseen image. Otherwise, when there is a clear separation between training of a (parametric) model and application to a new test image, the approach is considered to be atlas-free. On the other hand, atlas-based segmentation approaches often utilise a supervised classification technique after atlas alignment. We consider any method inferring prior knowledge of the anatomy by co-registering a pre-segmented atlas brain image to be atlas-based.

Most of the reviewed segmentation methods classify a voxel as to containing one (or a mixture) of the different brain tissues. Some atlas-based methods further divide the brain into multiple brain structures, and other regions of interest. This is generally not possible with an atlas-free method, with the exception of supervised classification using spatial context, because not all boundaries between segments need to be associated with a visible image feature such as an intensity gradient. Given our primary interest in segmenting and modelling the cerebral cortex, only a tissue classification is required.

4.1. Challenges of neonatal brain image segmentation

The problem of segmenting a brain MR image to identify elements in the image that contain a certain type of tissue, or correspond to a particular brain structure, is of fundamental interest in biomedical image analysis. Many methods have been developed for the tissue classification and structural whole brain segmentation for adult brain images. However, imaging the brain in early life, and even during gestation, has only become possible more recently. While the quality of the acquired 3D reconstructions of a neonatal brain scan keep improving, there are still a number of challenges to overcome. We summarise here the main challenges in neonatal brain image analysis with respect to the segmentation problem as noted before by Prastawa et al. (2005).

- Although MRI protocols for infant brain scans have a shorter scan time, motion artefacts still occur because the infant cannot be prevented to move. Motion correction of the acquired stack of 2D slices during the 3D volume reconstruction is thus needed. This reduces motion artefacts, but cannot eliminate these.

- WM myelination is not complete before the age of two years. Infant brains contain unmyelinated white matter (uWM) and myelinated white matter (mWM), where the boundaries of fully myelinated WM are often ambiguous (Rutherford, 2002).

- Due to different MR signal properties of brain tissues at the neonatal period, when compared to the mature brain, the (predominantly unmyelinated) WM is brighter than GM in T2w MR images and CSF is brighter than both WM and GM. This

leads to partial volume averaging at the pial interface of the cortex between cGM and CSF, with average intensity similar to WM. This is often the cause for voxel-wise misclassification of these partial volumes (Xue et al., 2007).

- The contrast between WM and GM is lower than for adult brains. This is due in part to the smaller size of the newborn's head, requiring it to be scanned at higher resolution, leading to higher noise levels. A shorter scan time to avoid sedation or restraining the infant is another cause for the lower contrast-to-noise ratio (CNR). This low CNR makes segmenting partial volume regions even more challenging.

4.2. Atlas-based brain image segmentation

Among the most successful techniques for brain segmentation that have emerged over the past decade are atlas-based methods, with methods using multiple atlases or a probabilistic atlas being considered state-of-the-art in brain image segmentation. These methods exploit prior knowledge of common brain architecture inferred from one or more brain atlases. The template brain image of an atlas is therefore registered with the subject brain image to be segmented. In case of a probabilistic atlas, which inherently encodes uncertainty of the exact spatial location of structures, affine registration may suffice, but generally at least a coarse deformable registration is required. Given the established correspondences between a subject and template brain image, the atlas labels (or prior probabilities, respectively) are mapped to the coordinates of the to-be-segmented image. This process is referred to as label propagation. The required accuracy of atlas alignment hereby depends on the subsequent refinement (or fusion) stage of each respective approach.

4.2.1. Brain image segmentation using a single atlas

Original atlas-based methods for adult brain segmentation have utilised a single manual segmentation of an individual brain (Christensen et al., 1994b; Collins et al., 1995). This approach is illustrated using neonatal brain images in figure 4.1. It has two major limitations, which hinder a robust automatic segmentation. Firstly, the atlas depicts a single anatomical instance only, although there exist substantial differences among human brains.

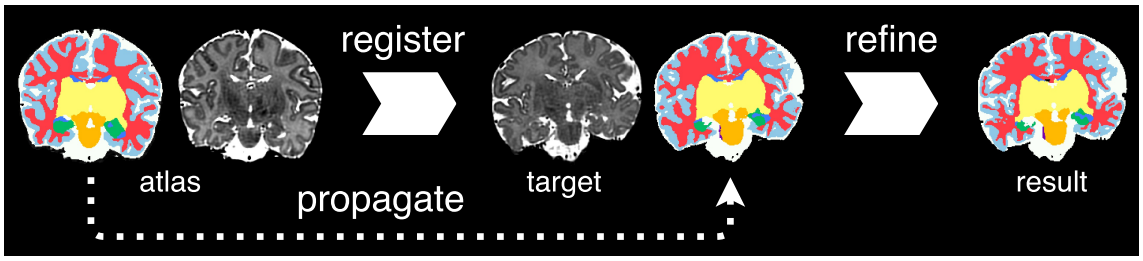


Figure 4.1.: Illustration of single-atlas-based segmentation.

This variability cannot be captured by a single atlas. Secondly, it relies on a robust and accurate propagation of the reference labels. This is in itself not an easy task due to the ill-posedness of the registration problem, which requires regularisation, and may not be able to identify corresponding anatomical points accurately. This is further complicated by aforementioned anatomical differences, such as in the folding of the cerebral cortex, especially during early brain development. For these reasons, single-atlas methods find nowadays rare application in brain image segmentation, but were instrumental in paving the way for the following atlas-based techniques.

A notable variant of the outlined single-atlas segmentation approach, originally applied to segment adult brain images, and later used by Weisenfeld et al. (2006a,b) to segment multiple brain images of preterm born neonates scanned at term-equivalent age for use in their automatic multi-atlas approach, is the semi-automatic method proposed by Warfield et al. (2000). In this semi-automatic method, the aligned atlas is used to locally adapt an intensity-based supervised classifier, namely k-Nearest Neighbour (kNN), using the Euclidean distance to the aligned reference structures. This single-atlas approach is less susceptible to aforementioned problems, because the classification is primarily based on the intensities of manually selected examples of the image to be segmented.

4.2.2. Brain image segmentation using multiple atlases

To overcome the limitations of single-atlas methods, Rohlfing et al. (2004) investigated different strategies for atlas-based segmentation using multiple atlases. It has further been shown by Heckemann et al. (2006) that fusing the information from multiple brain atlases, as illustrated in figure 4.2, achieves results superior to those based on a single atlas.

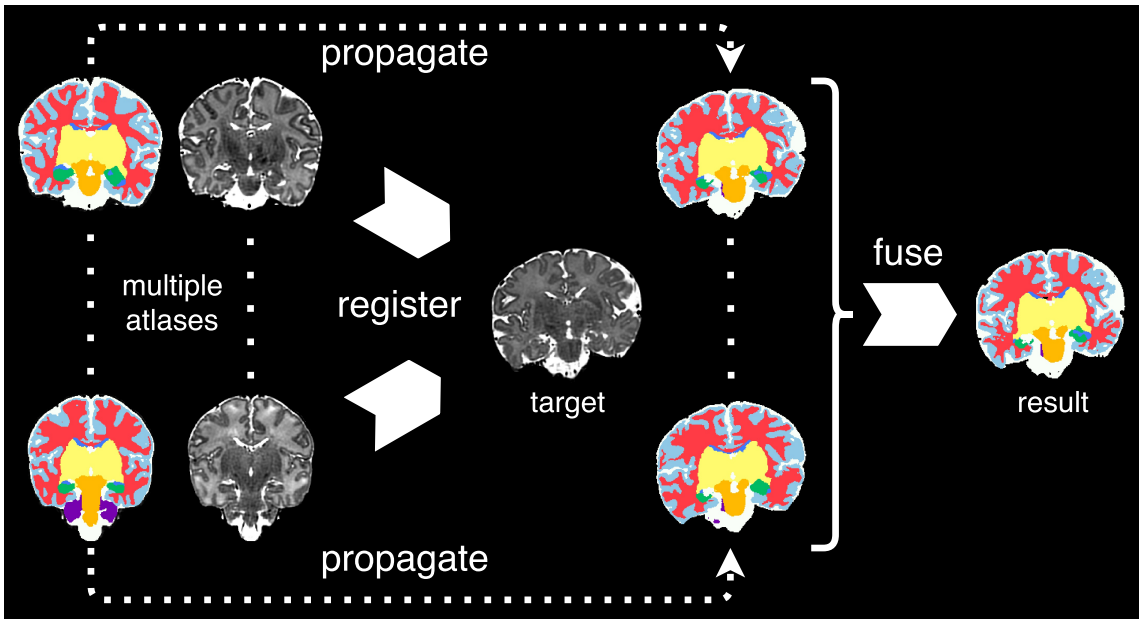


Figure 4.2.: Illustration of multi-atlas-based segmentation.

Atlas selection or non-uniform weighting, which gives more influence to atlases that better resemble the anatomy of the brain to be segmented and discards atlases with considerably misaligned anatomy, either due to registration errors or significant anatomical differences, improves the results (Aljabar et al., 2009). In these original multi-atlas segmentation approaches, the different deformed atlas labels were commonly combined to a consensus label using a voxel-wise majority vote. This label fusion approach counts the number of times each label is assigned to a specific voxel, and selects the one with highest frequency.

The performance of a multi-atlas segmentation depends on the quality of the atlases, the co-registration with the subject anatomy, and the technique used to combine the propagated reference segmentations, which further can mitigate some of the unavoidable registration errors. Various label fusion approaches besides majority voting have therefore been proposed in the literature (Iglesias and Sabuncu, 2015). The patch-based segmentation proposed by Coupé et al. (2011), for example, uses a non-local patch search without atlas deformation, where the spatial distance of patches is not considered by the weighting function. Instead, a patch similarity measure is used to weight the contribution of each patch label. Multiple matching patches from the same atlas can hereby contribute to the consensus label, reducing the number of required atlases. For more details and a broader

overview of current multi-atlas brain segmentation methods, including methods for the immature brain, and associated label fusion strategies, see Iglesias and Sabuncu (2015).

The neonatal brain segmentation method of Weisenfeld and Warfield (2009) adopts a unique label fusion approach based on the simultaneous truth and performance level estimation (STAPLE) algorithm (Warfield et al., 2004). For this, manually selected labelled prototypes of training samples in the atlas images, used to segment these images with the semi-automatic method of Warfield et al. (2000), are transformed to the subject image. For each atlas and tissue class, a non-parametric intensity distribution is estimated given the intensities at the corresponding transformed prototype locations. Multiple modalities can hereby be integrated, such as T1w, T2w, and proton density (PD) images. For each set of propagated atlas prototypes, a candidate segmentation is produced, given the respective non-parametric estimates. These are then fused using the STAPLE performance level of each segmentation that is estimated using the expectation maximisation (EM) algorithm. Training samples that are in clear disagreement with this fused segmentation are excluded, and the whole process repeated until convergence. Despite its widespread successful application, the soundness and usefulness of STAPLE results has recently been called into question by Van Leemput and Sabuncu (2014).

Anbeek et al. (2008, 2013) used a kNN classifier based on multi-modal intensity features and spatial information to segment a neonatal brain image into four (WM, cGM, dGM, and CSF) (Anbeek et al., 2008) or eight (uWM, mWM, cGM, dGM, CSF, ventricles, brainstem, and cerebellum) (Anbeek et al., 2013) classes, respectively. Co-registration of atlas and subject images is hereby needed because of the use of absolute voxel coordinates as spatial features. The kNN-based tissue probabilities, followed by an assignment of the label corresponding to the class with maximum probability, can therefore be viewed as a type of label fusion. Although this approach utilises a training set consisting of random samples drawn from each manually labelled atlas, only a single registration of the subject image to a pre-computed unbiased average template brain image is needed. A similar multi-stage kNN-based tissue segmentation approach using more extracted features, and a feature selection step, was proposed by Chita et al. (2013). This approach does not require co-registration of atlases because of spatial features that are relative to the brain

extraction mask. Only affine spatial normalisation to a reference brain image is needed such that all images have the same orientation. We thus consider it as atlas-free method.

Instead of kNN, Han (2013) performed multi-atlas label fusion and then employed a standard random forest (RF) (Breiman, 2001) for each structure to refine the delineation of structures in CT images. Specifically, they used a voxel-wise convex combination of the STAPLE and RF probabilities. A method for neonatal and infant brain tissue classification using a pre-trained RF without prior label fusion was proposed by Wang et al. (2015). This method only requires a reference brain image for affine spatial normalisation, if the brain images do not already have the same orientation. We therefore discuss it in section 4.4 together with other atlas-free supervised tissue classification methods.

A practical drawback of many multi-atlas segmentation approaches is the need for several registrations to segment a single brain image. Based on the number of atlases and the computational complexity of the registration algorithm, this can be a bottleneck in a brain image analysis pipeline. Although a pre-registration of the individual atlases to a common reference space is possible, such co-registration can introduce additional errors in the subject alignment of propagated atlases. A more significant drawback is the need for multiple annotated reference brain images to ensure a sufficient coverage of possible anatomical configurations (and pathologies), which can be labour-intense to obtain, and hence limit the number and diversity of the available atlases. Especially for segmenting images of the early developing brain, multi-atlas methods are not always feasible, due to the limited availability of imaging data and corresponding manual reference segmentations.

4.2.3. Brain image segmentation using a probabilistic atlas

Given the large anatomical differences during the period of early brain development, many atlas-based methods for segmenting images of the immature brain utilise a statistical model of image intensities in conjunction with a probabilistic atlas, which represents *a priori* knowledge about the spatial distribution of image labels, i.e., the different brain tissues, individual brain structures, or other ROIs. This approach is illustrated in figure 4.3, where tissue probability maps of the age appropriate neonatal atlas constructed by Kuklisova-Murgasova et al. (2011) are deformed to the target subject to infer this prior knowledge.

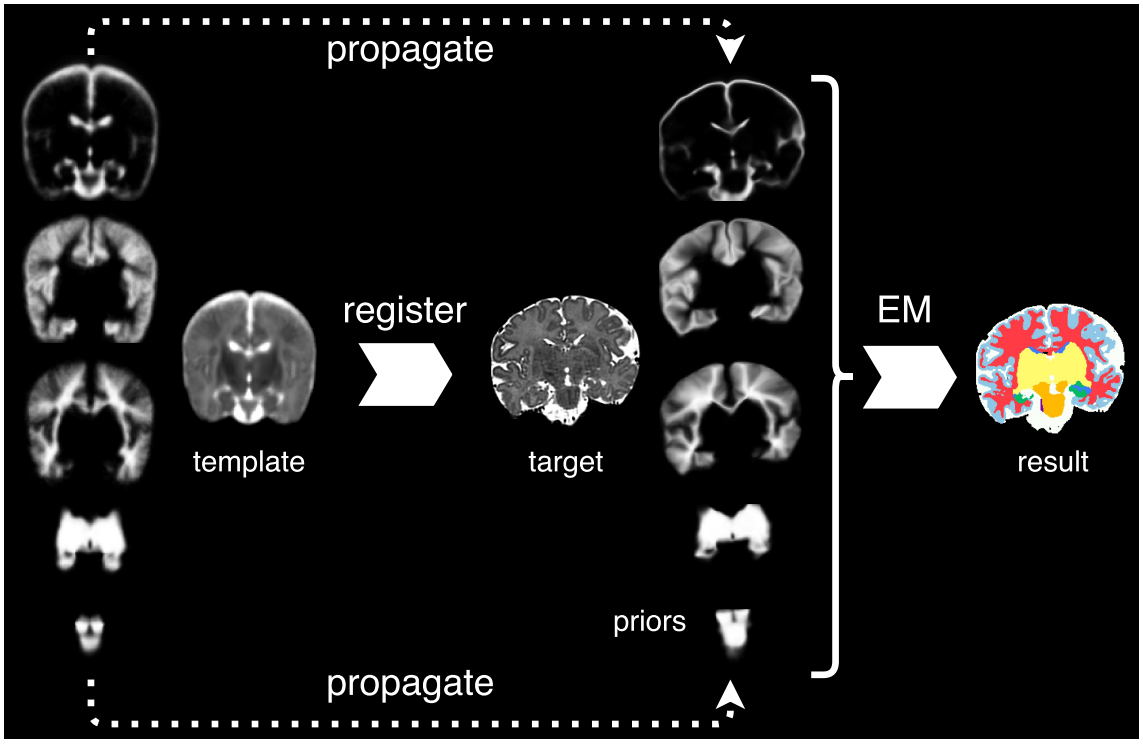


Figure 4.3.: Illustration of probabilistic atlas-based segmentation.

Kamber et al. (1995) were first to introduce a probabilistic model of spatially varying *a priori* probabilities for the segmentation of brain lesions. Such model can be obtained from a cohort of brain images that are representative for the studied population using one of the atlas construction techniques presented in chapter 3. For this, a sufficient number of segmented brain images is required, in order to capture enough anatomical variability within the studied population. The more images are used to construct the probabilistic atlas, the less influence have individual segmentation errors, such that a sufficiently accurate automatic segmentation approach can be used to segment the brain images from which the atlas is built.

An early approach combining a probabilistic atlas of brain structures with a voxel-wise classification was proposed by Collins et al. (1999), with the aim to improve the segmentation of the cortex. A supervised classification based on an artificial neural network (ANN) (Zijdenbos et al., 1996) is hereby combined with the original single-atlas propagation method of Collins et al. (1995), named ANIMAL[1]. Their ANN was previously

[1] http://www.bic.mni.mcgill.ca/ServicesSoftwareAdvancedImageProcessingTools/ANIMAL

employed for the automatic quantification of multiple sclerosis (MS) lesion volume (Zijdenbos et al., 1998) within the so-called intensity-normalized stereotaxic environment for classification of tissues (INSECT) framework. The voxel-wise ANN classification is refined using the atlas-based segmentation with spatially varying tissue priors inferred from the probabilistic ICBM-152 atlas (Mazziotta et al., 1995).

During the early period of life, a single probabilistic brain atlas cannot capture the full anatomical variability, and at the same time provide sufficient localisation of individual structures. Given a sufficient number of brain images from a wide age range, this problem is largely overcome by the use of a spatio-temporal probabilistic atlas, which provides age-specific prior information for each developmental phase. Moreover, techniques for a relaxation of the prior label probabilities (Shiee et al., 2011; Cardoso et al., 2011a,b, 2013) iteratively adapt the probabilistic atlas to the anatomy of the imaged subject. The use of a pre-constructed probabilistic atlas eliminates the need for multiple atlas-to-subject registrations, but comes with its own downsides, which may or may not impact the final segmentation, depending on the quality of the atlas, and how robust the probabilistic method is to atlas misalignment. A subject-independent probabilistic atlas construction requires an average brain template image to be computed, which should not be biased towards any of the individual atlas anatomies. The registration of such average brain template image with another brain image is generally less accurate than a separate registration of individual brain atlases, due to the reduced clarity of boundaries in an average brain template. And finally, a locally weighted atlas fusion, based on the similarity of each atlas within a local neighbourhood with the subject image to be segmented, can provide better subject-specific label probabilities (e.g. Ledig et al., 2012a). An adaptation of the atlas is also achieved by other techniques, such as the aforementioned relaxation of prior probabilities, or a combination of the inferred atlas probabilities with label probabilities based on image intensities, e.g., using an unsupervised k-means clustering (Xue et al., 2007; Makropoulos et al., 2014), or patch-based probabilities (Liu et al., 2014, 2016).

(last accessed: May 20, 2017)

Expectation maximisation

Probabilistic atlas-based methods combine the spatially varying prior probabilities of the atlas either with a statistical model of class conditional intensity distributions in a Bayesian framework, or use a supervised classifier trained on sets of intensities for each class, which are based on the inferred probabilities. The atlas priors further act as spatial constraints, to deal with overlapping intensity distributions. Because these methods classify a voxel largely based on its intensity, spatial intensity inhomogeneities observed in MR images have to be taken into account. This bias can be reduced in a pre-processing step using N4 (Tustison et al., 2010). The bias field estimation would benefit, however, from a tissue classification. Therefore, instead of two separate sequential stages, Wells et al. (1996) proposed a method where the tissue classification is interleaved with an estimation of a smoothly varying bias field. The EM (Dempster et al., 1977) algorithm was modified to estimate the maximum likelihood label assignment and the parameters of the bias field. The method of Wells et al. (1996) required pre-estimated class conditional intensity models based on the intensities of manually selected voxels. Van Leemput et al. (1999a) eliminated this manual step by using a probabilistic brain atlas. After spatial normalisation, i.e., co-registration of atlas and subject images, the tissue probabilities of the atlas are used to initialise the statistical intensity models, where commonly a Gaussian Mixture Model (GMM) is employed. This EM scheme was extended to also re-estimate the parameters of the GMM at each iteration. The tissue probabilities, referred to as anatomical (tissue) priors, furthermore spatially constrain the classification. A voxel with a low anatomical tissue prior is less likely assigned to this tissue class than a voxel with high anatomical tissue prior. Additional spatial smoothness of the segmentation is achieved by incorporating contextual information by means of a Markov random field (MRF), to favour assigning a voxel to the same class as the majority of its neighbours (Van Leemput et al., 1999b). This EM scheme was used by Habas et al. (2010b) to segment *in vivo* images of the fetal brain into skull, CSF, GM, WM, germinal matrix, and ventricles. They created a spatio-temporal atlas from manual delineations of the tissues in 14 images (Habas et al., 2010a), in order to cope with rapid structural and appearance changes during brain development.

Atlas-based methods are prone to produce erroneous results when the imaged anatomy is significantly different from the average anatomy of the population used to construct the atlas. Moreover, a population atlas contains high uncertainty about the presence of a certain structure in regions of high variability among subjects from which the atlas is constructed. In case of brain images, this is the case along the cortex, due to the large differences in cortical folding between individuals. Negative atlas bias can also result from an atlas representing the healthy population, while the subject image depicts an anatomy modified by pathology. Too strong tissue priors, which quickly drop to zero at tissue boundaries, can further bias the segmentation, and thereby overshadow real anatomical differences. A solution to this problem was presented by Shiee et al. (2011) for the segmentation of adult brain images with enlarged lateral ventricles (ventriculomegaly). It consists of adding a tissue prior relaxation to the iterative EM scheme, where the tissue priors are not simply modelled to be equal the anatomical priors inferred from the healthy population atlas, but as random variables drawn from a Dirichlet distribution, whose parameters are derived from this atlas. These random variables are not strictly speaking *a priori* knowledge, as they depend on the image data, but their role within the EM framework is that of prior tissue class probabilities. With each EM iteration, the anatomical priors are more and more adapted to reflect the actual subject anatomy. A similar prior relaxation was proposed about the same time by Cardoso et al. (2011a), where the prior tissue probabilities were originally updated only after the EM converged. In their following work on a segmentation method for infant brain images with structural abnormalities, such as ventriculomegaly, their anatomical prior relaxation is also done after each EM iteration (Cardoso et al., 2011b, 2013). This relaxation scheme is implemented in NiftySeg[2]. Melbourne et al. (2012) extended the approach of Cardoso et al. (2013) to also perform outlier rejection of intensity clusters, which have a large Mahalanobis distance from the estimated model. They participated in the NeoBrainS12 challenge.

A subject-specific atlas with naturally fewer anatomical differences can be constructed from longitudinal images taken from the same subject. Preterm born infants are commonly scanned at term-equivalent age, and sometimes again after one or two years for a follow

[2] https://sourceforge.net/projects/niftyseg/ (last accessed: May 20, 2017)

up assessment of brain development. At the age of about two years, the myelination of WM is nearly complete, and the inversion of WM and GM intensities reverted, with better tissue contrast in T1w images compared to the T2w image acquired at term-equivalent age (Rutherford, 2002). The method proposed by Shi et al. (2010) explicitly exploits the availability of such follow up scan to segment the T2w image of the neonatal brain. They construct a subject-specific probabilistic atlas by segmenting the T1w follow up image using the adaptive fuzzy C-means (AFCM) algorithm (Pham and Prince, 1999). Similar to the EM scheme of Wells et al. (1996), this adaptive version of the fuzzy C-means (FCM) clustering (Dunn, 1974) includes a simultaneous bias field estimation, to improve its robustness to MR intensity inhomogeneities. Faced with the problem of having to align the T1w follow up image with the T2w image at the neonatal period, both exhibiting very distinct MR appearances of the infant brain, Shi et al. (2010) adopted a joint segmentation-registration framework (Pohl et al., 2006) to align the images based on tissue distributions, as opposed to MR intensities, using HAMMER (Shen and Davatzikos, 2002).

One of the mentioned drawbacks of a probabilistic atlas is the uncertainty of tissue probabilities near the cortex. Shi et al. (2011a) therefore devised an infant brain segmentation method, where they construct a subject-specific probabilistic atlas from a set of pre-segmented images, weighting the individual atlases based on their cortical similarity with the to-be-segmented image. For this, a confidence map of quasi-probabilities for cGM is derived from the subject image using a Hessian filter. The eigenvalues of the Hessian matrix at a point of the T1w image are used to identify either line-, plane-, or blob-like structures, an approach commonly used for the detection of vessels (Frangi et al., 1998). The cGM confidence reflects how much the local image structure resembles a line or plane, respectively, given that the cortex has a highly folded sheet-like shape. The mutual information of a non-rigidly aligned atlas image with the T2w image is then evaluated for regions having a high confidence value of being part of the infant's cortex. A probabilistic atlas is constructed using these similarity weights, and the resulting anatomical priors are combined with the cGM confidence measure. The unified framework of Ashburner and Friston (2005), where tissue classification, bias field, and atlas registration are simultaneously estimated using an iterated conditional modes (ICM) (Besag, 1986) approach, is

then employed to segment the infant brain image with HAMMER (Shen and Davatzikos, 2002) as registration algorithm, and the tissue priors of the constructed atlas.

The unified framework of Ashburner and Friston (2005) was originally developed for adult brain image segmentation, and is implemented in the well-established SPM[3] software. This brain segmentation method is in particular used for VBM studies of the adult brain. SPM was applied by Wang et al. (2012) in conjunction with a customised probabilistic neonatal brain atlas (Kuklisova-Murgasova et al., 2011) in the NeoBrainS12 challenge, followed by connected component analysis to identify mislabelled partial volume. Another neonatal brain tissue segmentation based on SPM was developed by Beare et al. (2016). It uses mathematical morphology to adapt the atlas constructed by Kuklisova-Murgasova et al. (2011) to the anatomy of the neonatal brain to be segmented. We discuss this method in section 4.3, after briefly introducing the concepts behind morphological segmentation.

A processing pipeline for the segmentation of neonatal brain images using existing general purpose tools was presented by Wu and Avants (2012) as part of NeoBrainS12. The used tools are part of ANTs, namely N4 (Tustison et al., 2010) for bias correction, SyN (Avants et al., 2008) for atlas alignment, and Atropos (Avants et al., 2011) for tissue classification using the EM-MRF framework. In the neonatal processing pipeline, the MRF posterior probabilities were composed of the likelihood probability for the T1w and T2w voxel intensities, prior probability as smoothness constraint, and anatomical tissue priors.

A different approach combining the intensity-based posterior probabilities of Atropos with joint label fusion (Wang et al., 2013) was evaluated on the NeoBrainS12 data by Sanroma et al. (2016). It presents a generic ensemble learning technique for combining the probabilistic segmentations of multiple base methods, i.e., multi-atlas label fusion and EM-MRF posterior probabilities in the tested neonatal segmentation method. The posteriors of the joint label fusion were also used as anatomical priors in Atropos.

One of the difficulties in voxel-wise tissue classification of neonatal brain images is partial volume averaging. The mixing of tissues results in intensities at the cGM/CSF interface that are similar to WM. While the first-order neighbourhood MRF priors successfully remove isolated voxels that are labelled different from their neighbours, it is not sufficient

[3] http://www.fil.ion.ucl.ac.uk/spm/ (last accessed: May 20, 2017)

to remove mislabelled partial volume in neonatal brain images. This is because of the number of voxels on the CSF/GM boundary that can be misclassified. If two misclassified voxels are neighbouring each other, one can prohibit the other from being corrected by the MRF term. Different methods to deal with and correct these mislabelled partial volume voxels in neonatal brain image segmentations have been proposed. First to address this problem was Xue et al. (2007), using a knowledge-based driven approach within an unsupervised method for neonatal CSF, GM, and WM classification. This approach was adopted by Makropoulos et al. (2012, 2014) in an automatic atlas-based method for both tissue classification (Makropoulos et al., 2012) and structural whole brain segmentation (Makropoulos et al., 2014). Hereby, the MRF-based prior probabilities are modified at each EM iteration. If a voxel is currently classified as WM, and within its neighbourhood there are also CSF and cGM voxels, the MRF prior probability of this voxel for WM is reduced, and the difference in probability mass evenly redistributed among the CSF and cGM classes. Makropoulos et al. (2014) derived similar rules for other constellations, which commonly result in mislabelled partial volume in neonatal brain images. Another common approach to deal with partial volume is to introduce mixed tissue classes. After a first tissue classification without partial volume classes, to get good initial estimates of the Gaussian distribution parameters of each tissue, Cardoso et al. (2011a,b, 2013) perform a second EM classification with mixed partial volume classes. Ledig et al. (2012b), on the other hand, extended the MRF term to also take into account second-order neighbourhood information, i.e., the co-occurrence of three different tissue types within a local neighbourhood. Their MRF prior probability of a voxel to belong to the WM class is low, when there are voxels with high probability for cGM as well as CSF in its neighbourhood.

Non-parametric methods

A non-parametric atlas-based segmentation using a kNN classifier was proposed by Cocosco et al. (2003). The training set for each class is hereby selected by thresholding the co-registered anatomical priors. A pruning technique based on a minimum spanning tree (MST) is used for automatic rejection of misclassified training samples resulting from anatomical differences and atlas misalignment. This clustering technique is further used

by Prastawa et al. (2005) to separate unmyelinated from myelinated WM given T1w and T2w intensities of the combined WM tissue class of their atlas. WM samples with high image gradient are hereby excluded to avoid training samples close to tissue boundaries, and within the transition zone from uWM to mWM. Robust initial mean and covariance estimates for the Gaussian distributions of CSF, GM, uWM, and mWM are obtained using minimum covariance determinants (MCD) (Rousseeuw and Van Driessen, 1999). For a better separation of the overlapping intensity distributions, Prastawa et al. (2005) employ a second classification stage with non-parametric kernel density estimates, with outlier rejection based on the MST-pruning approach of Cocosco et al. (2003). This neonatal brain image segmentation method is available as open source software[4].

Another non-parametric method for neonatal brain images was proposed by Srhoj-Egekher et al. (2013) for NeoBrainS12. They construct a probabilistic atlas by averaging manual segmentations of multiple atlases after alignment with the subject image. In a first stage, a binary kNN classifier is used for each tissue to classify a voxel within a tissue-specific training mask derived from the atlas as either belonging to this class or not. At this stage, only voxels with high confidence and prior probability are labelled. After merging these tissue segmentation masks, the remaining unlabelled voxels, which occur primarily near tissue boundaries, are classified in a second stage using a naive Bayes classifier. For classification, T1w and T2w intensities, and the respective intensities after convolution of the images with Gaussian kernels and their derivatives at two different scales, are used as features. A principle component analysis is employed to obtain less correlated features.

Whole brain segmentation

So far discussed methods for neonatal brain image segmentation classify each voxel as belonging to one of the tissue classes, often including classes for brainstem, cerebellum, and skull. To our knowledge, Gousias et al. (2013) were first to segment the whole neonatal brain image into sub-cortical structures and various cerebral ROIs. The approach followed their work on whole brain segmentation of paediatric brain images (Gousias et al., 2010), and was made possible by the effort to manually delineate 18 sub-cortical structures and

[4] https://www.nitrc.org/projects/neoseg (last accessed: May 20, 2017)

32 cerebral ROIs in brain images of 5 healthy term-born and 15 preterm-born infants called ALBERTs, imaged between 36.6 and 44.9 weeks PMA (Gousias et al., 2012).

Makropoulos et al. (2014) improved upon this automatic multi-atlas neonatal brain image segmentation into 50 anatomical ROIs in several ways. Specifically, they adopted the probabilistic atlas-based EM scheme with simultaneous bias field estimation (Van Leemput et al., 1999b; Ledig et al., 2012a), divided the 50 ROIs of the ALBERTs into 87 sub-cortical, cortical, and cerebral ROIs, and modelled the intensity statistics of each ROI with a Gaussian. To obtain better initial estimates of the GMM mixing coefficients, the prior probabilities resulting from fusing the individually propagated ALBERTs are combined with subject-specific prior probabilities obtained through k-means clustering (Xue et al., 2007). These are further multiplied with MRF prior probabilities from first-order neighbour interactions as spatial smoothness constraints, including a knowledge-based partial volume correction (Xue et al., 2007). To deal with the fact that different cortical and cerebral ROIs have similar intensity distributions, because these are either made up of cGM or WM, a hierarchical mixture model is used, where all cortical/cerebral ROIs share the same Gaussian distribution parameters (Van Leemput et al., 2009). Due to strong overlap of intensity distributions of dGM regions, the likelihood of these structures is modelled by an ensemble of two models, label fusion and GMM. A weighting scheme favours the fused atlas priors over the GMM in homogeneous image regions, to avoid expansion of dGM regions into neighbouring areas. An extensive validation of the method was performed by Makropoulos et al. (2014) using T2w MR images of 234 mainly preterm infants, to verify its robustness at different ages of the developing neonatal brain. The method allows regional brain growth in neonates to be assessed in an automatic and reproducible way (Makropoulos et al., 2016), and is implemented in the developing region annotation with expectation maximisation (Draw-EM) package[5] of MIRTK. It has further been used by Wright et al. (2015) to reconstruct the WM/GM boundary of the cortex of fetuses at various gestational ages, for the construction of a spatio-temporal surface atlas.

[5] https://github.com/MIRTK/DrawEM (last accessed: May 20, 2017)

4.3. Unsupervised brain tissue segmentation

An unsupervised approach for classifying brain tissues in T1w images of the adult brain is implemented in the BrainSuite software[6]. It implements an entire image processing pipeline, including brain extraction, non-uniformity correction, tissue classification, graph-based topology correction of WM segmentation, and cortical surface reconstruction (Shattuck et al., 2001; Shattuck and Leahy, 2001, 2002). Although in general unsupervised, with empirical default parameters, it often requires manual parameter tuning to produce acceptable results, and is hence not fully unsupervised. Moreover, it is specific to adult T1w images and cannot be used to process T2w neonatal brain images.

A purely data-driven classification of the tissues in a neonatal brain image, without *a priori* information about the anatomy, is susceptible to noise and the challenges discussed in section 4.1. Consequently, only few methods segmenting an image of the neonatal brain without supervision have been proposed. It should be noted that semi-automatic methods are inherently supervised by means of the human input, with a possible exception for manual post-segmentation editing steps. An example is the LIGASE[7] method developed and used by Hill et al. (2010), where the user has to estimate the mean and width of the GM, WM, and CSF intensity distribution from a histogram, with additional manual editing of the reconstructed cortical surface in Caret[8]. As we are interested in fully automatic approaches, we excluded such semi-automatic methods from the discussion.

To compensate for the lack of neonatal brain atlases, Xue et al. (2007) proposed an unsupervised approach based on a EM-MRF approach similar to the respective atlas-based methods discussed in the previous section. Instead of anatomical atlas priors, they used a k-means clustering of the intensities to initialise the mixing coefficients of the GMM. A novelty of the work was the knowledge-based approach used to correct mixed CSF/GM partial volume that is often incorrectly labelled in neonatal images as WM.

An extension of the FCM clustering (Dunn, 1974), the AFCM algorithm (Pham and Prince, 1999), which is robust to spatial intensity inhomogeneities, found application in

[6] http://brainsuite.org/ (last accessed: May 20, 2017)
[7] http://brainvis.wustl.edu/LIGASE/ (last accessed: May 20, 2017)
[8] http://brainvis.wustl.edu/wiki/index.php/Caret:About (last accessed: May 20, 2017)

the method of Shi et al. (2010). As discussed in the previous section, Shi et al. (2010) used AFCM to segment the T1w image of the infant's brain at one or two years of age without supervision, in order to derive a subject-specific probabilistic atlas for the segmentation of the earlier neonatal brain scan.

Another interesting class of recent methods for neonatal brain image segmentation are those applying algorithms from mathematical morphology. Knowledge about common (neonatal) brain anatomy is hereby encoded via a combination of morphological filtering and seed-based segmentation. A type of morphological filtering is reconstruction by dilation, where the filter operates on the connected components of an image, with connectivity between neighbouring voxels based on intensity difference, in order to remove some image features without changing others. A method for morphological segmentation is the morphological watershed transformation. It is a region growing starting at the seed points (markers) based on image features such as the image intensities, or the output of an edge-detection or other image-enhancement filter.

Mangin et al. (1995) developed a method using homotopic dilation, homotopic erosion, and a homotopically deformable region (HDR) algorithm. The HDR algorithm deforms a binary image of an object via front propagation, where only simple points (Bertrand, 1994) are added, such that the initial topology of the object is preserved. It is used to identify the WM/GM interface as the minimum of a global energy, i.e., to segment the image into two regions, the interior and the exterior of WM, including sub-cortical structures. The energy function consists of a unary potential based on estimates of mean and standard deviation of WM and GM intensities, and a regularisation consisting of an Ising model to encourage a smooth object boundary. A 3D skeleton of the exterior is computed using a homotopic thinning algorithm. This skeleton is split into major sulci, the external brain surface, and inter-hemispheric fissure. The so identified simple surfaces form the nodes of an attributed relational graph representing the cortical topography. These processing steps are implemented by the publicly available Morphologist toolbox of BrainVISA[9].

A similar atlas-free approach used to segment the cortex in infant brain images from a wide range of ages with the HDR algorithm was presented by Leroy et al. (2011). The

[9] http://brainvisa.info/web/index.html (last accessed: May 20, 2017)

HDR algorithm is used to simultaneously deform two binary regions, one inside of the WM/GM interface and one outside of it. These coupled HDRs increase the robustness of the method towards partial volume effects at the GM/CSF interface, where intensities slightly outside the pial surface are similar to WM, resulting in partial image boundaries with an appearance similar to the WM/GM boundary. In an initial deformation stage, both surfaces are deformed until close to the cortical ribbon, which is identified using a morphological top hats filter (Serra, 1983). In a final deformation stage, the surfaces are further inflated/deflated simultaneously towards the WM/GM interface using the energy minimisation framework of Mangin et al. (1995), until these meet. The two surfaces thereby converge to the same boundary, which corresponds to the WM/GM interface.

A different approach for neonatal tissue segmentation involving filtering steps to enhance structures of interest at various stages, and seed-based segmentation with seeds informed by prior knowledge, was developed by Gui et al. (2011, 2012a). They participated with their morphological approach in the NeoBrainS12 challenge (Gui et al., 2012b). A difficulty is to make such approach robust to varying image quality and acquisition protocol. Even though these methods are very adaptable, it may be difficult to develop and stabilize the processing steps. A hybrid method for neonatal tissue segmentation, which combines the robustness of atlas-based methods with the adaptability of morphological operations, was more recently developed by Beare et al. (2016). Their atlas-based segmentation method utilises the unified segmentation (Ashburner and Friston, 2005) available in SPM, and is itself implemented as a freely available SPM toolbox called Morphologically Adaptive Neonatal Tissue Segmentation (MANTiS)[10]. The atlas employed by this toolbox is a modified version of the atlas constructed by Kuklisova-Murgasova et al. (2011).

4.4. Supervised brain tissue segmentation

Song et al. (2007) consider the segmentation task as a Bayesian inference problem in a MRF framework, with a data term derived from a spatially varying non-parametric model, and an intensity-based prior learned using a fuzzy non-linear support vector ma-

[10] http://developmentalimagingmcri.github.io/mantis (last accessed: May 20, 2017)

chine (SVM) (Burges, 1998). Their method does not require anatomical priors from a probabilistic atlas, but can readily make use of it when available. The pairwise potential of neighbouring voxels penalises differing label assignment and thereby regularises the segmentation. The particular non-parametric models used are Parzen window estimates with Gaussian kernels. The mean values of the kernels are from a random sample drawn from the intensity histogram, whereby a spatially localised sampling strategy results in intensities of voxels closer to the voxel to be classified to be drawn with higher chance than those of voxels that are farther away. The resulting MRF-based energy function is minimised using graph cuts (Boykov et al., 2001).

A multi-stage kNN tissue classification approach was proposed by Chita et al. (2013). In the first stage, voxels that can be assigned with high confidence to either WM, GM, or CSF are labelled. In the second stage, remaining voxels are classified considering also features based on the tissue probabilities corresponding to the first stage. Both of these stages use a separate two-class classification for each tissue type. Inconsistencies in the resulting tissue-specific segmentation are resolved in a third and final stage, which performs multi-class classification. Similar to Anbeek et al. (2013), the classification is based on both intensities and spatial features, whereby also a number of textual features are extracted, and a subset of all features is selected using a sequential floating forward search (Pudil et al., 1994). Unlike the spatial features used by Anbeek et al. (2013), the spatial features employed by Chita et al. (2013) are relative to the brain extraction mask of each image. Therefore, no spatial normalisation is needed before classification. The method was adapted and further evaluated in Moeskops et al. (2013). In this modified version, a SVM is used instead of a kNN classifier in the second stage. An advantage of this supervised approach is that the training images can be only partially annotated.

A different multi-stage classification of major tissues in infant brain images based on pre-trained RF multi-class classifiers and 3D Haar-like features (Viola and Jones, 2004) was proposed by Wang et al. (2015). The RF classifier of the first stage is trained not only on T1w and T2w images, but also FA images, in order to improve the classification for brain images acquired at an age of 6–8 months. At this age, before WM is fully myelinated, the distribution of WM and GM intensities are nearly identical for T1w and T2w images.

This renders classification very challenging, and the FA image provides valuable structural information to resolve ambiguities in tissue label assignment. At consecutive stages, the tissue probabilities of the previous stage are added as features. This auto-context scheme was initially proposed by Tu and Bai (2010), and applied to RF-based tissue classification in the atlas forest approach of Zikic et al. (2014).

The use of standard feed forward ANNs for tissue classification has already been investigated by Özkan et al. (1993) and Zijdenbos et al. (1998). Their application for brain image segmentation has been limited, however. Convolutional neural networks (CNNs), on the other hand, enjoy widespread attention in recent years.

First to propose a CNN for brain tissue segmentation of images of about 6–8 month old infants was Zhang et al. (2015). Similar to Wang et al. (2015), the CNN is trained on T1w, T2w, and FA images. A sufficient amount of training data was generated using iBEAT, an atlas-based deformable models approach discussed in the next section, followed by manual editing to correct mistakes. The network was trained on 2D image patches of one size for each voxel to be classified. Zhang et al. (2015) performed experiments with four different CNN architectures, and compared the results for eight infant brain images to those based on a linear SVM or RF classifier instead.

A CNN architecture for segmenting brain tissues in images from a wide range of ages, including the neonatal period and (ageing) adults, was presented by Moeskops et al. (2016). This network was trained separately for each set of images from a specific period of life using 2D in-plane image patches of three different sizes for each voxel to be classified. Patches in the out-of-plane direction or 3D patches were not used, because of the higher slice thickness of the acquired neonatal brain images. For neonatal brain tissue classification, the method was evaluated on the NeoBrainS12 challenge data and compared to other methods. The proposed CNN achieved better average Dice similarity coefficient (DSC) (Dice, 1945) for each of the individual test sets. It was recently outperformed by a slight margin by two other yet unpublished CNN approaches[11].

[11] On-line challenge team names "LRDE_LTCI" and "INFANT-INSIGHT" (last checked: Jun 20, 2017)

4.5. Deformable models of the cerebral cortex

Methods discussed thus far are based on a voxel-wise labelling of the brain image. A very different segmentation approach is that of deformable models. A deformable model actively adjusts its shape under the influence of external forces, pushing the model towards image features such as the edges of a structure, and internal forces, modelling the physical properties of elastic materials. Whereas voxel-wise labelling followed by explicit modelling of the boundaries of structures is a purely bottom-up approach to shape modelling, deformable models combine the bottom-up image-derived forces with a top-down modelling of the expected characteristics of the shape.

One class of deformable models are those, where the shape is a parametric function. First parametric models were used to extract contours in 2D images. These are known as active contours or snakes, respectively, a name inspired by the way the contours slither while minimising their energy (Kass et al., 1988). An early application in 3D was the generation of physically plausible animations in computer graphics (Terzopoulos and Fleischer, 1988). For an overview of early applications of parametric deformable models in medical image analysis, see the survey of McInerney and Terzopoulos (1996). A more recent introduction to deformable models in general is given by Xu et al. (2000) and Osher and Fedkiw (2003). A key property of parametric models is the predefined topology. In applications where the topology of the shape is not known, this is a disadvantage, because the topology of the parametric model has to be adapted explicitly when necessary.

Another type, which can change its topology naturally, are geometric models. The contour in 2D or surface in 3D is hereby given implicitly by the zero level set of a function defined over the image domain (Osher and Sethian, 1988). These methods are hence defined in an Eulerian domain. In contrast to this, parametric models correspond to a Lagrangian point of view. Besides being able to easily adapt its topology, numerical methods for solving the associated PDEs are generally more efficient than those for parametric models. In case of the cerebral cortex, which has a known sheet-like topology, not being able to control the topology of the model is a substantial limitation. This limitation is overcome by the topology-preserving level set method proposed by Han et al. (2001, 2003).

4.5.1. Parametric deformable models

In chapter 7, we present parametric deformable models for the reconstruction of the cortical surfaces of the neonatal brain from MRI data. Here, we provide a brief introduction to parametric deformable models as found in McInerney and Terzopoulos (1996) and Xu et al. (2000), before we summarise methods that have been proposed in the literature.

A parametric deformable model entails a continuous function

$$\boldsymbol{X} : \omega = [0,1]^{d-1} \rightarrow \Omega \tag{4.1}$$

where ω is the parametric domain of the model, and $\Omega \subset \mathbb{R}^d$ is the image domain of dimension d. The function $\boldsymbol{X}$ is a parametric curve in 2D, and a surface in 3D. It may be represented by a piecewise linear approximation, e.g., a triangular surface mesh. Parametric deformable models are therefore also referred to as explicit models, because of this explicit representation of the model's geometry.

Given an image $I : \Omega \rightarrow \mathbb{R}$ of the human brain, we want to find the function $\hat{\boldsymbol{X}}$, which best matches the cortical surface, while obeying to the internal constraints defined by the model. Two formulations of parametric deformable models are common. The first is based on an energy functional $\mathcal{E}(\boldsymbol{X})$, where the model minimising this energy is sought, i.e., $\hat{\boldsymbol{X}} = \arg\min_{\boldsymbol{X}} \mathcal{E}(\boldsymbol{X})$. Necessary conditions for a local minimum of $\mathcal{E}$ give rise to a gradient descent scheme based on a force balance type equation. The second formulation is directly based on the resulting formulation of $\boldsymbol{X}$ as a dynamic function over time, i.e.,

$$\boldsymbol{X} : \omega \times [0,1] \rightarrow \Omega \tag{4.2}$$

and a balance equation between internal model forces and external forces derived from the image I, together with partial derivatives describing the temporal dynamics of the model. While the energy formulation poses the problem as well understood minimisation problem, the dynamic formulation allows more flexibility in the design of the model, because only the forces acting on its discrete representation have to be defined, and a corresponding continuous and differentiable energy function may indeed not exist or be difficult to obtain.

Energy minimising formulation

In general, the energy functional of a deformable model is given by

$$\mathcal{E}(\boldsymbol{X}) = \mathcal{P}(\boldsymbol{X}, I) + \alpha \mathcal{S}(\boldsymbol{X}) \tag{4.3}$$

where the potential energy $\mathcal{P}$ measures how well the model fits the underlying imaging data, the internal energy $\mathcal{S}$ is proportional to the deflection and bending of the model, and $\alpha \in \mathbb{R}_{>0}$ balances data fidelity and model regularity. The snakes potential is

$$\mathcal{P}(\boldsymbol{X}, I) = \int_{\omega} P_I\big(\boldsymbol{X}(\boldsymbol{s})\big)\, d\boldsymbol{s} \tag{4.4}$$

where the image potential $P_I : \Omega \to \mathbb{R}$ is defined over the image domain. An example is the edge potential given by $P_I(\boldsymbol{x}) = -\|\nabla I_\sigma(\boldsymbol{x})\|^2$, where I_σ is the image after convolution with a Gaussian kernel with standard deviation σ. The internal energy can be written as

$$\mathcal{S}(\boldsymbol{X}) = \sum_{1 \le |\kappa| \le 2} \frac{1}{2} \int_{\omega} \beta_\kappa(\boldsymbol{s}) \left\| D^\kappa \boldsymbol{X}(\boldsymbol{s}) \right\|^2 d\boldsymbol{s} \tag{4.5}$$

where $\kappa \in \mathbb{R}^{d-1}_{\ge 0}$ is a multi-index, $|\kappa|$ denotes the ℓ_1-norm of κ, i.e., the sum of indices, and $D^\kappa \boldsymbol{X}(\boldsymbol{s})$ is a d-dimensional vector of partial derivatives of $\boldsymbol{X}(\boldsymbol{s})$ of order $|\kappa|$, i.e.,

$$D^\kappa \boldsymbol{X}(\boldsymbol{s}) = \frac{\partial^{|\kappa|} \boldsymbol{X}(\boldsymbol{s})}{\partial s_1^{\kappa_1} \cdots \partial s_{d-1}^{\kappa_{d-1}}} \tag{4.6}$$

The spatially adaptive weight of each regularisation term is given by $\beta_\kappa(\boldsymbol{s})$. The first order terms make the snake act like a membrane, whereas the second order terms correspond to the bending energy of a thin plate of metal (Kass et al., 1988).

When the contour is closed, the Gâteaux derivative after integration by parts is

$$\delta_{\boldsymbol{h}} \mathcal{S}(\boldsymbol{X}) = \sum_{1 \le |\kappa| \le 2} (-1)^{|\kappa|} \int_{\omega} \beta_\kappa(\boldsymbol{s}) \boldsymbol{h}^T(\boldsymbol{s}) D^{2\kappa} \boldsymbol{X}(\boldsymbol{s})\, d\boldsymbol{s} \tag{4.7}$$

where $\boldsymbol{h} : \omega \to \Omega$ is a perturbation of the model, i.e., the direction in which the derivative is computed. A necessary condition for a local minimum is that $\delta_{\boldsymbol{h}} \mathcal{S}(\boldsymbol{X}) = \boldsymbol{0}$ for all $\boldsymbol{h}$.

Hence, the minimiser $\hat{\boldsymbol{X}}$ satisfies the Euler-Lagrange PDE given by

$$\underbrace{\nabla P_I\left(\hat{\boldsymbol{X}}(\boldsymbol{s})\right)}_{\boldsymbol{f}_{\text{ext}}(\boldsymbol{s})} + \alpha \underbrace{\sum_{1\leq|\alpha|\leq 2} \beta_\kappa D^{2\kappa}\hat{\boldsymbol{X}}(\boldsymbol{s})}_{\boldsymbol{f}_{\text{int}}(\boldsymbol{s})} = \boldsymbol{0} \tag{4.8}$$

with internal and external forces acting on the model at point $\hat{\boldsymbol{X}}(\boldsymbol{s})$ given by $\boldsymbol{f}_{\text{int}}(\boldsymbol{s})$ and $\boldsymbol{f}_{\text{ext}}(\boldsymbol{s})$, respectively. To find a solution for which these forces are in equilibrium, i.e., have equal magnitude and point in opposite directions, the deformable model is made dynamic by introducing a time parameter $t \in [0, 1]$ as in equation (4.2). Setting the partial derivative with respect to t equal to the left-hand side of equation (4.8) gives

$$\frac{\partial \boldsymbol{X}(\boldsymbol{s}, t)}{\partial t} = \boldsymbol{f}_{\text{ext}}(\boldsymbol{s}, t) + \alpha \boldsymbol{f}_{\text{int}}(\boldsymbol{s}, t) \tag{4.9}$$

Given an initial estimate $\boldsymbol{X}(\boldsymbol{s}, 0) = \boldsymbol{X}_0(\boldsymbol{s})$, the stationary solution of this ODE, i.e., when the left-hand side is zero, thus is a minimiser of the model energy $\mathcal{E}(\boldsymbol{X})$.

Dynamic force formulation

We have seen that a deformable model minimises a potential energy, which is integrated over the Euclidean image domain, while at the same time minimising its internal energy. The final static model balances the external and internal forces corresponding to these energy terms. Instead of defining the deformable model based on its differentiable energy, it can be treated directly as a dynamic model by defining the forces acting on it from a Lagrangian point of view. Given a dynamic model as defined in equation (4.2) with a mass density $\mu(\boldsymbol{s})$ and damping density $\gamma(\boldsymbol{s})$, the Lagrange equation of motion is given by

$$\mu(\boldsymbol{s})\boldsymbol{a}(\boldsymbol{s}, t) + \gamma(\boldsymbol{s})\boldsymbol{m}(\boldsymbol{s}, t) + \alpha \boldsymbol{f}_{\text{int}}(\boldsymbol{s}, t) = -\boldsymbol{f}_{\text{ext}}(\boldsymbol{s}, t) \tag{4.10}$$

where $\boldsymbol{m}(\boldsymbol{s}, t) = \frac{\partial \boldsymbol{X}(\boldsymbol{s},t)}{\partial t}$ and $\boldsymbol{a}(\boldsymbol{s}, t) = \frac{\partial^2 \boldsymbol{X}(\boldsymbol{s},t)}{\partial t^2}$ are the instantaneous momentum and acceleration of the model at time t, respectively. When the dynamic model comes to rest, both momentum and acceleration are zero, i.e., $\boldsymbol{m} = \boldsymbol{a} = \boldsymbol{0}$, and we obtain the force balance condition $\alpha \boldsymbol{f}_{\text{int}} = -\boldsymbol{f}_{\text{ext}}$ corresponding to the Euler-Lagrange equation (4.8).

Parametric deformable models in medical image analysis

Parametric deformable models were pioneered by the snakes model of Kass et al. (1988), where a contour in a 2D image was modelled by a spline function, a variational formulation of its elastic properties, and an image potential attracting it towards image features. Specifically, the external force of the original snake model is an image potential, which is proportional to the magnitude of the image gradient. This has the drawback that the model must be initialised close to image edges. In a homogeneous image region, the gradient is zero, and no external force moves the contour within capture range of an image edge. In this case, only the elastic internal forces, which tend to shrink the contour on itself, are active. This issue is overcome for contours in 2D, and surfaces in 3D by a balloon force as proposed by Cohen and Cohen (1993). A balloon force, also referred to as pressure force, pushes the surface outwards as long as an image feature, e.g., the intensity itself, is within a specified range, and inwards otherwise. It can be designed to be inactive nearby an image edge, where the potential force is active. Another novelty introduced by Cohen and Cohen (1993) is a distance potential force, which is proportional to some measure of distance of the surface to image features such as edges. A problem of using a Chamfer or (minimum) Euclidean distance for the potential force is, that it fails to drive the surface into concavities, as demonstrated by Xu and Prince (1998). This renders the reconstruction of the cerebral cortex using such distance potential difficult, because it is highly folded. A good initialisation is needed for many deformable models, whose energy minimisation procedures are easily trapped in a local minimum of the cost function.

The gradient vector flow (GVF) proposed by Xu and Prince (1998) allows snakes to penetrate concave regions, overcoming some of the shortcomings of distance potentials. This external force is computed as a diffusion of the gradient vector field of the image, a binary segmentation, or an edge map, respectively. This increases its capture range, and achieves a force within concavities, that pushes the model deeper into these. Because the GVF cannot be written as potential energy function, the deformable model is formulated directly in terms of the force balance conditions, rather than a variational energy minimisation. The ability to formulate a deformable model based on the forces that act on it,

without the need for a precise definition of their anti-derivative, is a beneficial property, which greatly simplifies the modelling of new application-specific interactions.

Among the first deformable surfaces parametrised by a polyhedral mesh in 3D is the surface model proposed by Miller et al. (1991). Unlike previous methods for object modelling, where contours of adjacent slices of a volumetric dataset are extracted first individually, e.g., using the snakes model of Kass et al. (1988), and then stitched together to form a surface model, their method directly deforms a surface model in 3D. The constraint and external forces hereby act on each surface node individually. The model is initialised to a low resolution model, specifically an icosahedron with few nodes, and it is gradually refined using a global resampling, where each triangular face is subdivided into four by inserting a new model node at the centre of the triangle. This multi-resolution approach was inspired by the deformable template matching approach of Bajcsy and Kovacic (1989), a volumetric multi-resolution image registration method for elastically deforming an atlas image to match a given subject image. A curvature constraint, minimising the distance of a surface node from the centroid of adjacent nodes, keeps it from straying too far from its neighbours, and locally models the surface as an elastic sheet. To allow the model to deform into concavities, Miller et al. (1991) employed a deformation potential that searches along the surface normal direction for an image event, based on a simple shifted threshold operator. This operator is zero, when the intensity is below a given threshold corresponding to the lowest intensity of the inside of an object, and is otherwise proportional to the difference of the intensity and this threshold. The image feature search in normal direction overcomes the problem of a (minimum) distance potential, to not be able to drive the model into deep sulci. The deformation process iteratively moves each surface node in the direction of steepest descent along the cost surface, i.e., the direction opposite to the gradient of the cost function, with adaptive step size control. During this gradient descent, the polyhedral model preserves its initial closed topology.

Cootes et al. (1993, 1994) not only constrained the topology, but also the shape of the deforming model. A statistical shape model consisting of a mean shape and modes of variation is derived from a number of manually delineated and co-aligned contours of different examples of an object. The mean shape corresponds to the average model

node positions, whereas the modes of variation are derived from the eigenvectors of the covariance matrix of node positions. At each deformation step, a new best position is searched for each model node along the boundary normal direction. The set of new node positions are then projected onto the space of admissible shapes, as defined by the statistical shape model. Given the large biological variations in cortical folding, such shape constraints are potentially too strict for reconstructing the cortical surface, and are therefore generally not applied to this problem.

Sandor and Leahy (1997) circumvented the issue of accurately modelling the cortical folds by focusing on a coarse reconstruction of the outer brain surface. This is sufficient to identify the four lobes and major fissures. Also Davatzikos and Prince (1995) and Davatzikos and Bryan (1996) first extract a representation of the outer brain surface, and combine this surface model with a measure of sulcal depth at each surface node. They used the model of the outer cortical surface for spatial normalisation (Davatzikos, 1996).

Dale and Sereno (1993) used three flood fillings after WM segmentation to obtain a topologically correct close initialisation, and refine it using a shrink-wrapping procedure, with external forces pushing the surface inwards or outwards in normal direction, depending on the difference of the intensity at the surface node from the average GM intensity. The reconstructed cortical surface models were originally used to solve the electroencephalogram (EEG) or magnetoencephalography (MEG) inverse problem. In this application, electromagnetic signals are recorded with electrodes (for EEG) or magnetometers/gradiometers (for MEG) outside the skull. To be able to recover the set of currents inside the brain, that gave rise to the measured signals, cortical surface models were sought to be able to constrain this ill-posed inversion problem. The cortical surface reconstruction code of Dale and Sereno (1993) gave rise to FreeSurfer[12] (Fischl, 2012), a popular software for cortical surface reconstruction from T1w adult brain MR images. The original deformable surface model was extended in Dale et al. (1999) for the reconstruction of both WM/GM interface (WM surface) and GM/CSF interface (pial surface). This method still required manual editing to fix topological defects, but this step has later been automated (Fischl et al., 2001; Segonne et al., 2007). The surface model consists of an intensity-based data term

[12] https://surfer.nmr.mgh.harvard.edu/ (last accessed: May 20, 2017)

and two separately weighted spring forces, one in normal direction and the other force acting in the tangent plane. While the tangent spring forces redistribute surface nodes to regions where there is local expansion in order to maintain a regular sampling, the normal spring forces impose a smoothness constraint. The WM surface model is initialised close to the WM/GM interface using a previous tissue segmentation, and is refined using an external force to minimise the squared sum of differences between each surface node intensity and the mean WM intensity within a local neighbourhood. The pial surface is reconstructed from this WM surface by deforming it outwards, as long as the intensity at a surface node is above a determined global estimate of GM/CSF edge intensity. This assumes a good contrast between GM and CSF, as is common for T1w images of the mature brain. A triangle-triangle intersection test (Möller, 1997) is used to avoid self-intersections during the expansion. The surface reconstruction method of FreeSurfer has recently been improved for high-resolution T1w images with an image gradient-based surface refinement as part of the minimal processing pipeline of HCP (Glasser et al., 2013).

A different initialisation for the reconstruction of a topologically-correct mid-thickness surface of the cortex was proposed by Xu et al. (1997). This mid-thickness surface is halfway between the WM and pial surfaces. Their method uses FCM membership functions of WM, GM, and CSF. An initial surface tessellation is obtained by extracting the isosurface of the WM membership function for value 0.5. As long as this surface has incorrect Euler characteristic, a median filtering of the WM membership function is performed, and the isosurface is recomputed. The largest connected component of this topologically correct isosurface is used as initialisation of the deformable surface model. This initial model is deformed towards the peaks of the GM membership function using potential and pressure forces, to obtain a surface model of the mid-thickness layer of the cortex.

MacDonald et al. (1994) introduced a deformable polyhedral surface model, which jointly deforms multiple surfaces. They used this method to solve the problem of reconstructing the inner- and outer-cortical surface under the presence of partial volume, rendering the identification of especially the GM/CSF interface difficult (MacDonald et al., 1998). To avoid intersections, their model includes a self-proximity constraint between any pair of non-adjacent polygons. An additional constraint on the distance between closest

nodes of two surfaces encourages the surfaces to stay a certain distance apart. It is used to incorporate *a priori* knowledge of average cortical thickness. MacDonald et al. (1998) also utilise an external force, which searches for the closest image edge in surface normal direction, and allows it to penetrate narrow and deep sulci. Their method, named anatomic segmentation using proximities (ASP), and based on concentric surfaces with proximity links, is further detailed in MacDonald et al. (2000). While the cortical thickness constraint addresses the issue of partial volume within sulci, it has the drawback of biasing the cortical thickness measurements computed from the reconstructed surfaces. This issue has been raised and addressed later by Kim et al. (2005), who propose a new method for reconstructing the cortical surfaces referred to as constrained Laplacian ASP (CLASP). Instead of jointly deforming two concentric surfaces, CLASP adopts a sequential procedure, where the WM surface is reconstructed first, and subsequently expanded towards the pial surface. A classification method for detecting the partial volume fraction of CSF in deep sulci replaces the thickness constraints. The expansion of the WM surface for creating the pial surface is based on a Laplacian map between this surface and a skeletonised CSF map, similar to the cortical thickness estimation of Jones et al. (2000). The corresponding expansion paths prohibit self-intersections during the surface deformation. CLASP is used in CIVET[13] (Ad-Dab'bagh et al., 2006), an image processing pipeline for fully automated volumetric, corticometric, and morphometric analysis of adult T1w brain MR images, which is available as a service from the BIC at the MNI, McGill University. CIVET was only recently modified to accommodate the processing of T2w neonatal brain images. The new pipeline is called NEOCIVET (Kim et al., 2016).

A parametric deformable model was used by Wright et al. (2015) to reconstruct the WM/GM boundary of the fetal cortex. The so obtained WM surfaces were used to create a spatio-temporal surface atlas from surfaces reconstructed from T2w images of the brains of fetuses at different gestational ages. In their approach, the WM surface was reconstructed from a previous segmentation of the brain image into WM, cGM, CSF, lateral ventricles, dGM structures, brainstem, and cerebellum using the (probabilistic) atlas-based method of Makropoulos et al. (2014). After merging the posterior probabili-

[13] http://mcin-cnim.ca/technology/civet/ (last accessed: Jun 20, 2017)

ties of structures inside the WM/GM boundary, a distance transform (Maurer et al., 2003) was computed from the corresponding binary mask. The nodes of the initial, closed triangular surface mesh are deformed in normal direction relative to the minimum Euclidean distance to the segmentation boundary, whereby the magnitude of this force is clamped to an upper threshold, such that all nodes move coherently. After each step along the normal direction, a Laplacian smoothing of the surface mesh, with weights proportional to the inverse distance of neighbouring nodes, is performed. To allow the surface to dynamically expand and compress without affecting its vertex density and topology, long edges are split and short edges collapsed afterwards. The reconstructed WM surface closely follows the binary segmentation of WM and sub-cortical structures, with reported average node distance to the segmentation boundary of 0.1 mm. As common for cortical surface reconstruction methods that completely rely on a (fuzzy) tissue classification, it strongly depends on a good segmentation. But unlike the segmentation, the reconstructed WM surface model is guaranteed to have the correct topology.

4.5.2. Geometric deformable models

The surface of a geometric deformable model is implicitly given by a scalar function defined over the image domain. More specifically, the deforming surface corresponds to the zero level set of this function. A common choice of implicit function is the signed Euclidean distance. This allows a direct query of the Euclidean distance to the closest surface point from any other point. The geometric model evolves through a change of the implicit function. An advantage of such deformable model over parametric models is, that self-intersections are inherently prevented, eliminating the need for costly intersection checks, and that it can naturally change its topology without any need for reparametrisation. The automatic adaptation of the surface topology is however undesired when modelling the surface of a shape with known topology, such as the cortical surface. A geometric deformable model therefore either relies on a topologically correct underlying boundary map (segmentation), requires a post-processing to remove topological errors, or is combined with an update procedure, which explicitly preserves its initial topology.

A topology-preserving level set method for the reconstruction of the cortical surfaces

was developed by Han et al. (2001, 2003). First, AFCM membership functions for GM, WM, and CSF are computed, and the WM membership function filled in to include lateral ventricles and dGM structures. The topology of the WM segmentation is then corrected using a graph-based approach (Han et al., 2002). The boundary of this segmentation is homotopic to a sphere, and its tessellation yields the initial triangular mesh of the topology-preserving level set evolution. This level set evolution makes use of the fact, that the topology of the surface can only change, when the sign of the implicit function changes at a point. When the respective point is a simple point (Bertrand, 1994) of the binary object corresponding to non-positive function values, the sign change that is equivalent to removing this point from the binary object does not change the topology of this shape. The topology-preserving level set evolution hence ensures that the implicit function only changes sign for such simple points, and prevents a sign change otherwise. The initial tessellation of the WM surface is smoothed using potential and pressure forces based on the WM membership function. The mid-thickness and pial surfaces are then reconstructed starting with this WM surface. To deal with partial volume, the GM membership function is automatically edited beforehand, to produce evidence of CSF within narrow sulci. The basic approach is to locate the exterior skeleton of the WM segmentation, and to reduce the GM membership value on this skeleton, thus changing this particular GM into CSF (Han et al., 2004). After this automatic GM membership editing, the mid-thickness surface is obtained by evolving a geometric deformable model initialised to the WM surface, using an external GVF force based on the GM/CSF edge map, and a regional pressure force to keep the surface from moving outside of GM. The pial surface is then reconstructed by another geometric deformable model similar to the WM model, but based on the combined WM and GM membership functions. An implementation of this cortical reconstruction using implicit surface evolution (CRUISE) is freely available[14]. Instead of the original AFCM followed by graph-based topology correction, this software utilises the topology-preserving anatomical segmentation (TOADS) of Bazin and Pham (2008). The tissue segmentation preserves the topology of structures as defined by a topological atlas. The accuracy of this TOADS-CRUISE method was established by Tosun et al. (2006). The nested CRUISE

[14] http://www.nitrc.org/projects/toads-cruise/ (last accessed: May 20, 2017)

method was furthermore used by Xue et al. (2007) to reconstruct the cortical surfaces from T2w neonatal brain MR images using their atlas-free EM classification approach with knowledge-based partial volume correction.

Distance-based coupling of two deformable models for WM and pial surfaces was first proposed by Zeng et al. (1999) using a level set approach, where two separately embedded surfaces are simultaneously evolved. When the distance from one surface to the other is outside a normal range, it imposes a constraint on its interface propagation speed. The minimum distance from a point to a surface in an implicit surface model is directly given by the respective signed distance function. The used boundary indicator functions are based on the joined probability of the mean intensity in a region outside the surface to belong to the expected tissue outside, and the mean intensity of a region inside to belong to the expected tissue inside the target surface. This assumes known estimates of the likelihood functions of GM, WM, and CSF. The coupled level-sets method was formulated by Goldenberg et al. (2002) as energy minimisation problem. They used ratios of tissue probabilities instead of the product as boundary indicator functions, and a fast geodesic active contour approach to minimise this energy. Both methods start the coupled surfaces evolution with manually placed concentric spheres inside the WM. Similar to the parametric ASP method with cortical thickness constraint, the coupled level-sets approach has the potential of bias in cortical thickness measures.

Wang et al. (2011a) incorporated anatomical atlas priors into the image-based term of the coupled level sets approach. Three interrelated geometric deformable models are defined to subdivide (i.e., segment) the brain volume into four disjoint regions: WM, GM, CSF, and background. The joined energy of these models is based on Gaussian intensity distributions estimated for each tissue type within a local neighbourhood defined by a truncated spatial Gaussian kernel (Wang et al., 2009). This energy is minimised using gradient descent. To enforce a cortical thickness constraint, the speed functions of the WM/GM and GM/CSF boundaries are coupled. When the distance of one surface to the other is outside a specified distance range, it is either locally inflated or deflated, respectively, until the distance is within the allowed range, while the image-based force is disabled to not affect this corrective surface deformation. The cortical thickness constraint

is deactivated near sub-cortical structures, i.e., within a co-registered template mask enclosing the lateral ventricles and sub-cortical structures. The geometric deformable models are initialised in three stages, where the first two stages consist of the global minimisation of similar convex energy functions. The first stage separates CSF from both background and the combined WM and GM region. In the second stage, the WM and GM is divided into two regions near the WM/GM interface. Misclassified WM between CSF and GM is subsequently changed to either CSF or GM, based on the frequency of occurrence of each tissue within a local neighbourhood. The approach of Xue et al. (2007) is further adopted to correct misclassified CSF within narrow sulci, which has been labelled as WM instead.

Following the atlas-based approach of Shi et al. (2010), Wang et al. (2011b) extended their coupled level sets to include a longitudinal constraint for segmenting the neonatal image, given the tissue boundaries of the infant brain image acquired at two years of age. The images of the two time points are first segmented independently. The method of Wang et al. (2011a) is used to segment the neonatal image, whereas the infant brain image is segmented using AFCM. The tissue membership functions of the late time point image are used as subject-specific priors in the coupled level sets approach. The tissue boundaries are extracted from the AFCM segmentation using fast marching cubes (Sethian, 1999). After aligning the tissue segmentations of both time points using HAMMER, the longitudinal constraint is defined similar to the cortical thickness constraint. It limits how far the models can deviate from the corresponding late time point boundary. In later work, this approach has been further extended to utilise images from multiple modalities, namely T1w, T2w, and FA images, to segment each time point using the original coupled level sets approach, and to then refine the individual segmentations by coupling the level sets across time with the longitudinal constraint proposed previously. Different subject-specific atlas constructions providing anatomical priors for the coupled level sets were explored more recently. Atlases were fused based on the sparse representation coefficients given a dictionary consisting of local patches from co-registered atlases, either only from the T2w images (Wang et al., 2014b), or all available modalities and the tissue priors (Wang et al., 2014a). A number of these approaches are implemented in iBEAT[15] (Dai et al., 2013).

[15] http://www.nitrc.org/projects/ibeat (last accessed: May 20, 2017)

Part III.

Methods

Chapter 5.

Flexible, fast, and diffeomorphic registration of the neonatal brain

The Image Registration ToolKit (IRTK) (Rueckert et al., 1999; Schnabel et al., 2001) has been demonstrated to be still among the best performing medical image registration tools available (Klein et al., 2009a), despite the fact that it has been written more than 10 years prior to this study, and by now has reached an age of more than 15 years. The main use of the IRTK is for non-rigid registration of medical images using the FFD algorithm (Rueckert et al., 1999). Given the success of the FFD algorithm, we adopted it also for our work on modelling the anatomy of the developing brain. For this work, two properties of the original algorithm limit its applicability.

The first limitation is the relatively long runtime of IRTK, due to finite difference approximations of the energy gradient by in turn varying the parameters of each control point. This limitation is more of practical nature, but also prevents its use in routine clinical applications. The relatively long runtime of IRTK furthermore prolongs finding optimal parameters for a given registration task, and increases the computing resources or time required by a brain atlas construction based on deformable registration between each pair of images in a big dataset. A pairwise atlas construction, which benefits from the fast brain registration method presented in the following, is summarised in section 6.2.

The second limitation is that the original FFD algorithm does not ensure invertible mappings, which preserve the topology of developing brain structures as outlined in sec-

tion 2.2.4. The diffeomorphic variant proposed by Rueckert et al. (2006), where the total transformation is modelled as the composition of small diffeomorphic FFDs by restricting the maximum coefficient value, increases the memory required to store these compositions, and hinders a compact representation of an average FFD needed for atlas construction. Computing the average deformation of a set of composite FFDs generally results in a dense representation of the average displacement field. The interpolation of this dense displacement field at non-voxel positions then introduces deviations from the average displacements of the composite FFDs at this point. Inverting a FFD or dense displacement field also requires to numerically find the inverse mapping such that $\mathbf{\Phi}^{-1} \circ \mathbf{\Phi} = \mathbf{\Phi} \circ \mathbf{\Phi}^{-1} = \mathrm{Id}$.

In section 2.2.3, transformations which are modelled after the dynamics of a viscous fluid are summarised. These transformation models are given by the integration of small instantaneous velocity fields. This not only allows for large diffeomorphic deformations, but finding the inverse deformation corresponds to a simple exchange of integration bounds. Further advantages of these models for inter-subject registration under the presence of large biological variations are, that a mean deformation, which itself belongs to the group of diffeomorphic mappings, is given by the average velocity field, and that a symmetric and inverse consistent registration can be derived directly, without minimising an inverse consistency error (see section 2.4). In chapter 6, we make use of these properties for the construction of a spatio-temporal atlas of the developing brain.

It should be noted that aforementioned benefits are only given, when the respective model is indeed parametrised by a velocity field. In case of a greedy non-parametric approach, such as the viscous-fluid model (Christensen et al., 1996), diffeomorphic Demons (Vercauteren et al., 2007, 2009), and the greedy SyN algorithm (Avants et al., 2008), the intermediate velocity fields are not recorded. The result of these greedy registration methods is instead the final dense displacement field, which is commonly accompanied by a separate displacement field corresponding to the inverse transformation. The properties of a velocity based parametrisation, which are beneficial for the construction of mean shape models of the developing brain (see also chapter 6), and the study of brain development using the principles of CA introduced in section 1.3, are thus traded in for a more efficient registration, which does not require the costly space-time optimisation of LDDMM.

A further drawback of a greedy optimisation, even when the intermediate update fields are preserved, such as in case of the diffeomorphic FFD approach of Rueckert et al. (2006), is that it may lead to a suboptimal flow of diffeomorphisms in terms of the deviation from the shortest path connecting the identity mapping from the final deformation along the manifold of diffeomorphisms. When defining a mean deformation based on these velocity fields, a deformation corresponding to a longer path contributes more to the mean shape deformation than a flow that follows the geodesic. To illustrate this further, when the flow deforms an earlier time point of the developing brain to a later time point of the same anatomy, brain structures may implausibly expand, shrink, and expand again.

In this chapter, we first introduce the Medical Image Registration ToolKit (MIRTK). This modular software for medical image analysis is the result of our substantial refactoring and extension of IRTK. This new registration library improves upon the previously mentioned limitations of the original framework. Above all else, it implements registration methods that are based on parametric spline-based transformation models, and the minimisation of an energy function. A unique feature compared to other registration tools is the flexibility of the registration tool, which realises different algorithms based on a configurable simplified math expression of the energy function.

Afterwards, we consider two deformable registration methods for diffeomorphic registration of neonatal brain images implemented in MIRTK. The first approach is asymmetric and based on soft constraints of the Jacobian determinants of a classic FFD model as utilised by Rueckert et al. (2006) and Modat et al. (2012a). We then propose a symmetric method with strong guarantees of invertibility and inverse consistency for the unbiased deformable registration of neonatal brain images. This second approach is based on a cubic B-spline parametrisation of a SVF. This model was first used by Modat et al. (2012b) for the diffeomorphic registration of adult brain images. A parametrisation based on different RBFs with compact support, namely the Wendland kernels (Wendland, 1995), was used more recently also by Pai et al. (2015). The use of a parametric SVF model leads to an efficient algorithm that retains the benefits of a velocity based parametrisation for the study of brain development, while avoiding the considerably more costly optimisation of a non-stationary velocity field (Beg and Khan, 2007; De Craene et al., 2012).

In section 5.4, we assess the performance of the presented image registration approaches for the deformable alignment of imaged neonatal brain structures. We compare our proposed unbiased variant of the stationary velocity free-form deformation (SVFFD) algorithm to the related works of Modat et al. (2012b) and Lorenzi et al. (2013), and evaluate the performance for different alternative choices of numerical integration schemes and approximations of the similarity gradient on publicly available brain image datasets. To enable comparison to other evaluation studies that are based on adult brain images only, the methods are evaluated on both neonatal and adult brain images. The focus of our evaluation is on the assessment of the registration performance for neonatal brain images.

For the construction of brain atlases of the normal developing brain, we are especially interested in a computationally efficient registration, which produces diffeomorphic and inverse consistent spatial mappings between neonatal brain anatomies, without sacrificing anatomical overlap by enforcing overly smooth deformations. The SyN approach implemented in ANTs (Avants et al., 2008) has been shown by different researchers to be among the currently best performing registration methods for adult brain images (Klein et al., 2009a; Ou et al., 2014). Although it produces inverse consistent diffeomorphisms, it is excluded from our comparison because of the long runtime of up to ten hours (Tustison and Avants, 2013), and the greedy optimisation, which does not result in a (non-stationary) velocity field from which an average diffeomorphism can be computed. Further, due to the significant MR appearance changes during early brain development, methods based on SSD, such as LDDMM as proposed by Beg et al. (2005), and the symmetric log-domain Demons approach of Vercauteren et al. (2008), are unsuitable for the deformable registration of neonatal brain images, and therefore not considered in our evaluation.

Our results show that both displacement- and velocity-based parametric approaches achieve an overlap comparable to other state-of-the-art registration tools. MIRTK performs better than IRTK in terms of accuracy, inverse consistency, and runtime, with results similar to the very fast NiftyReg implementations of related approaches. Our proposed symmetric variant of the SVFFD algorithm achieves a completely unbiased alignment of neonatal brain structures with higher overlap than both NiftyReg and LCC LogDemons.

5.1. MIRTK: Medical Image Registration ToolKit

MIRTK builds on the data structures of IRTK. The registration component has, however, been newly implemented to facilitate an energy function oriented design based on analytic gradient calculations. An extended abstraction of the transformations, which allows for the delegation of model specific calculations through the concept of polymorphism, enabled the development of a universal registration tool. This is in contrast to IRTK, where each transformation model corresponds to a separate registration command specialized for this model. This new universal registration tool can not only register two images using different transformation models, but also (1) take advantage of multiple available modalities, (2) include point constraints such as landmarks, (3) separately or simultaneously register surface meshes, (4) combine different image dissimilarity measures and transformation constraints, (5) supports asymmetric and symmetric energy formulations, (6) and can perform longitudinal registration of a temporal image sequence.

The main features on which this flexible tool is based are highlighted in the following.

5.1.1. Modularity

Similar to other component-based toolboxes such as ITK and `elastix`, MIRTK can be extended with further transformation models, data matching terms, and transformation constraints, without having to re-implement components not specific to this feature. For example, the implementation of a new image dissimilarity need only provide a function for the evaluation of the criterion, and a function to compute its non-parametric gradient, i.e., the voxel-wise derivatives with respect to a change of the transformation. The components performing the resampling of the images given the current transformation and chosen interpolation scheme, and the evaluation of the gradient of this energy term with respect to the parameters of the selected transformation model, are the same for all image dissimilarity measures. More specifically, the calculation of the derivatives with respect to the parameters $\boldsymbol{c}$ is decoupled using the chain rule, i.e., $\nabla_{\boldsymbol{c}}\mathcal{D}_\star = (\nabla_{\boldsymbol{\Phi}}\mathcal{D}_\star)\,(\nabla_{\boldsymbol{c}}\boldsymbol{\Phi})$, where the image dissimilarity component for $\mathcal{D}_\star$ calculates $\nabla_{\boldsymbol{\Phi}}\mathcal{D}_\star$, and then delegates the remaining computation of $\nabla_{\boldsymbol{c}}\mathcal{D}_\star$ to the transformation model component.

5.1.2. Configurable energy function

As a novelty compared to other registration toolboxes, the objective of the registration tool is configured via a user-defined energy function based on a simplified math expression. This energy function is a sum of individual terms implemented by specific components. An energy term can either measure the dissimilarity of two images, determine the distance of point constraints, measure the distance of two surface meshes, or enforce a transformation constraint. For example, the expression corresponding to the original FFD algorithm is

```
Energy function = NMI(I(1), I(2) o T) + 0.001 BE(T)
```

where `T` denotes the FFD, `I(1)` and `I(2)` are the first and second input image, `I(2) o T` symbolises function composition, and `NMI` and `BE` correspond to the normalised mutual information (NMI) and bending energy (BE) terms. The registration tool therefore facilitates a number of registration tasks. Currently, energy terms for rigid and non-rigid single- or multi-channel registration based on the FFD and SVFFD model are supported. Example configurations are given in appendix C. The MIRTK can be extended with further components to realise more energy configurations.

5.1.3. Fast image registration

The runtime of MIRTK compares favourably to other FFD-based registration tools, as demonstrated in section 5.4. This efficiency can be attributed to the following characteristics: (1) the calculation of the energy gradient based on analytic derivatives as opposed to finite difference approximations, (2) images are cropped to avoid computations for background, (3) the transformation is only defined for the domain for which a local support region of the basis functions overlaps with image foreground, (4) the lower number of control points compared to the number of voxels, (5) a greedy line search, which is stopped when a predefined number of steps did not further reduce the energy value, and (6) a multi-threaded implementation of most components.

Besides the internal cropping of the images and early line search termination, most of these characteristics are shared with the highly optimised and fast NiftyReg implementation (Modat et al., 2010), which does however not feature a configurable software design.

5.2. Soft invertibility constraints

In the following, we briefly revisit the diffeomorphic FFD algorithm proposed by Rueckert et al. (2006) summarised in section 2.3 based on soft constraints of the Jacobian determinants of the displacements, with some of our modifications as implemented in MIRTK. We evaluate this FFD approach for neonatal brain image registration in section 5.4, and demonstrate that the addition of these soft constraints successfully reduces folding in the deformations. It may fail, however, to avoid these entirely without a considerable reduction in overlap of anatomical structures. This algorithm is furthermore based on the minimisation of an asymmetric energy function, and is therefore biased towards the target anatomy. The resulting transformations may thus not be inverse consistent, and invertibility is not guaranteed. Our unbiased diffeomorphic method is detailed afterwards.

The asymmetric energy function of the diffeomorphic FFD method is given by

$$\begin{aligned} E_{\text{FFD}}(I_0, I_1; \boldsymbol{u}_{\text{FFD}}) = \; & \mathcal{D}_\star\left(I_0, I_1 \circ \boldsymbol{\Phi}_{\text{FFD}}\right) + w_{\text{L}}\; \mathcal{S}_{\text{LE}}\left(\boldsymbol{u}_{\text{FFD}}\right) + w_{\text{B}}\; \mathcal{S}_{\text{BE}}\left(\boldsymbol{u}_{\text{FFD}}\right) \\ & + w_{\text{J}}\; \mathcal{P}_{\text{VP}}\left(\boldsymbol{\Phi}_{\text{FFD}}\right) + w_{\text{T}}\; \mathcal{P}_{\text{TP}}\left(\boldsymbol{\Phi}_{\text{FFD}}\right) \end{aligned} \tag{5.1}$$

where $I_{0,1}$ are the images to be registered, and $\boldsymbol{u}_{\text{FFD}}$ is the cubic B-spline function of local displacements of $\boldsymbol{\Phi}_{\text{FFD}}$ given by equation (2.71) with spline coefficients $\boldsymbol{c}$. In addition to (or instead of) the bending energy, $\mathcal{S}_{\text{BE}}$, and soft topology preservation constraint, $\mathcal{P}_{\text{TP}}$, the linear elastic potential, $\mathcal{S}_{\text{LE}}$, and local "volume preservation" term, $\mathcal{P}_{\text{VP}}$, which not only penalises small (and non-positive) Jacobian determinants, $J_{\boldsymbol{\Phi}} = |D_{\boldsymbol{x}}\boldsymbol{\Phi}|$, but also $J_{\boldsymbol{\Phi}} \gg 1$, can be employed, to enforce the required smoothness, and to reduce folding. The bending energy and $\mathcal{P}_{\text{VP}}$ are also utilised by Modat et al. (2012a) in a symmetric FFD approach, in addition to a soft inverse consistency constraint (see section 2.4). The soft constraints are weighted by the non-negative constants $w_\star$ relative to the image matching term. Of these, the original algorithm of Rueckert et al. (2006) utilised only the bending energy and soft topology preservation constraint, which penalises Jacobian determinants smaller than γ, where $0 < \gamma < 1$. The bending energy of the deformation is given by equation (2.32), whereas the definition of the linear elastic potential is given in equation (2.27).

The Jacobian-based penalty terms were previously defined in section 2.3, i.e.,

$$\begin{aligned}\mathcal{P}_{\mathrm{VP}}(\boldsymbol{\Phi}) &= \int_{\Omega} \left(\log\left(J_{\boldsymbol{\Phi}}(\boldsymbol{x})\right)\right)^2 \, d\boldsymbol{x} \\ \mathcal{P}_{TP}(\boldsymbol{\Phi}) &= \int_{\Omega} \mathbb{1}_{\{\boldsymbol{x}|J_{\boldsymbol{\Phi}}(\boldsymbol{x})<\gamma\}}(\boldsymbol{x}) \left(\frac{J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2}{\gamma^2} + \frac{\gamma^2}{J_{\boldsymbol{\Phi}}(\boldsymbol{x})^2} - 2\right) d\boldsymbol{x}\end{aligned} \tag{5.2}$$

where $\mathbb{1}_{\mathcal{X}}(\boldsymbol{x})$ is an indicator function that is one for points $\boldsymbol{x} \in \mathcal{X}$, and zero otherwise. It should be noted that $\mathcal{P}_{VP}$ is undefined for non-positive values. Similarly, $\mathcal{P}_{TP}$ is undefined for $J_{\boldsymbol{\Phi}} = 0$, and it has the undesired property that it may prevent already negative Jacobian determinants from being corrected for, since the penalty increases for negative values with decreasing magnitude. It attains its minimum for $J_{\boldsymbol{\Phi}} = \pm\gamma$.

Different implementations handle the cases of non-positive determinant values differently. An approach found in IRTK is to clamp the Jacobian determinant below a small threshold ϵ, e.g., $\epsilon = 10^{-6}$. This has the disadvantage that the value of the energy function does not change as long as the Jacobian determinant is below this threshold. This in turn may result in a premature termination of the iterative optimisation, because the total energy may not have improved when stepping along the gradient direction.

A different strategy to avoid non-positive Jacobian determinant values for the evaluation of $\mathcal{P}_{\mathrm{VP}}$ is taken by NiftyReg. After every gradient step, non-positive Jacobian determinant values are detected, and an iterative unfolding step, which aims to remove non-positive Jacobian determinant values, based on a separate maximisation of the determinant values less or equal than zero, is performed (Modat et al., 2012b). This strategy may fail to resolve all non-positive Jacobian determinant values, which in turn may result in undefined values (`NaN`) in the computations. This in consequence also results in a premature termination.

In MIRTK, we define both Jacobian-based penalty terms as piecewise functions, where the pieces above a small positive threshold ϵ are identical to previously defined penalties, but the penalty below ϵ is a linear continuation, which increases with decreasing Jacobian determinant. In order to obtain a $\mathbb{C}^1$-continuous penalty term, the slope of this linear penalty is equal to $\nabla_J \mathcal{P}_\star$ at $J = \epsilon$, and the intercept with the vertical axis at $J = 0$ is determined such that the two pieces below and above ϵ attain the same value for $J = \epsilon$, as shown in figure 5.1. Using $\mathcal{P}_{\mathrm{TP}}$ in addition to $\mathcal{P}_{\mathrm{VP}}$ avoids a separate unfolding step, and

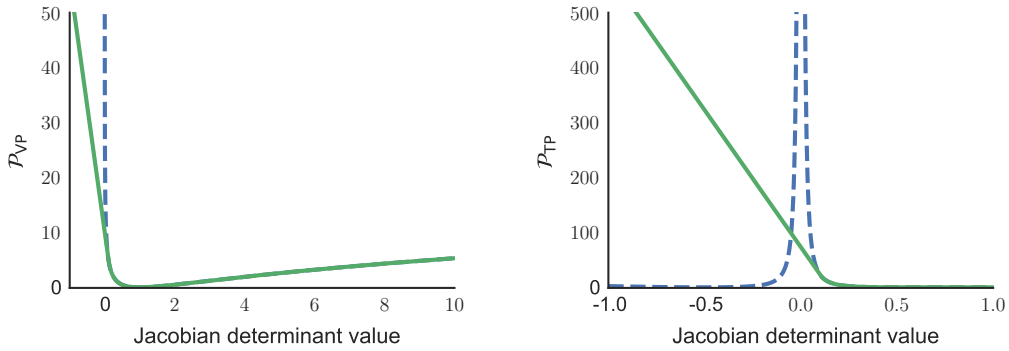


Figure 5.1.: Jacobian-based penalty functions. The linear continuation (solid, green) of the $\mathcal{P}_{\text{VP}}$ (left) and $\mathcal{P}_{\text{TP}}$ (right) penalty functions (dashed, blue) that are undefined for non-positive or zero values, respectively, results in a monotonically increasing penalty for increasingly negative determinant values. The used thresholds are $\epsilon = 0.1$ and $\gamma = 0.5$.

the resulting registration task is determined by a single well-defined energy minimisation.

5.3. Unbiased free-form deformation

As discussed in section 2.4, an asymmetric energy may lead to substantial bias towards the chosen target anatomy. This was noted first by Cachier and Rey (2000), and shown for hippocampal volume measurements by Yushkevich et al. (2010). This bias can be removed using a symmetric formulation. In case of the FFD model, this requires the simultaneous optimisation of forward and backward mappings, and the minimisation of an inverse consistency error (Modat et al., 2012b). A transformation model based on a SVF parametrisation has a number of advantages over the displacement-based FFD. The inverse is hereby given directly by negating the velocities. This enables symmetric formulations, which do not require a soft inverse consistency constraint. A SVF model results in an easier to implement and more efficient algorithm than symmetric approaches based on LDDMM such as geodesic shooting, which requires the costly solution of the EPDiff PDEs at each step (cf. section 2.2.3). We therefore adopt the parametric SVFFD model of Modat et al. (2012a), and propose a symmetric energy based on a single parametric spline function for unbiased brain image registration. Our proposed SVFFD algorithm produces optimally inverse consistent deformations, i.e., when the images are exchanged, only the sign of the velocities changes.

5.3.1. Parametric velocity field

The SVFFD model is given by a cubic B-spline function equivalent to the FFD model defined in section 2.2.4, where the B-spline function represents a SVF $\boldsymbol{v}(\boldsymbol{x})$ instead of a displacement field $\boldsymbol{u}(\boldsymbol{x})$. The diffeomorphic mapping that is generated by this SVF is given by the group exponential map introduced in section 2.2.3, i.e.,

$$\boldsymbol{\Phi}^t_{\text{SVFFD}} = \exp\left(t \underbrace{\sum_{l,m,n=0}^{3} w_l(\tilde{x}_1) w_m(\tilde{x}_2) w_n(\tilde{x}_3) \boldsymbol{c}_{\{i_1+l+n_1^c(i_2+m+n_2^c\,(i_3+n))\}}}_{=:\,\boldsymbol{v}(\boldsymbol{x})}\right) \tag{5.3}$$

Here, the exponent t is the scaling factor of the one-parameter subgroup of diffeomorphic mappings generated by the SVF, with $t = 1$ corresponding to the SVFFD, which aligns the registered anatomies, and $t = -1$ resulting in the inverse mapping. The local coordinates of the point $\boldsymbol{x} \in \Omega$ with respect to the control point lattice of size $n_\text{c} = \prod_{i=1}^{d} n_i^c$ are given by a rotation $\boldsymbol{R}$, translation $\boldsymbol{t}$, and anisotropic scaling with inverse scaling factors corresponding to the uniform spacing s_i of control points in the i-th dimension of the rotated lattice domain, i.e., $\hat{\boldsymbol{x}} = \text{diag}(s_1, \ldots, s_d)^{-1}\boldsymbol{R}\boldsymbol{x} + \boldsymbol{t}$. The coordinates of $\boldsymbol{x}$ relative to the edges of the cell of the lattice that it is contained in are then given by $\tilde{\boldsymbol{x}} = \hat{\boldsymbol{x}} - \lfloor \hat{\boldsymbol{x}} \rfloor$, and $\boldsymbol{i} = \lfloor \hat{\boldsymbol{x}} \rfloor - 1$ is the index of the "leftmost" control point, whose weight is non-zero. The weights $w_l(\tilde{x}_i)$ are

$$\begin{aligned} w_0(\tilde{x}_i) &= B_3(\tilde{x}_i + 1) \\ w_1(\tilde{x}_i) &= B_2(\tilde{x}_i + 0) \\ w_2(\tilde{x}_i) &= B_1(\tilde{x}_i - 1) \\ w_3(\tilde{x}_i) &= B_0(\tilde{x}_i - 2) \end{aligned} \tag{5.4}$$

with the polynomials $B_i(r)$ of the cubic B-spline function as defined in section 2.2.4.

The exponential map may be obtained using a numeric integration scheme to determine the solution of the stationary ODE given by equation (2.41), or using the more efficient SS approach given by equation (2.42). To further reduce the computational cost of the integration, a SS on the coarser control point lattice was proposed by Modat et al. (2012b).

This model inherits some of the benefits of the FFD model, namely an implicit smooth interpolation of the velocity field, and a parametric representation, which reduces the number of DOFs, and enables the analytic computation of derivatives. Note that even though the support of the basis functions is limited, a control point may still affect the deformation of a point outside its local support region. This can be explained by considering the trajectory of a particle that is subjected to the flow induced by the SVF. Even when this trajectory originates at a position outside the local support region of a given control point, it may enter it at a later time point $t > 0$.

In Modat et al. (2012b), it is noted that the SVF must be sufficiently smooth for the exponential to be a diffeomorphic mapping. Regularisation terms based on the Jacobian determinant of the SVF are therefore added. However, given the cubic B-spline parametrisation of the SVF, we can always find a maximum step length (or number of squaring steps) of the numerical integration, which guarantees positive Jacobian determinants of the scaled velocity field. The Jacobian determinant of $\boldsymbol{\Phi}^t_{\text{SVFFD}}$ at $\boldsymbol{x}$ is the product of Jacobian determinants along the trajectory of the particle originating at this point. By choosing the integration step such that the maximum scaled coefficient of the spline function is less than 0.4 times the minimum control point spacing, all factors are positive and thereby a diffeomorphic mapping is ensured (see section 2.3). This does not consider imprecisions of the numerical integration. When using an approximation scheme such as SS, and in particular the faster SS on a coarse control point lattice, folding and inconsistencies between $\exp(\boldsymbol{v})$ and $\exp(-\boldsymbol{v})$ may be observed, due to the interpolation errors resulting from the composition of discrete deformation fields.

However, an inverse mapping, which only deviates from the true inverse by aforementioned numerical integration errors, can always be computed. This is in contrast to a parametrisation of the spatial mapping based on displacements, where a numeric algorithm is needed to approximate the inverse either based on a scattered data interpolation or using a fixed point iteration. When the displacement field has regions of non-positive Jacobian determinants, regularisation may be needed to approximate a suitable inverse. Numerical errors caused by the integration of a SVF may be reduced by avoiding extreme Jacobian determinant values as suggested by Modat et al. (2012b).

5.3.2. Symmetric energy function

A SVF parametrisation of the deformation enables symmetric formulations, which do not require the numerical inversion or a simultaneous optimisation of an inverse. Similar to the consistent-midpoint-cost proposed by Beg and Khan (2007) for non-stationary LDDMM (see section 2.4), the energy function of the proposed SVFFD algorithm is

$$\begin{aligned} E_{\mathrm{SVFFD}}(I_0, I_1; \boldsymbol{v}) = \; & \mathcal{D}_\star \left(I_0 \circ \exp\left(-\boldsymbol{v}/2\right), I_1 \circ \exp\left(\boldsymbol{v}/2\right)\right) \\ & + w_{\mathrm{L}} \; \mathcal{S}_{\mathrm{LE}}\left(\boldsymbol{v}\right) + w_{\mathrm{B}} \; \mathcal{S}_{\mathrm{BE}}\left(\boldsymbol{v}\right) + \mathcal{P}_{\mathrm{JAC}}\left(\boldsymbol{v}\right) \end{aligned} \tag{5.5}$$

where $I_{0,1}$ are the images to be registered, and $\boldsymbol{v}$ is the SVF spline function of the diffeomorphic transformation $\boldsymbol{\Phi}_{\mathrm{SVFFD}}$. The image dissimilarity $\mathcal{D}_\star$ denotes one of the terms defined in section 2.1. In case of LNCC, it is equivalent to the data matching term of the non-parametric LCC LogDemons approach (Lorenzi et al., 2013). To explicitly enforce smoothness of the SVF, regularisation terms based on first and second order derivatives are employed. These are the elastic potential $\mathcal{S}_{\mathrm{LE}}(\boldsymbol{v}) = \mathrm{S}_{\mathrm{elastic}}(\boldsymbol{v})$ of the spline-based SVF similar to the non-parametric viscous-fluid model discussed in section 2.2.3, and the bending energy $\mathcal{S}_{\mathrm{BE}}(\boldsymbol{v})$ used in the classic FFD algorithm of Rueckert et al. (1999), but applied here to the SVF. In addition to these smoothness constraints, penalty terms based on the Jacobian determinants of the SVF can be added (see also section 2.3). The non-negative weights w_{L} and w_{B} determine the trade-off between the smoothness imposed on the SVF and the image matching term. A weighted sum of Jacobian-based penalty terms is denoted by $\mathcal{P}_{\mathrm{JAC}}(\boldsymbol{v})$. It should be noted that the bending energy and linear elasticity constraints are symmetric due to the squaring of partial derivatives, i.e., $\mathcal{S}_{\mathrm{BE}}\left(\boldsymbol{v}\right) = \mathcal{S}_{\mathrm{BE}}\left(-\boldsymbol{v}\right)$ and $\mathcal{S}_{\mathrm{LE}}\left(\boldsymbol{v}\right) = \mathcal{S}_{\mathrm{LE}}\left(-\boldsymbol{v}\right)$. The Jacobian-based penalties are not symmetric, however, and therefore must be evaluated for both the positive and negative SVF, i.e.,

$$\mathcal{P}_{\mathrm{JAC}}(\boldsymbol{v}) = \frac{w_{\mathrm{J}}}{2}\Big(\mathcal{P}_{\mathrm{VP}}(\mathrm{Id}+\boldsymbol{v}) + \mathcal{P}_{\mathrm{VP}}(\mathrm{Id}-\boldsymbol{v})\Big) + \frac{w_{\mathrm{T}}}{2}\Big(\mathcal{P}_{\mathrm{TP}}(\mathrm{Id}+\boldsymbol{v}) + \mathcal{P}_{\mathrm{TP}}(\mathrm{Id}-\boldsymbol{v})\Big) \tag{5.6}$$

where $\mathcal{P}_{\mathrm{VP}}$ and $\mathcal{P}_{\mathrm{TP}}$ are the penalty terms defined by equation (5.2), applied to the deformation corresponding to the instantaneous velocity field and its approximate inverse.

5.3.3. Energy minimisation

The energy given by equation (5.5) is minimised using Polak-Ribiére's variant of conjugate gradient descent (Press et al., 2007) with an inexact line search adopted from NiftyReg. The line search greedily steps along the gradient direction, where the length of the greedy steps is increased by 10% while the energy value decreases, or reduced by 50% otherwise. It is stopped when either a maximum number of iterations is exceeded, or the step length is outside a predefined interval. In addition to NiftyReg's implementation, this greedy line search is terminated prematurely to avoid unnecessary energy function evaluations, when a specified number of consecutive greedy steps did not reduce the energy value. The total length of each gradient step is equal the sum of the individual greedy steps.

The derivatives of the penalty terms with respect to the coefficients of the SVFFD are equivalent to the derivations for the displacement-based FFD model. The derivatives of $\mathcal{S}_{\mathrm{BE}}$ with respect to the spline coefficients are given in equation (B.19). A similar analytic derivation yields the gradient of $\mathcal{S}_{\mathrm{LE}}$ with respect to the coefficients $\boldsymbol{c}$ of $\boldsymbol{\Phi}_{\mathrm{SVFFD}}$.

The gradient vectors of the Jacobian penalty terms with respect to $\boldsymbol{c}$ are

$$\frac{\partial \mathcal{P}_\star}{\partial \boldsymbol{c}} = \int_\Omega Df\left(J_{\pm\boldsymbol{v}}(\boldsymbol{x})\right) \frac{\partial J_{\pm\boldsymbol{v}}(\boldsymbol{x})}{\partial \boldsymbol{c}}\, d\boldsymbol{x} \tag{5.7}$$

where

$$Df(J) = \begin{cases} \frac{2}{J}\log(J) & \text{for } \mathcal{P}_\star = \mathcal{P}_{\mathrm{VP}} \\ 2\left(\frac{J}{\gamma^2} - \frac{\gamma^2}{J^3}\right) & \text{for } \mathcal{P}_\star = \mathcal{P}_{\mathrm{TP}} \end{cases} \tag{5.8}$$

The partial derivatives $\frac{\partial J_{\pm\boldsymbol{v}}}{\partial c}$ of the determinants with respect to each control point coefficient are obtained using Jacobi's formula as in Modat et al. (2012b), i.e.,

$$\frac{\partial J_{\pm\boldsymbol{v}}}{\partial c_i} = \mathrm{tr}\left(\mathrm{adj}\left(\mathrm{Id} \pm D_{\boldsymbol{x}}\boldsymbol{v}\right)\left(\mathrm{Id} \pm \frac{\partial D_{\boldsymbol{x}}\boldsymbol{v}}{\partial c_i}\right)\right) \tag{5.9}$$

where the derivative of the Jacobian matrix $D_{\boldsymbol{x}}\boldsymbol{v}$ of the cubic B-spline function $\boldsymbol{v}$ given by equation (B.4) with respect to the spline coefficients is given in equation (B.8).

Using the chain rule, the derivatives of the image dissimilarity can be written as

$$\frac{\partial \mathcal{D}_\star}{\partial \boldsymbol{c}} = \int_\Omega \left(\frac{\partial \mathcal{D}_\star}{\partial I_0} \frac{\partial I_0}{\partial \boldsymbol{\Phi}_0} \frac{\partial \boldsymbol{\Phi}_0}{\partial \boldsymbol{v}} + \frac{\partial \mathcal{D}_\star}{\partial I_1} \frac{\partial I_1}{\partial \boldsymbol{\Phi}_1} \frac{\partial \boldsymbol{\Phi}_1}{\partial \boldsymbol{v}} \right) \frac{\partial \boldsymbol{v}}{\partial \boldsymbol{c}} \, d\boldsymbol{x} \tag{5.10}$$

where $\boldsymbol{\Phi}_0 = \exp(-\boldsymbol{v}/2)$, and $\boldsymbol{\Phi}_1 = \exp(\boldsymbol{v}/2)$. The first factor corresponds to the sum of voxel-wise non-parametric gradients of the image dissimilarity measure with respect to a change in the SVF for each reference point $\boldsymbol{x} \in \Omega$. The partial derivatives $\nabla_{I_{0,1}} \mathcal{D}_\star$ are given for NMI based on a Parzen window estimation of the joint probability distribution using cubic B-spline kernels in Modat et al. (2009), and for LNCC in Avants et al. (2008). The partial derivatives of the spline function $\boldsymbol{v}$ with respect to the spline coefficients is equivalent to the derivatives for the displacement-based FFD model.

The numerical approximation of the derivatives of the exponential maps $\nabla_{\boldsymbol{v}} \boldsymbol{\Phi}_{0,1}$ depends on the integration method. In MIRTK, we implemented different variants of explicit Runge-Kutta integration schemes for the computation of the exponential map and its derivatives. The formulas for the Euler integration are equivalent to the temporal diffeomorphic free-form deformation (TDFFD) algorithm of De Craene et al. (2012), which is a variant of LDDMM that uses a B-spline parametrisation of a non-stationary velocity field for longitudinal registration of cardiac images. While more accurate than SS, a Runge-Kutta integration for computing the voxel-wise energy gradient is too time consuming. A computation using SS is outlined in Ashburner (2007) and Modat et al. (2012b). Given a recursive definition of the exponential map by $\boldsymbol{\Phi}^{(k)} = \boldsymbol{\Phi}^{(k-1)} \circ \boldsymbol{\Phi}^{(k-1)}$, where k is the current number of in total n squaring steps, and $\boldsymbol{\Phi}^{(0)} = \mathrm{Id} + 2^{-n}\boldsymbol{v}$ (see also section 2.2.3), the partial derivatives with respect to the SVF are given by $\nabla_{\boldsymbol{v}} \boldsymbol{\Phi}^{(0)} = 2^{-n}$ and

$$\frac{\partial \boldsymbol{\Phi}^{(k)}(\boldsymbol{x})}{\partial \boldsymbol{v}} = \left(\frac{\partial \boldsymbol{\Phi}^{(k-1)}(\boldsymbol{y})}{\partial \boldsymbol{y}} \frac{\partial \boldsymbol{\Phi}^{(k-1)}(\boldsymbol{x})}{\partial \boldsymbol{v}} + \frac{\partial \boldsymbol{\Phi}^{(k-1)}(\boldsymbol{y})}{\partial \boldsymbol{v}} \right) \Bigg|_{\boldsymbol{y} = \boldsymbol{\Phi}^{(k-1)}(\boldsymbol{x})} \tag{5.11}$$

We instead propose a simpler approximation based on the assumption that $\nabla_{\boldsymbol{v}} \boldsymbol{\Phi}$ is small, and can be considered to be an instantaneous velocity field. The same is assumed for the update of the viscous fluid registration (Christensen et al., 1996) and diffeomorphic Demons (Vercauteren et al., 2007, 2008, 2009). The current SVFFD is composed with this velocity field using the BCH formula introduced in section 2.2.3, i.e.,

$$t_i \frac{\partial \mathcal{D}_\star}{\partial \boldsymbol{v}} = \boldsymbol{w}_i + \frac{1}{2}\left[t_i\boldsymbol{v}, \boldsymbol{w}_i\right] + \frac{1}{12}\left[t_i\boldsymbol{v}, \left[t_i\boldsymbol{v}, \boldsymbol{w}_i\right]\right] + \frac{1}{12}\left[\left[t_i\boldsymbol{v}, \boldsymbol{w}_i\right], \boldsymbol{w}_i\right] \tag{5.12}$$

where $t_0 = -0.5$, $t_1 = 0.5$, and $\boldsymbol{w}_i = \nabla_{\boldsymbol{\Phi}_i}\mathcal{D}_\star$. The first term of the formula is omitted to obtain the difference of the SVF before and after the update. This computation is similar to the update step of log-domain Demons (Vercauteren et al., 2008). By interpolating all vector fields (including the Lie bracket $[\cdot,\cdot]$) by cubic B-spline functions with control points defined on the same lattice as the parametric SVFFD, the approximated image dissimilarity gradient with respect to $\boldsymbol{c}$ is given directly by the spline coefficients of (5.12). Using the zeroth order approximation $\nabla_{\boldsymbol{v}}\mathcal{D}_\star = \frac{\boldsymbol{w}_i}{t_i}$ as in Lorenzi et al. (2013), the unbiased gradient of the symmetric energy function is given by the simplified equation

$$\frac{\partial \mathcal{D}_\star}{\partial \boldsymbol{c}} = 2\int_\Omega \left(\frac{\partial \mathcal{D}_\star}{\partial I_1}\frac{\partial I_1}{\partial \boldsymbol{\Phi}_1} - \frac{\partial \mathcal{D}_\star}{\partial I_0}\frac{\partial I_0}{\partial \boldsymbol{\Phi}_0}\right)\frac{\partial \boldsymbol{v}}{\partial \boldsymbol{c}}\, d\boldsymbol{x} \tag{5.13}$$

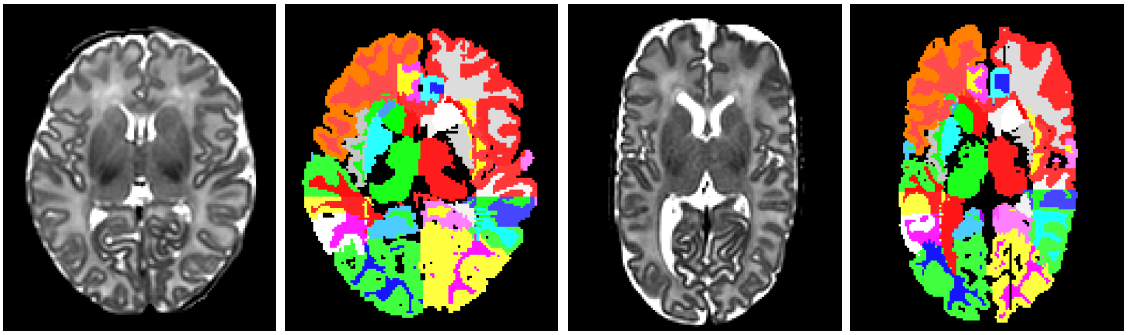
Different computations of the parametric gradient given the non-parametric similarity gradient have been proposed for the FFD algorithm. Because of the continuity of the cubic B-spline function, an analytic expression can be directly obtained, where

$$\frac{\partial \boldsymbol{v}}{\partial \boldsymbol{c}_i} = \boldsymbol{G}_{\text{FFD}}(\boldsymbol{x} - \boldsymbol{x}_i) \tag{5.14}$$

is the 3D kernel given by equation (2.67) centred at the i-th lattice point $\boldsymbol{x}_i$. The derivatives are zero outside the local support region of this B-spline kernel. The equations of the cubic B-spline FFD and its derivatives are summarised in more detail in appendix B. Tustison et al. (2009) suggested to instead find the coefficients, which approximate the similarity gradient in a least-squares sense, using the algorithm of Lee et al. (1997). Because of the original application to the modelling of geometric objects, it is referred to as "directly manipulated FFD" (DMFFD). An approximation based on separable convolution of the similarity gradient with the cubic B-spline weights was proposed by Modat et al. (2010) for a highly-parallel computation on the GPGPU. We implemented the exact gradient calculation and both approximations proposed for the FFD algorithm in MIRTK, and evaluated their influence on the accuracy and runtime of our SVFFD method.

5.4. Results

5.4.1. Brain image datasets

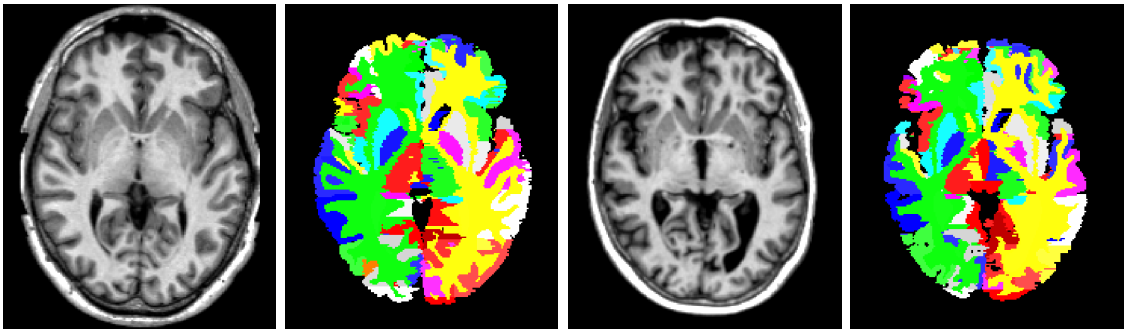


(a) ALBERT_01, 41 weeks GA (b) ALBERT_10, 40 weeks GA

Figure 5.2.: Exemplar axial slices of two brain scans from the ALBERTs dataset of (left) T2w image and (right) manual segmentation of 50 regions with 34 automatically separated cortical labels: (a) neonate scanned at 41 weeks GA, and (b) at 40 weeks GA.

ALBERTs[1] This dataset contains brain images of 15 premature-born and 5 term-born neonates, with manual annotations of 50 ROIs following the protocol defined by Gousias et al. (2012). An automatic probabilistic segmentation of cGM, which is included in the publicly available atlases provided by Makropoulos et al. (2014), has been used to split these ROIs into cortical and non-cortical labels. The resulting segmentations delineate 84 regions, of which 34 labels correspond to 17 cortical regions of left and right hemisphere, as shown in figure 5.2. The brain scans of the neonates with a mean age at scan of 40.7 weeks GA and a standard deviation of 2.0 weeks were acquired in a 3 T Philips Intera scanner. The T2w images were obtained with a fast spin echo (FSE) protocol using parameters: repetition time (TR) = 8000 ms, echo time (TE) = 160 ms, flip angle = 90°, and a 224x224 acquisition matrix with a field of view (FOV) of 22 cm. For each scan, 88–100 contiguous sagittal slices with a thickness of 2 mm were acquired. The reconstructed volumes were resampled to an isotropic voxel size of 0.86 mm, and oriented using the anterior commissure (AC)–posterior commissure (PC) line and the inter-hemispheric fissure. The images provided by Makropoulos et al. (2014) have been bias field corrected using N4 (Tustison et al., 2010). Extra-cranial intensities are zero.

[1] https://www.doc.ic.ac.uk/~am411/atlases-DrawEM.html (last accessed: Aug 17, 2017)



(a) OAS1_0061/1000, 20 years (b) OAS1_0113/1125, 83 years

Figure 5.3.: Exemplar axial slices of two brain scans from the OASIS dataset of (left) T1w image and (right) manual segmentation into 134 structures: (a) brain image of young adult scanned at the age of 20 years, (b) brain scan of older subject scanned at 83 years.

OASIS[2] This dataset consists of 30 adult brain images with manual annotations of 138 cortical and non-cortical brain structures used in the MICCAI 2012 Grand Challenge and Workshop on Multi-Atlas Labeling[3]. It contains de-faced T1w images from the "reliability" Open Access Series of Imaging Studies (OASIS) dataset (Marcus et al., 2007), whereas the manual annotations following publicly available protocols[4] were provided by Neuromorphometrics, Inc[5]. Small regions such as "vessel" and "cerebral exterior" are excluded from the evaluation. Of the remaining 134 structures, 98 ROIs correspond to 49 cortical structures of left and right hemisphere. The age at scan ranges from 18 to 90 years, with mean 34.3 years and a standard deviation of 20.7 years. Of the 30 participants, 21 were younger than 31 years at the time of scan. The provided images are motion-corrected averages of 3–4 repeated scans acquired in one session with isotropic voxel size of 1 mm. Brain masks were computed with PINCRAM (Heckemann et al., 2015), and bias field inhomogeneities corrected for using N4. Exemplar axial slices of two brain images and their manual segmentations are shown in figure 5.3.

[2] http://www.oasis-brains.org (last accessed: Aug 17, 2017)
[3] http://masiweb.vuse.vanderbilt.edu/workshop2012 (last accessed: Aug 17, 2017)
[4] http://braincolor.mindboggle.info/protocols/ (last accessed: Aug 17, 2017)
[5] http://www.neuromorphometrics.com (last accessed: Aug 17, 2017)

5.4.2. Registration evaluation

A repository containing the used REgistration PErformance Assessment Tools (REPEAT), including configuration files with parameters used for each evaluated registration software, is available publicly[6]. These use MIRTK utilities for pre-processing and quantification. Because the performance of registration methods depends on many implementation choices and parameters, all our source code and parameter settings for all compared methods are publicly available in an effort to maximise reproducibility.

For each dataset, the first image is selected as target for a spatial normalisation of all other images using first a rigid, and then an affine registration with 12 DOFs and NMI as similarity measure. Following the pre-processing of Ou et al. (2014), intensity differences between all images and the reference are reduced by a histogram matching (Nyul et al., 2000). It should be noted, however, that such pre-processing is not needed for NMI. As we encountered difficulties due to very high memory demand of one of the tools, the brain extracted images were cropped to smaller size by removing all background outside the axes-aligned minimal bounding box of the determined brain ROI. The so pre-processed images and their brain masks are the input of the different deformable registration methods.

For each method, several combinations of parameters were first explored individually using a subset consisting of the first five ALBERTs images. Information from previous adult brain registration evaluation studies (Klein et al., 2009a; Ou et al., 2014), and the documentation of each registration tool were used to guide this parameter exploration.

The quality of the pairwise inter-subject brain image registrations for a dataset consisting of n pre-processed brain scans is assessed using a number of measures. The most commonly used evaluation criterion, due to lack of ground truth deformations, is the mean pairwise overlap of annotated structures. To further quantify the similarity of intensities at each reference point across all co-registered images, as well as the invertibility and inverse consistency of the deformations, we also define voxel-wise measures, and average these within the brain ROI. Equivalent criteria were defined by the non-rigid registration evaluation project (NIREP)[7]. The specific measures used in our study are as follows.

[6] https://github.com/MIRTK/REPEAT (last accessed: Aug 17, 2017)
[7] http://www.nirep.org/registration_metrics (last accessed: Aug 17, 2017)

Mean overlap

The overlap measure quantifies the accuracy of the pairwise registrations using the available manual brain segmentations. Given the number of true positive (TP), true negative (TN), false positive (FP), and false negative (FN) label assignments for each structure, the mean Dice similarity coefficient (DSC) (Dice, 1945) of the l-th structure is given by

$$DSC(l) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{\substack{j=1 \\ j \neq i}}^{n} \frac{2\, TP_{ij}(l)}{2\, TP_{ij}(l) + FP_{ij}(l) + FN_{ij}(l)} \tag{5.15}$$

Standard deviation

The standard deviation quantifies how similar the intensity samples of all co-registered brain images are at each voxel of the selected target image, i.e.,

$$S(\boldsymbol{x}) = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \left(\tilde{I}_i(\boldsymbol{x}) - \bar{I}(\boldsymbol{x}) \right)^2} \tag{5.16}$$

where $\bar{I} = \frac{1}{n} \sum_{i=1}^{n} \tilde{I}_i$, and $\tilde{I}_i = \frac{1}{\sigma_i} (I_i - \mu_i)$ are the z-score normalised images with mean intensity μ_i and standard deviation σ_i of the intensities of the i-th co-registered image.

Jacobian determinant

For non-diffeomorphic transformation models, the number of target voxels with non-positive Jacobian determinant are recorded. The remaining values inside the brain ROI are log-transformed, and a robust range quantifying the degree of deformation is defined as the difference of 95th and 5th percentile as reported in Tustison and Avants (2013).

Inverse consistency error

The mean inverse consistency error (MICE) for target image I_i given the transformation $\boldsymbol{\Phi}_{ij}$ from I_i to I_j, and the reverse mapping $\boldsymbol{\Phi}_{ji}$ obtained with I_j as target, is given by

$$MICE_i(\boldsymbol{x}) = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^{n} \int_{\Omega} \left\| \boldsymbol{\Phi}_{ji} \circ \boldsymbol{\Phi}_{ij}(\boldsymbol{x}) - \boldsymbol{x} \right\|^2 d\boldsymbol{x} \tag{5.17}$$

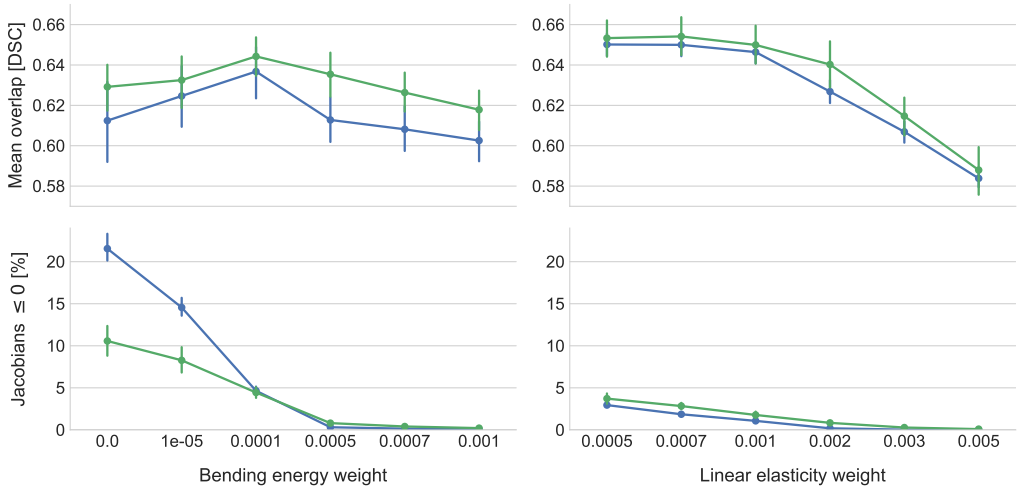


Figure 5.4.: Comparison of (top) mean overlap and (bottom) folding of 20 FFDs computed using (left) $\mathcal{S}_{\text{BE}}$ or (right) $\mathcal{S}_{\text{LE}}$ with (blue) LNCC or (green) NMI for 5 ALBERTs images.

5.4.3. Evaluation results

Image dissimilarity measure

The changing MR appearance of the developing brain during early life (see section 1.2) requires the use of a dissimilarity criterion, which is robust to these intensity variations. Of the similarity measures reviewed in section 2.1, the LNCC and NMI criteria are common choices for inter-subject brain registration. The influence of either MI or LNCC on the construction of fetal brain templates using SyN was investigated by Gholipour et al. (2014). The authors report sharper atlases for LNCC. In a first experiment, we compare the quality of inter-subject registration for NMI and LNCC in case of neonatal brain images and the FFD algorithm. The Jacobian penalties are not used, because for this experiment we are only interested in the relative performance for different similarity measures, even when folding is allowed to occur. The mean overlap of brain structures, and the percentage of non-positive Jacobian determinants inside the brain ROI, for 20 registrations using a subset of five ALBERTs images are reported in figure 5.4. Highest overlap is achieved with an average single-threaded runtime of 6.7 min with NMI using 64 bins compared to 52 min for LNCC with a box window of size 5 voxels[8]. We conclude that for the FFD algorithm, NMI is a better choice for neonatal brain image registration than LNCC.

[8]The runtime of LNCC could be improved by using separable convolutions with a Gaussian window.

Parametric gradient approximation

To assess the influence of the respective parametric similarity gradient approximations with respect to the cubic B-spline coefficients, given the voxel-wise non-parametric gradient, we computed all pairwise registrations between a subset of five ALBERTs images for each similarity gradient approximation technique:

- "Exact": Evaluation of the analytic gradient according to equation (5.14).
- "Convolution": Convolution with a 3D cubic B-spline kernel (Modat et al., 2010).
- "DMFFD": Cubic B-spline approximation of gradient field (Tustison et al., 2009).

In figure 5.5, the quantitative measures corresponding to the deformations obtained with the FFD algorithm, where NMI is used as image similarity measure and the linear elastic potential for regularisation, are compared for each respective gradient approximation. Equivalently, the results reported in figure 5.6 correspond to the diffeomorphisms computed with the SVFFD algorithm using the BCH-based gradient given by equation (5.13), where bending energy and linear elastic potential of the SVF are used for regularisation.

For both algorithms, we observe an increase in the range of log-Jacobian values for the DMFFD. This agrees with results reported by Tustison et al. (2009) for greedy SyN. In case of the FFD algorithm, a consequence of a more local deformation is an increase in folding, and a significantly higher inverse consistency error, while achieving only a slightly better mean overlap. A similar conclusion can be drawn from the results obtained with the SVFFD algorithm, where more extreme log-Jacobian values closer to zero can be observed. A more irregular SVF results in higher numerical errors in the exponential map, which can also be observed in figure 5.6 by an increased MICE. The comparison between exact parametric gradient evaluation and the convolution with a cubic B-spline kernel shows, that both yield very similar results, but the gradient approximation using a convolution filter is more efficient. We therefore use the convolution-based gradient approximation, which resembles the implementation in NiftyReg, for the remaining evaluation.

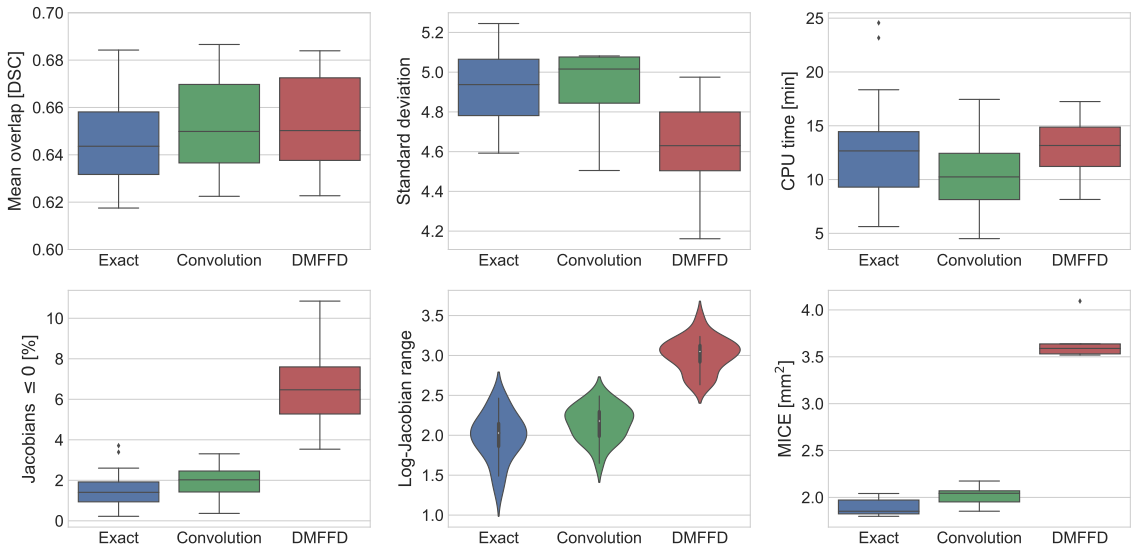


Figure 5.5.: Comparison of quantitative registration measures for exact parametric FFD energy gradient evaluation to two approximations, i.e., convolution with a cubic B-spline kernel (Modat et al., 2010), and the DMFFD approach (Tustison et al., 2009) using the spline approximation method of Lee et al. (1997). Mean values calculated from 20 pairwise registrations of 5 ALBERTs images using NMI as image similarity, and the linear elastic potential for regularisation. CPU time corresponds to a single-threaded execution.

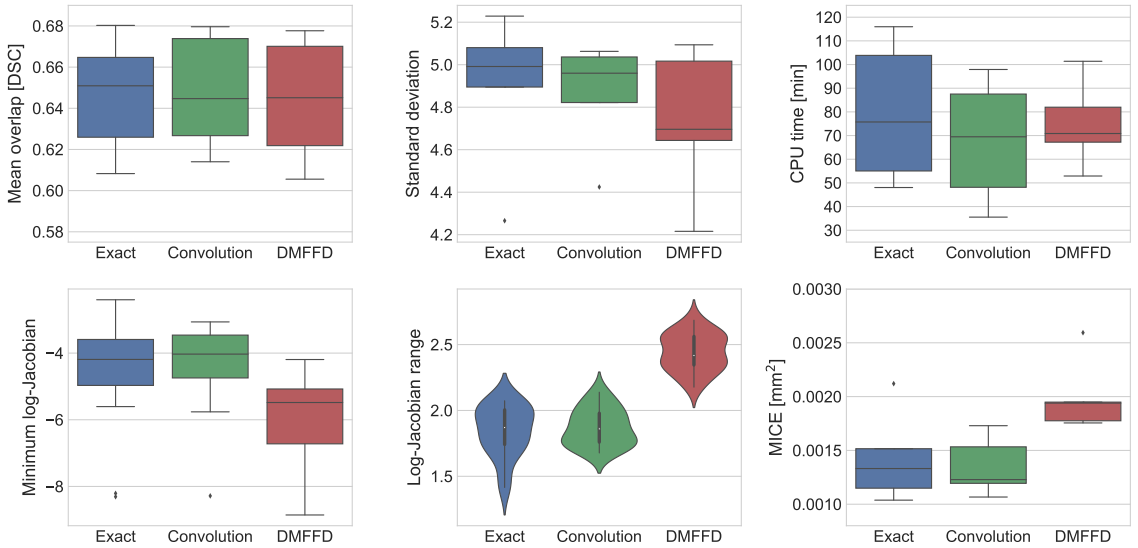


Figure 5.6.: Comparison of quantitative registration measures for exact parametric SVFFD energy gradient evaluation to two approximations, i.e., convolution with a cubic B-spline kernel (Modat et al., 2010), and the DMFFD approach (Tustison et al., 2009) using the approximation method of Lee et al. (1997). Mean values calculated from 20 pairwise registrations of 5 ALBERTs images using NMI, and bending energy and linear elastic potential of the SVF for regularisation. CPU time corresponds to a single-threaded execution.

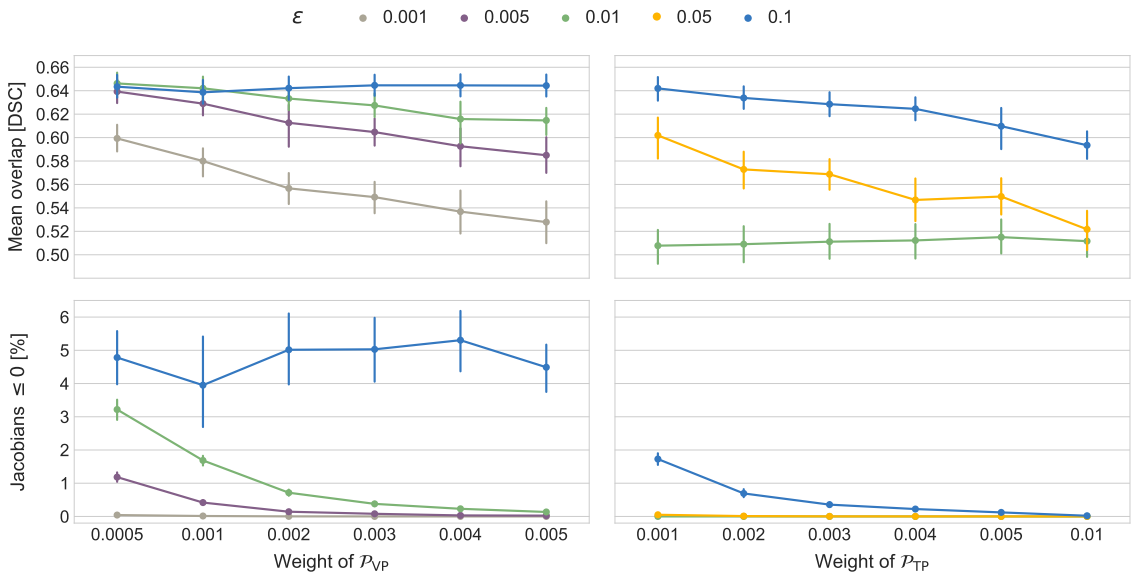


Figure 5.7.: Mean overlap and percentage of non-positive Jacobian determinant values for different (left) $\mathcal{P}_{\text{VP}}$ or (right) $\mathcal{P}_{\text{TP}}$ weights with $\gamma = 0.5$, and (hue; see legend above) different linear continuation thresholds ϵ. Mean values correspond to 20 registrations of 5 ALBERTs using the FFD algorithm with NMI and no additional smoothness constraints.

Soft invertibility constraints

In previous experiments, only the smoothness constraints $\mathcal{S}_{\text{BE}}$ and $\mathcal{S}_{\text{LE}}$ of the FFD energy have been utilised. As can be seen in figure 5.4, increasing the weight of the smoothness constraint noticeably reduces the number of target image points at which the Jacobian determinant becomes non-positive. However, increasing the smoothness may not prevent folding without a significant decrease in anatomical overlap. In the following, we assess the effectiveness of the Jacobian-based penalties $\mathcal{P}_{\text{VP}}$ and $\mathcal{P}_{\text{TP}}$ for diffeomorphic registration.

The influence of the linear continuation threshold ϵ is evaluated in figure 5.7 for 20 registrations of five ALBERTs using NMI and $w_{\text{L}} = w_{\text{B}} = 0$. A low ϵ value leads to a sudden increase of the penalty, which stops the registration before a sufficient overlap is achieved. When the threshold is too high in case of $\mathcal{P}_{\text{VP}}$, e.g., $\epsilon = 0.1$, the penalty becomes ineffective because it increases too slowly for decreasing determinants. A good balance is given for $\epsilon = 0.01$ in case of $\mathcal{P}_{\text{VP}}$, and $\epsilon = 0.1$ in case of $\mathcal{P}_{\text{TP}}$. These choices result in a smooth monotonically decreasing percentage of non-positive Jacobians with increasing penalty weight, without causing a rapid decline of the mean anatomical overlap.

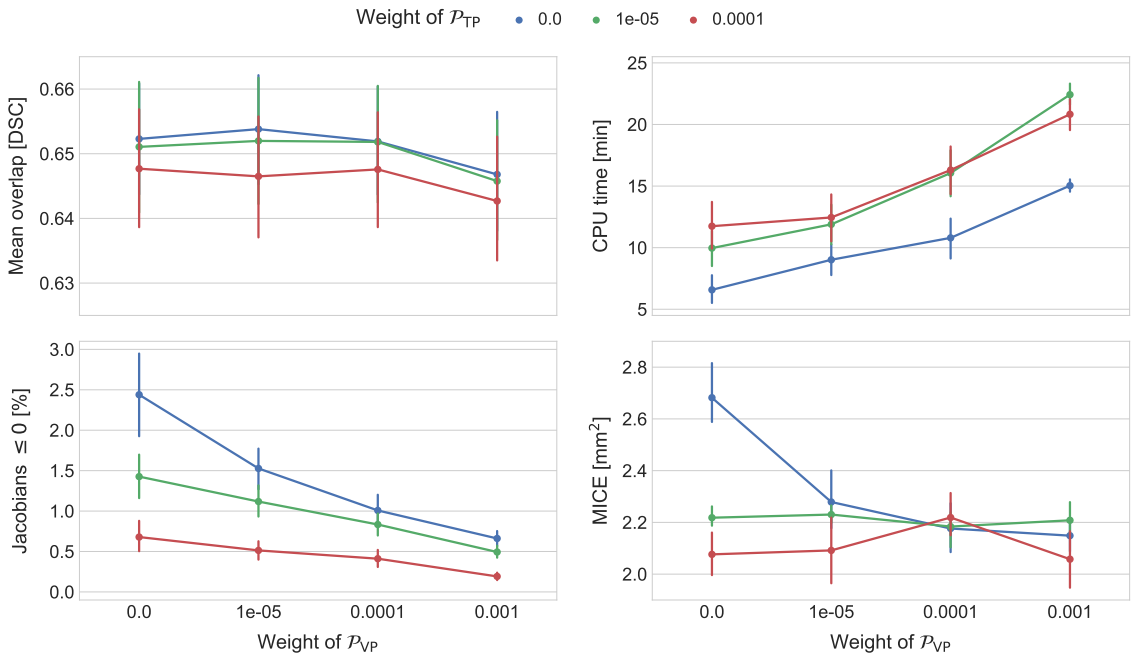


Figure 5.8.: Mean results of 20 pairwise registration of 5 ALBERTs using the FFD algorithm with NMI as image similarity, $w_L = 0.0005$, $w_B = 0.00001$, and different weights of (horizontal) $\mathcal{P}_{VP}$ with $\epsilon = 0.01$ and (hue; see legend above) $\mathcal{P}_{TP}$ with $\epsilon = 0.1$ and $\gamma = 0.5$.

The results for different weights of the Jacobian penalties using also the smoothness constraints are shown in figure 5.8. The occurrence of folding is noticeably reduced to below 0.5–1% with no or only a slight decrease in overlap. It can be observed, however, that the runtime increases significantly with increasing penalty due to slower convergence. We expect that an iterative unfolding as implemented in NiftyReg would counter this negative side-effect, by avoiding very strong penalty gradients at these points. In regions where the unfolding fails, our suggested linear continuation of the penalty terms takes effect to prevent undefined values. Combining both techniques is thus most promising to obtain diffeomorphic mappings with high certainty using the FFD algorithm.

For our work, we are however not only interested in a diffeomorphic registration of neonatal brain images, but also the earlier discussed advantageous of a velocity-based parametrisation. Having demonstrated the effectiveness of the soft Jacobian constraints in case of the FFD algorithm, we determine in the following if the proposed symmetric SVFFD method achieves similar anatomical overlap with guaranteed invertibility and inverse consistency.

Stationary velocity update

Different approximations of the velocity update field, i.e., the gradient of the exponential map with respect to the stationary velocities, and numerical integration schemes are summarised in section 5.3, i.e.,

- "Euler": Forward Euler integration (De Craene et al., 2012).
- "SS": Scaling and squaring on image lattice (Arsigny et al., 2006).
- "Fast SS": Scaling and squaring on control point lattice (Modat et al., 2012b).
- "BCH": Proposed approximation of composite SVFFD using BCH formula given by equation (5.12) similar to update step of LCC LogDemons (Lorenzi et al., 2013).

The results for 20 pairwise SVFFD registrations of five ALBERTs images using either forward Euler, dense, or sparse SS integration schemes to evaluate the exponential map with different number of integration steps are summarised in figure 5.9. The similarity gradient, with respect to the coefficients of the SVF, is hereby either computed using the respective integration scheme or the proposed zeroth order BCH approximation. The faster SS on the coarse control point lattice achieves only a sub-optimal overlap compared to the dense integration methods. It is further noted that the gradient computation based on the forward Euler integration is significantly more time consuming without resulting in a better anatomical overlap. The best results are given with the SS on the dense image lattice as integration method with 5–6 squaring steps. Moreover, using the proposed BCH update instead of the SS-based gradient evaluation (Modat et al., 2012b) significantly improves the overlap measures.

Comparison of registration methods

Previous experiments, based on a subset of five images of the ALBERTs dataset, were used to select optimal parameters for the FFD algorithm and proposed variant of the SVFFD method implemented in MIRTK. A similar parameter search using the same subset of neonatal brain images has been performed for `elastix` (Klein et al., 2010), NiftyReg (Modat et al., 2009, 2012b), and LCC LogDemons (Lorenzi et al., 2013).

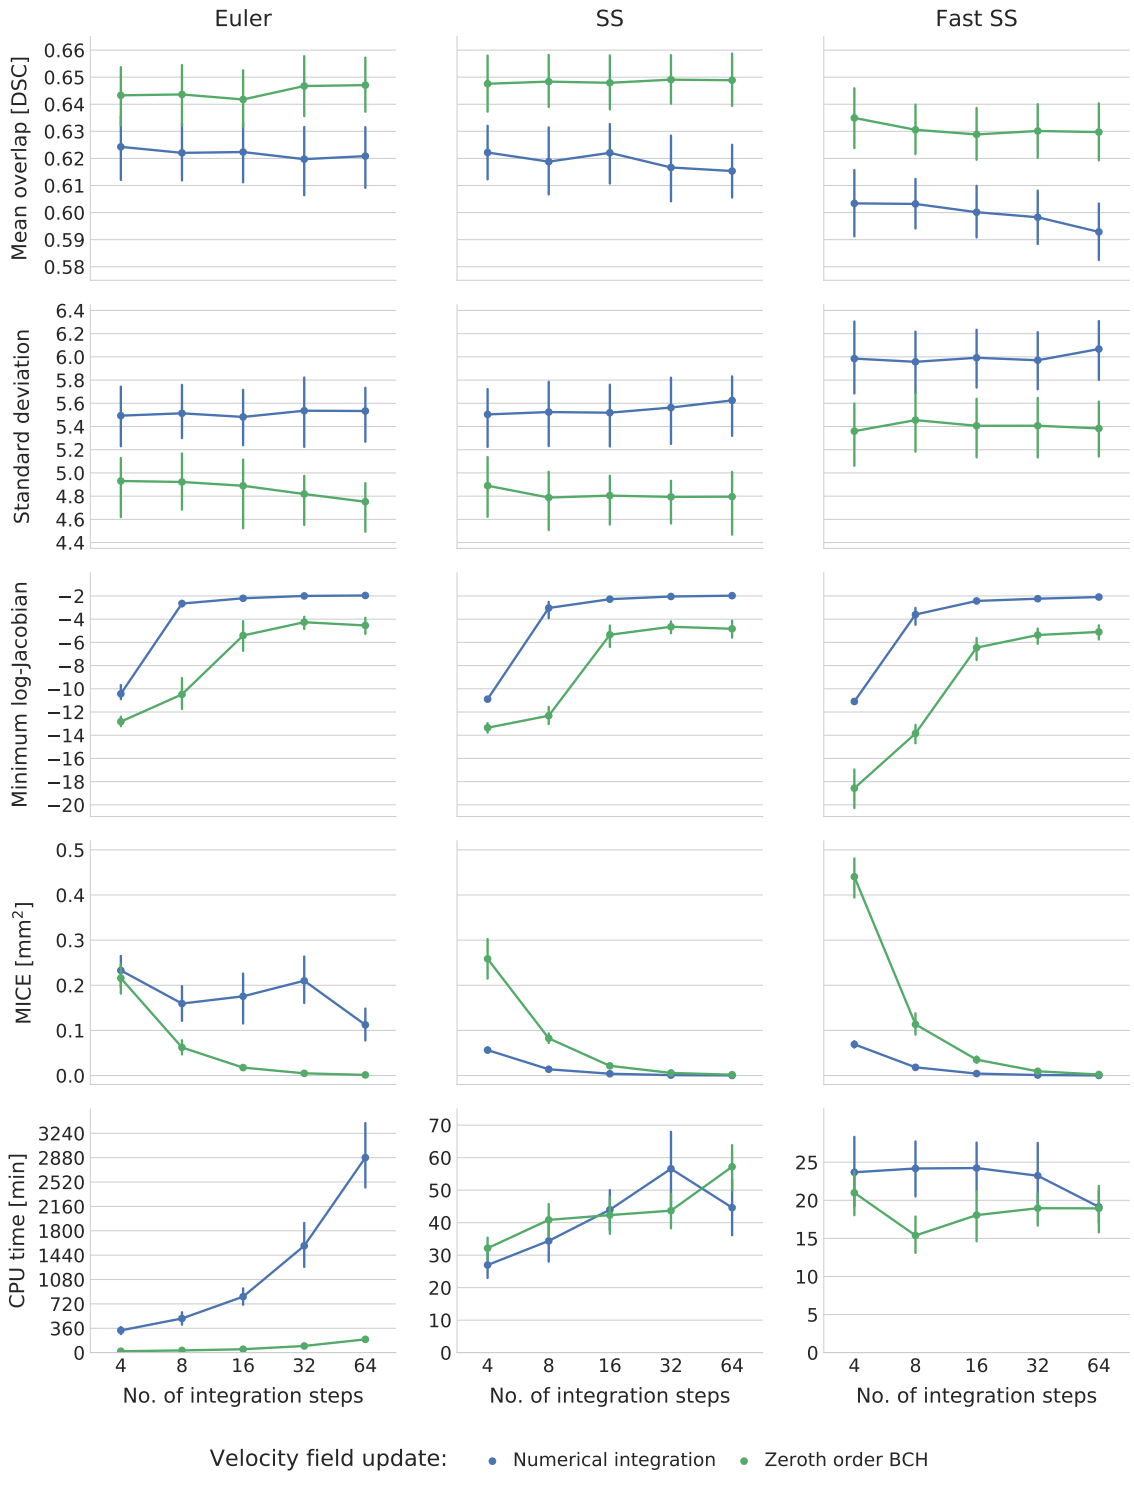


Figure 5.9.: Evaluation measures for 20 pairwise registrations of 5 ALBERTs images using the SVFFD algorithm with (columns) different numerical integration schemes with increasing number of steps, and (hue; see legend below) similarity gradient approximations.

In case of IRTK (Rueckert et al., 1999), the parameters reported in previous evaluation studies (Klein et al., 2009a; Ou et al., 2014) were used. These correspond to a steepest gradient descent based on finite difference approximations without bending energy regularisation. The use of a conjugate gradient descent significantly increases the runtime of IRTK and was not further investigated, because more recent and especially faster reference implementations of the FFD algorithm are given already by NiftyReg and `elastix`.

In case of `elastix`, only the bending energy in accordance with the original FFD algorithm of Rueckert et al. (1999) is used, because a linear elasticity constraint and Jacobian penalties are not implemented. A key feature of `elastix` is the adaptive stochastic gradient descent (Klein et al., 2009b), which drastically reduces the number of points at which the similarity gradient is evaluated at each iteration. Although the software was compiled with multi-threading enabled, it utilised only a single CPU. A better runtime is achieved without the slow bending energy evaluation of `elastix`, but with an increase in folding.

We used NMI as image similarity measure for all registrations except for LCC LogDemons, which is based on LNCC instead. The average results for all 380 pairwise registrations of the 20 ALBERTs images are summarised in table 5.1 and figure 5.10. We applied the same parameters for each method also to the 870 pairwise registrations of the 30 adult brain images of the OASIS dataset, and report the results in table 5.2 and figure 5.11.

With the selected parameters (cf. REPEAT project), all methods produce (almost) diffeomorphic mappings. Consistently good results are obtained with MIRTK and NiftyReg. Our proposed symmetric and inverse consistent SVFFD method with BCH-based gradient approximation achieves higher overlap than the F3D^2 method (Modat et al., 2012b) implemented in NiftyReg. Both SVFFD algorithms perform significantly better than LCC LogDemons, which also employs a BCH-based update of, however, a non-parametric SVF with a Demons type regularisation based on Gaussian filtering of the velocity field. This demonstrates the advantage of the SVFFD approach over log-domain Demons for diffeomorphic brain image registration. Of the methods based on the FFD algorithm, `elastix` achieves highest overlap for adult brain images, but significantly less good results for neonatal brain images, although the parameters were optimised for the latter images. The fastest registration is obtained with NiftyReg, closely followed by MIRTK.

	Mean overlap [DSC]	Jacobians ≤ 0 [%]	MICE [mm^2]	Runtime [min]
IRTK	0.543	0.011	1.34	26.5
Elastix	0.597	0.386	3.23	22.6
LCC LogDemons	0.608	0.000	0.0111	21.9
NiftyReg F3D	0.640	0.031	1.55	1.20
NiftyReg F3D^2	0.623	0.000	0.00197	5.86
MIRTK FFD	0.641	0.398	2.54	3.06
MIRTK SVFFD	0.637	0.000	0.00177	7.55

Table 5.1.: Results of different registration tools for ALBERTs dataset.

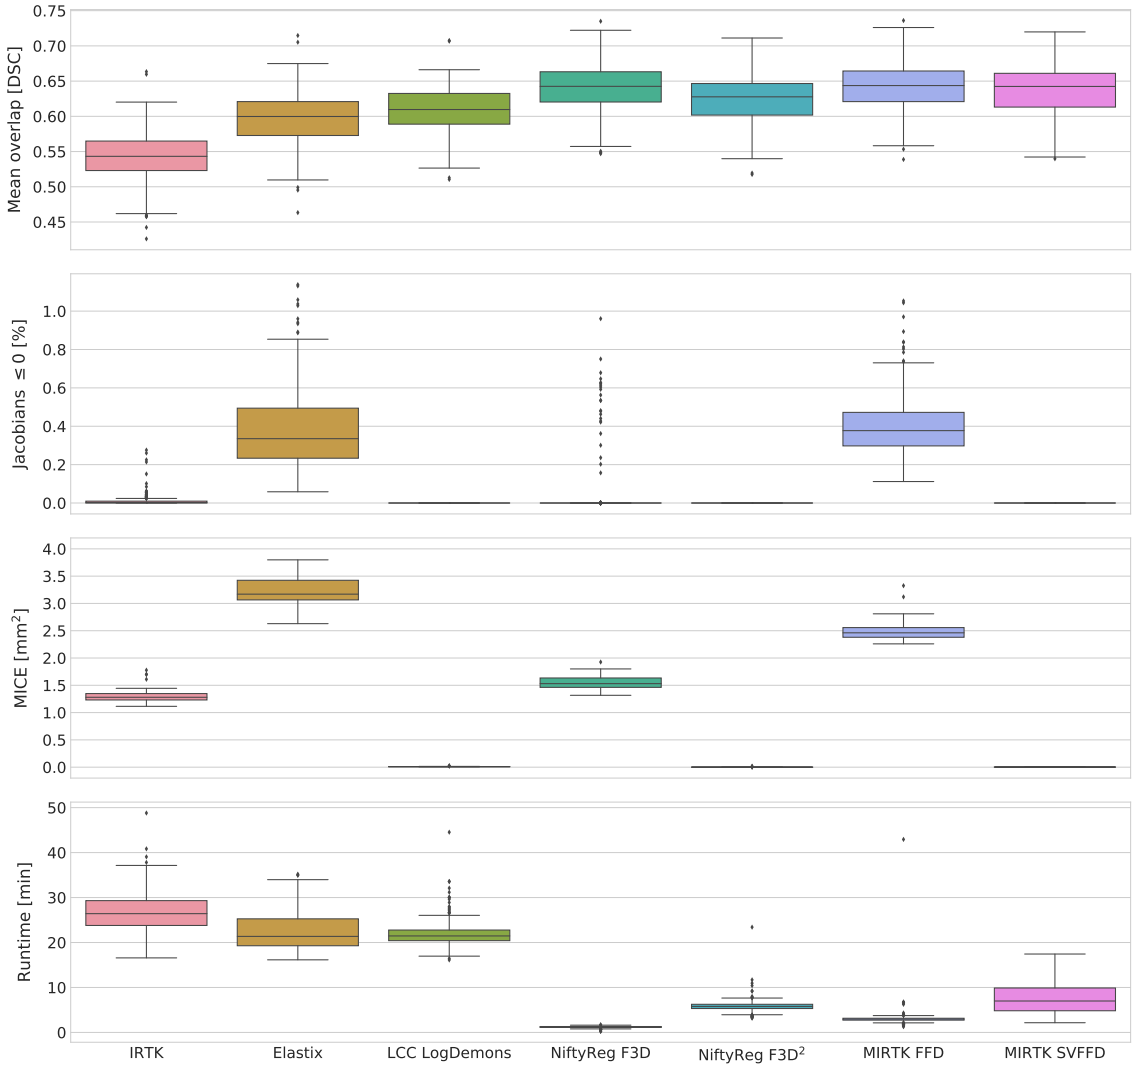


Figure 5.10.: Results of different registration tools for ALBERTs dataset.

	Mean overlap [DSC]	Jacobians ≤ 0 [%]	MICE [mm^2]	Runtime [min]
IRTK	0.604	0.100	1.48	78.0
Elastix	0.632	0.280	1.80	55.4
LCC LogDemons	0.556	0.000	0.242	50.4
NiftyReg F3D	0.625	0.240	13.5	2.66
NiftyReg F3D^2	0.628	0.000	0.403	17.0
MIRTK FFD	0.624	0.043	0.595	8.54
MIRTK SVFFD	0.645	0.000	0.00688	12.7

Table 5.2.: Results for OASIS dataset with parameters optimised for ALBERTs.

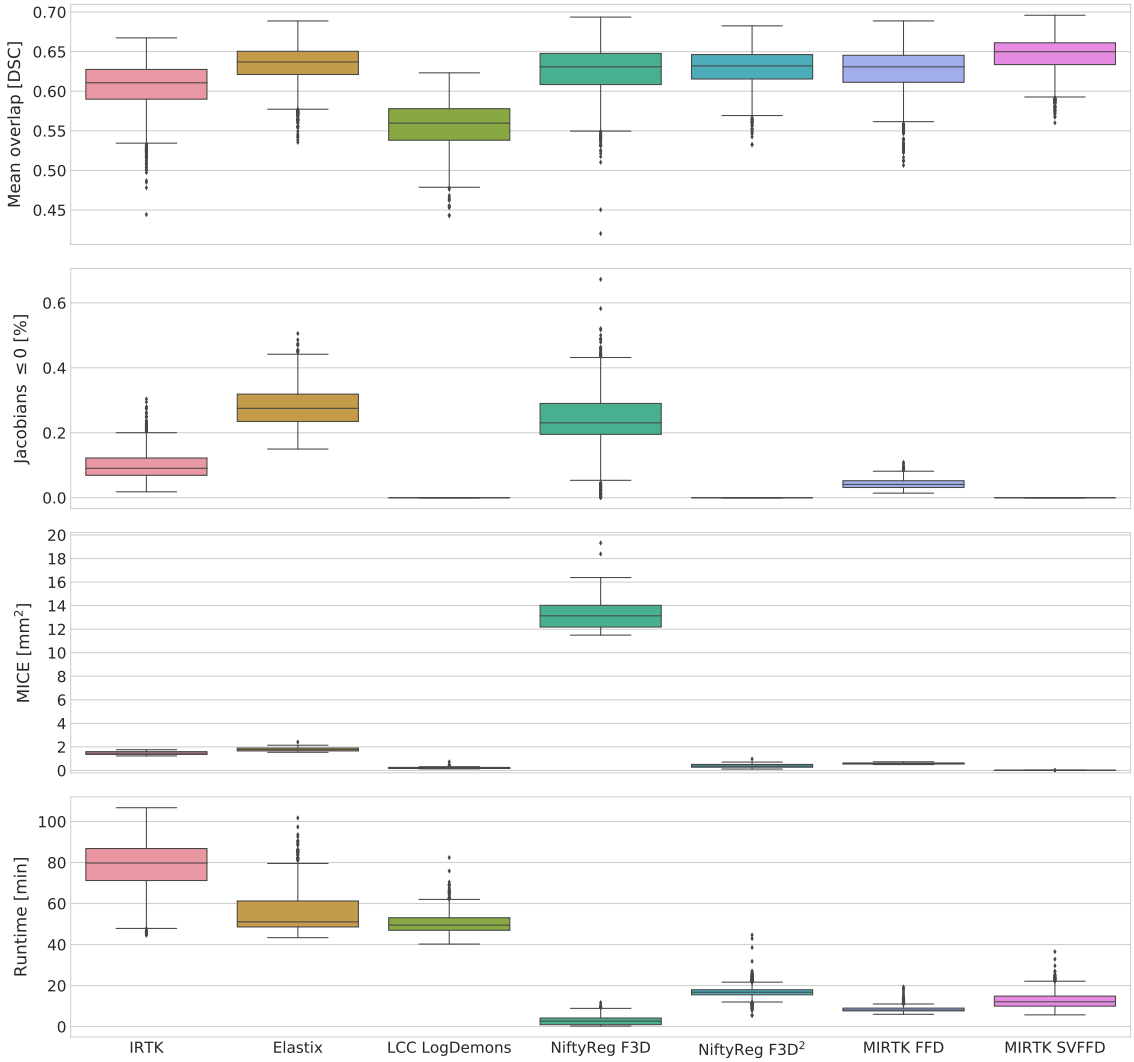


Figure 5.11.: Results for T1w adult brain images of OASIS dataset with parameters optimised for the registration of the T2w neonatal brain images of the ALBERTs dataset.

5.5. Conclusion

In this chapter, we evaluated and compared different approaches for the diffeomorphic registration of neonatal brain images based on cubic B-spline functions. To guarantee invertible, unbiased, and inverse consistent deformations, we proposed an extension of the SVFFD algorithm first motivated by Modat et al. (2012a). Our SVFFD algorithm is based on a symmetric energy formulation similar to the Demons approach of Lorenzi et al. (2013), but is not specific to a single image dissimilarity measure such as LNCC. It is the stationary equivalent to the consistent-midpoint-cost of LDDMM (Beg and Khan, 2007). This formulation is in contrast to the methods of Vercauteren et al. (2008, 2009) and Modat et al. (2012b), where the inverse transformation is given by a separate SVF, and thus requires an explicit symmetrisation of the forward and backward mappings. In our approach, the symmetry is inherent to the formulation based on a single SVFFD. The resulting algorithm is not biased towards the anatomy depicted by either of the two images. When the order of the images changes, only the sign of the SVFFD changes. This registration method is the foundation of the brain atlas construction techniques developed in the next chapter. The relevant configuration entries for the MIRTK `register` command, which realises the proposed SVFFD algorithm, are given in appendix C.1.

Chapter 6.

Spatio-temporal atlas of the normal developing brain

Medical image analysis has made significant progress over the past decades. With advances in acquiring high quality images of the developing human brain using MRI, analysing this data to understand brain development is rapidly becoming feasible (Kwon et al., 2014).

Many studies have shown that premature birth increases the risk of developing neuro-cognitive and -behavioural disorders (Rutherford, 2002; Blencowe et al., 2012; Sled and Nossin-Manor, 2013). A known difference to term-born infants is reduced cortical folding observed in premature-born infants at term-equivalent age (Hüppi et al., 1996; Inder et al., 2005; Kapellou et al., 2006). Further abnormalities are (cf. Kwon et al., 2014): (1) WM signal abnormalities such as lesions and diffuse excessive high signal intensities, (2) cGM signal abnormalities, and (3) cerebellar hemorrhages. The mechanisms of these alterations are yet unknown (Molnár and Rutherford, 2013). Studying the morphology and function of the brain during maturation, provides us not only with a better understanding of normal development, but may help identify causes for these factors.

The dHCP seeks to create a 4D connectome of early life. Understanding this connectome may provide insights into abnormal patterns of brain development. In this context, an atlas of normal brain development is important, as it facilitates spatial normalisation, VBM (Ashburner and Friston, 2000), DBM (Ashburner et al., 1998), and CA studies (Grenander and Miller, 1998; Thompson and Toga, 2002) (cf. section 1.3).

In the following, we first summarise a procedure in section 6.1 for obtaining an unbiased affine reference space for initial spatial normalisation of all brain scans. This global normalisation is followed by a spatio-temporal atlas construction based on deformable registration which reduces the misalignment of brain anatomies, and thereby increases the level of anatomical detail. For this, we propose two different approaches that both built on the unbiased, symmetric, and diffeomorphic SVFFD algorithm introduced in chapter 5.

In our first approach presented in section 6.2, we construct a spatio-temporal neonatal atlas from registrations between all pairs of images, and the Log-Euclidean mean of SVFs (Arsigny et al., 2006). It directly improves upon the method of Serag et al. (2012a), which is based on the FFD algorithm (Rueckert et al., 1999), and arithmetic mean of displacement fields. A longitudinal model of brain morphology is hereby obtained using Nadaraya–Watson kernel regression similar to Davis et al. (2007), with adaptive kernel size as proposed by Serag et al. (2012a).

A drawback of this approach is the quadratic computational complexity with respect to the number of images, which means that it is difficult to scale to large datasets. A more significant limitation is the reliance on accurate pairwise correspondences, and the dependence on average spatial correspondences, which can capture complex cortical folding. The less accurate the individual registrations are, the less detail is preserved. An approach which iteratively refines the transformations that relate each anatomy with the average space may be more robust towards residual misalignment.

Instead of expressing the spatio-temporal atlas in terms of the transformations relating each brain image to all others, in section 6.3 we derive a new atlas construction based on temporal kernel regression, the transformations relating each image to the atlas space, and their respective age-dependent Log-Euclidean mean. This novel approach jointly estimates mean shape and longitudinal change iteratively. Unlike other groupwise methods that were originally developed to construct single-time point atlases, our proposed method is based on a continuous recursive formulation of the spatio-temporal mean brain image with improved temporal consistency. A comparison of our methods to the atlas obtained with Serag et al.'s (2012a) technique demonstrates a marked increase in anatomical detail.

6.1. Unbiased global normalisation

Differences in the position and orientation of the anatomies depicted by the images from which the brain atlas is to be constructed are removed by an initial rigid alignment. Moreover, differences in brain size are also removed at this initial normalisation stage, such that the following template construction based on deformable registration can focus on local shape differences alone. The template brain images constructed from the globally normalised brain images may be rescaled afterwards such that the resulting spatio-temporal atlas includes global volume changes. This is discussed in more detail in section 6.2.3.

The global normalisation and age-dependent rescaling is similar to the atlas construction proposed by Kuklisova-Murgasova et al. (2011), where only affine registration is employed to create average brain images and tissue probability maps of the preterm born neonate. For this, the authors use a reference image that is rigidly aligned with the ICBM-152 adult brain template (Mazziotta et al., 2001). The neonatal reference brain image serves as common target for the affine image registrations. Bias towards the reference anatomy is subsequently accounted for by composing each affine transformation with the inverse of their Log-Euclidean mean (Arsigny et al., 2007). Rigid components reflecting only differences in scanner coordinate systems are hereby discarded. This removes the bias towards the global reference shape by mapping it to the barycentre of all observed anatomies. An age-dependent kernel regression is employed to shift this barycentre towards those anatomies that are more representative for a given PMA. In the work presented here, the kernel regression of the age-dependent rescaling of the brain images along an arbitrary set of axes is applied after an additional non-rigid deformation, and therefore discussed separately after we introduced the deformable atlas construction.

Under the assumption of ideal, and thereby inverse-consistent and transitive affine registrations, the Log-Euclidean mean is independent of the reference (Pennec and Arsigny, 2013) [1]. In practice, however, affine registration is more accurate for anatomies that are more similar to the reference anatomy. The choice of which of the two images is to remain

[1] A short video created by Miolane and Khanal illustrating this property is among the winners of the MICCAI Educational Challenge 2014. It can be viewed on-line at https://www.youtube.com/watch?v=XUz59y6HDEk (last accessed: 20th March 2017).

fixed, and which is being transformed and resampled, further effects the outcome of the registration. These registration errors result in inconsistent affine transformations that are neither inverse-consistent nor transitive (Woods et al., 1998). Consequently, some registration bias is retained when choosing a single reference for global normalisation.

In our atlas construction, we eliminate this remaining bias by computing all pairwise affine registrations, i.e., selecting each image in turn as reference. A mean transformation derived from all pairwise affine transformations, but different from the Log-Euclidean mean, was first employed by Woods et al. (1998) for the construction of an adult brain template from 22 brain MR images. Due to the lack of inverse-consistency, some residual global misalignment remains, after mapping each image by its corresponding Log-Euclidean mean transformation. We therefore repeat the global normalisation steps detailed below for another iteration. Alternatively, given how close the residual mean transformations are to the identity after just one iteration, the average intensity image of the thus aligned images may be used as single reference for the affine registrations at the second iteration, instead of computing all residual pairwise affine transformations. While the pairwise approach is considerably more costly in terms of computation time due to the quadratic rather than linear dependence of the number of registrations on the number of images, each affine registration takes only few seconds, and the added cost is low relative to a deformable template construction based on pairwise deformable registration. The benefits of a robust and fully unbiased affine pre-alignment of the images, which in turn impacts the subsequent deformable registrations, outweigh the added computational cost, especially when we consider that an atlas is generally constructed once and then re-used many times.

We now formalise our proposed iterative global normalisation. At the k-th iteration, the i-th transformation which maps image I_i to the global reference space is given by

$$\bar{\boldsymbol{A}}_i^{(k)} = \exp\left(\frac{1}{n}\sum_{j=1}^{n}\log \boldsymbol{A}_{ij}^{(k)}\right)\bar{\boldsymbol{A}}_i^{(k-1)} \text{ for } k>0, \text{ and } \bar{\boldsymbol{A}}_i^{(0)} = \boldsymbol{I}\,, \tag{6.1}$$

where $\boldsymbol{I}$ is the identity matrix, exp and log are the matrix exponential and corresponding logarithmic map (Alexa, 2002; Arsigny et al., 2007), and $\boldsymbol{A}_{ij}^{(k)}$ is the affine transformation obtained by registering the j-th to the i-th image after application of the mean transforma-

tions of the previous iteration, i.e., $I_i \circ \left(\bar{\boldsymbol{A}}_i^{(k-1)}\right)^{-1}$ and $I_j \circ \left(\bar{\boldsymbol{A}}_j^{(k-1)}\right)^{-1}$ [2]. The smoothing effect caused by interpolation errors when resampling the images in the reference space, is avoided through composition with the transformations $\boldsymbol{B}_i$ which map each image sampling space to the common world space. Specifically, the map relating a sampling point $\hat{\boldsymbol{x}}_i \in \mathbb{R}^d$ in the space of the i-th image with axes parallel to the sampling grid and scaled according to the image resolution, to its corresponding point $\hat{\boldsymbol{x}}_j$ in the j-th image is given by

$$\begin{aligned}\hat{\boldsymbol{x}}_j &= \boldsymbol{B}_j^{-1}\left(\bar{\boldsymbol{A}}_j^{(k-1)}\right)^{-1}\boldsymbol{A}_{ij}^{(k)}\bar{\boldsymbol{A}}_i^{(k-1)}\boldsymbol{B}_i\hat{\boldsymbol{x}}_i \\ &= \left(\bar{\boldsymbol{A}}_j^{(k-1)}\boldsymbol{B}_j\right)^{-1}\boldsymbol{A}_{ij}^{(k)}\left(\bar{\boldsymbol{A}}_i^{(k-1)}\boldsymbol{B}_i\right)\hat{\boldsymbol{x}}_i \\ &= \left(\boldsymbol{C}_j^{(k-1)}\right)^{-1}\boldsymbol{A}_{ij}^{(k)}\boldsymbol{C}_i^{(k-1)}\hat{\boldsymbol{x}}_i\end{aligned} \tag{6.2}$$

Registering images with affine pre-transformations $\bar{\boldsymbol{A}}_{\{i,j\}}^{(k-1)}$ is thus equivalent to registering images with the composite image to world maps denoted by $\boldsymbol{C}_{\{i,j\}}^{(k-1)}$. The registration algorithm finds the parameters of $\boldsymbol{A}_{ij}^{(k)}$ which maximise the NMI of the image pair. To provide a better initialisation and ensure a sufficient initial overlap, all images are first rigidly registered to a publicly available template brain image[3]. The corresponding rigid transformations then serve as the initial mean transformations $\bar{\boldsymbol{A}}_i^{(0)}$.

After global normalisation, the initial rigid alignment to a pre-existing template brain image and the unbiased population-specific global average brain image given by

$$\bar{I}_{\text{global}}(\boldsymbol{x}) = \frac{1}{n}\sum_{i=1}^{n}\frac{I_i\left(\boldsymbol{A}_i^{-1}\boldsymbol{x}\right) - \mu_i}{\sigma_i} \tag{6.3}$$

can be recomputed. Here, $\boldsymbol{A}_i = \bar{\boldsymbol{A}}_i^{(K)}$ is the final affine mean transformation after $K = 2$ iterations which maps image I_i to the global reference space, and μ_i and σ_i are the mean and standard deviation of the image intensities used to compensate for global differences of the intensity distributions. This rigid alignment only serves to help compare the resulting atlas to other templates defined in a commonly used reference space such as the adult MNI-305 and ICBM-152 templates, for example.

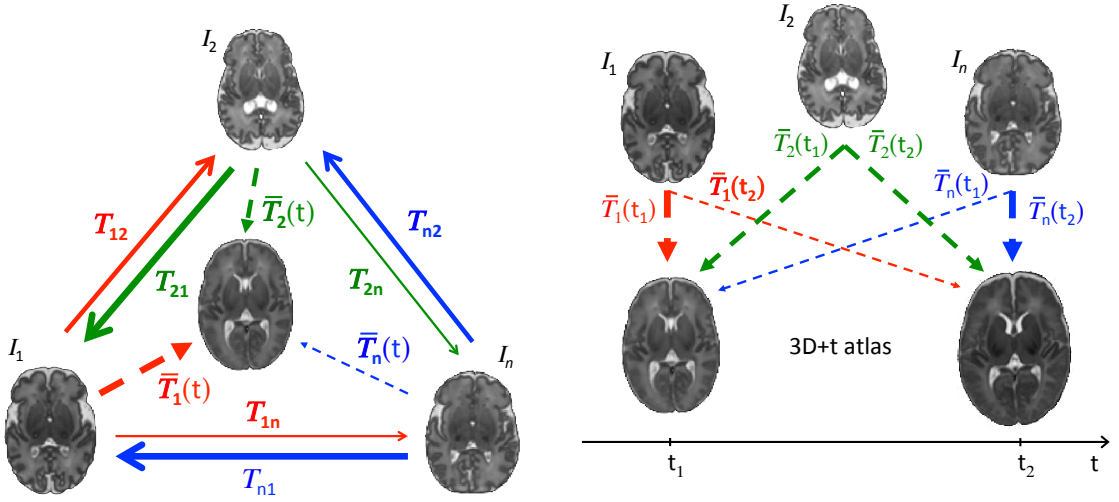
[2] Affine transformations and their 4x4 matrices are denoted by the same symbols (see also section 2.2.1).
[3] Available at http://brain-development.org/brain-atlases/ (last accessed: May 20, 2017).

6.2. Pairwise spatio-temporal atlas construction

We now present our first approach to the construction of an unbiased spatio-temporal brain atlas from individual observations at different developmental stages based on deformable registration between (all)[4] pairs of images. This neonatal atlas construction is a combination of the affine method of Kuklisova-Murgasova et al. (2011), which is based on the weighted Log-Euclidean mean of affine transformations, and the non-rigid approach by Serag et al. (2012a), which itself extends Seghers et al.'s (2004) method to 4D using temporal kernel regression with an adaptive kernel width. Our proposed construction results in template brain images of high anatomical detail comparable to previous non-rigid approaches, while overcoming in particular the inherent limitations of an atlas construction based on the arithmetic mean of displacement vector fields. In the following, we outline these limitations before presenting our proposed approach.

For the unbiased atlas construction from pairwise registrations (Seghers et al., 2004), a mean deformation is computed for each image, and its inverse is used to map the image intensities to the atlas domain. This is illustrated in figure 6.1a. However, neither the FFD model employed by Serag et al. (2012a) nor the arithmetic mean guarantee the required invertibility of the mean deformation. In order for a FFD to be guaranteed invertible, the displacements at control points have to be restricted to small deformations (Rueckert et al., 2006). This limits the application of the classic FFD for the diffeomorphic mapping of brain images with large biological variability as observed during early brain development across a wide age range. An atlas construction based on the inverse of the arithmetic mean deformation, such as the approach by Serag et al. (2012a), in practice relies on the smoothing of locally folded deformations as a result of the averaging of displacement vectors. Whereas the individual deformations between pairs of images may contain many points with non-positive Jacobian determinant, their arithmetic mean will contain considerably fewer such points when a large enough number of individual deformation fields (with a distinct spatial distribution of locally non-invertible mappings) are being averaged. However, no actual guarantees on the invertibility of the mean deformations are made.

[4] We will see that not all $n(n-1)$ registrations must be computed, but the complexity remains $O(n^2)$.



(a) Cross-sectional view (b) Longitudinal view

Figure 6.1.: Illustration of pairwise spatio-temporal atlas construction. The transformations $\bar{\boldsymbol{T}}_i(t)$ (dashed arrows) correspond to the Log-Euclidean mean of transformations $\boldsymbol{T}_{ij}$ between pairs of images (solid arrows). Kernel weights of the temporal regression are reflected in the width of the arrows of the corresponding transformations.

Another limitation results from the use of the arithmetic mean of displacement vectors. This mean deformation does not necessarily preserve the underlying deformation model of each pairwise deformation (Avants and Gee, 2004). Assuming the computed pairwise FFDs are invertible, i.e., elements of the group of diffeomorphic mappings with composition as group action, the addition of these displacement fields may not be an element of the group (Joshi et al., 2004). The topology of the shape which is deformed by this mean may thus not be preserved. The addition of displacement fields assumes an underlying linear vector space and is based on the small deformation assumption, where the addition is a valid approximation of their composition. In the large deformation setting, however, this approximation is no longer valid. In consequence, the interpolation of brain shapes given by the application of a linear combination of FFDs is not well defined under large deformations. A geometric definition of a mean deformation within the group of diffeomorphisms is hence desirable for constructing a template of mean brain shape.

Finally, we show that the individual mean deformations are inconsistent when the registrations between each pair of images are not inverse-consistent. More specifically, although each (arithmetic) mean deformation $\bar{\boldsymbol{T}}_i$ results in zero residual average displacement of each biased template image, $\hat{I}_i = I_i \circ \bar{\boldsymbol{T}}_i^{-1}$, to all observed images after mapping it to the atlas space, the average of the applied mean displacement, $\bar{\boldsymbol{d}} = \frac{1}{n} \sum_{i=1}^{n} \bar{\boldsymbol{T}}_i^{-1} - \mathrm{Id}$, on the other hand, is in general not zero when the pairwise transformations are not inverse-consistent, i.e., when $\boldsymbol{T}_{ij}^{-1} \neq \boldsymbol{T}_{ji}$. Inconsistent anatomies are thus averaged when creating an unbiased mean image from these n biased template images. Inverse-consistency is, however, in general not given for the transformations computed with the classic FFD algorithm (Rueckert et al., 1999) utilised by Serag et al. (2012a). Therefore, the resulting mean brain shape is not exactly at the centre of the observed anatomies. This shortcoming can be easily accounted for by composing each arithmetic mean deformation with the (weighted) average of these mean deformations. Neither Seghers et al. (2004) nor Serag et al. (2012a) have considered this corrective residual deformation, however, which compensates for the lack of inverse-consistency among the computed pairwise deformations.

In the following, we address each of the aforementioned limitations using the symmetric and inverse-consistent diffeomorphic SVFFD registration approach that we presented in the previous chapter. Application of this registration method both guarantees invertibility and inverse-consistency of the computed deformations between each pair of images. The symmetric energy formulation moreover reduces the number of required pairwise registrations by 50%. The inverse mappings are easily obtained by inverting the computed transformations. Given the SVFFDs between all pairs of images with non-zero weight, we can directly compute the Log-Euclidean mean deformation (Arsigny et al., 2006, 2007) instead of the arithmetic mean of displacement fields, without the need for numerical approximation of the group logarithm. In summary, the Log-Euclidean mean is a generalisation of the geometric mean of real numbers to the one-parameter subgroup of diffeomorphic mappings generated by SVFs. Utilising our SVFFD model, adopting the barycentric Log-Euclidean mean for both global normalisation as well as deformable atlas construction becomes feasible. This results in a consistent construction of a spatio-temporal mean shape atlas within the group of diffeomorphic mappings.

6.2.1. Log-Euclidean mean of SVFFD

Given the spatial correspondences between all pairs of n images obtained by symmetric pairwise SVFFD registration (incl. the identity for $i = j$), i.e., $\boldsymbol{x}_j = \boldsymbol{T}_{ij}(\boldsymbol{x}_i)$, where $\boldsymbol{T}_{ij} = \exp(\boldsymbol{v}_{ij})$, we seek for each image a mean transformation which maps it to the barycentre of all n observations (cf. dashed arrows in figure 6.1a). Specifically, given the definition of a distance measure between two diffeomorphic mappings generated by SVFs (see also section 2.2.3), the desired mean transformations are given by

$$\bar{\boldsymbol{T}}_i = \arg\min_{\boldsymbol{T}} \sum_{j=1}^{n} w_j \big(\operatorname{dist}(\boldsymbol{T}, \boldsymbol{T}_{ij})\big)^2 \tag{6.4}$$

where $\sum_{j=1}^{n} w_j = 1$, i.e., $w_j = 1/n$ in the uniform case. The distance of two deformations generated by SVFs $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$, respectively, is hereby defined as (Arsigny et al., 2006)

$$\operatorname{dist}(\boldsymbol{T}_1, \boldsymbol{T}_2) \coloneqq \|\log(\boldsymbol{T}_1) - \log(\boldsymbol{T}_2)\| = \|\boldsymbol{v}_1 - \boldsymbol{v}_2\| \tag{6.5}$$

The minimiser of (6.4) is the Log-Euclidean mean given by (Arsigny et al., 2006)

$$\bar{\boldsymbol{T}}_i = \exp\left(\sum_{j=1}^{n} w_j \log\left(\boldsymbol{T}_{ij}\right)\right) = \exp\left(\sum_{j=1}^{n} w_j \boldsymbol{v}_{ij}\right) = \exp\left(\bar{\boldsymbol{v}}_i\right) \tag{6.6}$$

where $\bar{\boldsymbol{v}}_i$ is the average SVF of those SVFFDs which map points of the i-th image to any other image. Because we place control points at the same positions of the global reference space for all computed SVFFDs, we can directly average the spline coefficients in order to obtain the corresponding Log-Euclidean mean SVFFD.

For the n SVFFDs given by (6.6) to map the images to the barycentre of the observed anatomies, the residual mean deformation given by $\exp(\sum_{i=1}^{n} w_i \bar{\boldsymbol{v}}_i)$ must be zero, i.e.,

$$\sum_{i=1}^{n} w_i \bar{\boldsymbol{v}}_i = \sum_{i,j=1}^{n} w_i w_j \boldsymbol{v}_{ij} = \sum_{i=1}^{n} \sum_{j=i}^{n} w_i w_j \left(\boldsymbol{v}_{ij} + \boldsymbol{v}_{ji}\right) = \boldsymbol{0} \tag{6.7}$$

The last equality is guaranteed only when the pairwise transformations are inverse-consistent, i.e., $\boldsymbol{v}_{ij} = -\boldsymbol{v}_{ji}$, which is ensured by our symmetric SVFFD registration method.

6.2.2. Age-dependent Log-Euclidean mean

Following the prior work of Kuklisova-Murgasova et al. (2011) and Serag et al. (2012a) on the construction of spatio-temporal neonatal brain atlases, we also utilise kernel regression to derive an age-dependent Log-Euclidean mean of diffeomorphic SVFFDs which map each observed brain anatomy to the common spatio-temporal atlas space. Such temporal kernel regression for spatio-temporal atlas construction was first proposed by Davis et al. (2007).

The constant weights w_j in (6.6) are hereby substituted by age-dependent weights

$$\hat{w}_j(t) = \frac{1}{\sigma_t\sqrt{2\pi}} \exp\left(\frac{-(t_j - t)^2}{2\sigma_t^2}\right) \tag{6.8}$$

with Gaussian temporal regression kernel centred at atlas time point t and corresponding standard deviation σ_t. The weights are truncated below a threshold to limit the global extent of the Gaussian and thereby reduce the number of required pairwise registrations.

The temporal extent of the kernel can be chosen to be either time-invariant, i.e., $\sigma_t = \sigma$, as in case of the kernel regression utilised by Kuklisova-Murgasova et al. (2011), or adapted in order to ensure a more equal number of brain images to be within the local support of each kernel as proposed by Serag et al. (2012a). An adaptive kernel width has been shown to be advantageous over a fixed kernel width for the construction of a spatio-temporal atlas of the early developing brain. This is due to the age at scan distribution of most neonatal datasets, where neonates are ideally imaged close to birth and most are born around term between 37 to 40 weeks gestation. Adapting the kernel width to be narrower at term combined with a data selection which eliminates sharp peaks in the age distribution, effectively smooths the age distribution and makes it more uniform. This results in temporally more consistent atlas templates given a more even distribution of observed anatomical samples across the considered age range.

The age-dependent Log-Euclidean mean transformation is now given by

$$\bar{\boldsymbol{T}}_i(t) = \exp\left(\sum_{j=1}^{n} w_j(t)\boldsymbol{v}_{ij}\right) \tag{6.9}$$

where t is the age associated with the co-domain of the subject-to-atlas coordinate maps,

and $w_j(t) = \hat{w}_j(t) / \sum_{k=1}^{n} \hat{w}_k(t)$ are the normalised regression weights. This weighted average shifts the barycentre of the individual deformations towards those samples which are closer in age to the respective atlas time point. Brain images from an earlier or later developmental phase have less or no influence on the mean deformation at this age. This is an important property given the large differences in brain shape at different ages, and the aim to construct an age-dependent brain atlas which better resembles the mean brain shape at each stage of brain development during early life.

6.2.3. Spatio-temporal mean intensity and brain shape

The atlas template of mean brain shape and intensity is estimated as the image which minimises the weighted sum of squared differences to the observations close to t, after mapping these into the respective spatio-temporal atlas coordinate system, i.e.,

$$\bar{I}_{\text{local}}(t) = \sum_{i=1}^{n} w_i(t) \tilde{I}_i(t) \tag{6.10}$$

where

$$\tilde{I}_i(t) = \frac{I_i \circ \boldsymbol{A}_i^{-1} \circ \bar{\boldsymbol{T}}_i(t)^{-1} - \mu_i}{\sigma_i} \tag{6.11}$$

is the i-th normalised intensity image mapped to the atlas space at time point t, and $w_i(t)$ are the normalised weights of the regression kernel. To counter the blurring effect, an edge enhancing filter based on convolution with the Laplacian is furthermore applied[5]. This spatio-temporal template image can be used for spatial normalisation of neonatal brain images, and for computing the deformations which relate a given brain image to this reference brain image for atlas-based segmentation, DBM, or statistical analysis in CA.

The resulting spatio-temporal brain atlas corresponds to the mean local shape and intensity of the population from which the atlas was constructed. Global differences in brain volume are not reflected in this template image as these have been compensated for during the global normalisation. For applications which analyse brain volume changes during early brain development, a brain atlas adjusted for average brain size over time can be created. These differences are encoded in the parameters of the affine transformations

[5] Using the Laplacian sharpening filter implemented in ITK which is also used by ANTs SyGN scripts.

computed during global normalisation. We can decompose these homogeneous coordinate transformations into a rigid transformation $\boldsymbol{R}_i$ followed by a scaling (and shearing) $\boldsymbol{S}_i$, s.t. $\boldsymbol{A}_i = \boldsymbol{S}_i\boldsymbol{R}_i$. A kernel regression of the thus obtained scaling parameters over time then recovers the transformation to be applied to the local shape and intensity template. The resulting template brain image not only reflects mean brain shape, but also average brain size at a given age. This spatio-temporal template image is given by

$$\bar{I}(t) = \sum_{i=1}^{n} w_i(t) \left(\tilde{I}_i(t) \circ \boldsymbol{S}(t) \right) \tag{6.12}$$

with age-dependent mean scaling (and shearing) of the spatial domain given by the regressed Log-Euclidean mean of the inverses of the non-rigid affine matrices $\boldsymbol{S}_i$, i.e.,

$$\boldsymbol{S}(t) = \exp\left(-\sum_{i=1}^{n} w_i(t) \log \boldsymbol{S}_i \right) \tag{6.13}$$

As noted before, without the intermediate Log-Euclidean mean deformations $\bar{\boldsymbol{T}}_i^{-1}(t)$, which compensate for local shape differences before the mapped observations are averaged, the obtained spatio-temporal atlas is similar to the atlas constructed with the method proposed by Kuklisova-Murgasova et al. (2011)[6]. Unlike the method proposed by Serag et al. (2012a), the here presented spatio-temporal atlas construction using deformable registration is a consistent extension of this affine atlas construction approach, where all mean transformations are the Log-Euclidean mean of diffeomorphic mappings. This has been made possible by the symmetric and inverse-consistent SVFFD registration method.

6.2.4. Anatomical atlas of mean brain morphology

The transformations computed for the template construction relate each point of the reconstructed brain scans to their respective anatomical point in the average space. Segmentations of brain tissues and structures can be transferred into the atlas space using these maps. Probabilistic and hard segmentations, which are consistent with the brain template, can be created thereof. Due to the lack of laborious manual segmentations, we

[6]Except for the differences in global normalisation discussed in section 6.1.

utilise an automatic atlas-based tissue (Makropoulos et al., 2012) and whole-brain segmentation (Makropoulos et al., 2014) to segment the images in our neonatal dataset. In this approach, multiple manually annotated neonatal brain images (Gousias et al., 2012) are non-rigidly registered to the to-be-segmented image. The 50 brain regions are subdivided further by Makropoulos et al. (2014) into cortical and non-cortical labels using the neonatal brain atlas of Serag et al. (2012a). The whole brain segmentation into 87 structures is obtained using an EM with MRF regularisation, and a partial volume correction which accounts for misclassifications of partial volume near the cortex (Xue et al., 2007).

Each posterior probability map $p_{l,i}$ of the segmentation of the class with label l in image I_i is transferred into the atlas space using the inverse subject-to-atlas transformations. The propagated segmentations are then averaged using the normalised temporal regression weights to obtain a mean probability map for the respective class, i.e.,

$$p_l(t) = \sum_{i=1}^{n} w_i(t)\tilde{p}_{l,i}(t) \tag{6.14}$$

with

$$\tilde{p}_{l,i}(t) = p_{l,i} \circ \boldsymbol{A}_i^{-1} \circ \bar{\boldsymbol{T}}_i(t)^{-1} \circ \boldsymbol{S}(t) \tag{6.15}$$

The more anatomical detail is preserved by the atlas construction, the steeper are the transitions of the mean class probabilities from one to zero. The mean probabilities are therefore in general less suitable for application in probabilistic atlas-based segmentation, because no spatial uncertainty is retained to reflect biological variability and they fail to balance an inaccurate atlas registration. This can be overcome by Gaussian smoothing of the probability maps when used for segmentation purposes and generally mandates the use of a prior relaxation technique (Shiee et al., 2011; Cardoso et al., 2011a,b, 2013).

Hard reference segmentations of the spatio-temporal atlas can be derived from the mean class probabilities, however, which may be useful for other applications such as the definition of regions of interest after spatial normalisation. For this, points in the atlas space are assigned the class label corresponding to the maximum mean probability, i.e.,

$$L(\boldsymbol{x}, t) = \arg_l \max p_l(\boldsymbol{x}, t) \tag{6.16}$$

6.2.5. Refinement of cortical correspondences

This first proposed atlas construction results in mean brain images for each age, which exhibit higher anatomical detail than the spatio-temporal neonatal atlas presented by Kuklisova-Murgasova et al. (2011), while overcoming the previously discussed limitations of the deformable atlas construction proposed by Serag et al. (2012a). However, such "single-shot" atlas construction relies on accurate registrations between any pair of brain images. Deformable brain image registration is an ill-posed problem with ambiguous and hard to determine anatomical correspondences, especially during early brain development with large variation of cortical folding, and thus substantial registration error remains. This decreases the markedness of cortical folds other than the major sulci. More cortical detail can be recovered by iterating the atlas construction. In case of an atlas construction based on registrations between all pairs of images, the same approach as used for the global normalisation could be applied. This adds, however, considerably to the computational burden, because an atlas construction based on deformable registration between all pairs of images is time-consuming already for a single iteration.

A much less computationally expensive refinement of the constructed unbiased spatio-temporal atlas for time points t_k can be achieved by registering each image to the respective atlas images $\bar{I}(t_k)$, and to compose these transformations with their respective Log-Euclidean mean similar to the affine method of Kuklisova-Murgasova et al. (2011), but with reference images given by the current unbiased template estimates at the considered discrete ages. A similar strategy has been successfully used by Gholipour et al. (2014, 2017) to build fetal brain image templates for different ages using greedy SyN (Avants et al., 2008), and the arithmetic mean of the obtained diffeomorphic displacement fields[7]. The initial estimates of mean intensity and brain shape are hereby given directly by the regressed affine spatio-temporal atlas without costly pairwise registrations between each pair of images. This iterative spatio-temporal atlas refinement successfully improves the anatomical detail of atlas templates at discrete time points, but has its drawbacks.

[7]Note that although SyN is a greedy LDDMM method (as discussed in section 2.2.3), the result of the greedy optimisation is a 3D displacement field rather than a 4D velocity field. We confirmed in personal correspondence with Ali Gholipour and by examining the publicly available ANTs source code that the arithmetic mean of displacements is used for atlas construction instead of integrating average velocities.

Firstly, the transformations relating each image to the atlas space at different time points are computed independent from each other, and therefore temporal consistency is not explicitly ensured. The registration of a given brain image to two nearby time points is only implicitly assumed to be temporally consistent due to the similarity of the respective template images to which it is registered. Because deformable registration is solved by local optimisation, which is easily trapped in local minima of the highly non-convex image dissimilarity cost function, even slight changes of the target image between adjacent time points may result in differences in the deformations that cannot be explained by the change in template shape alone.

Secondly, a continuous temporal function of mean intensity and brain shape requires several time-consuming minimisations of the registration energy for every time point as suggested by Gholipour et al. (2017). Alternatively, given a sufficient sampling of the spatio-temporal atlas space such that only small appearance and shape changes between discrete time points can be assumed, a temporal interpolation may be performed, e.g.,

$$\bar{I}(t) = \frac{\sum_{k=1}^{n} w_k(t)\bar{I}_k}{\sum_{k=1}^{n} w_k(t)} \tag{6.17}$$

with brain templates $\bar{I}_k$ constructed for time points t_k, and Gaussian weights

$$w_k(t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(t_k - t)^2}{2\sigma^2}\right) \tag{6.18}$$

A longitudinal deformation estimated from the image sequence may also be utilised.

Because an iterative construction has been shown to produce atlases of higher anatomical detail than a single iteration of the pairwise atlas construction, and given a favourable linear computational complexity which better scales for larger datasets, we re-examined our atlas construction, and derived a groupwise approach based on temporally consistent Log-Euclidean mean deformations using the SVFFD algorithm. This new groupwise atlas construction approach is presented in detail in the following section.

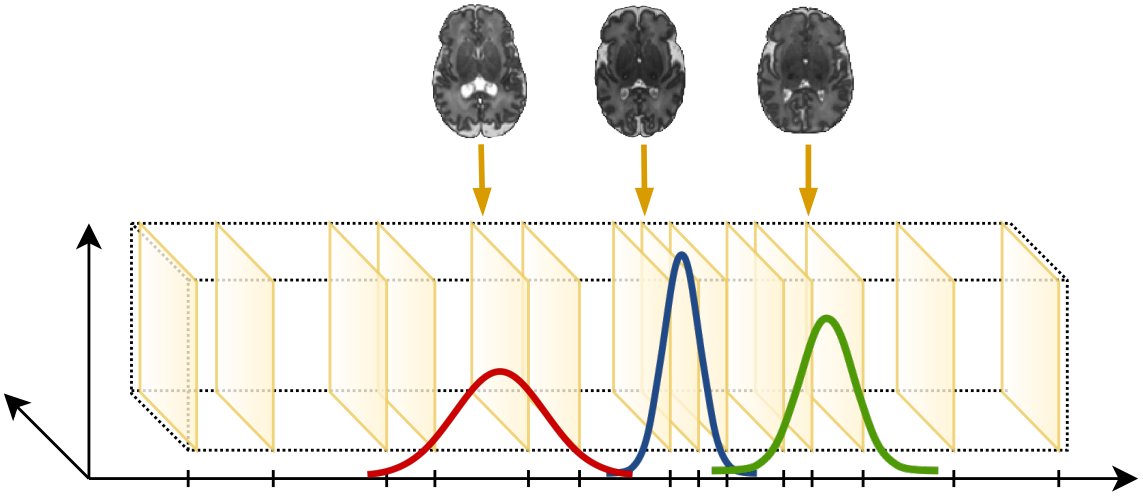


Figure 6.2.: Illustration of iterative spatio-temporal atlas construction with time corresponding to the horizontal axis, and cross-sectional image planes depicting the spatial domain of the atlas. Exemplary regression kernels are shown as curve overlays. Arrows from sample brain images to the atlas space illustrate the subject-to-atlas mappings.

6.3. Groupwise spatio-temporal atlas construction

Instead of expressing the spatio-temporal atlas in terms of the transformations relating each brain image to all others, we now derive a recursive continuous formulation based on temporal kernel regression, the transformations relating each image to the spatio-temporal atlas space, and their respective age-dependent Log-Euclidean mean. This definition of the spatio-temporal atlas is illustrated in figure 6.2. In this schematic drawing, the horizontal axis corresponds to the temporal domain and the image planes orthogonal to it depict the 3D space of the atlas. Each image plane pictures an atlas template to which a given sample brain image acquired at time t_i is registered to in order to infer the transformation which relates it to the atlas space. The temporal density of the observations is therefore illustrated by the spacing of the cross-sectional image planes along the time axis. The adaptation of the temporal regression kernel width to this age distribution is shown by three exemplary Gaussian curves overlayed over the diagram.

Once we defined the recursive mean intensity function in section 6.3.1, we modify this definition in section 6.3.2 to also incorporate an estimate of the longitudinal deformations of the atlas space from one time point to another. The steps which are iterated to construct a spatio-temporal atlas from n neonatal brain images of different ages are then summarised.

6.3.1. Recursive definition of spatio-temporal atlas

The mean intensity and brain shape template at iteration k is given by

$$\bar{I}_k(t) = \sum_{i=1}^{n} w_i(t) \left(\tilde{I}_i \circ \boldsymbol{T}_{k,i}^{-1}(t) \right) \tag{6.19}$$

where $\tilde{I}_i : \Omega_i \to \mathbb{R}$ is the i-th intensity normalised image after global alignment, i.e.,

$$\tilde{I}_i(\boldsymbol{x}) = \frac{I_i(\boldsymbol{A}_i^{-1}\boldsymbol{x}) - \mu_i}{\sigma_i} \tag{6.20}$$

and $\boldsymbol{T}_{k,i} : \Omega_i \times \mathbb{R} \to \Omega_t$ is the spatio-temporal mapping relating it to the atlas at time t.

This mapping can be modelled as composition of the inverse deformation $\boldsymbol{\phi}_{k,i}^{-1} : \Omega_i \to \Omega_{t_i}$ which locally aligns the image with the previous template at age t_i, and the residual mean deformation $\bar{\boldsymbol{\phi}}_k : \Omega_{t_i} \times \mathbb{R} \to \Omega_t$ regressed from the cross-sectional deformations, i.e.,

$$\boldsymbol{T}_{k,i}(t) = \bar{\boldsymbol{\phi}}_k\,(t) \circ \boldsymbol{\phi}_{k,i}^{-1} \tag{6.21}$$

where

$$\bar{\boldsymbol{\phi}}_k(t) = \exp\left(\sum_{i=1}^{n} w_i(t) \log \boldsymbol{\phi}_{k,i} \right) \tag{6.22}$$

is the age-dependent Log-Euclidean mean of atlas-to-subject deformations obtained through SVFFD registrations of all images to their respective age-matched average image of the previous iteration. The logarithmic maps are given directly by the computed SVFs.

More specifically, the transformation mapping the globally normalised image $\tilde{I}_i$ into the atlas space at time point t_i is given by

$$\boldsymbol{\phi}_{k,i} = \exp\left(\boldsymbol{v}_{k,i}\right) \tag{6.23}$$

where

$$\boldsymbol{v}_{k,i} = \arg\min_{\boldsymbol{v}} \mathcal{E}_{\text{svffd}}^{\text{sym}}(\boldsymbol{v}; \bar{I}_{k-1}(t_i), \tilde{I}_i) \tag{6.24}$$

is the (local) minimum of the cost function (5.5) of the symmetric SVFFD registration. Given the global normalisation, the initial deformation is the identity, i.e., $\boldsymbol{v}_{0,i} = \boldsymbol{0}$.

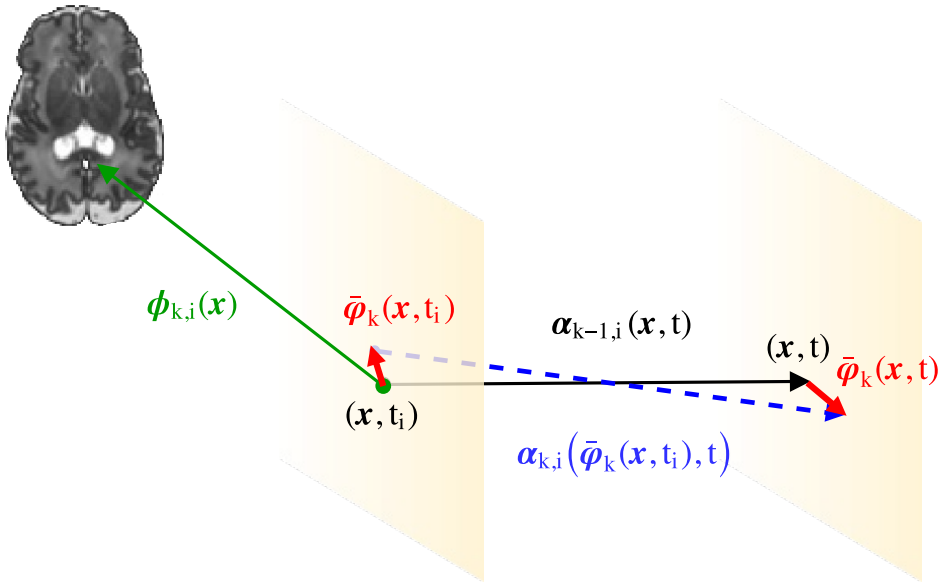


Figure 6.3.: Recursive estimation of mean growth via composition of space deformations. The cross-sectional SVFFD $\boldsymbol{\phi}_{k,i}$ (green arrow), computed by non-rigid registration of the i-th image to the atlas template at age t_i (left plane), is composed with the current longitudinal deformation $\boldsymbol{\alpha}_{k-1,i}$ (black arrow) to map the image to time t (right plane). This longitudinal estimate of mean growth evolves together with the spatial anatomical coordinates at the two time points which are deformed at iteration k by the age-specific mean deformations $\bar{\boldsymbol{\varphi}}_k$ (red arrows), resulting in the new estimate $\boldsymbol{\alpha}_{k,i}$ (blue arrow).

6.3.2. Joint estimation of mean shape and growth

In the previous section, the transformations which relate each brain image to different time points were considered to have the same spatial co-domain. This neglects the sequence of residual mean deformations $\bar{\boldsymbol{\phi}}_k$ that deform the atlas space at each time point differently for the mean shape to flow towards the weighted barycentre of the nearby observations at the respective age. As illustrated in figure 6.3, we can track the sequence of space deformations applied at each time t, and use these to derive a longitudinal coordinate map $\boldsymbol{\alpha}_{k,i}(t) : \Omega_{t_i} \to \Omega_t$ which relates an observed age t_i to time point t in order to correct for this anatomical mismatch in the co-domain, i.e.,

$$\begin{aligned} \boldsymbol{\alpha}_{k,i}(t) &= \bar{\boldsymbol{\varphi}}_k(t) \circ \boldsymbol{\alpha}_{k-1,i}(t) \circ \bar{\boldsymbol{\varphi}}_k^{-1}(t_i) \\ &= \left(\prod_{s=1}^{k} \bar{\boldsymbol{\varphi}}_s(t) \right) \left(\prod_{s=1}^{k} \bar{\boldsymbol{\varphi}}_s(t_i) \right)^{-1} \end{aligned} \tag{6.25}$$

where $\prod$ denotes a sequence of function compositions, and $\bar{\boldsymbol{\varphi}}$ are the adjusted Log-Euclidean mean transformations defined below. The non-recursive definition follows from the assumption that the initial longitudinal deformation between time t_i and any other time point is the identity transformation, i.e., $\boldsymbol{\alpha}_{0,i}(t) = \mathrm{Id}$. Because each transformation is the exponential map of a SVF, all compositions can be approximated using the BCH formula (Bossa et al., 2007), such that the composite transformation is another SVFFD.

The composition of the atlas-to-subject transformation defined in (6.23), which is obtained by registering the i-th image to the atlas template $\bar{I}_{k-1}(t_i)$, with the longitudinal deformation (6.25), yields an age-specific deformation $\boldsymbol{\varphi}_{k,i}(t) : \Omega_t \to \Omega_i$, i.e.,

$$\boldsymbol{\varphi}_{k,i}(t) = \boldsymbol{\phi}_{k,i} \circ \boldsymbol{\alpha}^{-1}_{k-1,i}(t) \tag{6.26}$$

For a set of discrete time points, these transformations may be found by separate registrations of the i-th image to the different atlas templates as discussed in section 6.2.5. However, obtaining it via composition with an estimate of the longitudinal atlas deformation instead, explicitly enforces temporal consistency and requires only n registrations.

The residual atlas deformation is now given by the Log-Euclidean mean of the age-adjusted atlas-to-subject transformations, i.e.,

$$\bar{\boldsymbol{\varphi}}_k(t) = \exp\left(\sum_{i=1}^{n} w_i(t) \log \boldsymbol{\varphi}_{k,i}(t)\right) \tag{6.27}$$

which are further used to redefine the total age-dependent deformations, i.e.,

$$\begin{aligned} \boldsymbol{T}_{k,i}(t) &= \bar{\boldsymbol{\varphi}}_k(t) \circ (\boldsymbol{\varphi}_{k,i}(t))^{-1} \\ &= \bar{\boldsymbol{\varphi}}_k(t) \circ \boldsymbol{\alpha}_{k-1,i}(t) \circ \boldsymbol{\phi}^{-1}_{k,i} \end{aligned} \tag{6.28}$$

This concludes the definition of the coordinate maps involved in this new recursive spatio-temporal atlas construction based on the diffeomorphic SVFFD model and the Log-Euclidean mean. As mentioned, all function compositions are efficiently approximated using the BCH formula, resulting in a compact representation of all coordinate maps by the spline coefficients of SVFFDs.

While the mean intensity and shape template of the spatio-temporal atlas can be evaluated for any time t, during the atlas construction, only those template images corresponding to the imaged time points t_i are needed. The result of the atlas construction is not only the spatio-temporal atlas of mean shape and intensity, and the transformations mapping each anatomy into the atlas space, but also an estimate of the longitudinal deformations corresponding to mean growth. It should be noted, though, that these deformations primarily reflect coarse anatomical changes, which is in part due to the approximations and the remaining misalignment after deformable registration, but also because finer details are the result of the temporal kernel regression of intensities at each atlas coordinate. The final mean shape is a combination of both, the sequence of Log-Euclidean mean deformations, and the regression of the intensities.

In summary, the atlas construction iterates the following steps for $k = \{1, ..., K\}$:

1. Generate $|\{t_i\}| \leq n$ template images $\bar{I}_{k-1}(t_i)$ given $\tilde{I}_i$ and $\boldsymbol{T}_{k-1,i}$ (6.19).
2. Compute $\boldsymbol{\phi}_{k,i} = \exp(\boldsymbol{v}_{k,i})$ which deform each image to match their template (6.24).
3. Compose longitudinal maps $\boldsymbol{\alpha}_{k-1,i}\, \forall t \in \{t_j | w_i(t_j) > 0\}$ given $\boldsymbol{\alpha}_{k-2,i}$ and $\bar{\boldsymbol{\varphi}}_{k-1}$ (6.25).
4. Compose cross-sectional maps $\boldsymbol{\phi}_{k,i}$ with $\boldsymbol{\alpha}_{k-1,i}\ \forall t \in \{t_j | w_i(t_j) > 0\}$ (6.26).
5. Compute Log-Euclidean means $\bar{\boldsymbol{\varphi}}_k$ for each observed time t_i (6.27).

After about 5 to 10 iterations (or convergence based on the residual mean deformations), template images of mean intensity and shape can be generated for a set of discrete time points together with the longitudinal deformations between consecutive templates.

These mean growth estimates may be used to initialise a longitudinal registration of the atlas time series, for example, using the TDFFD model (De Craene et al., 2012) with initial non-stationary velocities set equal to the piecewise SVFs of the SVFFDs. Such longitudinal registration is required due to aforementioned reasons, i.e., that the mean growth estimated by composition of mean atlas deformations only reveals coarse longitudinal changes. Structural changes resulting from the temporal regression of intensities are not considered by these deformations. A possible extension of our method is to integrate a refinement of the longitudinal coordinate maps at each iteration of the atlas construction.

6.4. Results

6.4.1. Image acquisition and reconstruction

The brain atlases presented in the following are constructed from MR images acquired by the dHCP[8]. The image data was collected at St. Thomas Hospital, London, on a Philips 3T scanner using a 32 channel dedicated neonatal head coil (Hughes et al., 2017). To reduce the effects of motion, T2w images were obtained using a turbo spin echo (TSE) sequence, acquired in two stacks of 2D slices in sagittal and axial planes, using parameters: TR = 12 s, TE = 156 ms, SENSE factor 2.11 (axial) and 2.58 (sagittal). Overlapping slices with resolution $0.8 \times 0.8 \times 1.6$ mm^3 were acquired. During motion corrected reconstruction (Cordero-Grande et al., 2016; Kuklisova-Murgasova et al., 2012), the images are resampled to an isotropic voxel size of 0.5 mm. T1w images were acquired using an inversion recovery TSE sequence at the same resolutions with TR = 4.8 s, TE = 8.7 ms, SENSE factor 2.26 (axial) and 2.66 (sagittal). Babies were not sedated, and instead were encouraged to sleep naturally in the scanner. All images were reviewed by a paediatric neuroradiologist and checked for possible abnormalities.

6.4.2. Preprocessing and whole brain segmentation

The reconstructed images were preprocessed using a structural image processing pipeline designed for dHCP. The preprocessing includes a brain extraction (Smith, 2002) and bias field correction (Tustison et al., 2010). The bias corrected MR images are segmented into 87 subcortical, and left and right brain hemisphere structures using a modified version of the developing region annotation with expectation maximisation (Draw-EM) algorithm developed by Makropoulos et al. (2014). This improved whole brain segmentation is implemented in the publicly available Draw-EM package[9]. For details regarding these processing steps, see the analysis pipeline summary of Makropoulos et al. (2017b).

[8] http://www.developingconnectome.org/ (last accessed: Aug 12, 2017)
[9] https://github.com/MIRTK/DrawEM/ (last accessed: May 20, 2017)

6.4.3. Brain atlas evaluation criteria

A number of measures were used to assess the quality of the brain registrations and the image alignment in the common average space, after applying the respective mean deformations. These mean deformations are either an arithmetic mean of displacement vectors in case of the classic FFD model, or the proposed Log-Euclidean mean of the diffeomorphic SVFFD model. A more accurate alignment directly translates into sharper brain templates. These evaluation criteria are consequently used to select the optimal registration parameters, and to compare brain atlases built with different approaches. First, we define the per-voxel/-structure measures, before detailing their mean calculation.

Gradient magnitude

The magnitude of the gradient of the image of mean intensities in the atlas space, i.e.,

$$M_k(\boldsymbol{x}) = \frac{1}{2h} \left\| \bar{I}_{\text{local}}(\boldsymbol{x} + h, t_k) - \bar{I}_{\text{local}}(\boldsymbol{x} - h, t_k) \right\| \tag{6.29}$$

where the gradient is approximated using a central difference scheme with h equal to the voxel size, is proportional to the sharpness of image edges. When boundaries of anatomical structures with distinct MR appearance are aligned well, such as in particular the pial and inner-cortical surface, then we can observe a strong gradient at this boundary. This evaluation criterion has also been used by Gholipour et al. (2014).

To quantify specifically the sharpness of the cortex in the mean brain images, we also assess the gradient magnitude of the cGM probability map of the constructed atlas.

Standard deviation

This criterion quantifies how similar the normalised intensity samples of the n brain images are at a specific atlas coordinate. It is the well known standard deviation given by

$$S_k(\boldsymbol{x}) = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \left(\tilde{I}_i(\boldsymbol{x}, t_k) - \bar{I}_{\text{local}}(\boldsymbol{x}, t_k) \right)^2} \tag{6.30}$$

The lower this value, the lower is the standard error of the mean intensity.

Intensity entropy

This information theoretic measure quantifies the uncertainty of the mean intensity estimate. In our context, it serves as an index of the diversity of individual intensity samples. The voxel-wise entropy at point $\boldsymbol{x}$ of the atlas at time point t_k is given by

$$H_k(\boldsymbol{x}) = -\sum_{j=1}^{b} \left(\frac{h_j(\boldsymbol{x}, t_k)}{n} \right) \log \left(\frac{h_j(\boldsymbol{x}, t_k)}{n} \right) \tag{6.31}$$

where $b = 100$ is the number of bins of the histogram $\boldsymbol{h}(\boldsymbol{x}, t_k) \in \mathbb{N}_0^b$ used to estimate the probability distribution of normalised image intensities at spatio-temporal atlas coordinate $(\boldsymbol{x}, t_k)$, and $h_j(\boldsymbol{x}, t_k)$ is number of intensity values which fall into the j-th histogram bin.

Label entropy

A segmentation-based voxel-wise entropy measure for evaluation of a spatial normalisation procedure was proposed by Robbins et al. (2004). In this case, the entropy of the labels assigned to a specific point in the standard space is evaluated. When all labels are in agreement, i.e., the atlas label for a given point is uniquely determined, the entropy is zero. It attains its highest value when all labels are equally likely, i.e., when each brain segmentation suggests a different label for a given atlas coordinate. The label entropy at point $\boldsymbol{x}$ of the atlas at time point t_k is given by

$$LE_k(\boldsymbol{x}) = -\sum_{l \in \mathcal{L}} \left(\frac{h_l(\boldsymbol{x}, t_k)}{n} \right) \log \left(\frac{h_l(\boldsymbol{x}, t_k)}{n} \right) \tag{6.32}$$

where $h_l(\boldsymbol{x}, t_k)$ is the number of times label l is propagated from one of the brain segmentations by their respective image to atlas transformations using nearest neighbour interpolation, and $\mathcal{L}$ denotes the set of integral labels, i.e., either the set of Draw-EM tissue classes $\mathcal{L}_\mathrm{T} = \{1, \ldots, 9\}$ or the set of Draw-EM brain structure labels $\mathcal{L}_\mathrm{S} = \{1, \ldots, 87\}$. We refer to the entropy of the tissue class labels as tissue (class) entropy, and the entropy of the whole brain structure labels as structure (label) entropy, respectively.

Pairwise overlap

This overlap measure quantifies the accuracy of the pairwise registrations using the automatic whole brain segmentation. Given the number of true positive (TP), true negative (TN), false positive (FP), and false negative (FN) label assignments for each structure and pair of segmentations of the i-th and j-th images propagated to the atlas at time point t_k, the mean Dice similarity coefficient (DSC) (Dice, 1945) of the l-th structure is given by

$$DSC_k(l) = \frac{1}{n^2} \sum_{i,j=1}^{n} \frac{2\ TP_{ijk}(l)}{2\ TP_{ijk}(l) + FP_{ijk}(l) + FN_{ijk}(l)} \tag{6.33}$$

Temporal consistency

To quantify the effect of the choice of temporal kernel width, we define a measure of temporal consistency (TC). The TC of structure l of a spatio-temporal atlas sampled at n_t ordered time points t_k, i.e., $t_k < t_{k+1}\ \forall k \in [1, n_\mathrm{t} - 1]$, is the generalised DSC given by

$$TC(l) = \sum_{k=1}^{n_\mathrm{t}-1} \frac{2|\boldsymbol{p}_l(t_k) \cdot \boldsymbol{p}_l(t_{k+1})|}{\|\boldsymbol{p}_l(t_k)\|^2 + \|\boldsymbol{p}_l(t_{k+1})\|^2} \tag{6.34}$$

where the scalar product of two probabilistic segmentations is given by

$$\boldsymbol{p}_l(t_1) \cdot \boldsymbol{p}_l(t_2) = \int_\Omega p_l(\boldsymbol{x}, t_1) p_l(\boldsymbol{x}, t_2)\, d\boldsymbol{x} \tag{6.35}$$

and the squared ℓ_2-norm of each probabilistic segmentation is $\|\boldsymbol{p}_l(t)\|^2 = \boldsymbol{p}_l(t) \cdot \boldsymbol{p}_l(t)$.

Average of spatial measures

The spatial measures are averaged for each discrete atlas time point t_k using the probabilistic segmentation of the atlas into nine Draw-EM tissue classes. Given a spatial atlas quality measure $Q_k(\boldsymbol{x})$, the average for tissue class $c \in \mathcal{L}_\mathrm{T} = \{1, \ldots, 9\}$ is given by

$$\mu_k[Q](c) = \frac{\int_\Omega p_c(\boldsymbol{x}, t_k) Q_k(\boldsymbol{x})\, d\boldsymbol{x}}{\int_\Omega p_c(\boldsymbol{x}, t_k)\, d\boldsymbol{x}} \tag{6.36}$$

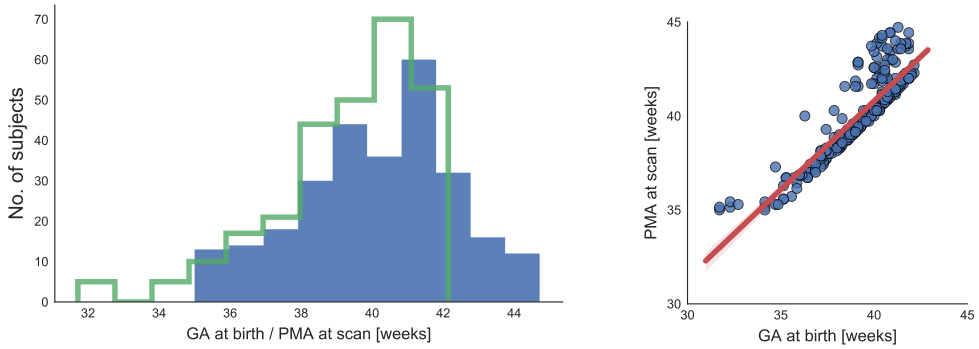


Figure 6.4.: Distribution of GA at birth (left, steps) and PMA at scan (left, bars). The scatter plot on the right is overlayed with the regression line mapping from GA at birth to PMA at scan. The slope close to 1 shows that most neonates were scanned shortly after birth, while the line intercept of 0.8 weeks corresponds to the average time to scan.

6.4.4. Selection of neonatal brain images

Out of 495 processed brain MR scans of 474 subjects, we selected images of 275 neonates. This set of brain images was obtained via a process of exclusion based on (a) a manual image quality score between 1 (poor) and 4 (good), excluding images with score 1 or 2, (b) the availability of reconstructed and motion corrected T1w and T2w images, (c) the exclusion of follow up scans, (d) a maximum time to scan of 4 weeks, and (e) a minimum age at scan of 35 weeks PMA. The age distributions are shown in figure 6.4, while individual ages are listed in appendix D.1.

6.4.5. Selection of model parameters

Before we compute all registrations required by the pairwise atlas construction, we determine suitable parameters for the deformable registration based on the defined evaluation criteria using a small subset of all images. For this, we selected three random subsets of images from 10 subjects each: (a) 5 neonates scanned at 36 weeks, and 5 neonates scanned at 38 weeks, (b) 5 neonates scanned at 42 weeks, and 5 neonates scanned at 44 weeks, (c) 3 of the 5 neonates scanned at 38 weeks, 3 of the 5 neonates scanned at 42 weeks, and 4 neonates scanned at 40 weeks. Given these image sets, we construct a brain template for each of the three age groups without temporal kernel regression. The defined quality measures are first evaluated for each age group, and then averaged across the subsets.

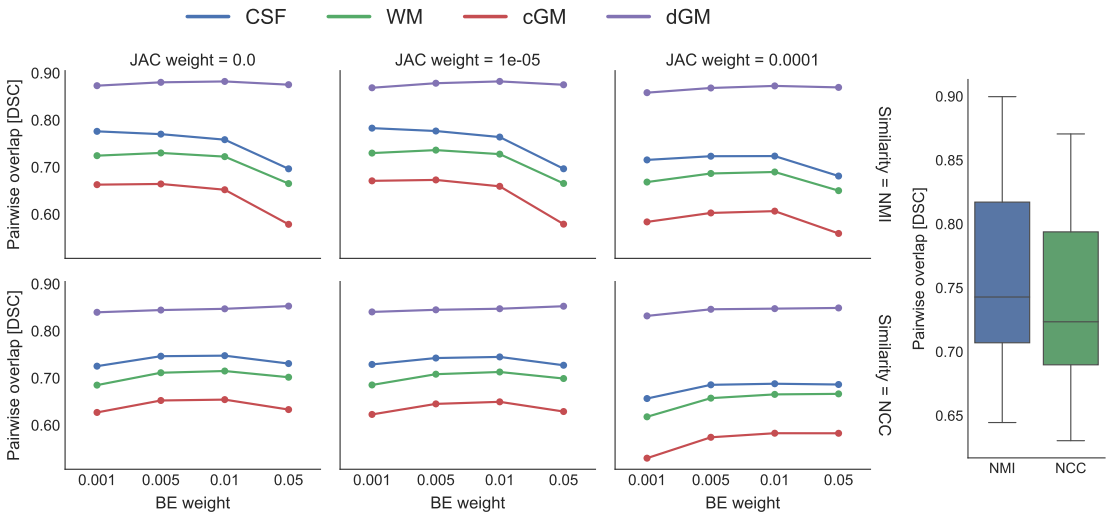


Figure 6.5.: Mean DSC of SVFFD registrations for different dissimilarity measures and regularisation weights. Mean overlap of brain structures, averaged separately for each tissue, is consistently higher for NMI compared to LNCC. On the right, the DSCs for different BE and JAC weights for which both dissimilarity measures achieved best overlaps are compared, i.e., BE weight of 0.001 or 0.005, and JAC weight of 0 or 10^{-5}, respectively.

Image dissimilarity measure

First, we compared the pairwise registration results for two dissimilarity measures commonly used for neonatal brain image registration, namely NMI and LNCC. In the work of Gholipour et al. (2014, 2017) on constructing a fetal brain atlas using SyN, the latter has been found to result in slightly sharper templates. Figure 6.5 plots the mean DSC of our SVFFD algorithm for different weights of the bending energy (BE) and Jacobian-based penalty (JAC) terms[10]. It can be observed that in case of MIRTK and our neonatal dataset, the overlap obtained with NMI is consistently higher, while at the same time requiring considerably less computation time (on average less than 40% of CPU time). These results correspond to a window size of 7x7x7 voxels for LNCC and 64 histogram bins used to estimate the joint probability distribution for the computation of NMI. We confirmed the same results for different LNCC window sizes and number of NMI bins.

Given the better performance and lower computational cost of NMI, we used the NMI image dissimilarity measure for the pairwise neonatal brain image registrations.

[10]Note that the results presented in this chapter are based on an earlier version of our SVFFD algorithm summarised in section 5.3, where only $\mathcal{S}_{\text{BE}}$ (BE) and $\mathcal{P}_{\text{VP}}$ (JAC) regularisation terms are used.

Control point spacing and regularisation weights

Next, we performed a grid search on the parameters regularising the smoothness of the pairwise transformations, namely the control point spacing and the weights of the BE and JAC terms of the energy function. The mean DSC for combinations of parameter values achieving the highest overlaps are shown in figure 6.6. Based on the average overlap of the brain structures for different registration parameters, we limit the parameter values further and examine the mean of the evaluation criteria for these more narrow ranges in figure 6.7. More specifically, we evaluated the quality measures for a control point spacing of 2.0 and 2.5 mm, a BE weight of 0.001, 0.005, and 0.01, and a JAC weight of 0 and 0.00001, respectively. This pre-selection of narrow parameter ranges is done so the mean values over the remaining two parameters, when examining the influence of a single parameter, are not negatively impacted by a sub-optimal choice of these other parameters. The best results using the SVFFD registration of MIRTK are obtained with a control point spacing of 2.0 mm, a BE weight of 0.005, and a JAC weight of 0.00001.

For a comparison of the proposed atlas construction based on the Log-Euclidean mean of diffeomorphic SVFFDs to the reference atlas constructed using the arithmetic mean of FFDs, we perform a separate grid search for the parameters of the classic FFD algorithm. In case of the FFD model, the JAC term corresponds to an undesired volume preservation and its weight is therefore set to zero. The average segmentation overlap after pairwise FFD registration for different control point spacings and BE weights is shown in figure 6.8.

Additionally, because the FFD model does not ensure the required invertibility of the mean deformations, the number of negative Jacobian determinants of the corresponding arithmetic mean FFDs are summarised in figures 6.9 and 6.10. Note that the individual pairwise FFDs may have considerably more non-invertible points with negative Jacobian determinant, but we only need to be able to invert the smoother mean deformation fields. Given the observed numbers of negative Jacobian determinants and corresponding segmentation overlap measures, we restrict the registration parameters to a control point spacing of 2.0 and 2.5 mm, respectively, and a BE weight of either 0.001 or 0.005. The FFD parameter selection is then based on the atlas quality measures for these values.

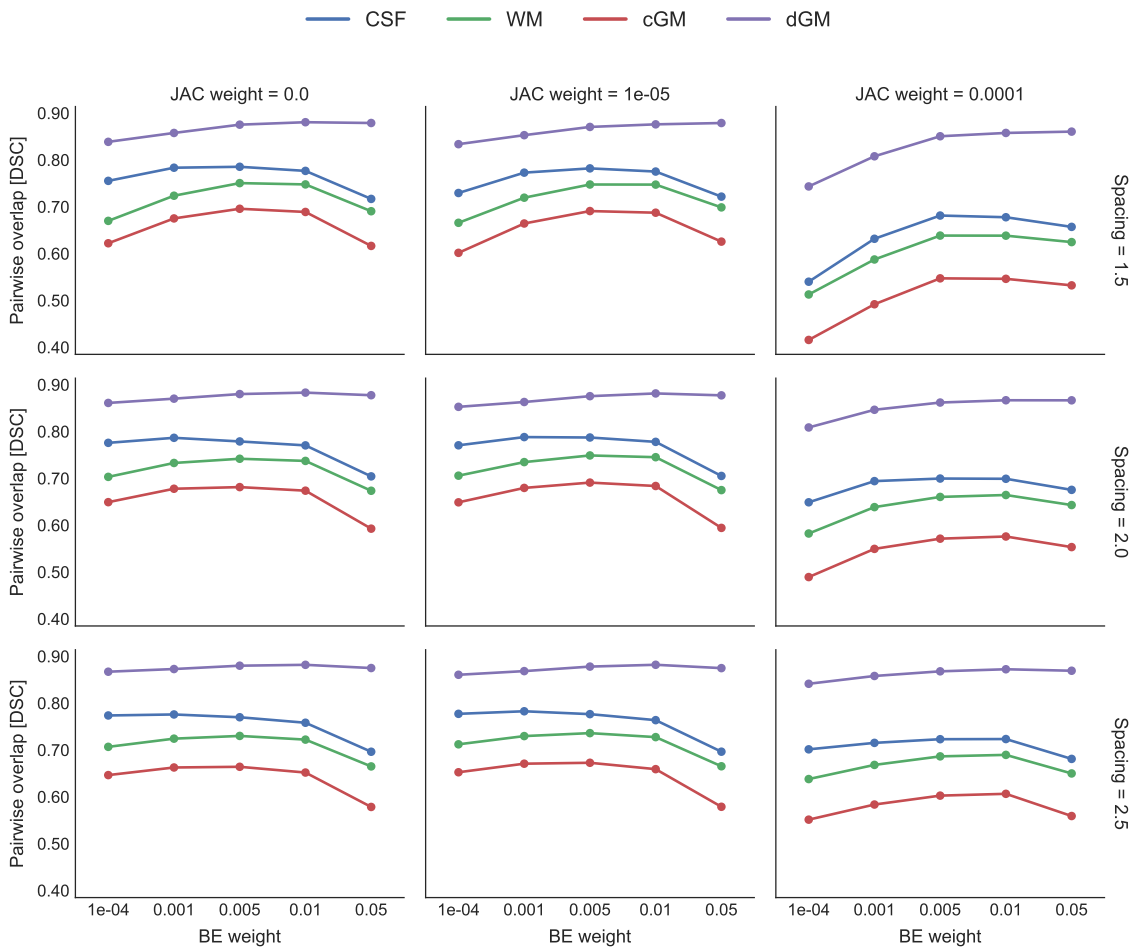


Figure 6.6.: Mean DSC for different combinations of SVFFD parameters.

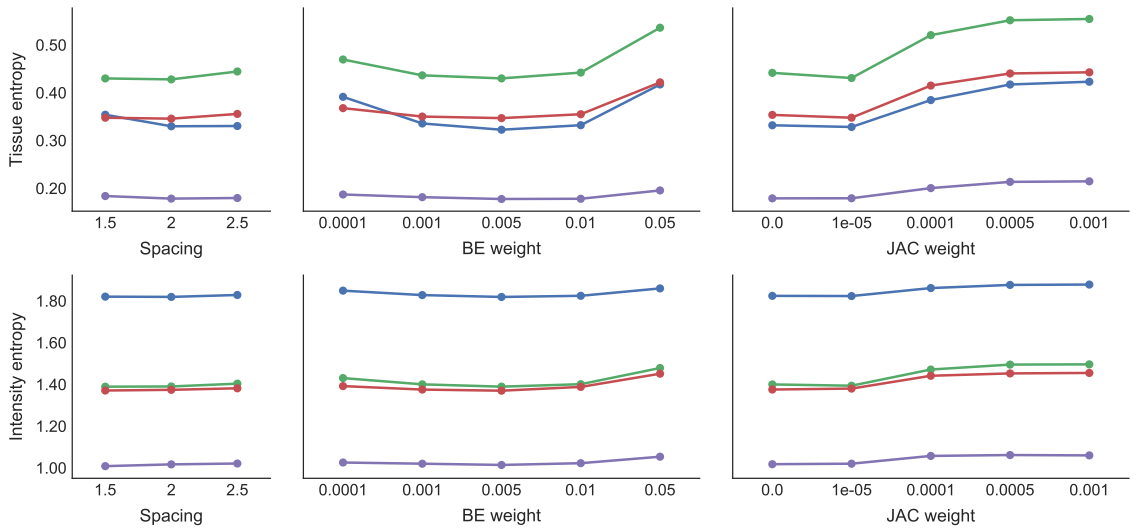


Figure 6.7.: Mean evaluation measures for SVFFD parameters selected after grid search.

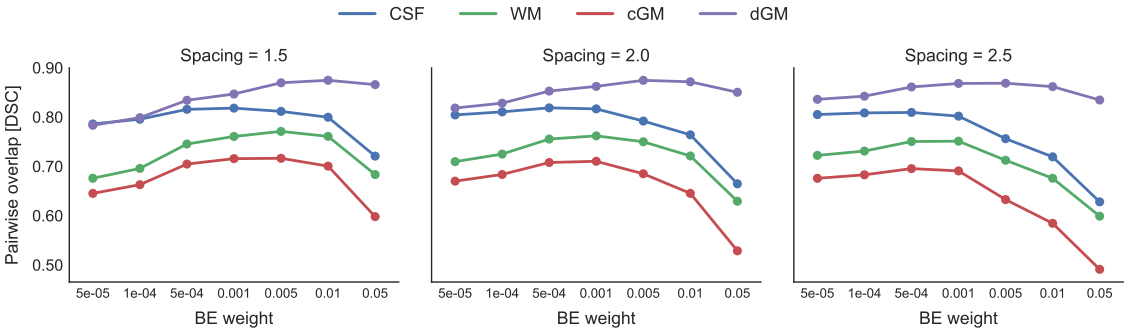


Figure 6.8.: Mean DSC corresponding to different combinations of FFD parameters.

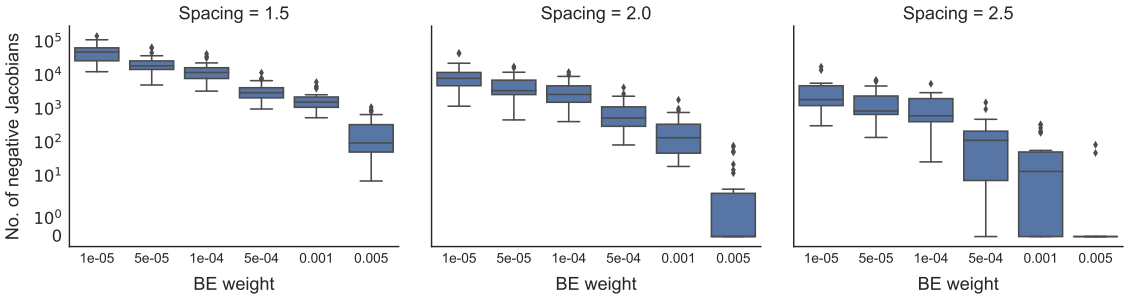


Figure 6.9.: Number of negative Jacobian determinants of arithmetic mean FFDs.

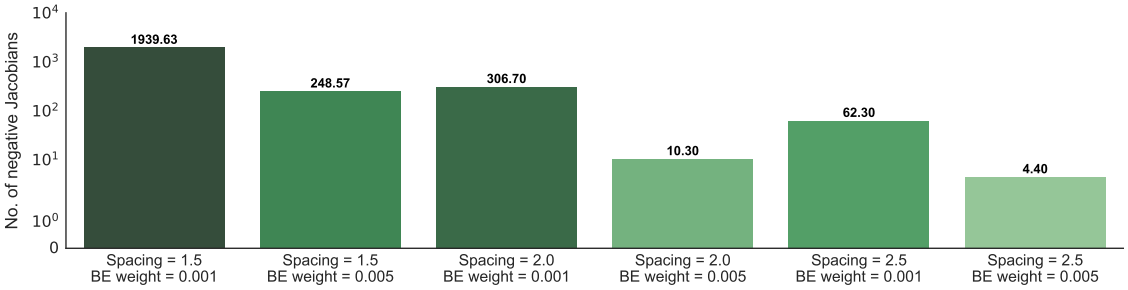


Figure 6.10.: Number of negative Jacobian determinants for selected FFD parameters.

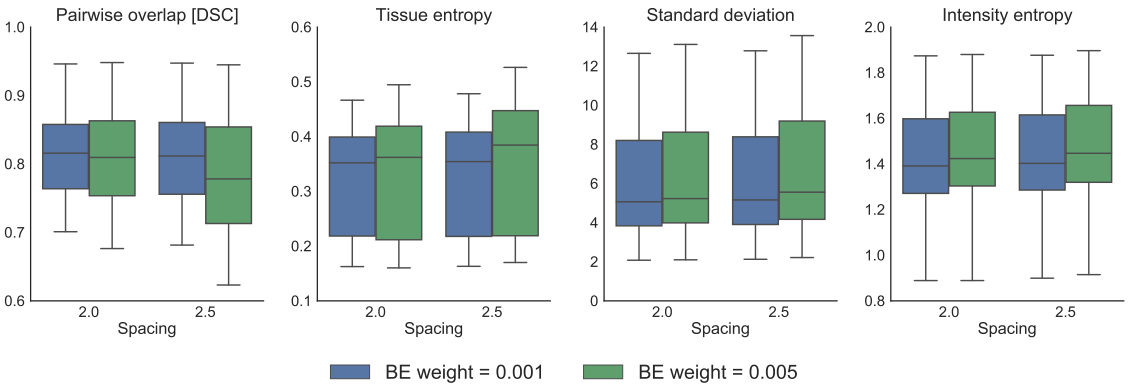


Figure 6.11.: Evaluation criteria for selected FFD parameters.

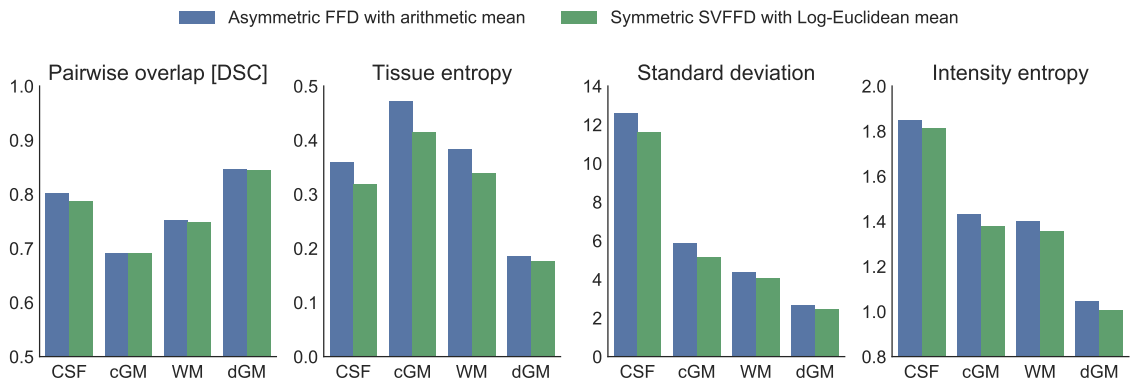


Figure 6.12.: Comparison of evaluation measures for the different pairwise atlas constructions based on the subsets of in total 30 neonatal subjects used for parameter selection. These preliminary results suggest that, although overlap between image pairs is similar, the template created with the SVFFD model and Log-Euclidean mean is slightly sharper.

A comparison of the atlas quality measures for all four combinations of these parameters in figure 6.11 reveals that the best combination of evaluated parameter values for the images in the dHCP dataset is either a control point spacing of 2.0 mm combined with a BE weight of 0.005, or a control point spacing of 2.5 mm with a BE weight of 0.001. Which of these two choices is considered best depends on whether a lower number of Jacobian determinants is preferred over a slightly lower label overlap and consistency. Here, we constructed a reference atlas with the arithmetic mean of FFDs using a control point spacing of 2.5 mm and a BE weight of 0.001. This choice corresponds to slightly better performance according to the evaluation criteria based on which we compare the different approaches. At the same time, the number of negative Jacobian determinants still enables a robust approximation of inverse deformations using the Newton-Raphson method.

An initial comparison of the evaluation measures for the FFD and SVFFD models corresponding to the selected optimal parameters, and evaluated for the three subsets of in total 30 neonatal subjects scanned at different PMAs, is given in figure 6.12. This preliminary comparison of the evaluation measures for atlases constructed without temporal kernel regression suggests that both registration methods achieve similar segmentation overlap, but that the diffeomorphic SVFFD registration combined with the Log-Euclidean mean produces better atlas templates, i.e., suggesting an advantage of our inverse-consistent and diffeomorphic atlas construction over the method proposed by Serag et al. (2012a).

Standard deviation of temporal regression kernel

For the construction of mean images at weekly intervals of the age range from 36 to 44 weeks PMA, for which a sufficient number of brain scans is included in our dataset, we use an algorithm similar to Serag et al. (2012a) to determine an age-specific standard deviation, σ_t, of the temporal regression kernel for each age, $t_k \in \{36, 37, 38, 39, 40, 41, 42, 43, 44\}$. This kernel width selection differs from the algorithm proposed by Serag et al. (2012a) in that no stopping criterion based on the difference in brain volumes is considered, because our goal is a uniform distribution of brain samples at the equally spaced discrete time points. We first specify a constant kernel width, σ, and determine the median number, m, of images with non-zero kernel weight for each atlas age t_k. We then iteratively increase or decrease each kernel width σ_{t_k} by a small value until the number of images is equal m.

Using this algorithm, we determined adaptive kernel widths for different target σ values, and constructed a separate spatio-temporal atlas for each corresponding set of kernel widths listed in table 6.1, and visualised in figure 6.14. The TC of different tissue classes for each constructed atlas is shown in figure 6.13. As expected, it is observed that TC improves with increasing temporal kernel widths. When $\sigma \to \infty$, the atlas becomes constant in time, and thus has maximum TC, but fails to model temporal shape differences. Reducing the kernel width results in more distinct and sharper mean brain shapes for each time point, with fewer brain images contributing to each template image. Therefore, we also compare the average pairwise overlap of spatially normalised Draw-EM structures and the mean voxel-wise standard deviation of T2w image intensity samples in figure 6.15. We expect an atlas with a consistent quality across time to achieve a similar DSC value for each time point, without sacrificing a substantial reduction of overall quality. But we should also consider the higher anatomical detail at later time points due to increased cortical folding.

A target σ value below half a week results in insufficient overlap to ensure temporal smoothness of the mean brain images at weekly intervals. A target kernel width of 0.75 to 1 week produces temporally smooth atlases with noticeable temporal shape differences. After visual inspection, we found that a good compromise for our dataset is obtained for target $\sigma = 1$. The kernel width σ_t for any time point is obtained by linear interpolation.

$\boldsymbol{\sigma}$	0.25	0.50	0.75	1.00
$\boldsymbol{m}$	44	101	143	196
$\boldsymbol{\sigma_{36}}$	0.51	0.99	1.31	1.64
$\boldsymbol{\sigma_{37}}$	0.35	0.69	0.99	1.31
$\boldsymbol{\sigma_{38}}$	0.25	0.50	0.75	1.00
$\boldsymbol{\sigma_{39}}$	0.14	0.39	0.56	0.77
$\boldsymbol{\sigma_{40}}$	0.14	0.31	0.43	0.64
$\boldsymbol{\sigma_{41}}$	0.11	0.27	0.43	0.73
$\boldsymbol{\sigma_{42}}$	0.11	0.35	0.56	0.86
$\boldsymbol{\sigma_{43}}$	0.27	0.51	0.77	1.18
$\boldsymbol{\sigma_{44}}$	0.51	0.82	1.09	1.50

Table 6.1.: Adaptive kernel widths.

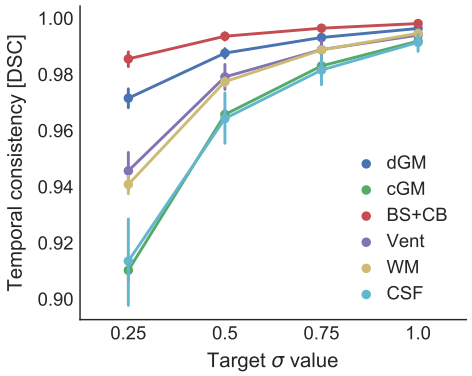


Figure 6.13.: TC for different target σ values.

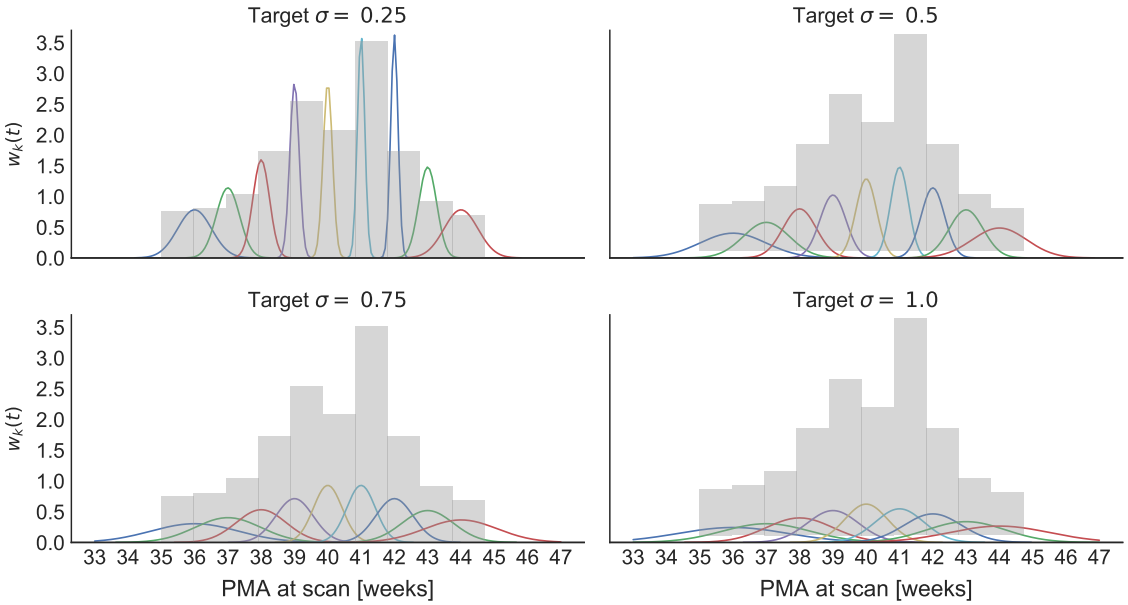


Figure 6.14.: Temporal regression kernels for different ages and target σ values. The respective standard deviations of the shown regression kernels centred at the discrete time points t_k are listed in table 6.1. The curves are overlayed on the PMA at scan histogram.

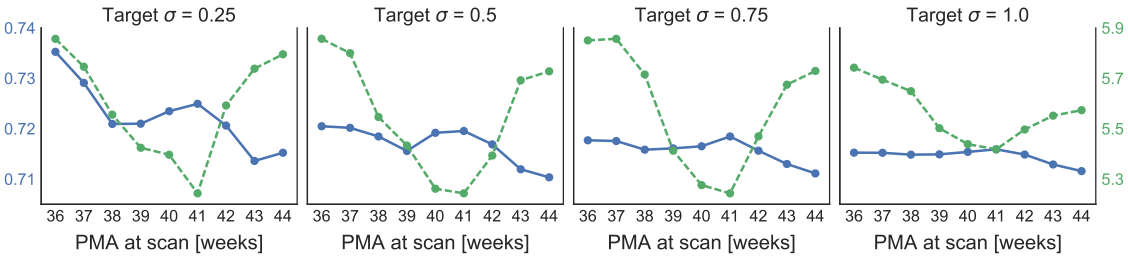


Figure 6.15.: Mean quality measures for discrete time points of atlas constructed with pairwise approach and different target σ values. The left vertical axis corresponds to the blue curve of the pairwise overlap (DSC), whereas the right vertical axis corresponds to the dotted green curve depicting the mean voxel-wise standard deviation of T2w intensities.

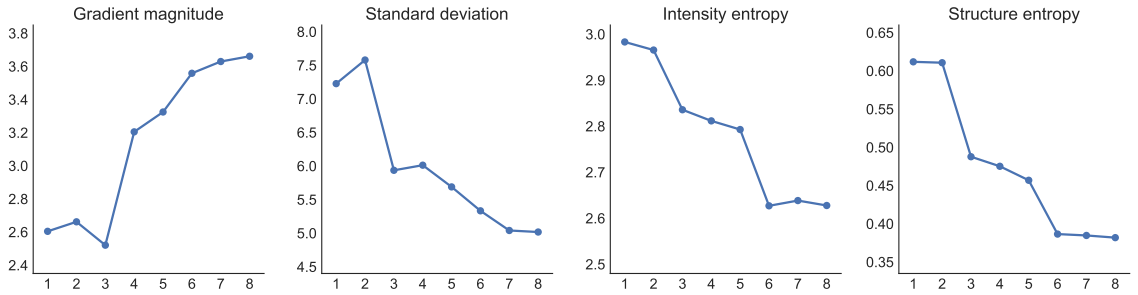


Figure 6.16.: Convergence of atlas quality measures in case of groupwise construction. The horizontal axis corresponds to number of completed iterations, k. The jumps from iteration $k = 2$ to $k = 3$, and from $k = 5$ to $k = 6$ are the result of halving the minimum control point spacing of the SVFFD at the highest image resolution level between iterations.

Parameters for groupwise atlas construction

Having determined the registration parameters for constructing a spatio-temporal neonatal atlas using direct registrations between all pairs of images, we performed another parameter search for the proposed groupwise approach based on previously determined parameter value ranges. This was done, because at each iteration of the groupwise approach, the target image is the current average brain image, which has a different appearance than an actual brain scan. For this groupwise parameter search, we constructed a single time point atlas at 40 weeks PMA from a subset of 64 brain images of term-born neonates. The SVFFD parameters chosen for the pairwise atlas construction were found to be suitable also for the groupwise method. Moreover, due to the iterative refinement, the approach demonstrated to be more robust towards variations in registration parameters. However, the initial target images for the pairwise registrations after affine alignment only are considerable noisy due to a significant amount of residual anatomical misalignment. We therefore constrain the deformations at early iterations more strongly using a sparser control point placement, and progressively half the control point spacing of the SVFFD after a few iterations. For the first two iterations, a control point spacing of 8 mm at the highest image resolution level is used. At iterations three to five, the control point spacing is 4 mm, whereas it is 2 mm for the final iterations. We found that the quality of the atlas did not substantially improve after 8 iterations. The convergence of a number of the atlas quality measures for the groupwise construction is shown in figure 6.16.

method	description
affine	Atlas after our proposed unbiased global normalisation.
reference1	Atlas construction of Serag et al. (2012a) using arithmetic mean of FFDs.
reference2	Atlas constructed with sum of mean FFDs equal to zero.
pairwise	Proposed atlas construction based on Log-Euclidean mean of pairwise SVFFDs.
refined	The "pairwise" atlas with 10 iterations of image to atlas deformation refinement.
groupwise	Proposed temporally consistent groupwise atlas construction.

Table 6.2.: Summary of compared spatio-temporal atlas construction methods.

6.4.6. Comparison of atlas construction methods

With the selected parameters, we constructed a spatio-temporal neonatal atlas for weekly intervals from 36 to 44 weeks PMA using the discussed construction methods. We compared the atlases constructed with our proposed approaches based on the SVFFD algorithm and Log-Euclidean mean deformations to the method of Serag et al. (2012a), which uses the classic FFD algorithm and arithmetic mean of displacements. Moreover, we applied the non-inverse-consistency correction noted in section 6.2 to this reference method. Specifically, the sum of inverse mean displacement fields is not ensured to be zero by the original approach. We therefore compose these inverse mean transformations with their arithmetic mean. This modified reference approach is referred to as "reference2", whereas the original method of Serag et al. (2012a) is "reference1". Our proposed pairwise and groupwise atlas construction methods are referred to as "pairwise" and "groupwise", respectively. Additionally, we perform an iterative refinement of the template brain images at discrete time points constructed with our pairwise approach as discussed in section 6.2.5. This approach is comparable to the groupwise fetal atlas construction of Gholipour et al. (2017) using our diffeomorphic SVFFD algorithm and the Log-Euclidean mean instead of SyGN. We compare this "refined" spatio-temporal atlas to our proposed "groupwise" approach. We also report these measures for the "affine" atlas after unbiased global alignment of brain images without shape deformation as a baseline of the evaluation criteria. A summary of the compared construction methods is given in table 6.2.

The mean of the voxel-wise measures averaged for each method over the entire brain masks separately for each discrete time point is reported in table 6.3. Box plots of these measures visualising also the temporal variance of the average measures are shown in

method	affine	reference1	reference2	pairwise	refined	groupwise
Gradient magnitude	1.764	2.837	2.868	3.096	3.648	**3.661**
Standard deviation	9.503	6.591	6.436	5.563	**4.821**	5.017
Intensity entropy	3.038	2.933	2.915	2.811	2.693	**2.628**
Tissue entropy	0.721	0.459	0.445	0.384	0.311	**0.297**
Structure entropy	0.825	0.526	0.513	0.454	0.389	**0.382**
cGM gradient magnitude	0.041	0.074	0.075	0.081	0.102	**0.108**

Table 6.3.: Overall mean quality measures of constructed atlases.

figure 6.17, whereas plots of individual average measures over time are shown in figure 6.18. The latter visualisation highlights the temporal smoothness even for average measures with a larger interquartile range, and explains this larger variance of average measures by the difference between early and late atlas time points.

As expected, a number of the evaluation measures are strongly correlated. This is especially the case for the voxel-wise label entropy of tissue classes and individual brain structures, because both are based on the automatic Draw-EM segmentations. We therefore focus on the structure entropy as it is based on a more fine granular parcellation. In regards to the intensity based evaluation criteria, the measures of intensity variance, namely standard deviation and entropy, demonstrate a similar relative performance of the construction methods. A comparison of these voxel-wise measures averaged separately for each tissue class in figure 6.19 shows that the apparent worse performance of the "groupwise" method when compared to the iteratively "refined" pairwise atlas is caused by a higher standard deviation of intensities for CSF. This is explained by the fact that this tissue class is at the boundary of the brain mask at which the NMI similarity criterion may not be evaluated during the registration as a result of only partial overlap of the foreground regions. An improved handling of partial boundary overlap during the registration would likely improve the within CSF measures. Of higher interest is however the atlas quality for brain tissues. The average measures for cGM, WM, and dGM for each constructed spatio-temporal atlas are compared in figure 6.20. It can be observed that all methods perform similarly well for deep brain structures, as these require less local deformation to be aligned, but that our proposed methods improve the quality of the atlas noticeably within WM and especially cGM regions. Our pairwise atlas construction using the Log-

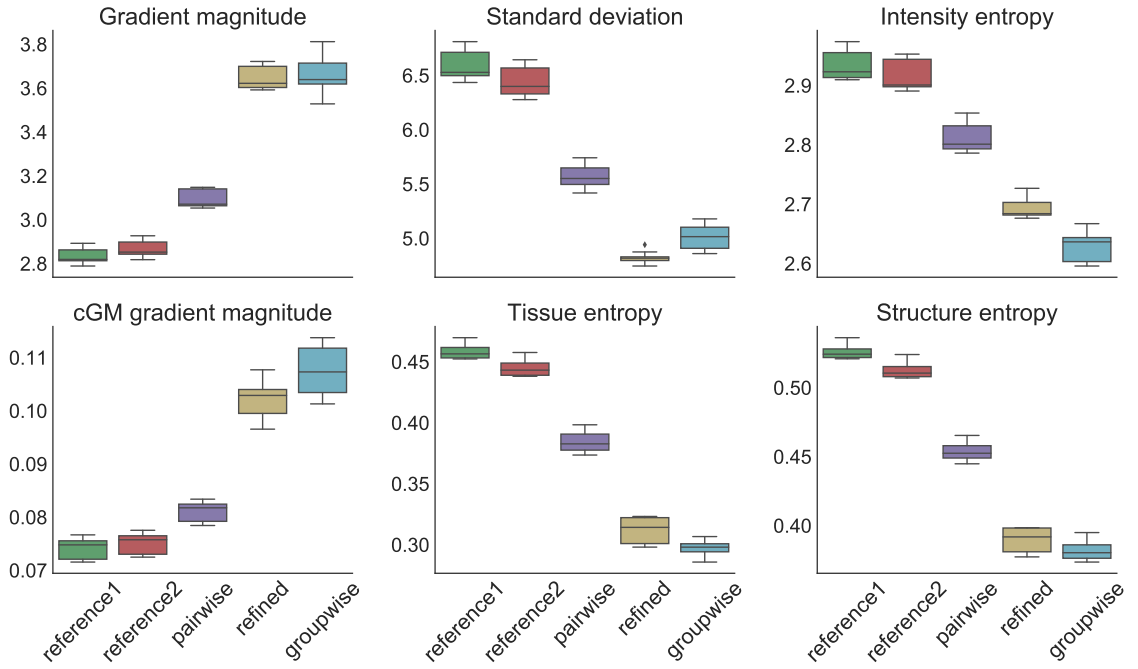


Figure 6.17.: Overall mean quality measures of spatio-temporal atlases with local spatial normalisation. The lower and upper bounds of the boxes correspond to the 25th and 75th percentile, respectively, of the nine values obtained by averaging the voxel-wise measures within the ROI containing all brain tissues for each discrete atlas time point. The height of each box reflects the variance of these average values across the ages. The mean of the average values corresponds to the centreline of the box (not shown). The median average value is shown as horizontal line inside each box.

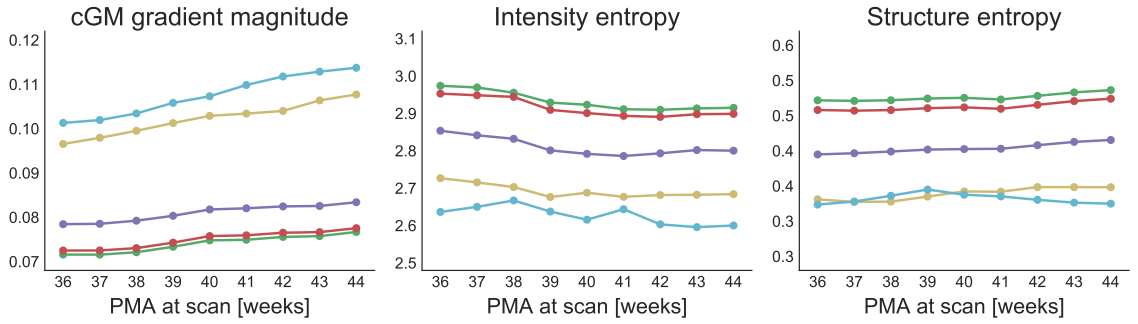


Figure 6.18.: Overall mean quality measures of constructed atlases over PMA at scan. Individual average values for each time point for selected evaluation measures shown in figure 6.17 with both wide (cGM gradient magnitude) and narrow (structure entropy) interquartile range. A wider interquartile range is due to stronger differences between early and late time points, but the atlas quality overall varies smoothly in time. A stronger average cGM gradient magnitude at later time points is explained by more cortical folding and the corresponding increase in cortical surface area.

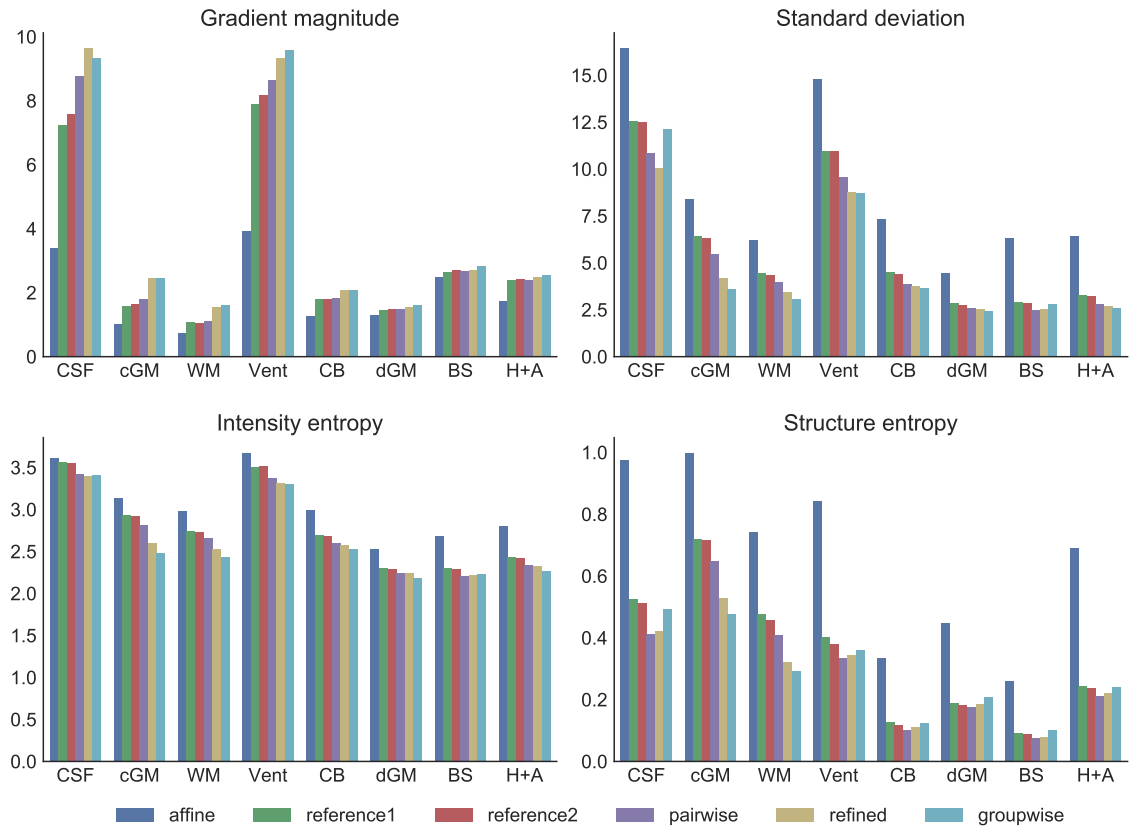


Figure 6.19.: Measures averaged separately for each tissue class: cerobrospinal fluid (CSF), cortical grey matter (cGM), white matter (WM), lateral ventricles (Vent), cerebellum (CB), deep grey matter (dGM), brain stem (BS), and hippocampi and amygdalae (H+A).

Euclidean mean of inverse-consistent diffeomorphic mappings is more consistent near the cortex than the approach proposed by Serag et al. (2012a). This is in part because we are able to more locally align cortical features with a lower control point spacing of the SVFFD which does not introduce singularities with negative Jacobian determinant. In contrast, the FFD model requires stronger regularisation in order to preserve the topology of the brain images. The iterative atlas constructions furthermore clearly outperform the pairwise methods due to the iterative compensation of residual anatomical misalignment. Our proposed "groupwise" spatio-temporal atlas construction performs slightly better than the individual iterative refinement of atlas templates at discrete time points. The TC measures shown in figure 6.21 suggest that our groupwise approach also results in more temporally consistent atlas templates, a difference which will become clearer after the visual assessment presented in the next section.

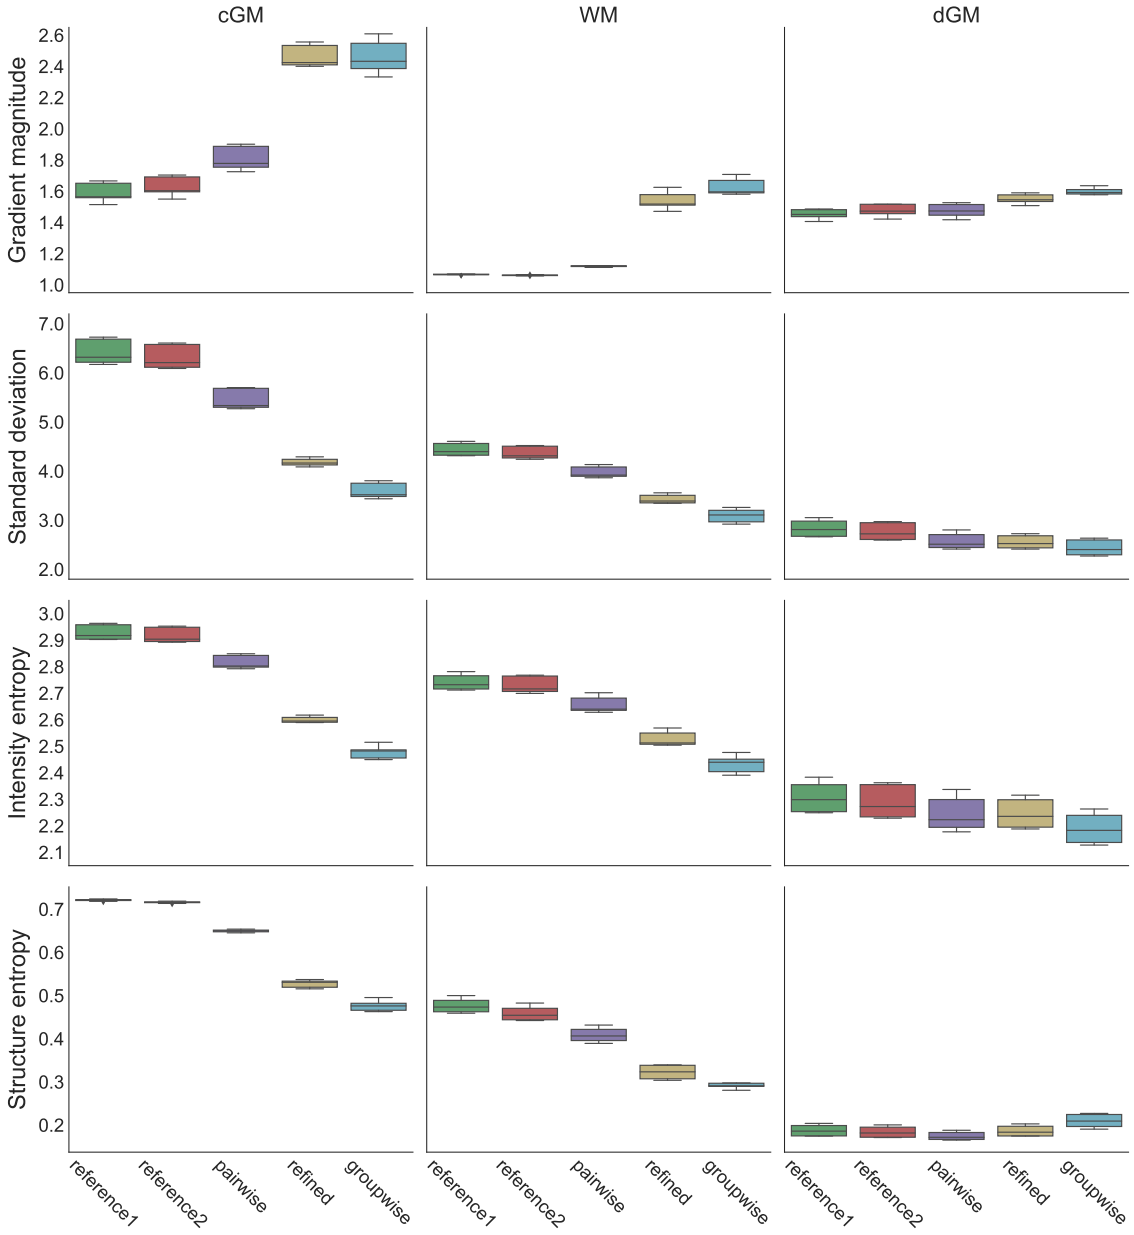


Figure 6.20.: Atlas quality measures averaged separately for cGM, WM, and dGM.

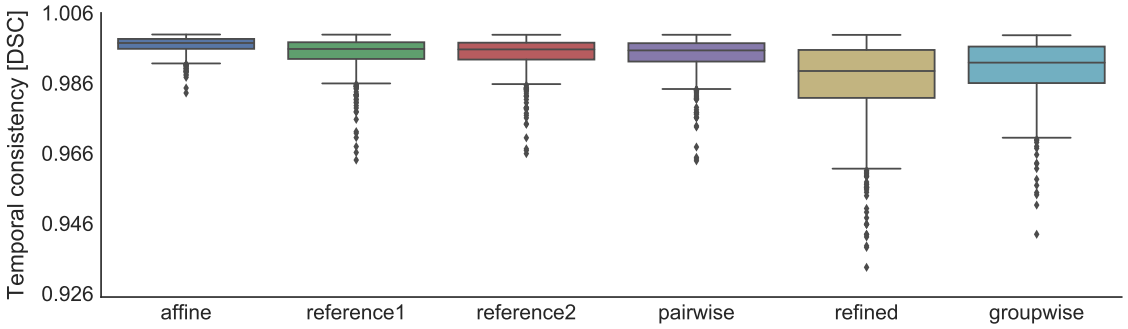


Figure 6.21.: Mean temporal consistency of atlas brain structures.

6.4.7. Visual assessment of neonatal atlases

Having observed improvements in the defined quantitative measures, we further assess the quality of the constructed atlases visually. Four time points at a regular interval of two weeks of the atlases constructed with the reference method of Serag et al. (2012a), our derived pairwise approach using the SVFFD model and Log-Euclidean mean instead of the arithmetic mean of FFDs, the template images of the pairwise method after 10 iterations of individual subject-to-atlas deformation refinement, and images generated by the proposed groupwise approach after 8 iterations are compared side by side in figure 6.22. To facilitate a direct comparison of local shape differences, these images depict the globally normalised templates, i.e., excluding the mean scaling reflecting a global change in brain size over time. Most noticeable are the sharper cortical details achieved by the iterative approaches[11]. Differences in the atlases obtained with either of the two pairwise or iterative techniques, respectively, are more subtle, and less clear by inspecting a single cross-section.

Slightly clearer cortical boundaries may be observed in the atlas created with our pairwise method when compared to the reference approach, in particular at the frontal lobe. These differences are more apparent when examining coloured renders of the voxel-wise intensity and tissue class entropy measures in figure 6.23. Our pairwise method achieves spatially more anatomical consistency in the alignment of the individual brain scans in the atlas space, but both methods produce atlases of high quality. An advantage of our approach is a consistent modelling of mean shape by topology preserving deformations.

The improved temporal consistency of the atlas constructed with our groupwise approach, as opposed to an individual iterative refinement of each time point, is highlighted in the detail views of smaller ROIs in figures 6.24 to 6.26. When brain scans are registered individually to each template, developing cortical structures may be missing at certain time points or some folding patterns recede instead of deepen further as expected for normal development. This is not observed in the temporal sequences generated by our groupwise method, which consistently maps each brain scan to each individual time point considering the longitudinal inter-atlas deformations.

[11] Note that the same Laplacian sharpening is applied to all images.

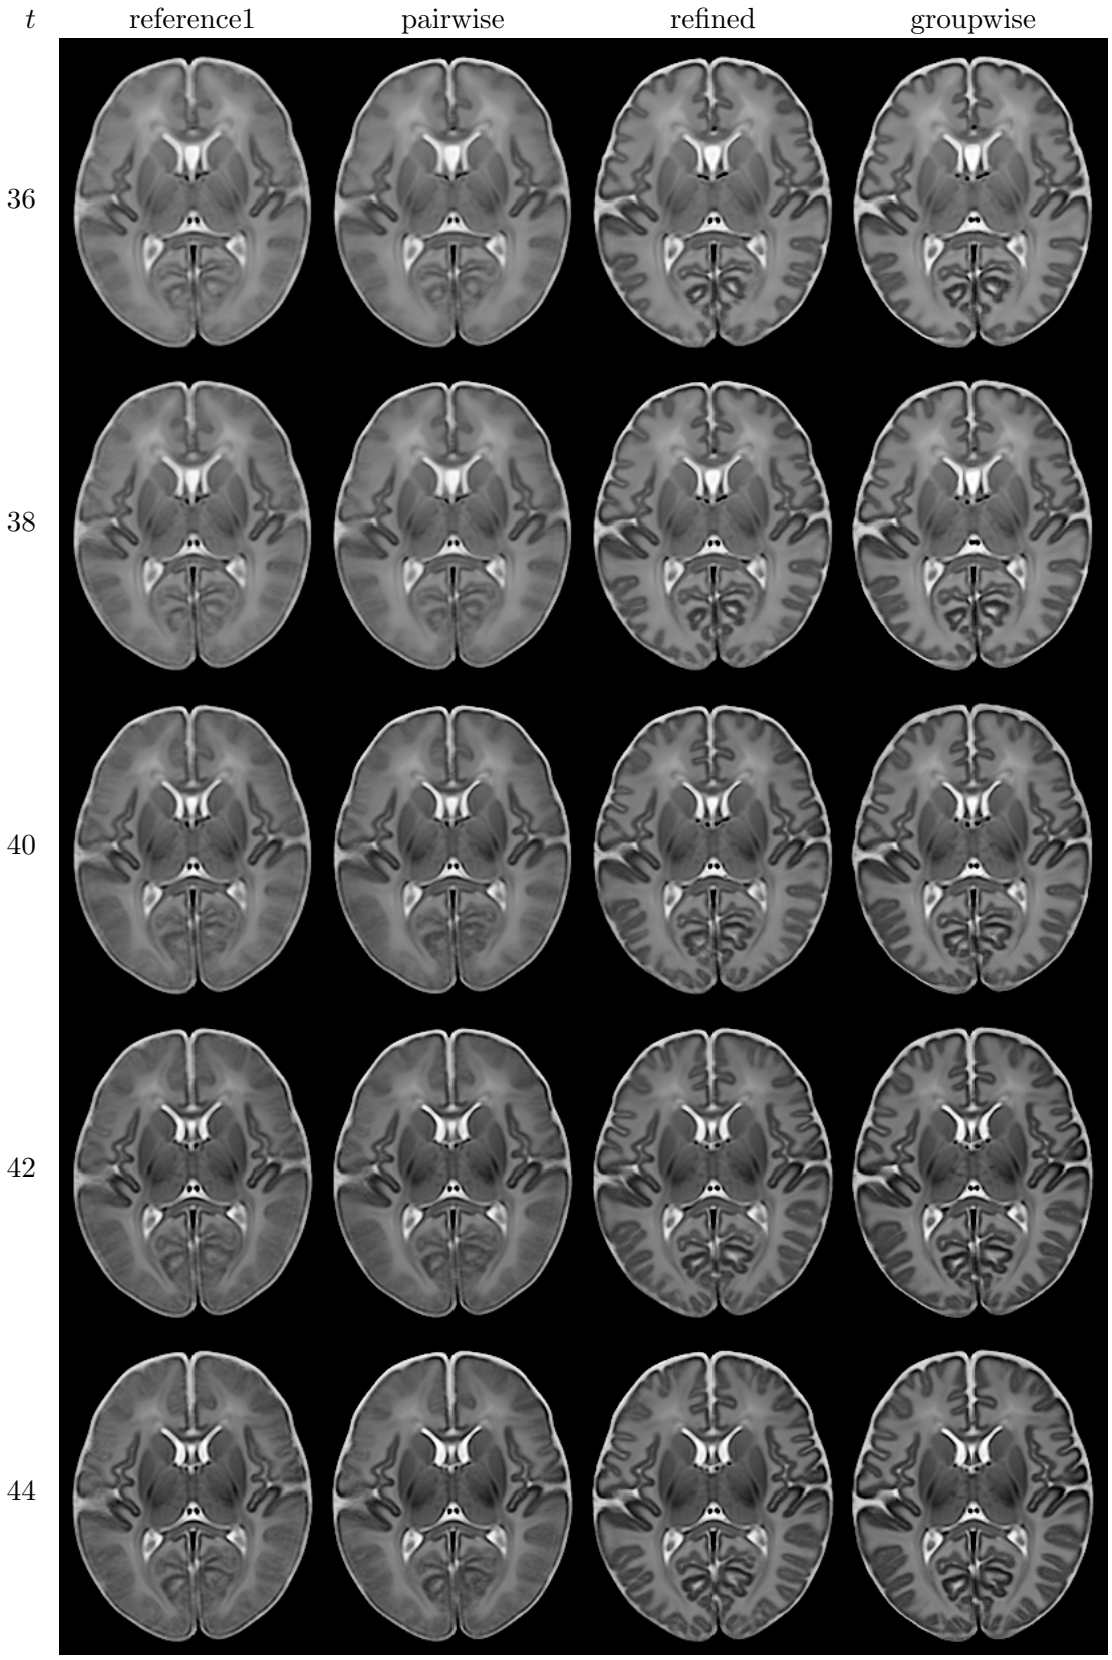


Figure 6.22.: Axial mid-sections of globally normalised spatio-temporal atlases.

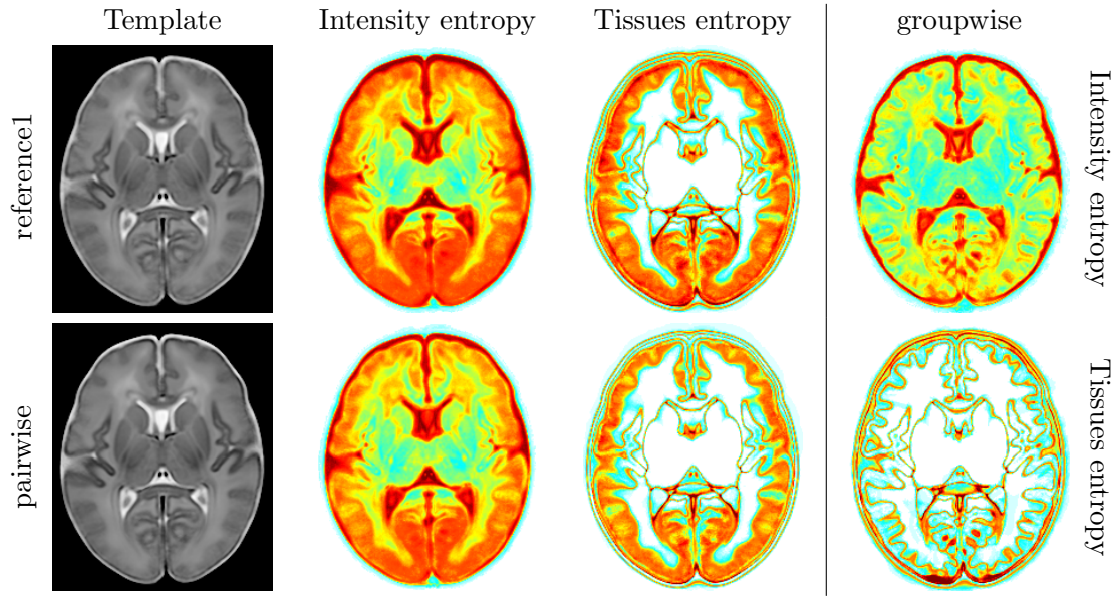


Figure 6.23.: Template at 36 weeks of spatio-temporal atlases constructed using the pairwise approach (top) using the FFD algorithm with arithmetic mean of displacements, and (bottom) our method using the SVFFD algorithm with Log-Euclidean mean of diffeomorphic mappings. Subtle differences (left) of the mean intensity at the cortex of the frontal lobe can be observed; (middle left) colour mappings of voxel-wise intensity and (middle right) tissue entropy reveal further cortical differences between the methods; (right) the voxel-wise entropy measures of the groupwise approach are presented here for comparison.

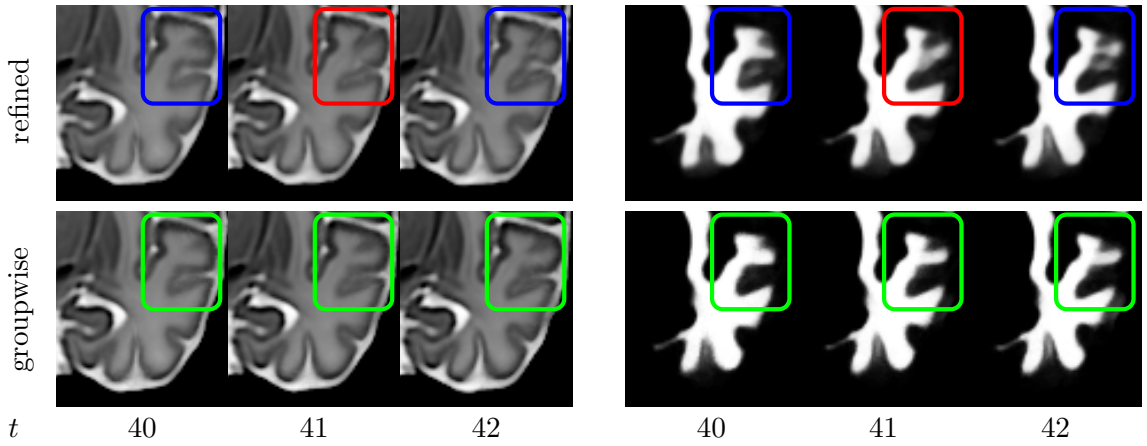


Figure 6.24.: Detail view of coronal cross-sections of spatio-temporal atlases constructed (top) with the pairwise approach followed by iterative refinement, and (bottom) our groupwise method. The three leftmost columns depict consecutive time points of the mean intensity template, whereas the three rightmost columns show the corresponding cGM probability maps. A temporal inconsistency of the marked major gyrus of the left temporal lobe is noticeable at week 41 (red box), when the time points were refined independently. Our groupwise method better preserves a consistent gyrus over time (green boxes).

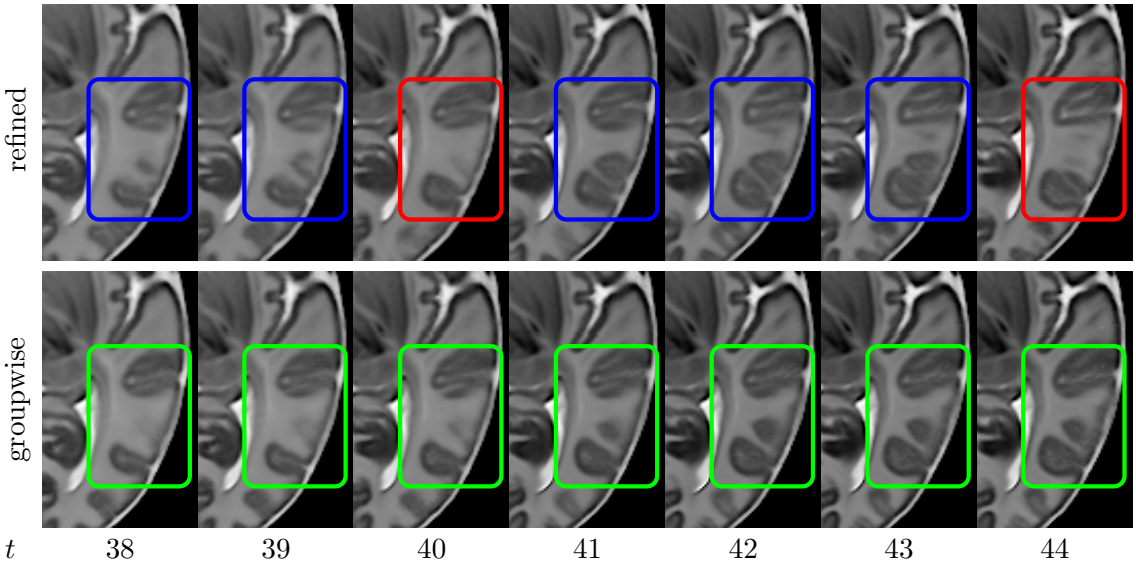


Figure 6.25.: Detail view of axial cross-sections of spatio-temporal atlases constructed (top) with the pairwise approach followed by iterative refinement, and (bottom) our groupwise method. Cortical structures in the left temporal lobe develop, but vanish at week 40 and 44 again in case of the individual refinement of time points (top, red boxes). By contrast, the sequence generated by our groupwise method shows a consistent formation of a sulcus.

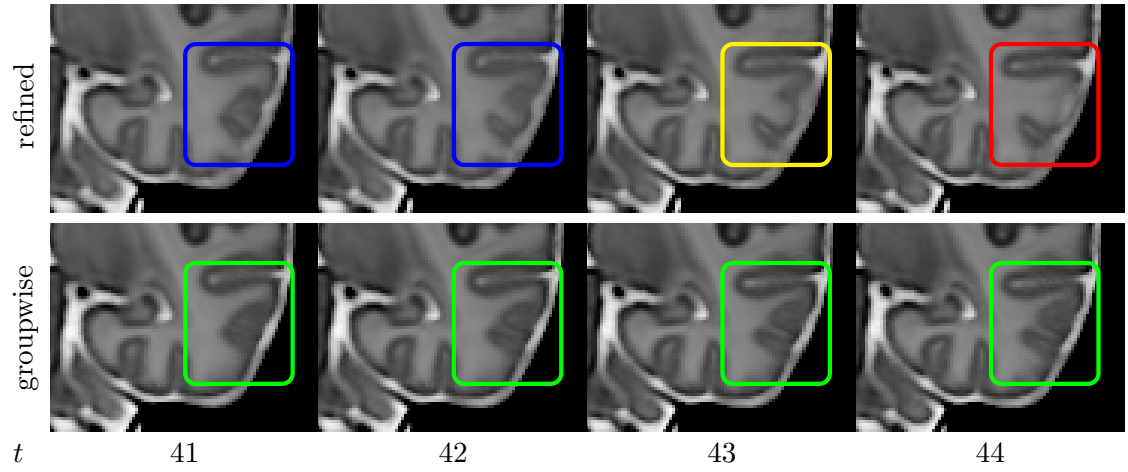


Figure 6.26.: Detail view of coronal cross-sections of spatio-temporal atlases constructed (top) with the pairwise approach followed by iterative refinement, and (bottom) our groupwise method. In the sequence obtained by individual iterative refinement of the atlas at each time point, a forming sulcus vanishes between week 43 (yellow box) and 44 (red box). Our groupwise approach, on the other hand, generates a temporal sequence which is consistent with our expectation of increasing cortical folding.

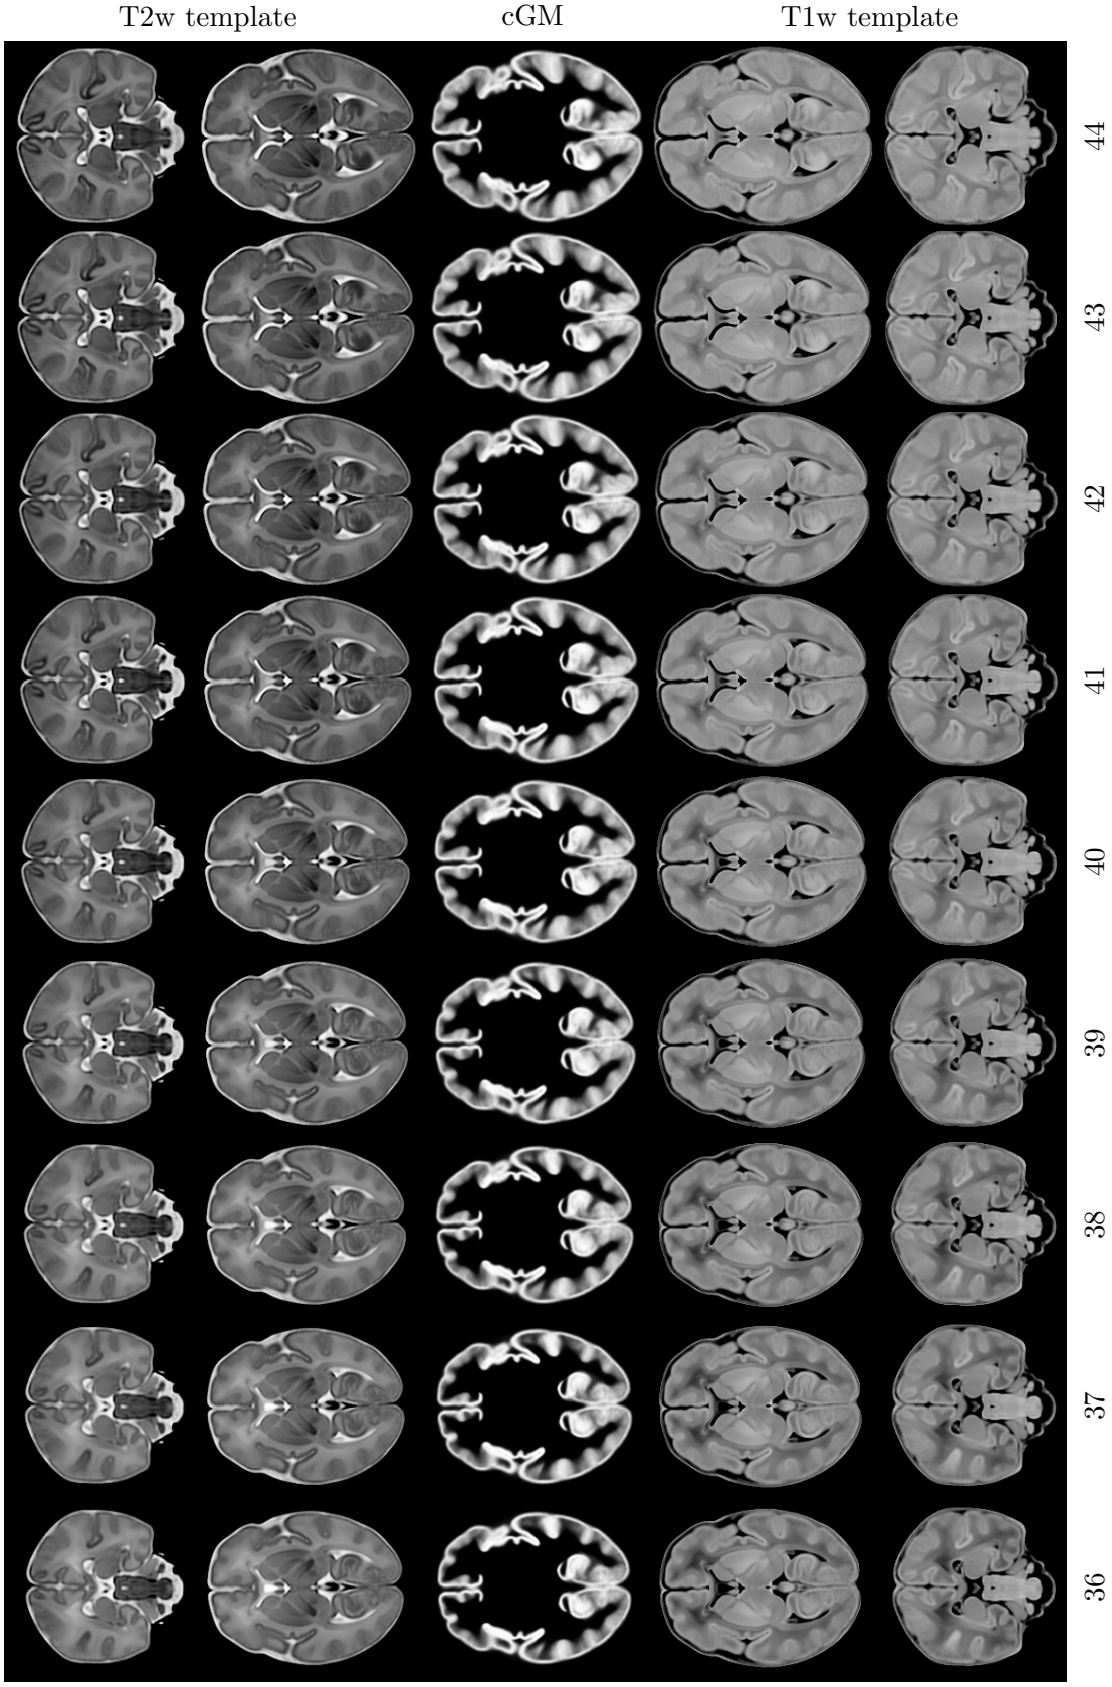


Figure 6.27.: Spatio-temporal atlas constructed with the proposed pairwise approach.

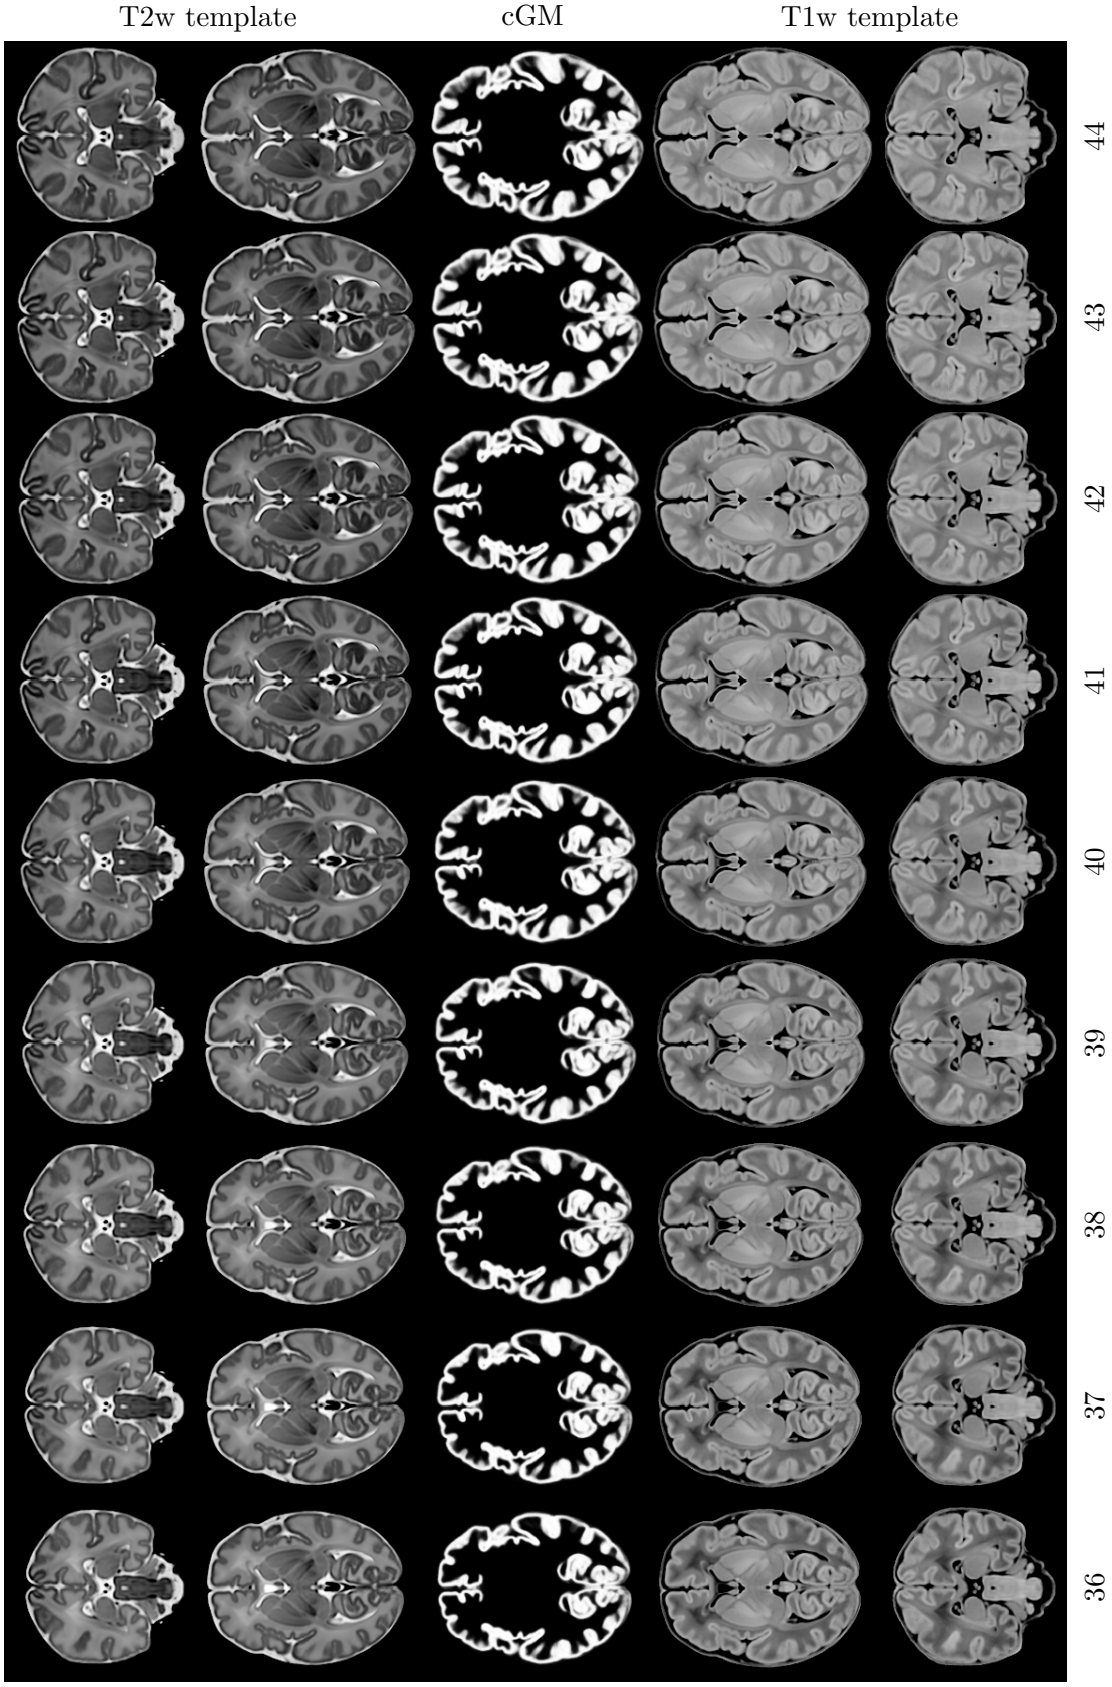
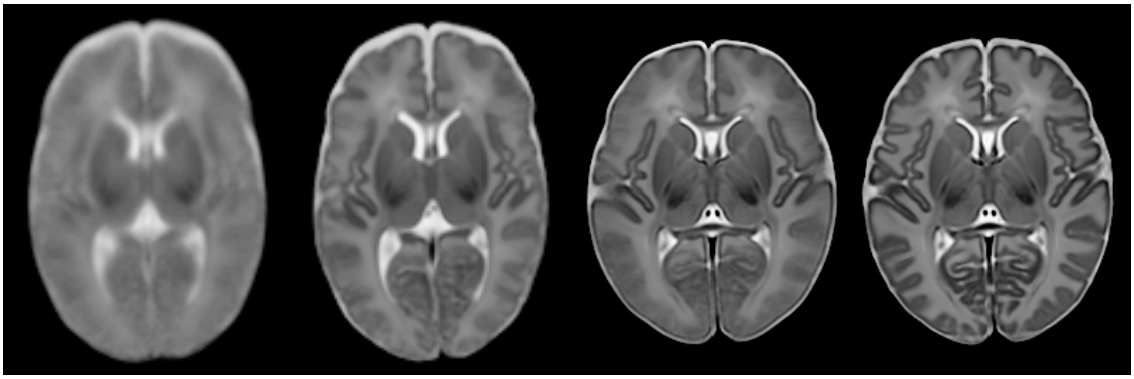


Figure 6.28.: Spatio-temporal atlas constructed with the proposed groupwise approach.



(a) (b) (c) (d)

Figure 6.29.: Qualitative comparison of mean T2w image at 40 weeks of (a) atlas created by Kuklisova-Murgasova et al. (2011) from 142 preterm born neonates using Log-Euclidean mean of affine transformations with constant kernel width of 2 weeks, (b) atlas created by Serag et al. (2012a) from 204 preterm born neonates using arithmetic mean of pairwise FFDs with adaptive kernel widths, (c) our pairwise atlas constructed from 275 term-born neonates using the Log-Euclidean mean of SVFFDs also with adaptive kernel widths, and (d) the atlas constructed from the same brain images using our groupwise approach.

T1w and T2w mean shape and intensity images at weekly intervals for the age range from 36 to 44 weeks PMA that were generated with the proposed pairwise technique are displayed in figure 6.27, along with the spatial tissue probability map for cGM. Unlike previously shown images, these are scaled relative to the mean brain size considering the mean scaling of the affine transformations computed during the initial global spatial normalisation. Cross-sections of the respective images of the atlas constructed with the proposed groupwise method are shown in figure 6.28. The clear markedness of cortical folds in the atlas constructed with the groupwise approach is reflected in the sharpness of the spatial cGM probability distribution.

6.4.8. Qualitative comparison to previous atlases

Previously, we compared the atlases constructed with our proposed methods based on the SVFFD algorithm to corresponding atlases constructed according to previous techniques proposed by Kuklisova-Murgasova et al. (2011) and Serag et al. (2012a) from the same set of neonatal images, and an improved registration using MIRTK. In figure 6.29 we give a comparison to the spatio-temporal neonatal brain atlases pre-dating our work, that were made publicly available by the original authors of these methods. Noticeable is the

significant improvement in image quality of the new atlases. The neonatal brain images acquired by dHCP are of remarkably higher quality than previous datasets of preterm born neonates that were available for the construction of these previous atlases. It should further be noted that at each time point of our new atlases 196 anatomies are being averaged[12], whereas the mean images created by Serag et al. (2012a) are the average of 15-19 individual brain scans (Serag et al., 2012c). This explains the lower cortical detail of our pairwise atlas despite the higher image quality. High image quality and marked cortical folds are presented by the atlas generated by our groupwise method.

6.5. Conclusion

We presented two approaches for the construction of a spatio-temporal atlas of the developing human brain. Our first approach improves upon a previous method developed by our research group. Unlike this initial technique using the original FFD algorithm implemented in IRTK, and the arithmetic mean of displacement fields as previously proposed by Seghers et al. (2004), we utilise our open source implementation of the SVFFD algorithm in MIRTK and the Log-Euclidean mean. This ensures topology preserving mappings and inverse consistency of the computed spatial mappings between any pair of brain images.

A disadvantage of such pairwise approach to atlas construction is, however, the quadratic computational complexity with respect to the number of brain images. A more significant limitation is the reliance on accurate pairwise correspondences. The less accurate the individual registrations, the less cortical detail is preserved. An approach which iteratively refines the transformations that relate each image with the average space are more robust towards residual anatomical misalignment. We have thus derived a new groupwise method which demonstrates better temporal consistency when compared to other extensions of iterative methods to the creation of a spatio-temporal atlas.

We conclude with a remark by Sled and Nossin-Manor (2013), that a difficulty in brain atlas formation is the normal variability among individuals, and the inherent uncertainty of the GA. In time, these are mitigated by larger datasets acquired shortly after birth.

[12]It should be noted, however, that the weight of many of these may be relatively low (see figure 6.15).

Chapter 7.

Deformable models of the neonatal cortex

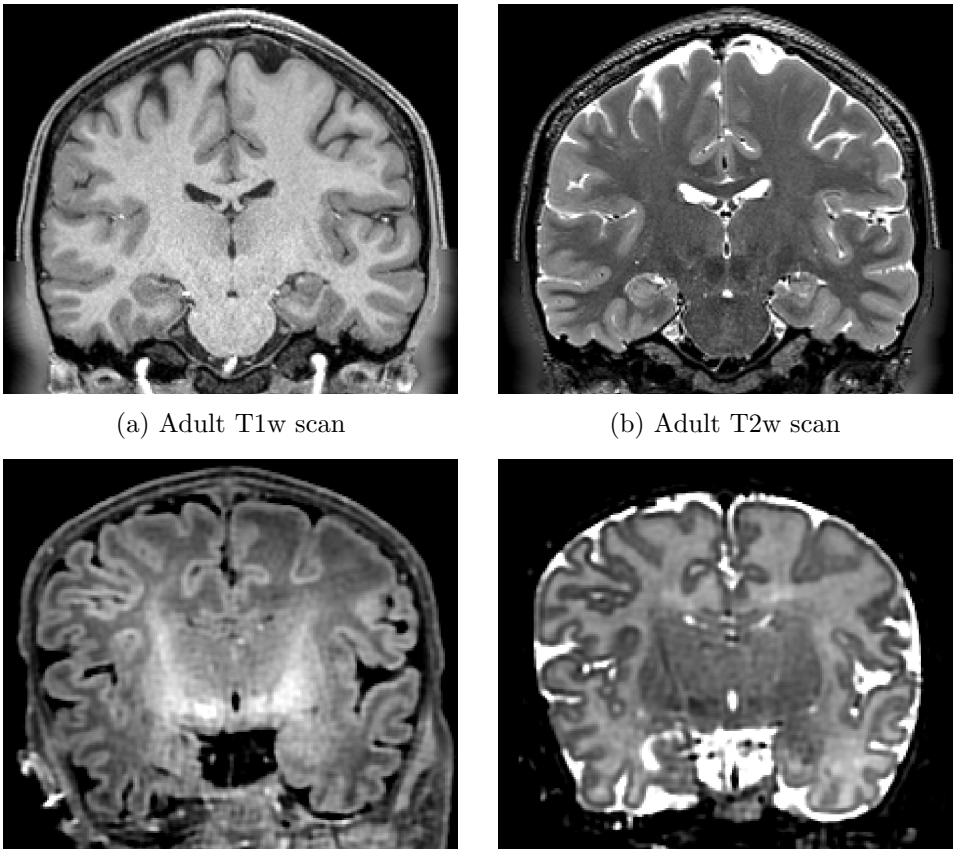
In chapter 6, a method for the unbiased construction of age-specific volumetric brain atlases is presented, where the shape of structures is given as an implicit function of either the template brain image intensities or the probabilities of a given structure to be present at a spatial atlas coordinate. Such representation is well suited for volumetric image analysis including, for example, spatial normalisation as required by VBM (Ashburner and Friston, 2000), and atlas-based segmentation (Makropoulos et al., 2014). One major drawback of most volumetric atlases for surface-based analysis is, however, the lack of detail prevalent near the cortex due to the vast biological variability of cortical folding, and the difficulty of accurately matching cortices when aligning brain images from which the atlas is constructed. This is a common limitation of volumetric image registration approaches. It is caused by the locally ambiguous appearance and shape of cortical folds, which cannot be distinguished well by a local similarity term.

In brain morphometry, such as the study of cortical folding (Shimony et al., 2016), an accurate delineation of the cortex and average surface templates that are representative for a population are desired. Explicit cortex models are a common prerequisite of methods for the analysis of functional and diffusion-weighted MRI brain scans (Glasser et al., 2013). Amongst other applications, these models are needed for surface matching (Yeo et al., 2010; Robinson et al., 2014; Lombaert et al., 2015; Wright et al., 2015), cortical parcellation to define functional units (Yeo et al., 2010; Parisot et al., 2016), and the mapping and visualisation of functional data (Fischl et al., 1999; Van Essen et al., 2001).

There exist a number of quite mature and freely available software tools for the reconstruction and analysis of the cortex in adults. The two most popular and first available tools for surface-based brain image analysis are FreeSurfer (Dale et al., 1999; Fischl, 2012) and Caret (Van Essen et al., 2001; Van Essen, 2012). Both find application in processing and analysis pipelines of the adult HCP (Glasser et al., 2013), and are integrated in Connectome Workbench[1]. Other widely used and comprehensive toolboxes are BrainSuite (Shattuck and Leahy, 2002) and BrainVISA (Rivière et al., 2009). The Morphologist[2] toolbox of BrainVISA, recently revised by Fischer et al. (2012), reconstructs 3D models of the boundary surfaces of the cortex from T1w images of the mature brain. Also the TOADS-CRUISE[3] plug-ins for JIST (Lucas et al., 2010) provide tools for adult brain segmentation and cortical modelling (Han et al., 2004; Huo et al., 2016).

In T1w images of the adult brain, WM is brighter than GM, and the intensity of CSF is below the mean intensity of GM as seen in figure 7.1a. Tools developed for adult brain image analysis often rely on this relation between tissues and their MR signal properties. Because of the higher water content, which alters the MR signal properties, along with other developmental processes such as myelination, the tissue contrast of the imaged neonatal brain is significantly different, and WM is mostly unmyelinated (uWM). While in adult brain imaging, a T1w MR sequence generates images well suited for structural analysis, T2w images exhibit higher tissue contrast in case of neonates. Exemplar sections of T1w and T2w scans at 40 weeks PMA are shown in figures 7.1c and 7.1d. The intensity of CSF is higher in neonatal T2w images than the mean uWM intensity, violating a common assumption of software tools developed for adult T1w images. As a result, these tools can not be applied to brain MR images of neonates and infants. For this reason, approaches to reconstruct the neonatal cortex usually rely on a previous segmentation of the neonatal brain tissues using a tissue classification specifically designed for neonatal brain images (Xue et al., 2007; Dubois et al., 2008; Hill et al., 2010). The accuracy of the resulting cortical surface models is hereby limited by the quality of the tissue classification.

[1] https://www.humanconnectome.org/software/connectome-workbench.html (last accessed: May 20, 2017)
[2] http://brainvisa.info/web/morphologist.html (last accessed: May 20, 2017)
[3] https://www.nitrc.org/projects/toads-cruise/ (last accessed: May 20, 2017)



(a) Adult T1w scan (b) Adult T2w scan

(c) Neonatal T1w scan (d) Neonatal T2w scan

Figure 7.1.: Illustration of different tissue contrast in adult and neonatal brain images. Blurred boundaries in (a, b) are caused by "defacing" used to anonymise patient data.

Segmentations of neonatal brain images often contain substantial errors due to the considerably poorer image contrast compared to adult brain images because of shorter acquisition times and infant motion. Partial volume, the increased water content of cerebral tissues, and unmyelinated WM all complicate a voxel-based classification (Xue et al., 2007). This results in CSF mislabelled as uWM or cGM as shown later in figure 7.9.

In this chapter, we present a method that we developed specifically for the reconstruction of the cortical surfaces of the developing human brain at the neonatal period based on non-self-intersecting deformable surfaces. Similar to aforementioned approaches, we also employ a previous brain segmentation. Unlike these approaches, however, errors in the tissue segmentation are subsequently accounted for by a local image-based refinement given the original bias corrected T2w image (and T1w image if available). The pial surface is then found by deforming the white matter surface outwards, also using an

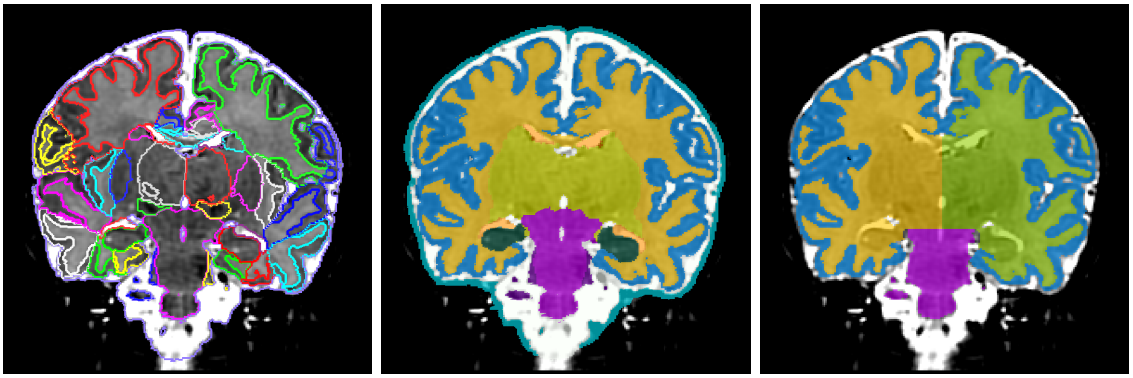
image-based force that is independent of the initial tissue classification. This avoids the usual overestimation of the thickness of the cGM segment due to the difficulty of separating gyri that are separated only by narrow sulci. We demonstrate that the final white matter and pial surfaces reconstructed with our proposed method define the boundaries of the neonatal cortex more accurately than the employed neonatal brain segmentation. The resulting surface meshes are further guaranteed to not intersect and subdivide the brain volume into disjoint regions. The proposed method generates topologically correct surfaces which facilitate both a flattening and spherical mapping of the cortex.

7.1. Segmentation of the brain volume

The neonatal cortical surface reconstruction builds upon a previous brain extraction (Smith, 2002), bias field correction (Tustison et al., 2010), and a segmentation of the neonatal brain into WM, GM, and CSF, along with a segmentation of dGM structures, the lateral ventricles, brainstem (BS) and cerebellum (CB). For this we employ the developing region annotation with expectation maximisation algorithm of Makropoulos et al. (2014) as implemented in the publicly available Draw-EM package[4] of the MIRTK. Draw-EM is a state-of-the-art atlas-based segmentation of neonatal brain scans with EM refinement, which finds application in the dHCP (Makropoulos et al., 2017b). A coronal view of a typical Draw-EM tissue segmentation into WM, cGM, dGM, CSF, lateral ventricles, hippocampi and amygdala, BS, and CB is shown in figure 7.2b. The corresponding full segmentation of the brain into 87 structures is shown in figure 7.2a.

Given the Draw-EM labels of right and left structures, we determine two orthogonal cutting planes. The first cutting plane separates the structures of the right and left cerebral hemispheres. It is found by PCA of the covariance matrix of the coordinates of voxels which have a face neighbour belonging to the other hemisphere. The plane is defined by a point on the plane corresponding to the centre of all inter-hemisphere voxels, and the eigenvector of the covariance matrix corresponding to the eigenvalue with smallest absolute value as normal vector, i.e., the direction of least variation in inter-hemisphere

[4] https://github.com/MIRTK/DrawEM/ (last accessed: May 20, 2017)



(a) Draw-EM structures (b) Draw-EM tissues (c) Disjoint regions

Figure 7.2.: Coronal sections of exemplar Draw-EM segmentation of the neonatal brain into (a) 87 structures, and (b) WM, cGM, dGM, CSF, lateral ventricles, hippocampi and amygdala, BS, and CB; (c) regions of right/left cerebrum, cortex, and BS+CB derived from this, while introducing a planar medial cut and separating the stem from the cerebrum.

voxel position. The normal vector of the second plane is defined to be orthogonal to the normal of the medial cutting plane. With the plane normal fixed, we then find a plane distance from the origin such that the plane cuts through the superior part of the BS, separating BS and CB from the cerebrum. For this, we search above and below the centre of voxels classified as either BS or dGM with a face neighbour belonging to the respective opposite class. Given a test plane, we extract the curve corresponding to the intersection of this plane with the BS segmentation. When this curve is not closed, the test plane is discarded. Of all possible closed cuts within the search distance, the cut closest to the BS/dGM which only intersects voxels classified as BS is selected. This ensures that no cerebral tissue is separated from the volumes of the two hemispheres.

Morphological closing of the dGM structures is then performed to include voxels labelled as CSF and intra-cranial background within the sub-cortical neighbourhood where the cortical surface cannot be defined as the interface between WM and cGM. This step is optional and was done here to facilitate future work where brain volumes should be partitioned into right/left cerebrum, cortex, and BS plus CB without requiring an accurate segmentation of sub-cortical structures, while ensuring at the same time that all brain structures are included within the reconstructed surfaces. The core idea of this future work is to utilise volumetric mapping of each disjoint region combined with a surface matching technique in order to improve volumetric brain image registration.

7.2. Reconstruction of the cortical surfaces

7.2.1. Deformable surface model

The individual models developed for the reconstruction of the neonatal cortex detailed in the following sections all employ a deformable surface that is governed by a simplified version of the Lagrange equations of motion in discrete form, which are integrated using a numerical scheme. More specifically, the final node positions adapted to the image data are obtained by integrating the differential equation using a simple forward Euler scheme,

$$\boldsymbol{x}_i(t + dt) = \boldsymbol{x}_i(t) + \boldsymbol{f}_i(t)dt \tag{7.1}$$

where $\boldsymbol{x}_i(t)$ is the position of surface node i at time point t, dt is the Euler integration step length, and $\boldsymbol{f}_i(t)$ is the weighted sum of the external and internal forces of the respective deformable surface model acting on node i as defined later. The Euler integration proceeds until either a maximum number of steps is exceeded or the internal and external forces for all nodes are in equilibrium, i.e., $\forall i,\ \|\boldsymbol{f}_i(t)\| < \epsilon$, where ϵ is a small number. For an introduction to deformable models based on dynamic force equations, see section 4.5.1.

Normalisation of force magnitude

The magnitude of external forces is in the following force terms normalised from a range $[a, b]$ to $[0, 1]$ before multiplying by the positive weight of the respective force. The auxiliary function used for this normalisation is the S-shaped membership function

$$\mathrm{smf}(x, a, b) = \begin{cases} 0, & \text{if } x \leq a \\ 2\left(\frac{x-a}{b-a}\right)^2, & \text{if } a < x \leq \frac{a+b}{2} \\ 1 - 2\left(\frac{x-b}{b-a}\right)^2, & \text{if } \frac{a+b}{2} < x < b \\ 1, & \text{if } x \geq b \end{cases} \tag{7.2}$$

plotted in figure 7.3, where x is a value to be normalised, e.g., the Euclidean distance to the target position. This normalisation is done for two reasons. Firstly, to limit the influence of x values outside $[a, b]$ to achieve a more uniform propagation of the surface model, and

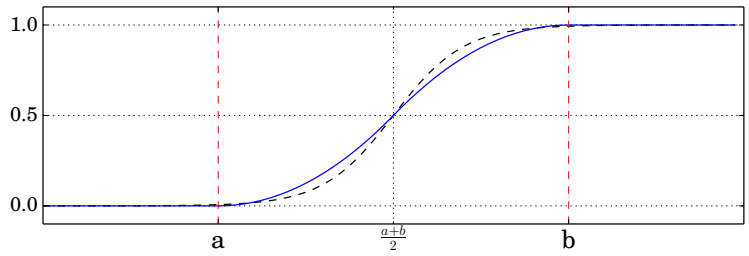
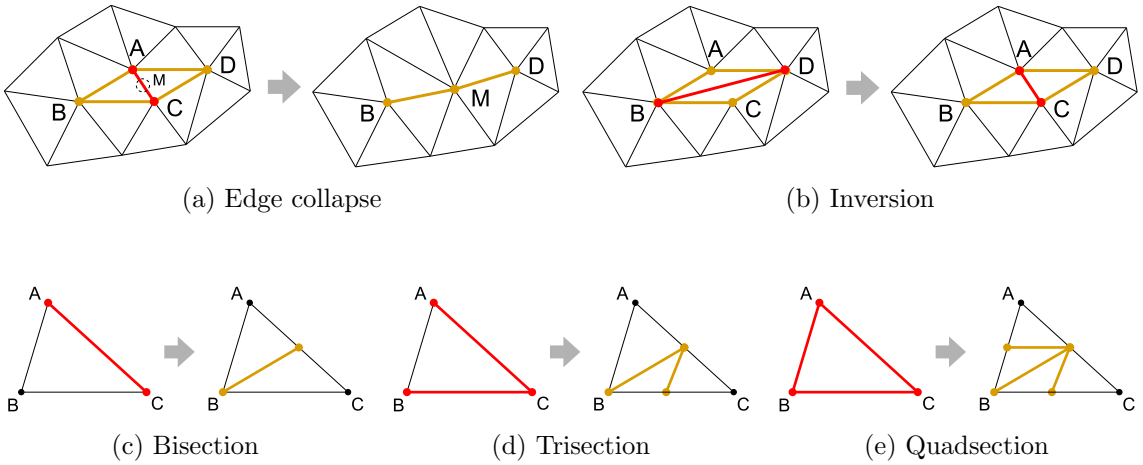


Figure 7.3.: S-shaped membership function (blue curve) used to normalise forces, whereas the dashed curve depicts the tanh function $y = (1 + \tanh(x))/2$ for $a = -2.5$ and $b = 2.5$.

secondly to allow better control in the relative weighting of each individual surface force with respect to the other internal and external forces acting on the deformable surface. The lower bound a is therefore generally set to zero and the upper bound b for example to the 95th percentile of all x values which gives a robust estimate of the maximum x value, and thereby reducing the influence of outlier values. From a different point of view, Cohen and Cohen (1993) describe such force magnitude normalisation as locally adaptive integration time step such that the surface evolves with the same speed everywhere. A normalisation using the tanh function instead of equation (7.2) has also been utilised by Dale and Sereno (1993) in early work on FreeSurfer. This function is more costly to compute, however, and has similar shape as seen by the dashed curve in figure 7.3.

Local adaptive remeshing

Due to the piecewise linear approximation of the deformable surfaces, where forces only act on the nodes of the discrete surface mesh, it is important to regularly adapt the sampling of the deformable surface to allow it to locally contract and expand. We therefore perform after each step topology preserving local adaptive remeshing operations consisting of edge collapse, triangle inversion, and subdivisions as detailed by Park et al. (2001) and schematically illustrated in figure 7.4. This keeps the average distance between adjacent nodes within a predefined range and thus achieves a consistent sampling. In case of the deformable pial surface model of section 7.2.4, no triangle inversion is done. The reason for this is that it can replace an edge running along the length of a sulcus by an edge running across the sulcus, which in turn can lead to unwanted filling in of the sulcus.



(a) Edge collapse (b) Inversion

(c) Bisection (d) Trisection (e) Quadsection

Figure 7.4.: Illustration of local adaptive remeshing operations.

Non-self-intersection constraint

Given an initial surface mesh at $t = 0$ with node positions $\boldsymbol{x}_i$ such that no self-intersections occur, we ensure that this non-self-intersection condition also holds after each integration step. For that, the force $\boldsymbol{f}_i$ is at each step, and for each node individually, either halved a pre-set maximum number of times or set to zero if the resulting surface mesh would otherwise self-intersect after the deformation. This is repeated until no more self-intersections are introduced by the surface evolution at the current time step.

7.2.2. Initial genus-0 surface model

In the first stage of the surface reconstruction, we seek to obtain a good initial starting point for the extraction of a topologically correct white matter surface. For this we utilise the available brain segmentation after mapping the voxel labels to disjoint segments of right and left cerebrum as described in section 7.1, while excluding cortical labels. This segmentation of the cerebrum may contain local misclassifications and consequently has generally not the correct topology. A direct extraction of the boundary using, for example, the marching cubes algorithm (Lorensen and Cline, 1987) is therefore prohibited. Instead, we compute for each brain hemisphere the convex hull of the boundary of its cerebrum half. This convex hull is a closed genus-0 surface with no "holes", i.e., has spherical topology.

The triangular mesh corresponding to the convex hull of the segmentation boundary is

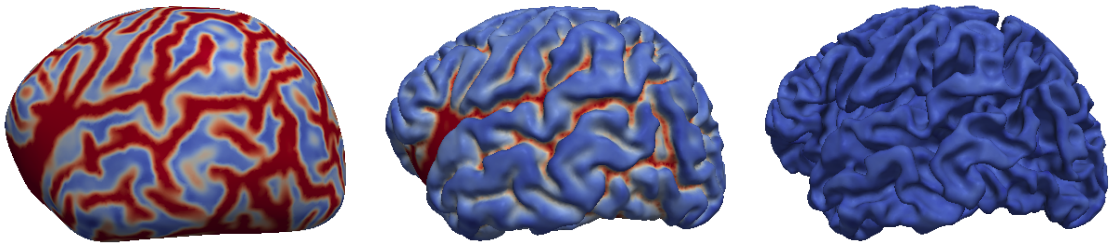


Figure 7.5.: Left convex hull deformed towards WM segmentation. Colours depict segmentation boundary distance in normal direction, where red is $\geq 5\,\text{mm}$ and blue $\leq 0\,\text{mm}$: (left) initial convex hull, (middle) after 10 iterations, and (right) after 120 iterations.

then deformed inwards, as illustrated in figure 7.5, towards the zero level set of the implicit surface given by the (signed) Euclidean distance transform (Maurer et al., 2003) by

$$\boldsymbol{f}_i^{init}(t) = w_d \boldsymbol{f}_i^s(t) + w_b \boldsymbol{f}_i^b(t) + w_g \boldsymbol{f}_i^g(t) + w_r \boldsymbol{f}_i^r(t) \tag{7.3}$$

where $w_{\{d,b,g,r\}}$ are positive weights. The external force $\boldsymbol{f}^s$ attracts the model towards the segmentation boundary, while the remaining forces are internal forces: a bending force $\boldsymbol{f}^b$ to minimise curvature, a Gauss curvature based spring force $\boldsymbol{f}^g$ to remove creases, and a node repulsion force $\boldsymbol{f}^r$ to reduce self-collisions. These terms are detailed in the following.

Joining the hemispheres

At this initial stage, the two hemispheres are processed separately in order to allow the two deformable surfaces to freely advance into the space of the medial longitudinal fissure. Intersections between the hemispheres that might occur at the introduced medial cut, where no cortex is present, are removed again by joining the surfaces of the two brain hemispheres at this medial cut to form a single closed genus-0 surface as illustrated in figure 7.6. This surface mesh is moreover merged with the surface of the union of BS and CB segments at the previously introduced BS cut. In doing so, we ensure that individual surfaces do not intersect each other during and after the consecutive deformations of this combined, non-self-intersecting surface. Triangles of the so obtained initial white matter surface mesh are assigned labels identifying them as part of right cerebrum, left cerebrum, or BS and CB. During all merging steps, we make sure to preserve the spherical topology.

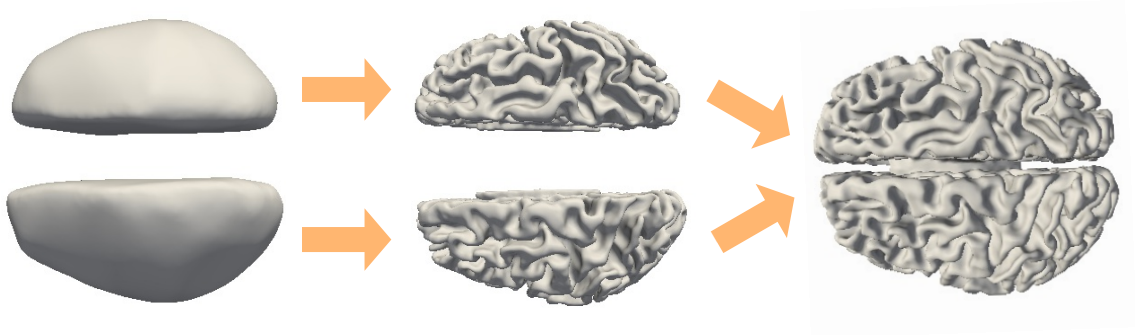


Figure 7.6.: Illustration of initial white matter surface assembly.

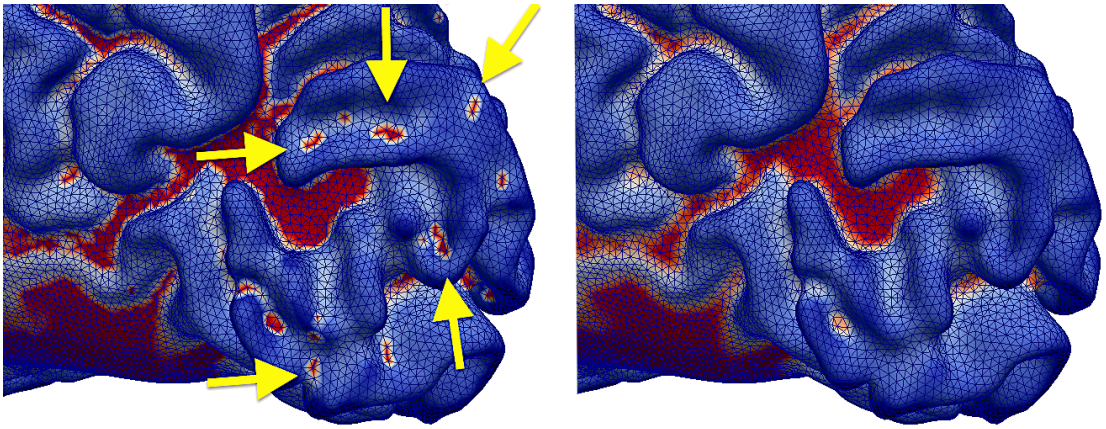
Implicit surface distance force

The initial convex hull of each brain hemisphere is deformed towards the boundary of the segmentation by the external force given by

$$\begin{aligned} \boldsymbol{f}_i^s &= m_i^s \times \boldsymbol{n}_i \\ m_i^s &= \operatorname{sgn} d_i^s \times \operatorname{smf}(|d_i^s|, 0, d_{max}^s) \end{aligned} \tag{7.4}$$

where $\boldsymbol{n}_i$ is the surface normal at node i, and m_i^s the signed magnitude given the distance, d_i^s, of the surface node to the segmentation boundary in inwards (negative) or outwards (positive) surface normal direction. The S-shaped membership function serves in this force term to limit its magnitude to be less or equal d_{max}^s. We set this upper distance threshold equal the 95th percentile, $d_{95\%}^s$, of the absolute values of all node distances, $|d_i^s|$, whereby a lower threshold d_ϵ^s is used to avoid amplification of the external force when the surface model is everywhere close to the segmentation boundary, i.e., $d_{max}^s = \max(d_\epsilon^s, d_{95\%}^s)$.

It has been shown for active contours by Xu and Prince (1998) that a force based on the minimum distance may fail to drive the contour into concavities such as the sulci of the human cortex. A minimum distance potential was originally proposed by Cohen and Cohen (1993) for the evolution of an active contour for the extraction of object boundaries. We therefore compute the distance of a surface node to the implicit surface of the segmentation in surface normal direction instead, by finding the closest zero crossing of the corresponding Euclidean distance transform along a ray cast in both inwards and



(a) Boundary distance (b) Filtered distance

Figure 7.7.: Distance of WM segmentation boundary in normal direction (a) before clustering based hole filling, and (b) after the hole filling and smoothing, where nodes coloured red have distance $\geq 5\,\text{mm}$, and blue nodes a distance $\leq 0\,\text{mm}$. A number of small holes in the segmentation that were accounted for are indicated in (a) by yellow arrows.

outwards surface normal directions. The absolute minimum surface distance at ray points is thereby used to adapt the step length for finding the zero crossing. To reduce the influence of small irregularities in the segmentation boundary, a Laplacian smoothing of surface distances is performed with weights

$$w_{ij} = \begin{cases} \max\left(0, \langle \boldsymbol{n}_i, \boldsymbol{n}_j \rangle\right), & \text{if } A_{ij} = 1 \\ 0, & \text{otherwise} \end{cases} \tag{7.5}$$

where A_{ij} is the node adjacency matrix of the surface mesh.

Implicit surface hole filling

A novelty of our distance potential force is a filtering of the computed surface distances, d^s, to account for small holes in the segmentation boundary. Such small holes manifest themselves in the surface distance map as small clusters of supposedly distant points as seen in figure 7.7a. The distance values of these clusters are consequently filled in to avoid the surface to deform into them. Specifically, surface clusters are formed by growing regions starting at a seed node k that has not been assigned to a cluster yet s.t. $\forall i, |d^s_k| \geq |d^s_i|$. The region growing stops when either the surface distance drops below a predefined threshold

or the angle made up by the surface normal vectors of the seed node and the propagating front node exceeds a specified threshold. This process is repeated until no more seed nodes have an absolute surface distance $|d_i^s| > d_{thres}^s$. Clusters which are elongated, e.g., corresponding to sulcal banks, or have an excessively large radius are then removed. The segmentation boundary distance of nodes belonging to the remaining clusters is set to zero, followed by a few Laplacian smoothing iterations, whereby only modified distance values are being updated. The effect of this filtering is shown in figure 7.7b.

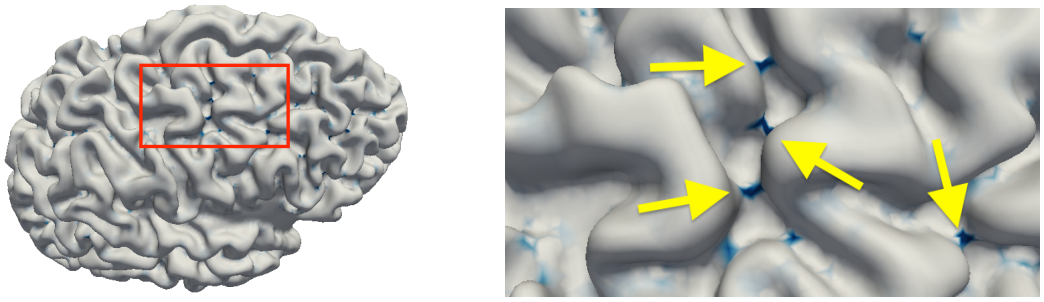
Bending force

Smoothness of the surface is encouraged by the bending force, $\boldsymbol{f}^b$, which minimises the surface curvature by attracting a node towards the barycentre of its adjacent nodes, $\boldsymbol{x}_{ij}$, for which $A_{ij} = 1$, i.e., $\bar{\boldsymbol{x}}_i = \frac{1}{N_i}\sum_{j=1}^{N_i} \boldsymbol{x}_{ij}$, where $N_i = \sum_{j=0}^{N} A_{ij}$ is the number of nodes adjacent to node i. This force was also used by Lachaud and Montanvert (1999) and Park et al. (2001). It is given by

$$\boldsymbol{f}_i^b(t) = \bar{\boldsymbol{x}}_i - \boldsymbol{x}_i - \frac{1}{N_i}\sum_{j=1}^{N_i} (\bar{\boldsymbol{x}}_{ij} - \boldsymbol{x}_{ij}) \tag{7.6}$$

Gauss curvature force

In addition to the bending force, we further define a new spring force, $\boldsymbol{f}^g$, that avoids the creation of unwanted creases perpendicular to the length of a sulcus while the surface deforms into a narrow sulcus of the cortex. This is shown in figure 7.8 for the deforming surface of the right hemisphere with negative Gauss curvature mapped to blue colouring. Similar to the bending force, this spring force attracts nodes to the barycentre of either all or a subset of adjacent nodes, whereby the magnitude of this new spring force is scaled to be proportional to the Gauss curvature of the deformable surface. Nodes with negative Gauss curvature as found in the "centre" of a hyperbolic paraboloid, are attracted only by adjacent nodes either below or above the tangent plane. This smooths out small dents along a gyrus and in particular thin creases bridging the gap between two gyri that are not usually observed in a human cortex.



(a) Deforming right hemisphere surface. (b) Zoomed view of creases.

Figure 7.8.: Small creases with negative Gauss curvature forming during the deformation.

The Gauss curvature based spring force is defined as

$$\boldsymbol{f}_i^g = m_i^g \times \frac{\bar{\boldsymbol{f}}_i^g}{\|\bar{\boldsymbol{f}}_i^g\|} \tag{7.7}$$

with

$$\bar{\boldsymbol{f}}_i^g = \frac{\sum_{j=1}^{N_i} \delta_{ij}^g \times (\boldsymbol{x}_{ij} - \boldsymbol{x}_i)}{\sum_{j=1}^{N_i} \delta_{ij}^g} \tag{7.8}$$

where δ_{ij}^g is an indicator function which, for negative Gauss curvature K_i, is 1 only for points on one side of the tangent plane at node i. Which side of the tangent plane depends on the sign of the external force $\boldsymbol{f}_i^s$ that is positive for nodes inside the implicit target surface and negative outside of it. Hence, points with a negative Gauss curvature outside the segmentation are forced inwards, while points with a negative Gauss curvature inside the segmentation are forced outwards. The indicator function is thus defined as

$$\delta_{ij}^g = \begin{cases} \max\left(0, \operatorname{sgn} m_i^s \times \operatorname{sgn}\langle \boldsymbol{x}_{ij} - \boldsymbol{x}_i, \boldsymbol{n}_i \rangle\right), & \text{if } K_i < 0 \\ 1, & \text{if } K_i \geq 0 \end{cases} \tag{7.9}$$

The magnitude of this spring force is $m_i^g = \operatorname{smf}(|K_i|, K_{min}, K_{max})$, where K_{min} and K_{max} are lower and upper Gauss curvature thresholds. The force is zero for $|K_i| < K_{min}$.

Node repulsion force

The internal force, $\boldsymbol{f}^r$, reduces the occurrence of self-intersections that have to be resolved after each integration step by having nodes repel each other. Its magnitude is proportional

to the pairwise distance of the surface mesh nodes as well as the dot product of their surface normal vectors. This results in nodes belonging to close opposing triangles to repel each other stronger than nodes of triangles belonging to the same surface patch.

The node repulsion force is defined as

$$\boldsymbol{f}_i^r = \sum_{j=1}^{N} m_{ij}^r \times \frac{\boldsymbol{x}_i - \boldsymbol{x}_j}{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|} \tag{7.10}$$

with pairwise node repulsion magnitude and direction given by

$$m_{ij}^r = \begin{cases} s_{ij}^r \left(\frac{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|}{r} - 1 \right)^2, & \text{if } 0 < \|\boldsymbol{x}_i - \boldsymbol{x}_j\| < r \\ 0, & \text{otherwise} \end{cases} \tag{7.11}$$

where the factor s_{ij}^r is proportional to the angle made up by the surface normal vectors of the two repelling nodes. This factor is zero when the angle is less or equal 90°, i.e.,

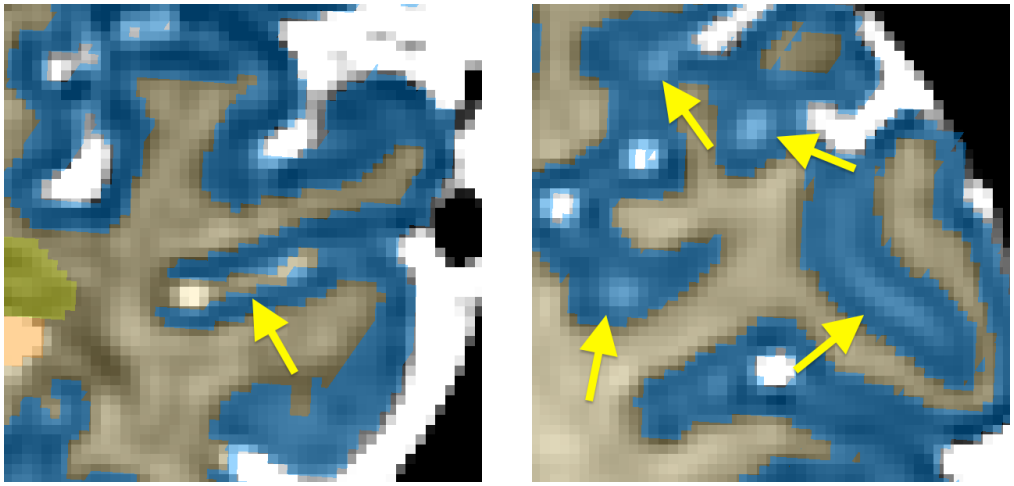
$$s_{ij}^r = \mathrm{smf}(-\langle \boldsymbol{n}_i, \boldsymbol{n}_j \rangle, 0, 1) \tag{7.12}$$

The terms of the sum in equation (7.10) are non-zero only for neighbouring nodes j within radius r of surface node i. We use an Octree (Meagher, 1982) to speed up the computation.

7.2.3. White matter surface model

Draw-EM occasionally assigns CSF voxels within sulci incorrect WM or GM labels due to partial volume or averaging which takes place during the motion correction. The wrongly labelled voxels, such as those highlighted by an arrow in figure 7.9a, cause the initial white matter surface to not deform completely into an affected sulcus. Thus, we use a deformable model based only on the intensities of the bias corrected images to extract the final white matter surface. The initial configuration of this deformable model is the topologically correct surface obtained from the WM segmentation.

The internal forces of this deformable white matter surface model are identical to the initial deformable model, but the segmentation boundary distance force $\boldsymbol{f}_i^s$ is replaced by a novel force $\boldsymbol{f}_i^w$ which pushes the surface node i towards its closest WM/cGM edge in



(a) CSF mislabelled as WM (b) CSF mislabelled as cGM

Figure 7.9.: Common Draw-EM tissue segmentation errors.

surface normal direction, i.e.,

$$\boldsymbol{f}_i^{wm}(t) = w_d \boldsymbol{f}_i^{w}(t) + w_b \boldsymbol{f}_i^{b}(t) + w_g \boldsymbol{f}_i^{g}(t) + w_r \boldsymbol{f}_i^{r}(t) \tag{7.13}$$

where the segmentation boundary distance is replaced by the edge distance, d_i^w, i.e.,

$$\begin{aligned} \boldsymbol{f}_i^w &= m_i^w \times \boldsymbol{n}_i \\ m_i^w &= \operatorname{sgn} d_i^w \times \operatorname{smf}(|d_i^w|, 0, d_{max}^w) \\ d_{max}^w &= \max(d_\epsilon^w, d_{95\%}^w) \ . \end{aligned} \tag{7.14}$$

WM/cGM interface distance

The image edge distance, d_i^w, is found by analysing the 1D intensity profile and directional derivative sampled at equally spaced normal ray points. Both the intensity and image derivative are sampled from the discrete image using a cubic B-spline interpolation kernel and its derivative, respectively. Examples of these intensity profiles are shown in the left column of figure 7.10.

We first identify minima, maxima, and inflection points of the intensity profile. It can be observed that a white matter surface edge occurs between a maximum with WM intensity, followed by a minimum with cGM intensity (left/right of green vertical line). Starting at

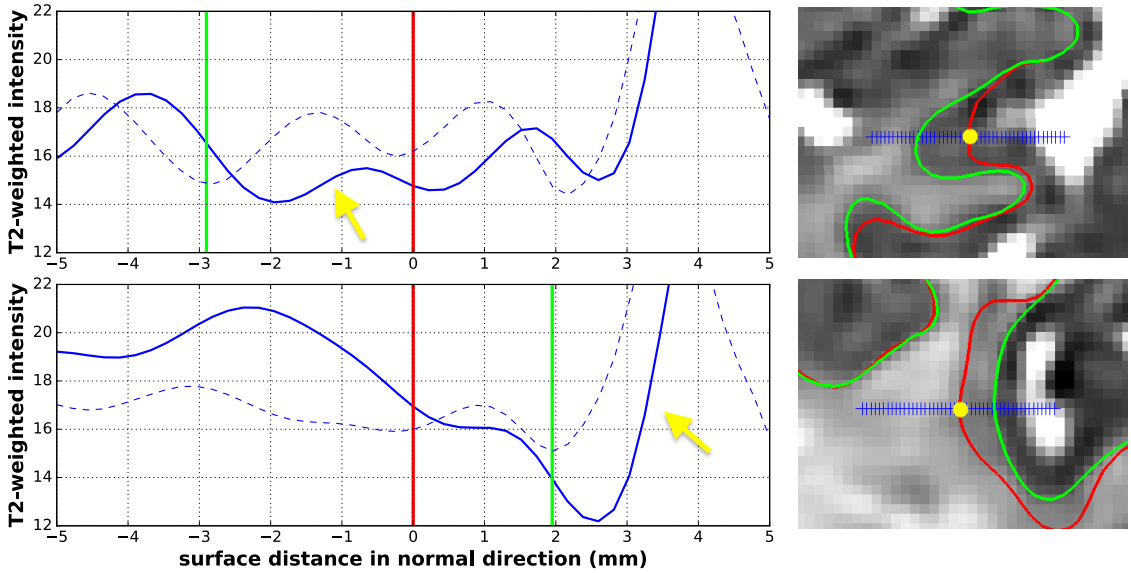
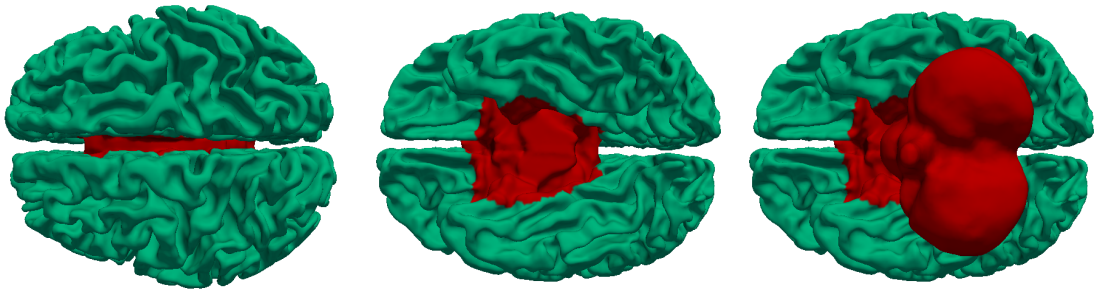


Figure 7.10.: Intensity (left, blue curve) and derivative (left, dashed blue curve) sampled from the image (right) at points marked by blue crosses. A yellow dot (right) marks the node on the initial surface (red). The green contour depicts the final white matter surface. The yellow arrows on the left mark the corresponding cGM/CSF edge of the pial surface.

the ray centre (red vertical line), we search for such image edge both inwards and outwards from the surface node until either a WM/cGM edge candidate is found, a maximum search depth is exceeded, or the ray intersects the surface. The outwards search is stopped when an opposing cGM/WM or cGM/CSF edge is encountered. If a suitable edge was found either only inwards or outwards of the deformable surface, it is selected. Otherwise, the combined probability of both maximum and minimum intensities to belong to the WM and cGM class, respectively, are used to choose one of the two candidate edges. To encourage the surface to deform into deep sulci and skip CSF voxels encountered inside the current surface, the inwards edge is preferred whenever a cGM/CSF edge is found inside. These edges are identified by a minimum in the intensity profile followed by a maximum with an intensity above the WM mean intensity plus five times the standard deviation of WM intensities. The image edge distance d_i^w is then equal the signed distance to the strongest minimum of the directional intensity derivative within the distance range of the encountered maximum and minimum edge intensity. When no edge is found, $d_i^w = 0$. A median filtering followed by Laplacian smoothing of edge distances mitigates errors caused by incorrectly identified edges due to noise and other image artefacts.



(a) Dorsal view (b) Ventral view, no BS+CB (c) Ventral view

Figure 7.11.: Initial white matter surface model with non-cortical triangles coloured red. The classification of surface elements as either cortical or non-cortical is based on the distance to the cGM segmentation. During the intensity-based white matter surface reconstruction, only cortical surface nodes (green) are deformed.

In order to avoid wrong edges near the lateral ventricles, where CSF is adjacent to cGM, and to further not be mislead by WM/dGM edges of non-cortical structures, we use the Draw-EM segmentation of dGM to adjust the intensity thresholds accordingly deep inside the cerebrum. Moreover, we detect transitions from uWM to darker mWM, which can in the T2w images at term predominantly be observed in the posterior periventricular WM (Rutherford, 2002), or from dGM to cGM, respectively, and choose the further outside edge, as seen in the second row of figure 7.10.

Exclusion of non-cortical nodes

Because WM/cGM edges can only be found for nodes near a WM/cGM interface but not for nodes near non-cortical structures such as the corpus callosum, the hippocampi, and amygdala, we keep the initial positions of these nodes by setting the node force $\boldsymbol{f}_i^{wm} = 0$. For that, we label nodes of the initial surface as either cortical or non-cortical by searching within a small radius for a voxel labelled as cGM. When none is found, the node is assigned a cortex mask value of zero, and one otherwise. We then erode this mask defined on the surface to ensure that only nodes near a segmentation boundary where both WM and cGM labels are present are deformed. An example of such cortex mask is depicted in figure 7.11, where nodes of triangles coloured red have a zero mask value and therefore no force is applied to these nodes during the extraction of the WM/cGM interface.

Multi-scale integration scheme

In addition to the filtering of image edge distances, which locally smooths the magnitude of the image-edge distance force, the external force of adjacent nodes is further averaged to allow entire regions to move coherently in the appropriate direction. The number of averaging steps, n_{avg}, is halved after convergence of the surface deformation, followed by a new Euler integration. The surface nodes are thus allowed to move more freely each time the deformation process is restarted. An averaging of node forces with subsequent fewer iterations has previously been proposed for the flattening and inflation of surfaces of the adult cortex in FreeSurfer (Fischl et al., 1999). At each scale, convergence is reached when the percentage of active nodes falls below a pre-defined threshold. The status of a node changes from active to passive when it has not moved more than a negligible small distance for a certain number of past update steps, e.g., due to a collision with an opposing surface element. The complete multi-scale integration scheme for the extraction of the white matter surface is summarised in figure 7.12, including the enforcement of the hard non-self-intersection constraint discussed earlier.

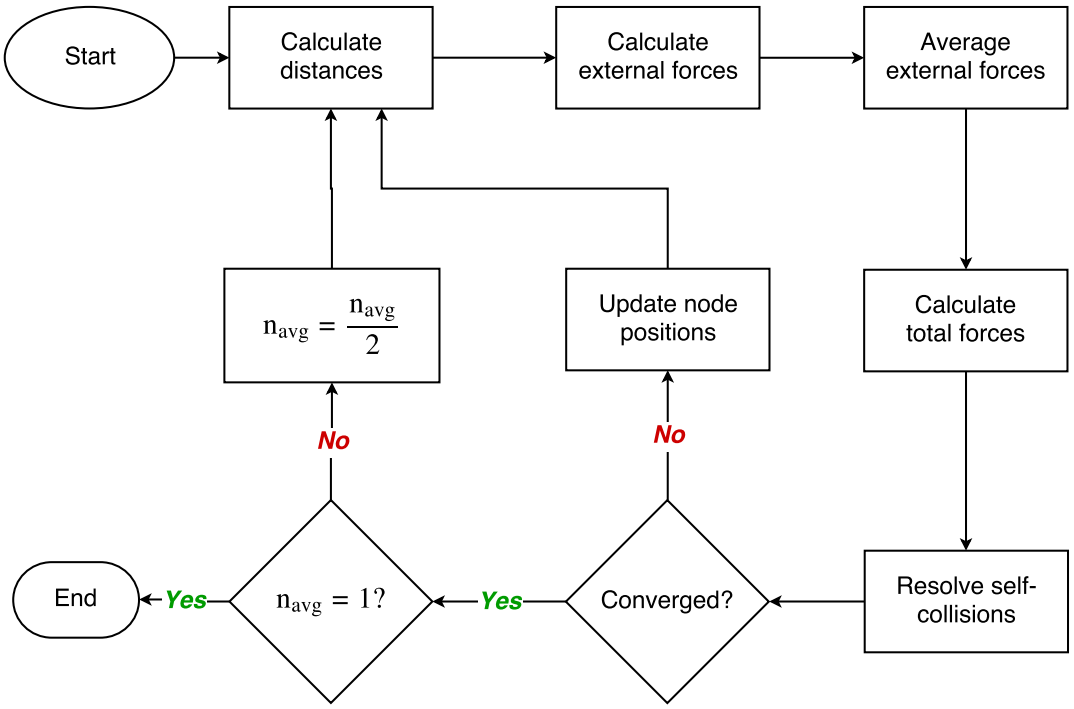


Figure 7.12.: Flow chart of white matter surface model integration scheme.

7.2.4. Pial surface model

We have so far reconstructed a surface model of the WM/cGM interface with spherical topology. Due to partial volume effects, a neonatal brain tissue segmentation often overestimates cortical thickness and mislabels CSF as cGM (see figure 7.9b). Another challenge is to find the boundary between adjacent gyri which either touch each other or are separated only by a narrow gap where no CSF is evident in the images. The latter is caused by partial volume owing to the low image resolution as well as the motion correction needed to reconstruct a MR volume without motion artefacts. Therefore, the pial surface is derived from the reconstructed white matter surface by deforming this surface outwards towards the cGM/CSF interface without considering the tissue segmentation.

We duplicate the cortical nodes (coloured green in figure 7.11) of the white matter surface mesh such that the resulting mesh consists of right and left white matter, right and left pial, and non-cortical surface nodes and faces. Both, the white matter and the pial surface faces form together with the non-cortical faces a closed genus-0 mesh each. When deforming only pial nodes while keeping the position of white matter and non-cortical surface nodes, the non-self-intersection constraint prevents the pial surface to intersect the white matter and non-cortical faces such that the surfaces partition the MR volume.

In a first deformation step, we move the pial surface nodes outwards in surface normal direction for a few iterations up to a small maximum distance from the white matter surface with non-self-intersection checks disabled. This is done so the initial pial surface nodes are strictly outside the white matter surface as otherwise the non-self-intersection constraint prevents the nodes to move outwards when duplicate faces are found to be intersecting each other. After this initial deformation, we use a deformation model identical to the white matter surface model but with an external force which, starting from the already reconstructed white matter surface, searches for the closest cGM/CSF image edge instead. These edges correspond to a positive derivative of T2w image intensities in surface normal direction as shown by yellow arrows in figure 7.10. When no such cGM/CSF edge is found, e.g., within a narrow sulcus due to partial volume, opposing gyri expand towards each other until stopped by the non-self-intersection constraint.

7.3. Results

7.3.1. Datasets

The image data used in the experiments is from the first batch collected for the dHCP. A summary of the acquisition and reconstruction parameters is given in section 6.4.1. In total, 492 brain MR scans of 453 neonates (dHCP-n492), including 38 follow up scans of premature born neonates, have been processed with a pipeline developed for dHCP. This pipeline is described in detail in Makropoulos et al. (2017b) and includes the brain extraction, bias field correction, and Draw-EM segmentation outlined in section 7.1. The segmentation is followed by our cortical surface reconstruction. Using this pipeline, we obtained models of the cortical surfaces for all 492 brain scans. Of these, scans of 43 term-born neonates (dHCP-n43) were selected independent from this work for a preliminary open access release based on a number of manual and quantitative measures computed from the structural, functional, and diffusion-weighted images. The age distribution of the selected subjects was further ensured to reflect the distribution of the ages at birth and the ages at scan of all imaged neonates. The distribution of GA at birth and PMA at scan of the both datasets are shown in figure 7.13. The IDs of the scans and individual ages are listed in appendix D.2. A subset consisting of 10 of the 492 structural brain scans (dHCP-n10), which is disjoint from the dHCP-n43 set of images used for manual evaluation, has been used for the development of the presented surface reconstruction.

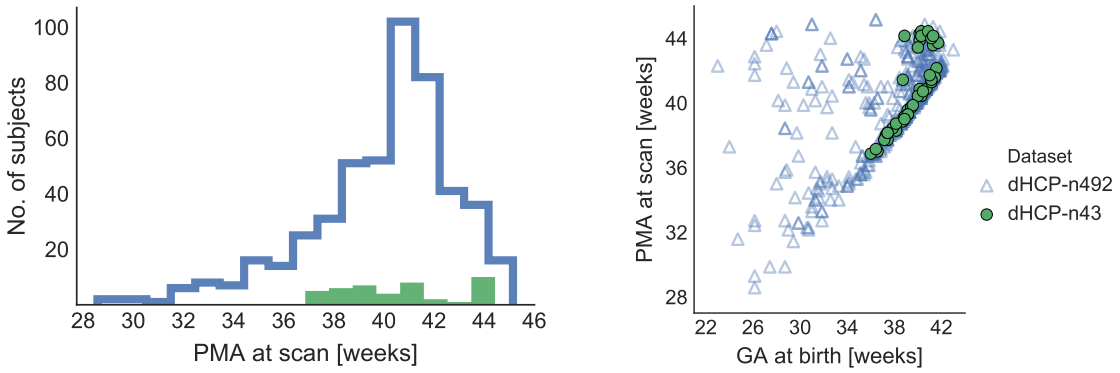


Figure 7.13.: Distribution of PMA at scan (left) for all 492 images, and the subset of 43. The right plot shows the correlation between GA at birth and PMA at scan for both sets.

7.3.2. Manual evaluation

A difficulty in evaluating the accuracy by which the cortical models depict the true anatomy is the lack of ground truth surfaces and segmentations. Imaging the developing brain has further only been made possible more recently and therefore no real state-of-the-art reference methods for the analysis of these brain images exist for comparison. A manual brain segmentation is elaborate and too time consuming for a larger cohort. Moreover, identifying accurate boundaries between tissues from 2D sections is difficult at best, and impossible at worst. For a quantitative evaluation, we performed a manual scoring and comparison of the reconstructed surfaces for the dHCP-n43 subset, as these are included in the first data release. The goal of this evaluation was to establish whether the proposed method yields white matter surfaces which can overcome some of the prominent Draw-EM segmentation errors, i.e., whether the image-based white matter surfaces depict the WM/cGM interface more accurately than the obtained tissue segmentations. For this purpose, we developed a custom evaluation tool which presents an expert with a 2D render of a T2w image patch intersected with the reconstructed surface model. In order to reduce the burden on the human rater and the time needed to assess the results, we further implemented an approach to automatically pre-select ROIs where the proposed white matter surface model differs noticeably from the underlying tissue segmentation. This is to focus the attention on surface patches where there are differences between the two boundaries. In the following we detail this automatic ROI selection and 2D screenshot generation. Afterwards, we briefly describe the graphical interface used by the expert. The implementation of the evaluation tools is made available for reference as open source[5].

Region of interest selection

An ROI used for evaluation is a 3D sub-volume extracted from the T2w image, where the span of this sub-volume is set to 50mm in each dimension. The voxel at the centres of the ROIs is selected as follows. Given a white matter surface mesh of the WM segmentation boundary used to reconstruct the initial genus-0 surface and the final white matter surface

[5] https://github.com/MIRTK/CorticalSurfaceEvaluation (v3.0; last accessed: May 20, 2017)

obtained with the proposed method, we compute for each node of the surface the Euclidean distance to the closest point on the segmentation boundary. We then do the same for each node of the surface mesh which was directly obtained from the segmentation, where the minimum distances to the proposed white matter surface model is computed. Afterwards, we establish links of closest nodes between the two surface meshes to construct a single mesh with only one connected component. This is done to make the following surface node clustering independent of the order of the two input surface meshes. With the two surfaces linked by their closest nodes, we cluster the nodes of the joined mesh based on the previously computed node distances using a region growing approach. This region growing starts at seed nodes with highest distance values. Adjacent nodes of a propagating front node are added to the cluster as long as it exceeds a predefined distance threshold. The average positions, $\bar{\boldsymbol{x}}_k$, of nodes belonging to the k-th cluster are the respective centre points of the chosen ROIs. To limit the number of ROIs per brain scan, we sort the clusters by decreasing number of cluster nodes and keep only the top K clusters. This selection strategy favours ROIs where the reconstructed white matter surface differs most from the segmentation boundary. To avoid bias in the evaluation, we augment the set of ROIs to include also ROIs with centres at spatially uniform distributed and randomly selected nodes. When a ROI overlaps by more than a specified volume ratio with an already selected ROI, it is skipped and the next ROI is chosen instead. For each brain scan, we generate 20 ROIs using the described approach with at least 5 randomly selected.

Optimal orthogonal viewing plane

To enable a quick visual assessment of each selected ROI, we further choose a single orthogonal slice orientation at which to view the 3D sub-volume with the extracted 2D slice centred at the ROI centre voxel. For this we intersect both the segmentation boundary and the reconstructed white matter surface with three orthogonal slices with axial, coronal, and sagittal orientation. The orthogonal view corresponding to the highest Hausdorff distance of the two surface contours that are visible in the respective 2D renders is the one selected for manual assessment.

Manual scoring

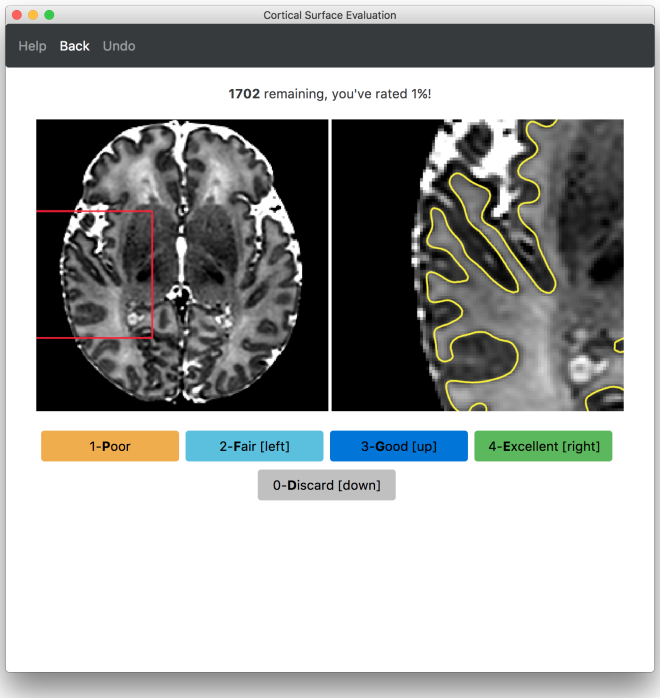
A graphical user interface (GUI) shown in figure 7.14a was used by an expert rater to assign a single score to the surface contour overlaid on the T2w image patch. We chose an even number of quality scores to encourage the expert to provide a tendency towards better or worse, and thereby avoid an overwhelmingly neutral opinion. The number of scores was further limited to four quality scores as it is otherwise difficult for a human rater to make consistent choices. Specifically, the scores available to the expert were:

1. **Poor**: Contour substantially deviates from tissue boundary; missing gyri.
2. **Fair**: Contour close, but with some more obvious mistakes.
3. **Good**: Contour mainly correct, with minor mistakes.
4. **Excellent**: Contour follows the tissue boundary reasonably well.

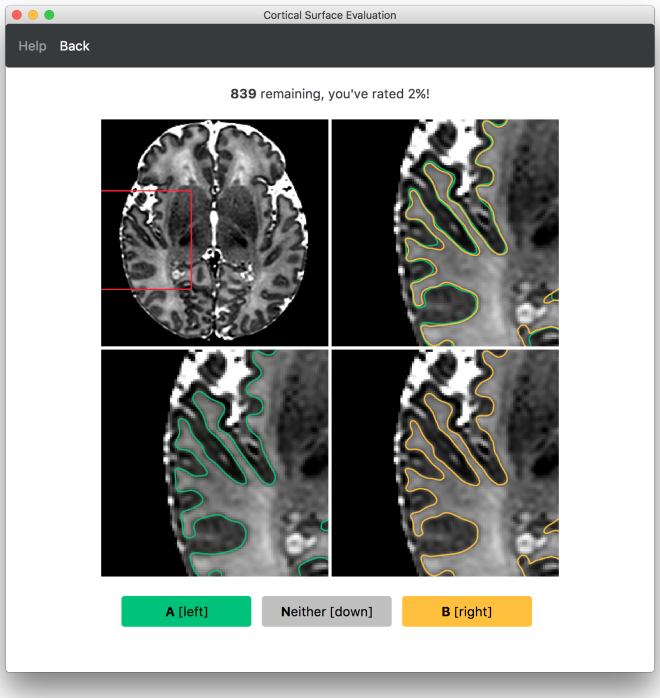
As the 2D renders presented to the rater were selected automatically and some randomly, these may sometimes not be sufficient to enable an assessment of the quality of the cortical surface, e.g., at the boundary of the image volume or a sagittal view through the medial longitudinal fissure. The expert can therefore also choose to discard a presented ROI screenshot instead of assigning a score. All 2D renders of the same image patch are then excluded, including those used for direct comparison of two surfaces.

Manual comparison

A similar GUI shown in figure 7.14b was used by an expert rater to directly compare two cortical surface meshes reconstructed by different algorithms. The task of the expert is to choose the surface contour which more accurately depicts the true tissue boundary. When both surfaces represent the tissue boundary similarly well or contain both comparably significant errors, the expert is encouraged to choose the option “Neither” instead. The assignment of each of the two cortical surface contours to one of the option labels “A” or “B” with corresponding colours green and orange, respectively, is hereby random for each ROI. This avoids bias of the expert towards either one of the two choices and not to predominantly choose either “A” or “B” when uncertain, but requires the expert to sufficiently assess the visible contours before entering their choice.



(a) Manual scoring interface



(b) Manual comparison interface

Figure 7.14.: Graphical interface of custom cortical surface evaluation application.

7.3.3. Expert opinion

In the following we present the results of the manual evaluation of the white matter surfaces of the dHCP-n43 subset carried out by two experts. We compare the results for our method with the approach used by Wright et al. (2015). The approach of Wright et al. (2015) is to deform an initial bounding sphere towards the WM segmentation in normal direction, with force proportional to the minimum distance. The resulting surface has spherical topology and closely follows the segmentation boundary. It is therefore considered as surrogate of a topological correct Draw-EM segmentation.

The automatic ROIs selection method detailed in section 7.3.2 resulted in a set of 860 pre-rendered screenshots for each of the two surfaces for manual scoring, plus one set with both overlaid together for direct comparison. Figure 7.15 summarises the percentage of individually rated contours that have been assigned a common quality score. It can be noted that the proposed method achieves more (48%) "Excellent" results, and fewer (14%) "Fair" results than the reference method (34% "Excellent"; 18% "Fair"), while the percentage of "Poor" results is below 2% for both. The reference model is given a "Good" score in 46% of cases, whereas our model is considered of the same quality only in 36% of the rated contours. However, from the increase in "Excellent" scores we conclude that our approach successfully improves the cortical surface which in case of the reference model is only considered "Good", or "Fair" even. Another observation is that although the expert "AM" is overall more positive about the quality of both models, the relative improvements are similar according to both experts. Regarding intra-rater agreement, an expert assigned a higher score to the opposite model than its counterpart in 28% of cases.

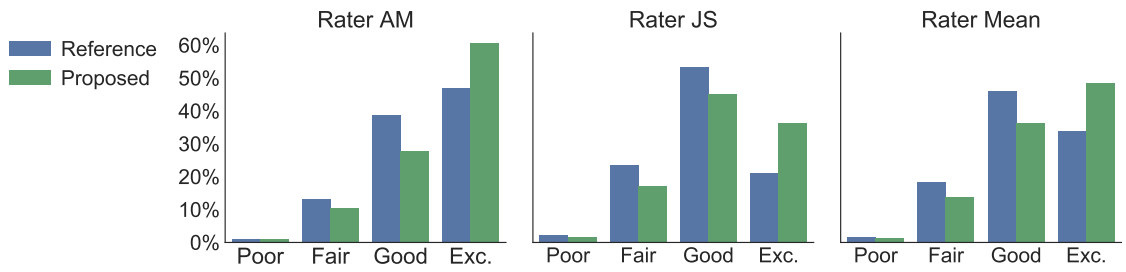


Figure 7.15.: Distribution of expert scores for white matter surface reconstructed either by (blue) the method used by Wright et al. (2015), or (green) our proposed approach.

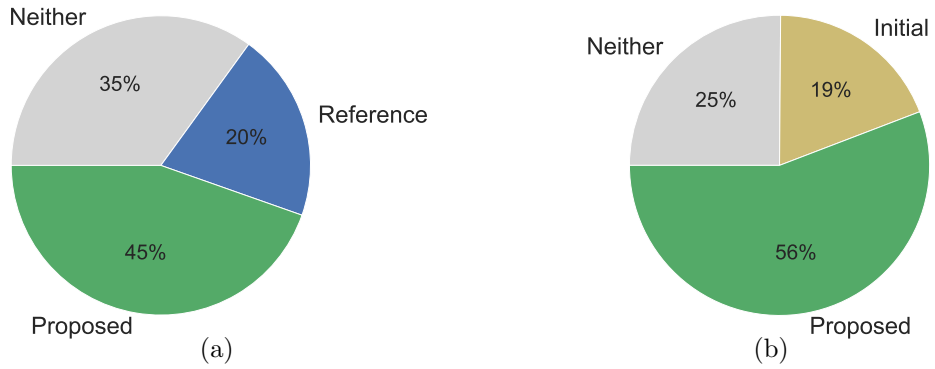


Figure 7.16.: Comparison of proposed white matter surface model to (a) reference model, and (b) initial segmentation-based surface obtained before intensity-based refinement, with different colours for each model. The ratio for which neither is considered better is gray.

When comparing our model to the reference for the same ROI, our model is given a higher score more than twice as often, while neither model is assigned a higher score than the other for 35% of all ROIs. This result is summarised in figure 7.16a. Similarly, the choices of the experts comparing our final white matter surface after image-based refinement directly to the initial segmentation-based surface with clustering-based hole filling of the segmentation boundary distances are presented in figure 7.16b. In 56% of comparisons, the image-edge based refinement has improved the cortical surface, while in 19% of cases the initial model is considered to follow the inner cortical boundary more closely than our final model after image-based refinement. In the remaining 25% of cases, the segmentation and thereby the initial surface has high accuracy already such that neither an improvement nor an impairment of the surface quality is observed.

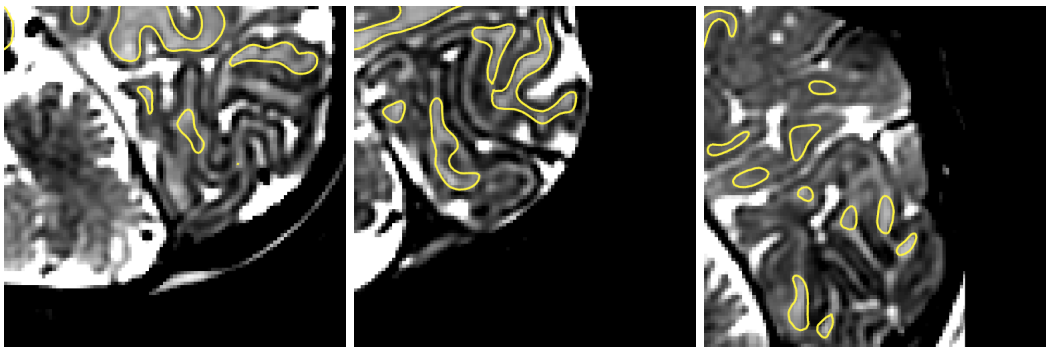
These findings based on the opinion of two independent experts suggest that our proposed white matter surface reconstruction produces better 3D models of the WM/cGM interface of the neonatal cortex than the Draw-EM segmentation. The comparison of segmentation-based and intensity-based models provides convincing evidence that our WM/cGM image edge distance force successfully corrects a considerable number of errors of the tissue classification. Finally, it should be noted that most of the low quality scores are due to smaller gyri not being reflected in the WM segmentation, but cGM only. As we have not explicitly addressed the problem of missing WM in this work, our model is not able to correct these mistakes as seen in figures 7.17a and 7.17b.

7.3.4. Qualitative assessment

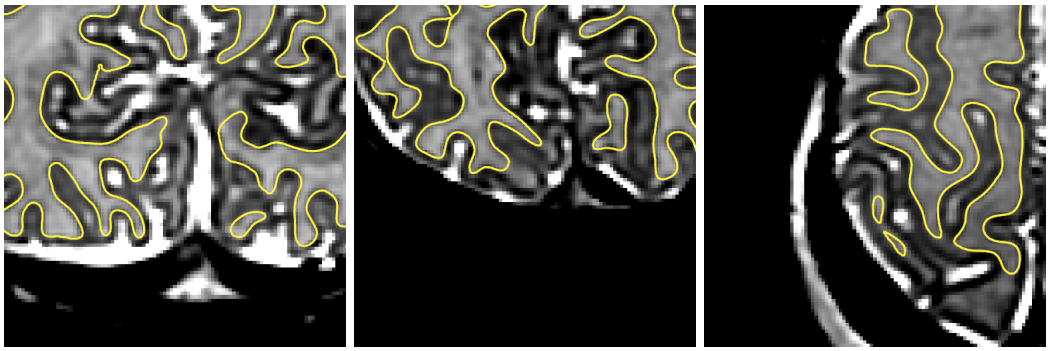
For a first visual assessment of the accuracy of the white matter surface model and to give examples of what quality score has been assigned by an expert to which contours, we show in figure 7.17 a random selection of screenshots which were rated by the experts individually, grouped by their respective quality score.

The reconstructed white matter and pial surface models from the brain scan of one subject are shown in figures 7.18a to 7.18d. These are intersected with a coronal slice of the T2w image in figure 7.18e. This image is overlaid with the WM (red), cGM (blue), and CSF (green) segmentation. It can be observed that the white matter surface follows the WM/cGM interface also where CSF has been mislabelled as WM (yellow boxes). The pial surface further depicts the cGM/CSF boundary more truthfully then the cGM segmentation where CSF appears dark due to partial volume (white boxes). Figures 7.18f and 7.18g show two close up views of the surface meshes of the initial surface model (left) and the final white matter surface (right), demonstrating the effect of the image edge-based refinement of the initial segmentation-based white matter surface. Specifically, figure 7.18f shows an incorrect gyrus within the superior temporal sulcus resulting from misclassified CSF which was successfully "flattened" by our deformable white matter surface model. Figure 7.18g shows a case where a large hole in the segmentation of the central sulcus resulted in a narrow initial surface, stopped only from self-intersecting by the non-self-intersection constraint, and is filled after the image edge-based refinement.

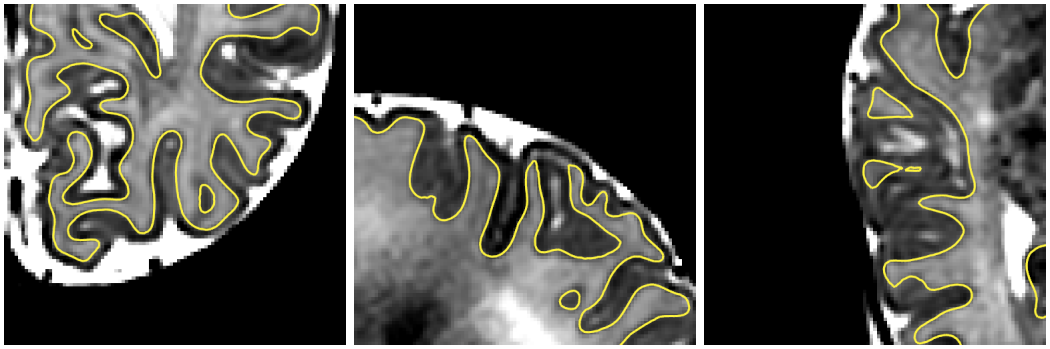
Coronal views of the brain images of nine other subjects, selected from each one-week wide bin of the histogram of PMA at scan within the range from 35.71 to 43.43 weeks, are displayed in figure 7.19, overlaid with the contours of the reconstructed white matter and pial surfaces intersected by the coronal image plane. It can be observed that the surfaces follow the boundaries of the cortex well in these slices. It must be noted though that the overall accuracy of the surface models cannot be assessed from a single slice view. We inspected multiple slices and found no major errors, except of expected mistakes of mainly missing WM because these errors are not currently addressed by our cortical surface models. These limitations are discussed in section 7.4.



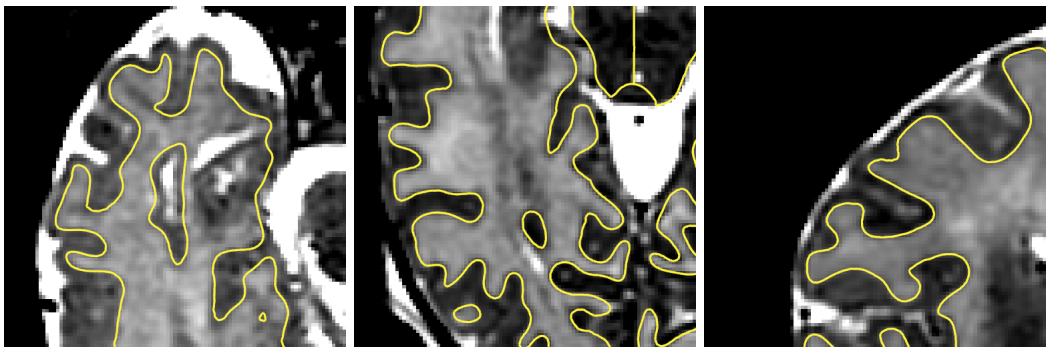
(a) “Poor”



(b) “Fair”

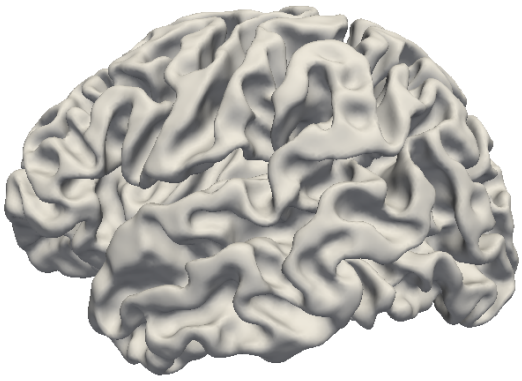


(c) “Good”

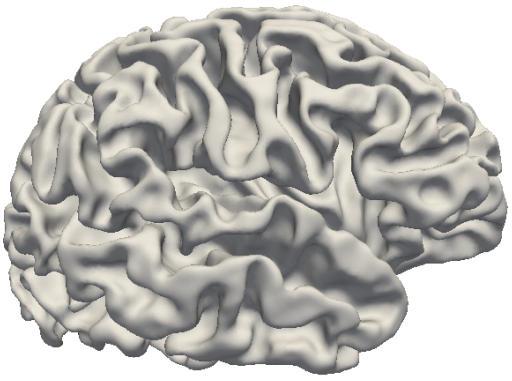


(d) “Excellent”

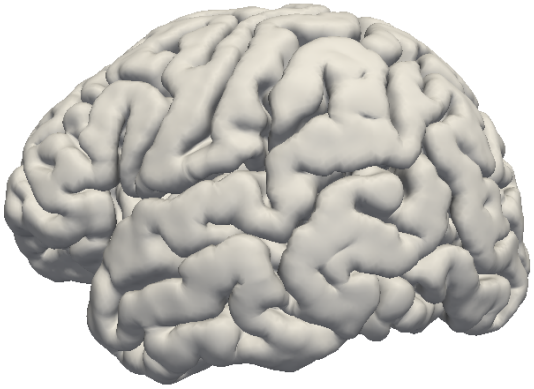
Figure 7.17.: Exemplar image patches overlaid with our white matter surface contours and scored by an expert as either (a) “Poor”, (b) “Fair”, (c) “Good”, or (d) “Excellent”.



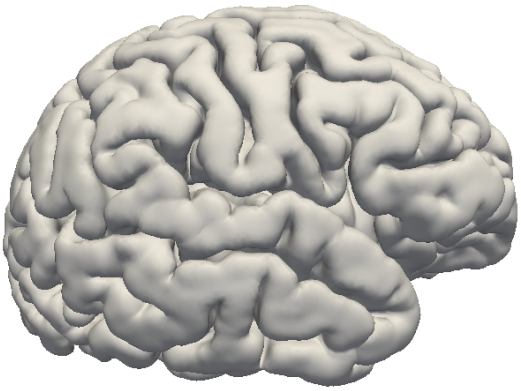
(a) Left white matter surface



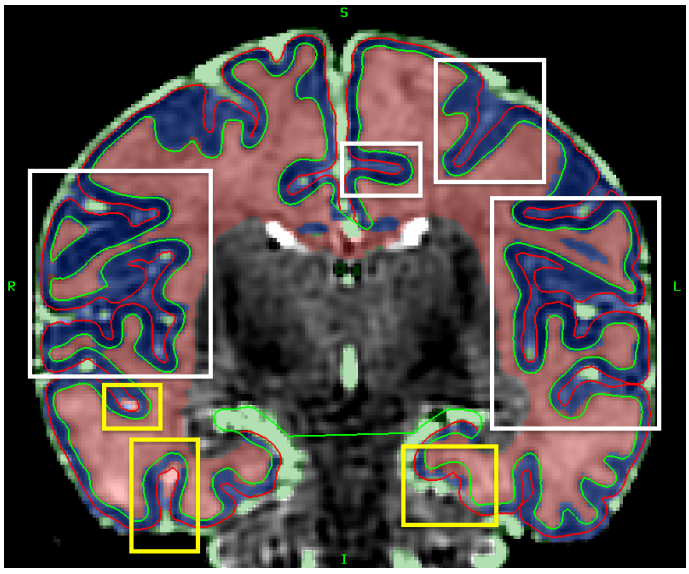
(b) Right white matter surface



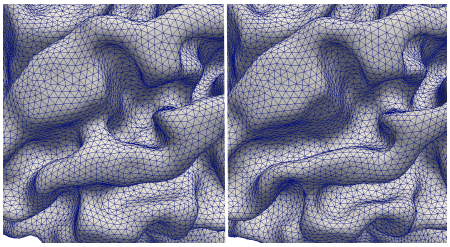
(c) Left pial surface



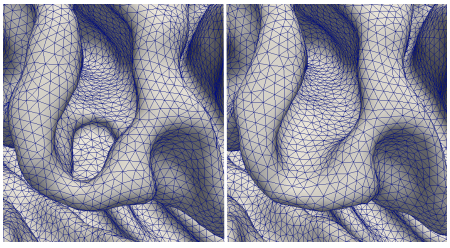
(d) Right pial surface



(e) Illustration of corrected segmentation errors



(f) Corrected sulcal bank



(g) Filled in hole

Figure 7.18.: Qualitative results of cortical surface reconstruction for one brain scan.

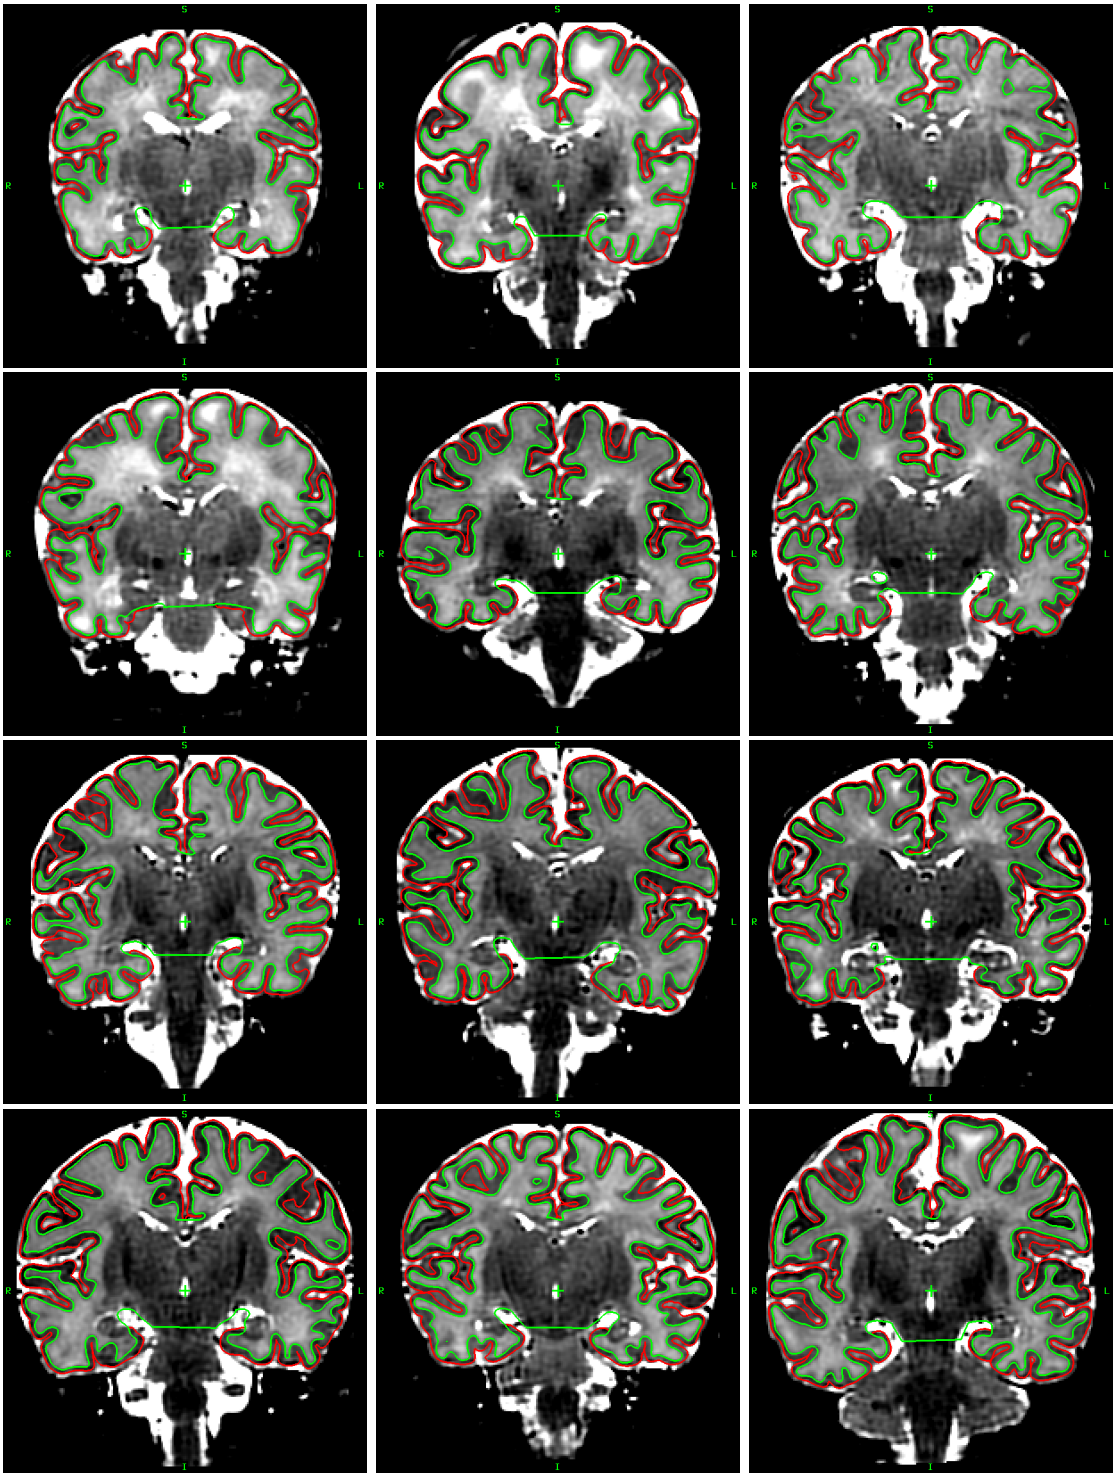


Figure 7.19.: T2w image slices of 12 MRI scans of neonatal brains of different subjects intersected with reconstructed surfaces. From left to right, and top to bottom, the age at scan difference between neonates is less or equal 1 week. The PMA at scan of the youngest brain is 35.7 weeks (top left), whereas the oldest was scanned at 43.4 weeks (bottom right).

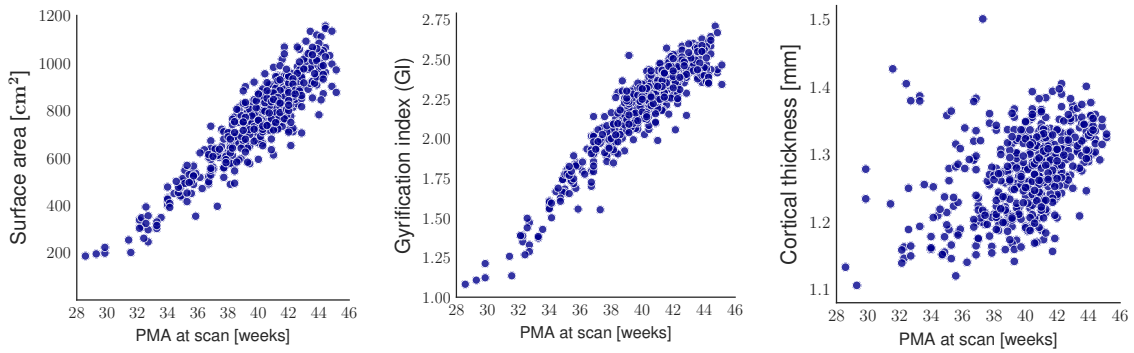


Figure 7.20.: Mean cortical surface area, GI, and thickness with increasing age.

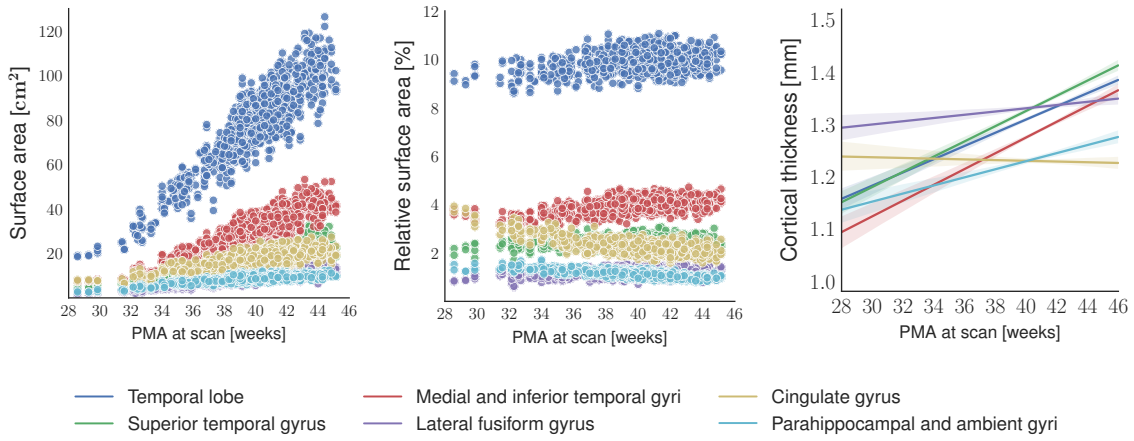


Figure 7.21.: Mean cortical measures for bilateral ROIs with increasing age.

7.3.5. Quantitative measures

Given the cortical surface models obtained with our proposed reconstruction technique, we measured the surface area, thickness, and gyrification index (GI) of the cortex in the dHCP-n492 cohort. Figures 7.20 and 7.21 report these automatically obtained measures with increasing PMA at scan. The cortical ROIs in figure 7.21 are derived from the cortical labels of the volumetric Draw-EM segmentation.

Moeskops et al. (2013) measured the cortical thickness of manually segmented early preterm brains with values around 1-1.2 mm. The thickness in our neonatal cohort, mainly consisting of term-equivalent neonatal brains, was found by our method to be on average around 1.4 mm, with 25th and 75th percentiles of 1.33 and 1.53 mm, respectively. Authors of other automatic cortex segmentation methods report thickness values of 1.2-1.4 mm (Xue et al., 2007) and 1.6 mm (Makropoulos et al., 2016). It should be noted that

the latter is based on the Draw-EM tissue segmentation. As expected, considering partial volume resulting in CSF mislabelled as cGM, the Draw-EM based thickness measures are higher than our results, likely overestimating cortical thickness. The range of our measurements is in better agreement with the values reported by other authors.

Gyrification was manually measured in post-mortem brains by Armstrong et al. (1995). The reported GI values range from 2.0 at around 36 weeks to 2.3-2.8 for different parts of the cortex at around 44 weeks PMA. The GI values obtained from our cortical surface models similarly range from 2.1 at 36 weeks to 2.6-2.9 at 44 weeks PMA.

7.4. Conclusion

We presented a new method for reconstructing the neonatal cortex utilising deformable models of spherical topology. Common errors in the Draw-EM segmentation are hereby corrected for by an edge-based refinement which identifies nearby image edges of the WM/cGM interface. We published the source code of our implementation in the MIRTK Deformable package[6]. The surfaces obtained with this method were found by experts to be superior to the initial Draw-EM segmentation when it comes to accurately delineating the cortical anatomy. Current limitations of the models (see also section 8.2), which do not account for WM that has not been correctly classified as such, need to be addressed in future work by adopting an improved tissue segmentation and a supervised machine learning method to train a regressor of the WM/cGM image edge distance given local image context relative to the orientation and shape of the deforming surface.

[6] https://github.com/MIRTK/Deformable (last accessed: May 20, 2017)

Part IV.

Conclusion

Chapter 8.

Summary and future work

Recent advances in imaging the living brain during gestation and the first weeks of life, enable the study of human brain maturation, and allow us to derive models of normal brain development. These computational models provide us with a reference through which anatomies can be compared, to distinguish normal variation from possible impairments, caused for example by premature birth. In this thesis, the focus has been on the development of methods used to obtain volumetric templates and surface models of the neonatal brain shortly after birth. Volumetric atlases of the human brain are invaluable for unbiased spatial normalisation, where individual brain scans are registered to a reference space, the study of mean shape of primarily sub-cortical structures such as the hippocampi, and the delineation of ROIs for VBM and to measure their volumes for use as imaging biomarkers in disease classification. Accurate models of the cerebral cortex constitute rich complementary representations of the human brain. These surface models are essential for the study of cortical brain morphometry, connectivity, and function. A consistent parcellation of the cortical surfaces extracted from structural MRI is the foundation, from which maps of structural and functional connectivity between grey matter regions can be generated using diffusion-weighted and functional MRI. This is prominently exemplified by the frameworks proposed through the adult Human Connectome Project (HCP)[1] led by researchers at Washington University, University of Minnesota, and Oxford University (Glasser et al., 2013). With recent advances in imaging the immature human brain, the

[1] http://humanconnectome.org/ (last accessed: Jun 20, 2017)

Developing Human Connectome Project (dHCP)[2] led by King's College London, Imperial College London and Oxford University, aims to create the first connectome spanning the period of early life. A prerequisite towards achieving this goal is an accurate segmentation and reconstruction of 3D models of the cerebral cortex of the developing brain.

Accordingly, this thesis has presented new methods for the unbiased construction of spatio-temporal atlases of the developing brain, and the delineation of cortical boundaries in T2w images. Our specific achievements are summarised in the following. Limitations of the presented approaches and directions for future work are given thereafter.

8.1. Achievements

Evaluation of fast diffeomorphic neonatal brain image registration methods. A prerequisite for the brain atlas construction is a registration, which produces spatial mappings that preserve the topology of anatomies. Its parametrisation should further facilitate a definition of a mean deformation that is not biased towards any of the observations. Because average models are more representative for the studied population, when constructed from a larger number of observations, a faster method that reduces the time needed to construct a template is beneficial. Different methods for diffeomorphic registration of adult brain images have been proposed (see chapter 2). In chapter 5, we evaluated the accuracy, inverse consistency, and runtime of efficient methods for the registration of neonatal brain images. We showed that the parametric SVFFD model, first introduced by Modat et al. (2012b), can be used for fast registration of neonatal brain images with guaranteed invertibility. This model enables the definition of a mean, which itself belongs to the one-parameter subgroup of diffeomorphisms generated by SVFs. We proposed a symmetric variant of the SVFFD method based on a single parametric spline function, which is not biased towards either of the anatomies depicted in the images. We obtained an improved anatomical alignment compared to the original method of Modat et al. (2012b) with a gradient approximation similar to LogDemons using the BCH formula.

[2] http://www.developingconnectome.org/ (last accessed: Jun 20, 2017)

Two approaches for the unbiased construction of spatio-temporal brain atlases. Having established the performance of our proposed unbiased SVFFD algorithm for the diffeomorphic registration of neonatal brain images, we derived two related approaches for the construction of a spatio-temporal brain template of the developing brain in chapter 6.

Our first approach is an extension of previous methods developed in our research group (Kuklisova-Murgasova et al., 2011; Serag et al., 2012a). It overcomes the inherent limitations of the method proposed by Serag et al. (2012a) based on the original FFD algorithm (Rueckert et al., 1999), as discussed in section 6.2. This method is based on registrations between all pairs of images in the dataset. A particular shortcoming of the previous approach is that these pairwise registrations are not guaranteed to preserve the topology of the anatomies. Moreover, as shown in section 6.2.1, the individual arithmetic mean FFDs, used to map each individual brain anatomy to the atlas space, are not consistent when the pairwise registrations are not inverse consistent. Inverse consistency is, however, not given by the used asymmetric registration. These methodological shortcomings of the original pairwise atlas construction were resolved using our unbiased, symmetric, and diffeomorphic SVFFD algorithm.

We then proposed in section 6.3 a new groupwise spatio-temporal atlas construction similar to other such techniques, but based on the SVFFD algorithm and Log-Euclidean mean (Arsigny et al., 2006). This iterative approach to atlas construction significantly reduces the computational burden imposed by the previous pairwise approach. Unlike other groupwise methods that were originally developed to construct single time point atlases, our proposed method is based on a continuous recursive formulation of the spatio-temporal mean brain image. We demonstrated that our joint estimation of mean shape and (coarse) temporal growth results in an improved temporal consistency, when compared to an approach, which separately registers each brain image to individual time points.

The spatio-temporal atlases constructed from 275 brain images of term-born neonates acquired by dHCP show a marked increase in the level of visible anatomical detail compared to publicly available atlases. Our proposed groupwise atlas construction results in a spatio-temporal brain template well suited for spatial normalisation, VBM, and DBM, and can serve as deformable template in CA for the study of early brain development.

Deformable parametric surface models of the neonatal cortex. In chapter 7, we proposed a new method for the reconstruction of the neonatal cortex from T2w images. The appearance of the developing brain in MR images is significantly different from adult brains (see section 1.2). This prohibits the application of existing tools for the reconstruction of the adult cortex such as FreeSurfer. Methods for the reconstruction of the neonatal cortex often rely completely on an accurate voxel-wise brain (tissue) segmentation. Segmentations of neonatal brain images still contain substantial errors, however, due to the considerably poorer image contrast compared to adult brain images, owing to the shorter acquisition time and infant motion. More importantly, partial volume, the increased water content of cerebral tissues, and unmyelinated WM all complicate a voxel-wise classification (see section 4.1). This results in CSF being mislabelled as WM or cGM. A surface reconstruction based on the segmentation alone is thus susceptible to these mistakes.

Our method presented in chapter 7 was developed for the reconstruction of the cortical surfaces of the developing human brain at the neonatal period using non-self-intersecting parametric deformable surfaces. Similar to aforementioned approaches, we also employ a previous brain segmentation. Unlike these approaches, however, errors in the segmentation are subsequently accounted for by a local image-based refinement given the original bias corrected T2w (and if available T1w) intensities. More specifically, our cortical surface reconstruction is based on a Draw-EM segmentation of the neonatal brain, and was specifically developed to counter misclassifications of CSF that were commonly observed. This was achieved by introducing an external deformable surface force that is based on the intensity profiles along the surface normal direction. Signed distances of image edges corresponding to the inner- or outer-cortical interfaces are determined based on these intensity profiles at each step of the deformable surface propagation. We showed that the cortical surfaces reconstructed with our deformable models follow the true cortical boundaries more closely than models that were reconstructed using the segmentation alone.

A complete pipeline for the segmentation and cortical reconstruction of the neonatal brain using our approach was made publicly available[3]. It finds application for the processing of the structural MR images in the dHCP.

[3] https://github.com/DevelopingHCP/structural-pipeline (last accessed: 28 Sep 2017)

8.2. Limitations and future work

Estimating an accurate longitudinal template deformation

Our groupwise spatio-temporal brain atlas construction proposed in chapter 6, jointly estimates the mean shape and longitudinal correspondences by approximating the composition of SVFFDs using the BCH formula. This approximation of the longitudinal change does not preserve high frequency components in the composite deformations. Moreover, because the mean shape is not only determined by the spatial correspondences relating each anatomy to the atlas space, but also the regressed mean intensities, these longitudinal deformations do not fully capture the anatomical correspondences of different atlas time points. In other words, when a template at one time point is deformed to a later time point using the longitudinal growth model estimated by our approach, the deformed template will not perfectly match the template anatomy estimated at this later time point.

An accurate longitudinal deformation, which is consistent with the spatio-temporal template anatomy, would provide valuable complementary information to the mean shape and intensity model. For example, these longitudinal deformations can be utilised for registering developing brains at different developmental stages, where each brain image is registered to its age-equivalent template. The correspondences between the two brain images are then given by the composition of these cross-sectional deformations with the longitudinal template deformation (Serag et al., 2012b). During the groupwise atlas construction, an additional longitudinal registration could be employed to refine the current estimate of a growth model, for example using the TDFFD algorithm (De Craene et al., 2012). The spline function of the non-stationary velocity field of this longitudinal registration can be initialised based on the piecewise constant SVFs of the longitudinal SVFFDs estimated by our presented approach.

Robustness of cortical surface reconstruction

We have shown in section 7.3, that our proposed white matter surface model depicts the WM/cGM interface more accurately than the Draw-EM segmentation. In particular, it is robust to segmentation errors, where CSF is mislabelled as WM. This is due to our

novel WM/cGM image edge force, which has been specifically designed to counter these misclassifications. Another type of segmentation error, which is often observed in the posterior medial cerebrum, is WM that is incorrectly classified as cGM because of the hypo-intensities in these regions. In this work, we have not attempted to address these types of errors, where entire cortical folds are not segmented. Consequently, our cortical surface models cannot correct for these mistakes, which are the main reason for a "Poor" or "Fair" expert quality score. This is especially apparent in the screenshots shown in figure 7.17a, and would best be addressed by an improved tissue segmentation. A balloon force as used by Park et al. (2001), which expands the deformable white matter surface while inside WM, may further help to push the surface into missing folds and closer to the edges that our WM/cGM image edge force is looking for.

The presented WM/cGM image edge distance computation is based on hard-coded rules to identify the WM/cGM image edge along the surface normal direction given the corresponding 1D intensity profile. These hard-coded decisions favour an inwards deformation of the surface to counter misclassified CSF within sulci. Given enough ground truth cortical surfaces or WM segmentations, respectively, a supervised machine learning approach promises to generalize better than the current hard-coded distance computation. Given the improved segmentation of the cortex with our surface models, we plan in future work to manually correct the remaining errors in order to generate this training data. When utilising also spatial information, such machine learning based edge detection promises to also correct for some of the false negative WM classification errors.

Finally, a quantitative evaluation of the surface reconstruction given ground truth delineations of the cortical interfaces would provide a clearer indication of the quality of the reconstructed surfaces. For this work, such labour intense annotations have not been available. We expect these to be of great value, however, not only for quality control, but also aforementioned possibilities of applying supervised machine learning techniques to further improve the accuracy of the reconstructed surfaces, and to also eliminate the remaining segmentation errors that have not been addressed in this work.

Bibliography

Ad-Dab'bagh, Y., Einarson, D., Lyttelton, O., Muehlboeck, J.-S., Mok, K., Ivanov, O., Vincent, R. D., Lepage, C., Lerch, J., Fombonne, E., Evans, A. C., 2006. The CIVET image-processing environment: a fully automated comprehensive pipeline for anatomical neuroimaging research. In: Corbetta, M. (ed.), *Proceedings of 12th Annual Meeting of the Organization for Human Brain Mapping (OHBM)*. Florence, Italy, p. S45.

Alexa, M., 2002. Linear combination of transformations. In: *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*. San Antonio, Texas, USA, pp. 380–7. doi:10.1145/566570.566592

Alexander, B., Murray, A. L., Loh, W. Y., Matthews, L. G., Adamson, C., Beare, R., Chen, J., Kelly, C. E., Rees, S., Warfield, S. K., Anderson, P. J., Doyle, L. W., Spittle, A. J., Cheong, J. L., Seal, M. L., Thompson, D. K., 2016. A new neonatal cortical and subcortical brain atlas: the Melbourne Children's Regional Infant Brain (M-CRIB) atlas. *NeuroImage* 147 (Supplement C), 841–51. doi:10.1016/j.neuroimage.2016.09.068

Aljabar, P., Heckemann, R. a., Hammers, A., Hajnal, J. V., Rueckert, D., 2009. Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy. *NeuroImage* 46 (3), 726–38. doi:10.1016/j.neuroimage.2009.02.018

Altaye, M., Holland, S. K., Wilke, M., Gaser, C., 2008. Infant brain probability templates for MRI segmentation and normalization. *NeuroImage* 43 (4), 721–30. doi:10.1016/j.neuroimage.2008.07.060

Amit, Y., 1994. A nonlinear variational problem for image matching. *SIAM Journal on Scientific Computing* 15 (1), 207–24. doi:10.1137/0915014

Anbeek, P., Išgum, I., Kooij, B. J. M. V., Mol, C. P., Kersbergen, K. J., Viergever, M. A., Vries, L. S. D., Benders, M. J. N. L., 2013. Automatic segmentation of eight tissue classes in neonatal brain MRI. *PLoS One* 8 (12), e81895. doi:10.1371/journal.pone.0081895

Anbeek, P., Vincken, K. L., Groenendaal, F., Koeman, A., van Osch, M. J. P., van der Grond, J., 2008. Probabilistic brain tissue segmentation in neonatal magnetic resonance imaging. *Pediatric Research* 63 (2), 158–63. doi:10.1203/PDR.0b013e31815ed071

Andersson, J. L. R., Jenkinson, M., Smith, S., 2007. Non-linear registration, aka spatial normalisation. Technical Report TR07JA2, FMRIB Centre, Oxford, UK.

Ardekani, B. A., Guckemus, S., Bachman, A., Hoptman, M. J., Wojtaszek, M., Nierenberg, J., 2005. Quantitative comparison of algorithms for inter-subject registration of 3D volumetric brain MRI scans. *Journal of Neuroscience Methods* 142 (1), 67–76. doi:10.1016/j.jneumeth.2004.07.014

Armstrong, E., Schleicher, A., Omran, H., Curtis, M., Zilles, K., 1995. The ontogeny of human gyrification. *Cerebral Cortex* 5 (1), 56–63.

Arsigny, V., Commowick, O., Ayache, N., Pennec, X., 2009. A fast and Log-Euclidean polyaffine framework for locally linear registration. *Journal of Mathematical Imaging and Vision* 33 (2), 222–38. doi:10.1007/s10851-008-0135-9

Arsigny, V., Commowick, O., Pennec, X., Ayache, N., 2006. A Log-Euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.), *Proceedings of 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 924–31. doi:10.1007/11866565_113

Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2007. Geometric means in a novel vector space structure on symmetric positivedefinite matrices. *SIAM Journal on Matrix Analysis and Applications* 29 (1), 328–47. doi:10.1137/050637996

Arsigny, V., Pennec, X., Ayache, N., 2005. Polyrigid and polyaffine transformations: a novel geometrical tool to deal with non-rigid deformations - application

to the registration of histological slices. *Medical Image Analysis* 9 (6), 507–23. doi:10.1016/j.media.2005.04.001

Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. *NeuroImage* 38 (1), 95–113. doi:10.1016/j.neuroimage.2007.07.007

Ashburner, J., 2012. SPM: A history. *NeuroImage* 62, 791–800. doi:10.1016/j.neuroimage.2011.10.025

Ashburner, J., Friston, K. J., 1999. Nonlinear spatial normalization using basis functions. *Human Brain Mapping* 7 (4), 254–66. doi:10.1002/(SICI)1097-0193(1999)7:4<254::AID-HBM4>3.0.CO;2-G

Ashburner, J., Friston, K. J., 2000. Voxel-based morphometry–the methods. *NeuroImage* 11 (6), 805–21. doi:10.1006/nimg.2000.0582

Ashburner, J., Friston, K. J., 2005. Unified segmentation. *NeuroImage* 26 (3), 839–51. doi:10.1016/j.neuroimage.2005.02.018

Ashburner, J., Friston, K. J., 2011. Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. *NeuroImage* 55 (3), 954–67. doi:10.1016/j.neuroimage.2010.12.049

Ashburner, J., Hutton, C., Frackowiak, R., Johnsrude, I., Price, C., Friston, K., 1998. Identifying global anatomical differences: Deformation-based morphometry. *Human Brain Mapping* 6 (5-6), 348–57. doi:10.1002/(SICI)1097-0193(1998)6:5/6<348::AID-HBM4>3.0.CO;2-P

Avants, B. B., Epstein, C. L., Grossman, M., Gee, J. C., 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis* 12 (1), 26–41. doi:10.1016/j.media.2007.06.004

Avants, B. B., Gee, J. C., 2004. Shape averaging with diffeomorphic flows for atlas creation. In: *Proceedings of 2nd IEEE International Symposium on Biomedical Imaging (ISBI)*. vol. 1. pp. 595–8. doi:10.1109/ISBI.2004.1398608

Avants, B. B., Tustison, N. J., Wu, J., Cook, P. A., Gee, J. C., 2011. An open source multivariate framework for n-tissue segmentation with evaluation on public data. *Neuroinformatics* 9 (4), 381–400. doi:10.1007/s12021-011-9109-y

Avants, B. B., Yushkevich, P., Pluta, J., Minkoff, D., Korczykowski, M., Detre, J., Gee, J. C., 2010. The optimal template effect in hippocampus studies of diseased populations. *NeuroImage* 49 (3), 2457–66. doi:10.1016/j.neuroimage.2009.09.062

Bai, J., Abdul-Rahman, M. F., Rifkin-Graboi, A., Chong, Y. S., Kwek, K., Saw, S. M., Godfrey, K. M., Gluckman, P. D., Fortier, M. V., Meaney, M. J., Qiu, A., 2012. Population differences in brain morphology and microstructure among Chinese, Malay, and Indian neonates. *PLoS One* 7 (10). doi:10.1371/journal.pone.0047816

Bajcsy, R., Kovacic, S., 1989. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing* 46 (1), 1–21. doi:10.1016/S0734-189X(89)80014-3

Bajcsy, R., Lieberson, R., Reivich, M., 1983. A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. *Journal of Computer Assisted Tomography* 7 (4), 618–25.

Ball, G., Srinivasan, L., Aljabar, P., Counsell, S. J., Durighel, G., Hajnal, J. V., Rutherford, M. a., Edwards, a. D., 2013. Development of cortical microstructure in the preterm human brain. *Proceedings of the National Academy of Sciences* 110 (23), 9541–6. doi:10.1073/pnas.1301652110

Baluwala, H. Y., Saddi, K. A., Schnabel, J. A., 2010. Log-unbiased elastic image registration with spatial constraint for 3D CT lung images. In: *Proceedings of Medical Image Analysis and Understanding*. Warwick, UK, pp. 5–9. doi:10.13140/2.1.2970.3682

Baluwala, H. Y., Saddi, K. A., Schnabel, J. A., 2011. Non-rigid chest image registration with preservation of topology and rigid structures. *Annals of the British Machine Vision Association (BMVA)* 2011 (4), 1–15.

Bates, A., Schuh, A., Williams, B., Lanier, M., McConnell, K., Loew, W., Fleck, R., Woods, J., Dumoulin, C., Amin, R., Apr 2017. Using dynamic MRI to create moving

boundary conditions for CFD: determining the causes of upper airway motion in sleep apnea. In: *Proceedings of 25th Annual Meeting of International Society for Magnetic Resonance in Medicine (ISMRM)*. Honolulu, HI, USA.

Battin, M. R., Maalouf, E. F., Counsell, S. J., Herlihy, a. H., Rutherford, M. a., Azzopardi, D., Edwards, a. D., 1998. Magnetic resonance imaging of the brain in very preterm infants: visualization of the germinal matrix, early myelination, and cortical folding. *Pediatrics* 101 (6), 957–62. doi:10.1542/peds.101.6.957

Bazin, P.-L., Pham, D. L., 2008. Homeomorphic brain image segmentation with topological and statistical atlases. *Medical Image Analysis* 12 (5), 616–25. doi:10.1016/j.media.2008.06.008

Beare, R. J., Chen, J., Kelly, C. E., Alexopoulos, D., Smyser, C. D., Rogers, C. E., Loh, W. Y., Matthews, L. G., Cheong, J. L. Y., Spittle, A. J., Anderson, P. J., Doyle, L. W., Inder, T. E., Seal, M. L., Thompson, D. K., 2016. Neonatal brain tissue classification with morphological adaptation and unified segmentation. *Frontiers in Neuroinformatics* 10, 1–17. doi:10.3389/fninf.2016.00012

Beg, M. F., Khan, A., 2007. Symmetric data attachment terms for large deformation image registration. *IEEE Transactions on Medical Imaging* 26 (9), 1179–89. doi:10.1109/TMI.2007.898813

Beg, M. F., Miller, M. I., Trouvé, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision* 61 (2), 139–57. doi:10.1023/B:VISI.0000043755.93987.aa

Benkarim, O. M., Sanroma, G., Zimmer, V. A., Muñoz-Moreno, E., Hahner, N., Eixarch, E., Camara, O., González Ballester, M. A., Piella, G., 2017. Toward the automatic quantification of in utero brain development in 3D structural MRI: A review. *Human Brain Mapping* 35 (4), 2772–87. doi:10.1002/hbm.23536

Bertrand, G., 1994. Simple points, topological numbers and geodesic neighborhoods in cubic grids. *Pattern Recognition Letters* 15 (10), 1003–11. doi:10.1016/0167-8655(94)90032-9

Besag, J., 1986. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)* 48 (3), 259–302. url: http://www.jstor.org/stable/2345426

Bhatia, K., Hajnal, J., Puri, B., Edwards, A., Rueckert, D., 2004. Consistent groupwise non-rigid registration for atlas construction. In: *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI)*. vol. 1. pp. 908–11. doi:10.1109/ISBI.2004.1398686

Bhatia, K. K., Aljabar, P., Boardman, J. P., Srinivasan, L., Murgasova, M., Counsell, S. J., Rutherford, M. a., Hajnal, J., Edwards, a. D., Rueckert, D., 2007. Groupwise combined segmentation and registration for atlas construction. In: Ayache, N., Ourselin, S., Maeder, A. (eds.), *Proceedings of 10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 532–40. doi:10.1007/978-3-540-75757-3_65

Bihan, D. L., 2014. Looking Inside the Brain: The Power of Neuroimaging. Princeton University Press, Princeton, NJ, USA.

Blencowe, H., Cousens, S., Oestergaard, M. Z., Chou, D., Moller, A. B., Narwal, R., Adler, A., Vera Garcia, C., Rohde, S., Say, L., Lawn, J. E., 2012. National, regional, and worldwide estimates of preterm birth rates in the year 2010 with time trends since 1990 for selected countries: A systematic analysis and implications. *The Lancet* 379 (9832), 2162–2172. doi:10.1016/S0140-6736(12)60820-4

Bookstein, F. L., 1989. Principal warps thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11 (6), 567–85.

Bossa, M., Hernandez, M., Olmos, S., 2007. Contributions to 3D diffeomorphic atlas estimation: application to brain images. In: Ayache, N., Ourselin, S., Maeder, A. (eds.), *Proceedings of 10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 667–74. doi:10.1007/978-3-540-75757-3_81

Bossa, M., Olmos, S., 2008. A new algorithm for the computation of the group logarithm of diffeomorphisms. In: *Proceedings of 1st MICCAI Workshop on Mathematical Foundations of Computational Anatomy*. pp. 1–11.

Bossa, M., Zacur, E., Olmos, S., 2008. Algorithms for computing the group exponential of diffeomorphisms: Performance evaluation. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. pp. 1–8. doi:10.1109/CVPRW.2008.4563005

Boykov, Y., Veksler, O., Zabih, R., 2001. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (11), 1222–39. doi:10.1109/34.969114

Breiman, L., 2001. Random forests. *Machine Learning* 45 (1), 5–32. doi:10.1023/A:1010933404324

Bro-Nielsen, M., Gramkow, C., 1996. Fast fluid registration of medical images. In: Höhne, K. H., Kikinis, R. (eds.), *Proceedings of 4th International Conference on Visualization in Biomedical Computing (VBC)*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 265–76. doi:10.1007/BFb0046964

Broit, C., 1981. Optimal registration of deformed images. Ph.D. thesis, University of Pennsylvania, Philadelphia.

Burges, C. J., 1998. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery* 2 (2), 121–67. doi:10.1023/A:1009715923555

Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N., 2003. Iconic feature based nonrigid registration: the PASHA algorithm. *Computer Vision and Image Understanding* 89 (2-3), 272–98. doi:10.1016/S1077-3142(03)00002-X

Cachier, P., Pennec, X., Ayache, N., 1999. Fast non rigid matching by gradient descent: Study and improvements of the "Demons" algorithm. Technical Report RR-3706, INRIA. url: https://hal.inria.fr/inria-00072962

Cachier, P., Rey, D., 2000. Symmetrization of the non-rigid registration problem using inversion-invariant energies: Application to multiple sclerosis. In: Delp, S. L., DiGoia, A. M., Jaramaz, B. (eds.), *Proceedings of 3rd International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 472–81. doi:10.1007/978-3-540-40899-4_48

Cahill, N. D., Noble, J. A., Hawkes, D. J., 2009. Demons algorithms for fluid and curvature registration. In: *Proceedings of 6th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 730–33. doi:10.1109/ISBI.2009.5193151

Cardoso, M. J., Clarkson, M. J., Ridgway, G. R., Modat, M., Fox, N. C., Ourselin, S., 2011a. LoAd: A locally adaptive cortical segmentation algorithm. *NeuroImage* 56 (3), 1386–97. doi:10.1016/j.neuroimage.2011.02.013

Cardoso, M. J., Melbourne, A., Kendall, G. S., Modat, M., Hagmann, C. F., Robertson, N. J., Marlow, N., Ourselin, S., 2011b. Adaptive neonate brain segmentation. In: *Proceedings of 14th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 378–86.

Cardoso, M. J., Melbourne, A., Kendall, G. S., Modat, M., Robertson, N. J., Marlow, N., Ourselin, S., 2013. AdaPT: An adaptive preterm segmentation algorithm for neonatal brain MRI. *NeuroImage* 65, 97–108. doi:10.1016/j.neuroimage.2012.08.009

Ceritoglu, C., Oishi, K., Li, X., Chou, M. C., Younes, L., Albert, M., Lyketsos, C., van Zijl, P. C. M., Miller, M. I., Mori, S., 2009. Multi-contrast large deformation dif-

feomorphic metric mapping for diffusion tensor imaging. *NeuroImage* 47 (2), 618–27. doi:10.1016/j.neuroimage.2009.04.057

Chefd'Hotel, C., Hermosillo, G., Faugeras, O., 2002. Flows of diffeomorphisms for multi-modal image registration. In: *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 753–56. doi:10.1109/ISBI.2002.1029367

Chita, S. M., Benders, M., Moeskops, P., Kersbergen, K. J., Viergever, M. A., Isgum, I., 2013. Automatic segmentation of the preterm neonatal brain with MRI using supervised classification. In: *Proceedings of Society of Photographic Instrumentation Engineers (SPIE)*. vol. 8669. pp. 1–6. doi:10.1117/12.2006753

Choi, Y., Lee, S., 2000. Injectivity conditions of 2D and 3D uniform cubic B-spline functions. *Graphical Models* 62 (6), 411–27. doi:10.1006/gmod.2000.0531

Christensen, G., Joshi, S., Miller, M., 1997. Volumetric transformation of brain anatomy. *IEEE Transactions on Medical Imaging* 16 (6), 864–877. doi:10.1109/42.650882

Christensen, G. E., 1994. Deformable shape models for anatomy. Ph.D. thesis, Washington University.

Christensen, G. E., Johnson, H. J., 2001. Consistent image registration. *IEEE Transactions on Medical Imaging* 20 (7), 568–82. doi:10.1109/42.932742

Christensen, G. E., Miller, M. I., Vannier, M., 1994a. A 3D deformable magnetic resonance textbook based on elasticity. In: *Proceedings of the American Association for Artificial Intelligence, Symposium: Applications of Computer Vision in Medical Image Processing.*

Christensen, G. E., Rabbitt, R. D., Miller, M. I., 1994b. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine & Biology* 39, 609–18.

Christensen, G. E., Rabbitt, R. D., Miller, M. I., 1996. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing* 5 (10), 1435–47. doi:10.1109/83.536892

Clouchoux, C., Kudelski, D., Gholipour, A., Warfield, S. K., Viseur, S., Bouyssi-Kobar, M., Mari, J. L., Evans, A. C., Du Plessis, A. J., Limperopoulos, C., 2012. Quantitative in vivo MRI measurement of cortical development in the fetus. *Brain Structure and Function* 217 (1), 127–139. doi:10.1007/s00429-011-0325-x

Cocosco, C. A., Zijdenbos, A. P., Evans, A. C., 2003. A fully automatic and robust brain MRI tissue classification method. *Medical Image Analysis* 7 (4), 513–27. doi:10.1016/S1361-8415(03)00037-9

Cohen, L. D., Cohen, I., 1993. Finite-element methods for active contour models and balloons for 2-D and 3-D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15 (11), 1131–47. doi:10.1109/34.244675

Collignon, A., Vandermeulen, D., Suetens, P., Marchal, G., 1995. 3D multi-modality medical image registration using feature space clustering. In: *Proceedings of 1st International Conference on Computer Vision, Virtual Reality and Robotics in Medicine*. pp. 195–204. doi:10.1007/978-3-540-49197-2_22

Collins, D. L., Holmes, C. J., Peters, T. M., Evans, A. C., 1995. Automatic 3-D model-based neuroanatomical segmentation. *Human Brain Mapping* 3 (3), 190–208. doi:10.1002/hbm.460030304

Collins, D. L., Zijdenbos, A. P., Baar, W. F. C., Evans, A. C., Baare, W. F. C., Evans, A. C., 1999. ANIMAL+INSECT: Improved cortical structure segmentation. In: *Proceedings of 16th International Conference on Information Processing in Medical Imaging*. pp. 210–23. doi:10.1007/3-540-48714-X

Cootes, T. F., Hill, A., Taylor, C. J., Haslam, J., 1993. The use of active shape models for locating structures in medical images. In: Barrett, H. H., Gmitro, A. F. (eds.), *Proceedings of 13th International Conference on Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 33–47. doi:10.1007/BFb0013779

Cootes, T. F., Hill, A., Taylor, C. J., Haslam, J., 1994. Use of active shape models for locating structures in medical images. *Image and Vision Computing* 12 (6), 355–65. doi:10.1016/0262-8856(94)90060-4

Cordero-Grande, L., Teixeira, R. P. A., Hughes, E. J., Hutter, J., Price, A. N., Hajnal, J., 2016. Sensitivity encoding for aligned multishot magnetic resonance reconstruction. *IEEE IEEE Transactions on Computational Imaging* 2 (3), 266–80. doi:10.1109/TCI.2016.2557069

Counsell, S. J., Maalouf, E. F., Fletcher, A. M., Duggan, P., Battin, M., Lewis, H. J., Herlihy, A. H., Edwards, A. D., Bydder, G. M., Rutherford, M. A., 2002. MR imaging assessment of myelination in the very preterm brain. *American Journal of Neuroradiology* 23 (5), 872–81.

Coupé, P., Manjón, J. V., Fonov, V., Pruessner, J., Robles, M., Collins, D. L., 2011. Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage* 54 (2), 940–54. doi:10.1016/j.neuroimage.2010.09.018

Dai, Y., Shi, F., Wang, L., Wu, G., Shen, D., 2013. iBEAT: A toolbox for infant brain magnetic resonance image processing. *Neuroinformatics* 11 (2), 211–25. doi:10.1007/s12021-012-9164-z

Dale, A., Fischl, B., Sereno, M., 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage* 9 (2), 179–94. doi:10.1006/nimg.1998.0395

Dale, A. M., Sereno, M. I., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: A linear approach. *Journal of Cognitive Neuroscience* 5 (2), 162–76. doi:10.1162/jocn.1993.5.2.162

Davatzikos, C., 1996. Spatial normalization of 3d brain images using deformable models. *Journal of Computer Assisted Tomography* 20 (4), 656–65.

Davatzikos, C., Bryan, N., 1996. Using a deformable surface model to obtain a shape representation of the cortex. *IEEE Transactions on Medical Imaging* 15 (6), 785–95. doi:10.1109/42.544496

Davatzikos, C., Prince, J. L., 1995. An active contour model for mapping the cortex. *IEEE Transactions on Medical Imaging* 14 (1), 65–80. doi:10.1109/42.370403

Davis, B. C., Fletcher, P. T., Bullitt, E., Joshi, S., 2007. Population shape regression from random design data. In: *Proceedings of 11th IEEE International Conference on Computer Vision (ICCV)*. pp. 1–7. doi:10.1109/ICCV.2007.4408977

Davis, M. H., Khotanzad, A., Flamig, D. P., Harms, S. E., 1997. A physics-based coordinate transformation for 3-D image matching. *IEEE Transactions on Medical Imaging* 16, 317–28. doi:10.1109/42.585766

de Boor, C., 2001. A Practical Guide to Splines - Revised Edition. Springer-Verlag, New York.

De Craene, M., Piella, G., Camara, O., Duchateau, N., Silva, E., Doltra, A., D'hooge, J., Brugada, J., Sitges, M., Frangi, A. F., 2012. Temporal diffeomorphic free-form deformation: application to motion and strain estimation from 3d echocardiography. *Medical Image Analysis* 16 (2), 427–50. doi:10.1016/j.media.2011.10.006

Dempster, A. P., Laird, N. M., Rubin, D. B., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39 (1), 1–38. url: http://www.jstor.org/stable/2984875

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., Killiany, R. J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31 (3), 968–80. doi:10.1016/j.neuroimage.2006.01.021

Di Stefano, L., Mattoccia, S., Mola, M., 2003. An efficient algorithm for exhaustive template matching based on normalized cross correlation. In: *Proceedings of 12th International Conference on Image Analysis and Processing*. pp. 322–327. doi:10.1109/ICIAP.2003.1234070

Dice, L. R. ., 1945. Measures of the amount of ecologic association between species. *Ecology* 26 (3), 297–302. doi:10.2307/1932409

Dittrich, E., Riklin Raviv, T., Kasprian, G., Donner, R., Brugger, P. C., Prayer, D., Langs, G., 2014. A spatio-temporal latent atlas for semi-supervised learning of fetal brain segmentations and morphological age estimation. *Medical Image Analysis* 18 (1), 9–21. doi:10.1016/j.media.2013.08.004

Dubois, J., Benders, M., Cachia, A., Lazeyras, F., Ha-Vinh Leuchter, R., Sizonenko, S. V., Borradori-Tolsa, C., Mangin, J. F., Hüppi, P. S., 2008. Mapping the early cortical folding process in the preterm newborn brain. *Cerebral Cortex* 18 (6), 1444–54. doi:10.1093/cercor/bhm180

Dunn, J. C., 1974. A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters. *Journal of Cybernetics* 3 (3), 32–57. doi:10.1080/01969727308546046

Dupuis, P., Grenander, U., Miller, M. I., 1998. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics* LVI (3), 587–600.

Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N., 2009. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: *Proceedings of 12th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 297–304.

Edwards, P. J., Hill, D. L., Little, J. a., Hawkes, D. J., 1998. A three-component deformation model for image-guided surgery. *Medical Image Analysis* 2 (4), 355–67. doi:10.1016/S1361-8415(98)80016-9

Evans, A., Collins, D., Mills, S., Brown, E., Kelly, R., Peters, T., 1993. 3D statistical neuroanatomical models from 305 MRI volumes. In: *IEEE Conference Record Nuclear Science Symposium and Medical Imaging Conference*. vol. 3. pp. 1813–17. doi:10.1109/NSSMIC.1993.373602

Evans, A. C., Dai, W., Collins, L., Neelin, L., Marrett, S., 1991. Warping of a computereized 3-D atlas to match brain image volumes for quantitative neuroenetomical and functional analysis. In: *Proceedings of Society of Photographic Instrumentation Engineers (SPIE)*. pp. 236–46. doi:10.1117/12.45221

Faugeras, O., Hebert, M., 1986. The representation, recognition, and locating of 3-D objects. *International Journal of Robotics Research* 5 (3), 27–52. doi:10.1177/027836498600500302

Feng, W., Reeves, S. J., Denney, T. S., Lloyd, S., Dell'Italia, L., Gupta, H., 2009. A new consistent image registration formulation with a B-spline deformation model. In: *Proceedings of 6th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 979–82. doi:10.1109/ISBI.2009.5193218

Fernández, V., Llinares-Benadero, C., Borrell, V., 2016. Cerebral cortex expansion and folding: what have we learned? *EMBO Journal* 35 (10), 1–24. doi:10.15252/embj.201593701

Fischer, B., Modersitzki, J., 2002. Fast diffusion registration. In: Nashed, M. Z., Scherzer, O. (eds.), *Inverse Problems, Image Analysis, and Medical Imaging*. vol. 313. American Mathematical Society, Providence, RI, USA, pp. 117–127. doi:10.1090/conm/313/05372

Fischer, B., Modersitzki, J., 2003. Curvature based image registration. *Journal of Mathematical Imaging and Vision* 18 (1), 81–5. doi:10.1023/A:1021897212261

Fischer, B., Modersitzki, J., 2004. A unified approach to fast image registration and a new curvature based registration technique. *Linear Algebra and its Applications* 380, 107–24. doi:10.1016/j.laa.2003.10.021

Fischer, C., Operto, G., Laguitton, S., Perrot, M., Denghien, I., Rivire, D., Mangin, J.-F., 2012. Morphologist 2012: the new morphological pipeline of BrainVISA. In: *Proceedings of Human Brain Mapping*.

Fischl, B., 2012. FreeSurfer. *NeuroImage* 62, 774–81. doi:10.1016/j.neuroimage.2012.01.021

Fischl, B., Liu, A., Dale, A. M., 2001. Automated manifold surgery: Constructing geometrically accurate and topologically correct models of the human cerebral cortex. *IEEE Transactions on Medical Imaging* 20 (1), 70–80. doi:10.1109/42.906426

Fischl, B., Sereno, M., Dale, A., 1999. Cortical surface-based analysis: II. inflation, flattening, and a surface-based coordinate system. *NeuroImage* 9 (2), 195–207. doi:10.1006/nimg.1998.0396

Flifla, J., Dillenseger, J. L., Haigron, P., Coatrieux, J. L., 1992. Quaternion based registration method in medical imaging. In: *Proceedings of the 1992 International Biomedical Engineering Days*. pp. 161–65. doi:10.1109/IBED.1992.247105

Fonov, V., Evans, A. C., Botteron, K., Almli, C. R., McKinstry, R. C., Collins, D. L., 2011. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* 54 (1), 313–27. doi:10.1016/j.neuroimage.2010.07.033

Fornefett, M., Rohr, K., Stiehl, H. S., 2001. Radial basis functions with compact support for elastic registration of medical images. *Image and Vision Computing* 19 (1-2), 87–96. doi:10.1016/S0262-8856(00)00057-3

Frangi, A. F., Niessen, W. J., Vincken, K. L., Viergever, M. a., 1998. Multiscale vessel enhancement filtering. In: Wells, W. M., Colchester, A., Delp, S. (eds.), *Proceedings of 1st International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 130–37. doi:10.1007/BFb0056195

Gee, J. C., Bajcsy, R. K., 1999. Elastic matching: Continuum mechanical and probabilistic analysis. In: Toga, A. W. (ed.), *Brain Warping*. Academic Press, San Diego, ch. 11, pp. 183–197.

Ghadimi, S., Mohtasebi, M., Abrishami Moghaddam, H., Grebe, R., Gity, M., Wallois, F., 2017. A neonatal bimodal MR-CT head template. *PLoS One* 12 (1), e0166112. doi:10.1371/journal.pone.0166112

Gholipour, A., Limperopoulos, C., Clancy, S., Clouchoux, C., Akhondi-asl, A., Estroff, J. A., Warfield, S. K., 2014. Construction of a deformable spatiotemporal MRI atlas of the fetal brain: Evaluation of similarity metrics and deformation models. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.), *Proceedings of 17th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer International Publishing, pp. 292–9. doi:10.1007/978-3-319-10470-6_37

Gholipour, A., Rollins, C. K., Velasco-Annis, C., Ouaalam, A., Akhondi-Asl, A., Afacan, O., Ortinau, C. M., Clancy, S., Limperopoulos, C., Yang, E., Estroff, J. A., Warfield, S. K., 2017. A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth. *Scientific Reports* 7 (1), 476. doi:10.1038/s41598-017-00525-w

Glasser, M., Sotiropoulos, S., Wilson, J., Coalson, T., Fischl, B., Andersson, J., Xu, J., Jbabdi, S., Webster, M., Polimeni, J., Van Essen, D., Jenkinson, M., 2013. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* 80, 105–24. doi:10.1016/j.neuroimage.2013.04.127

Glocker, B., Sotiras, A., Komodakis, N., Paragios, N., 2011. Deformable medical image registration: setting the state of the art with discrete methods. *Annual Review of Biomedical Engineering* 13, 219–44. doi:10.1146/annurev-bioeng-071910-124649

Goldenberg, R., Kimmel, R., Rivlin, E., Rudzsky, M., 2002. Cortex segmentation - a fast variational geometric approach. In: *Proceedings of IEEE Workshop on Variational and Level Set Methods in Computer Vision*. pp. 127–33. doi:10.1109/VLSM.2001.938891

Gousias, I. S., Edwards, a. D., Rutherford, M. a., Counsell, S. J., Hajnal, J. V., Rueckert, D., Hammers, A., 2012. Magnetic resonance imaging of the newborn brain: Manual segmentation of labelled atlases in term-born and preterm infants. *NeuroImage* 62 (3), 1499–509. doi:10.1016/j.neuroimage.2012.05.083

Gousias, I. S., Hammers, A., Counsell, S. J., Srinivasan, L., Rutherford, M. a., Heckemann, R. a., Hajnal, J. V., Rueckert, D., Edwards, a. D., 2013. Magnetic resonance imaging of

the newborn brain: automatic segmentation of brain images into 50 anatomical regions. *PLoS One* 8 (4), e59990. doi:10.1371/journal.pone.0059990

Gousias, I. S., Hammers, A., Heckemann, R. A., Counsell, S. J., Dyet, L. E., Boardman, J. P., Edwards, A. D., Rueckert, D., 2010. Atlas selection strategy for automatic segmentation of pediatric brain MRIs into 83 ROIs. In: *Proceedings of IEEE International Conference on Imaging Systems and Techniques (IST)*. pp. 290–93. doi:10.1109/IST.2010.5548493

Gousias, I. S., Rueckert, D., Heckemann, R. A., Dyet, L. E., Boardman, J. P., Edwards, A. D., Hammers, A., 2008. Automatic segmentation of brain MRIs of 2-year-olds into 83 regions of interest. *NeuroImage* 40 (2), 672–84. doi:10.1016/j.neuroimage.2007.11.034

Grabner, G., Janke, A. L., Budge, M. M., Smith, D., Pruessner, J., Collins, D. L., 2006. Symmetric atlasing and model based segmentation: An application to the hippocampus in older adults. In: Larsen, R., Nielsen, M., Sporring, J. (eds.), *Proceedings of 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 58–66. doi:10.1007/11866763_8

Grenander, U., Miller, M. I., 1998. Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics* 56 (4), 617–94.

Gu, S., Meng, X., Sciurba, F. C., Ma, H., Leader, J., Kaminski, N., Gur, D., Pu, J., 2014. Bidirectional elastic image registration using B-spline affine transformation. *Computerized Medical Imaging and Graphics* 38 (4), 306–14. doi:10.1016/j.compmedimag.2014.01.002

Gui, L., Lisowski, R., Faundez, T., Hüppi, P. S., Lazeyras, F., Kocher, M., 2011. Automatic segmentation of newborn brain MRI using mathematical morphology. In: *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 2026–30. doi:10.1109/ISBI.2011.5872810

Gui, L., Lisowski, R., Faundez, T., Hüppi, P. S., Lazeyras, F., Kocher, M., 2012a. Morphology-driven automatic segmentation of MR images of the neonatal brain. *Medical Image Analysis* 16 (8), 1565–79. doi:10.1016/j.media.2012.07.006

Gui, L., Lisowski, R., Faundez, T., Petra, S. H., 2012b. Morphology-based segmentation of newborn brain MR images. In: *MICCAI NeoBrainS12 Challenge*. p. 8.

Guimond, A., Meunier, J., Thirion, J., 1998. Automatic computation of average brain models. In: Wells, W. M., Colchester, A., Delp, S. (eds.), *Proceedings of 1st International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 631–40. doi:10.1007/BFb0056249

Guimond, A., Meunier, J., Thirion, J.-P., 2000. Average brain models: A convergence study. *Computer Vision and Image Understanding* 77 (2), 192–210. doi:10.1006/cviu.1999.0815

Ha, L. K., Krüger, J., Fletcher, P. T., Joshi, S., Silva, C. T., 2009. Fast parallel unbiased diffeomorphic atlas construction on multi-graphics processing units. In: *Proceedings of the 9th Eurographics Conference on Parallel Graphics and Visualization (EG PGV)*. pp. 41–8. doi:10.2312/EGPGV/EGPGV09/041-048

Habas, P. a., Kim, K., Corbett-Detig, J. M., Rousseau, F., Glenn, O. a., Barkovich, a. J., Studholme, C., 2010a. A spatiotemporal atlas of MR intensity, tissue probability and shape of the fetal brain with application to segmentation. *NeuroImage* 53 (2), 460–70. doi:10.1016/j.neuroimage.2010.06.054

Habas, P. A., Kim, K., Rousseau, F., Glenn, O. A., Barkovich, A. J., Studholme, C., 2010b. Atlas-based segmentation of developing tissues in the human brain with quantitative validation in young fetuses. *Human Brain Mapping* 31, 1348–58. doi:10.1002/hbm.20935

Habas, P. a., Scott, J. a., Roosta, A., Rajagopalan, V., Kim, K., Rousseau, F., Barkovich, a. J., Glenn, O. a., Studholme, C., 2012. Early folding patterns and asymmetries of the normal human brain detected from in utero MRI. *Cerebral Cortex* 22 (1), 13–25. doi:10.1093/cercor/bhr053

Han, X., 2013. Learning-boosted label fusion for multi-atlas auto-segmentation. In: Wu, G., Zhang, D., Shen, D., Yan, P., Suzuki, K., Wang, F. (eds.), *Proceedings of 4th International Workshop on Machine Learning in Medical Imaging, held in conjunction with MICCAI 2013*. Springer International Publishing, pp. 17–24. doi:10.1007/978-3-319-02267-3_3

Han, X., Pham, D. L., Tosun, D., Rettmann, M. E., Xu, C., Prince, J. L., 2004. CRUISE: Cortical reconstruction using implicit surface evolution. *NeuroImage* 23 (3), 997–1012. doi:10.1016/j.neuroimage.2004.06.043

Han, X., Xu, C., Braga-Neto, U., Prince, J. L., 2002. Topology correction in brain cortex segmentation using a multiscale, graph-based algorithm. *IEEE Transactions on Medical Imaging* 21 (2), 109–21. doi:10.1109/42.993130

Han, X., Xu, C., Prince, J. L., 2003. Topology preserving level set method for geometric deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (6), 755–68. doi:10.1109/TPAMI.2003.1201824

Han, X., Xu, C., Tosun, D., Prince, J. L., 2001. Cortical surface reconstruction using a topology preserving geometric deformable model. In: *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*. pp. 213–20. doi:10.1109/MMBIA.2001.991736

He, J., Christensen, G. E., 2003. Large deformation inverse consistent elastic image registration. In: Taylor, C. J., Noble, J. A. (eds.), *Proceedings of 18th International Conference on Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 438–49. doi:10.1007/978-3-540-45087-0_37

Heckemann, R. a., Hajnal, J. V., Aljabar, P., Rueckert, D., Hammers, A., 2006. Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33 (1), 115–26. doi:10.1016/j.neuroimage.2006.05.061

Heckemann, R. A., Ledig, C., Gray, K. R., Aljabar, P., Rueckert, D., Hajnal, J. V., Hammers, A., 2015. Brain extraction using label propagation and group agreement: Pincram. *PLoS One* 10 (7), 1–18. doi:10.1371/journal.pone.0129211

Henn, S., 2006. A full curvature based algorithm for image registration. *Journal of Mathematical Imaging and Vision* 24 (2), 195–208. doi:10.1007/s10851-005-3621-3

Hermosillo, G., Chefd'Hotel, C., Faugeras, O., 2001. A variational approach to multimodal image matching. In: *Proceedings IEEE Workshop on Variational and Level Set Methods in Computer Vision*. pp. 21–28. doi:10.1109/VLSM.2001.938877

Hermosillo, G., Faugeras, O., 2001. Dense image matching with global and local statistical criteria: a variational approach. In: *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. vol. 1. pp. I–73–I–78. doi:10.1109/CVPR.2001.990458

Hernandez, M., 2014. Gauss-Newton inspired preconditioned optimization in large deformation diffeomorphic metric mapping. *Physics in Medicine & Biology* 59 (20), 6085–115. doi:10.1088/0031-9155/59/20/6085

Hernandez, M., Bossa, M. N., Olmos, S., 2007. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. In: *Proceedings of 11th IEEE International Conference on Computer Vision (ICCV)*. pp. 1–8. doi:10.1109/ICCV.2007.4409126

Hernandez, M., Bossa, M. N., Olmos, S., 2009. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. *International Journal of Computer Vision* 85 (3), 291–306. doi:10.1007/s11263-009-0219-z

Hernandez, M., Olmos, S., Pennec, X., 2008. Comparing algorithms for diffeomorphic registration: Stationary LDDMM and diffeomorphic demons. In: Pennec, X. (ed.), *Proceedings of 2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy*. pp. 24–35.

Hestenes, M., Stiefel, E., 1952. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards* 49 (6), 409.

Hill, J., Dierker, D., Neil, J., Inder, T., Knutsen, A., Harwell, J., Coalson, T., Van Essen, D., 2010. A surface-based analysis of hemispheric asymmetries and folding of cerebral cortex in term-born human infants. *Journal of Neuroscience* 30 (6), 2268–76. doi:10.1523/JNEUROSCI.4682-09.2010

Holmes, C. J., Hoge, R., Collins, L., Woods, R., Toga, A. W., Evans, A. C., 1998. Enhancement of MR images using registration for signal averaging. *Journal of Computer Assisted Tomography* 22 (2), 324–33. doi:10.1097/00004728-199803000-00032

Horn, B. K., Schunck, B. G., 1981. Determining optical flow. *Artificial Intelligence* 17 (1-3), 185–203. doi:10.1016/0004-3702(81)90024-2

Hua-mei Chen, Varshney, P., 2003. Mutual information-based CT-MR brain image registration using generalized partial volume joint histogram estimation. *IEEE Transactions on Medical Imaging* 22 (9), 1111–19. doi:10.1109/TMI.2003.816949

Hughes, E. J., Winchman, T., Padormo, F., Teixeira, R., Wurie, J., Sharma, M., Fox, M., Hutter, J., Cordero-Grande, L., Price, A. N., Allsop, J., Bueno-Conde, J., Tusor, N., Arichi, T., Edwards, A. D., Rutherford, M. A., Counsell, S. J., Hajnal, J. V., 2017. A dedicated neonatal brain imaging system. *Magnetic Resonance in Medicine* 78 (2), 794–804. doi:10.1002/mrm.26462

Huo, Y., Plassard, A. J., Carass, A., Resnick, S. M., Pham, D. L., Prince, J. L., Landman, B. A., 2016. Consistent cortical reconstruction and multi-atlas brain segmentation. *NeuroImage* 138, 197–210. doi:10.1016/j.neuroimage.2016.05.030

Hüppi, P. S., Maier, S. E., Peled, S., Zientara, G. P., Barnes, P. D., Jolesz, F. A., Volpe, J. J., 1998. Microstructural development of human newborn cerebral white matter assessed in vivo by diffusion tensor magnetic resonance imaging. *Pediatric Research* 44 (4), 584–90. doi:10.1203/00006450-199810000-00019

Hüppi, P. S., Schuknecht, B., Boesch, C., Bossi, E., Felblinger, J., Fusch, C., Herschkowitz, N., 1996. Structural and neurobehavioral delay in postnatal brain development of preterm infants. *Pediatric Research* 39 (5), 895–901. doi:10.1203/00006450-199605000-00026

Iglesias, J. E., Sabuncu, M. R., 2015. Multi-atlas segmentation of biomedical images: A survey. *Medical Image Analysis* 24 (1), 205–19. doi:10.1016/j.media.2015.06.012

Im, K., Guimaraes, A., Kim, Y., Cottrill, E., Gagoski, B., Rollins, C., Ortinau, C., Yang, E., Grant, P., 2017. Quantitative folding pattern analysis of early primary sulci in human fetuses with brain abnormalities. *American Journal of Neuroradiology* 38 (7), 1449–55. doi:10.3174/ajnr.A5217

Inder, T. E., Warfield, S. K., Wang, H., Hüppi, P. S., Volpe, J. J., feb 2005. Abnormal cerebral structure is present at term in premature infants. *Pediatrics* 115 (2), 286 LP – 294.

Išgum, I., Benders, M. J. N. L., Avants, B., Cardoso, M. J., Counsell, S. J., Gomez, E. F., Gui, L., Huppi, P. S., Kersbergen, K. J., Makropoulos, A., Melbourne, A., Moeskops, P., Mol, C. P., Kuklisova-Murgasova, M., Rueckert, D., Schnabel, J. A., Srhoj-Egekher, V., Wu, J., Wang, S., de Vries, L. S., Viergever, M. A., 2015. Evaluation of automatic neonatal brain segmentation algorithms: The NeoBrainS12 challenge. *Medical Image Analysis* 20 (1), 135–51. doi:10.1016/j.media.2014.11.001

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., Smith, S. M., 2012. FSL. *NeuroImage* 62, 782–90. doi:10.1016/j.neuroimage.2011.09.015

Jiang, S., Xue, H., Glover, A., Rutherford, M. A., Rueckert, D., Hajnal, J. V., 2007. MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR): Application to fetal, neonatal, and adult brain studies. *IEEE Transactions on Medical Imaging* 26 (7), 967–80. doi:10.1109/TMI.2007.895456

Jones, S. E., Buchbinder, B. R., Aharon, I., 2000. Three-dimensional mapping of cortical thickness using laplace's equation. *Human Brain Mapping* 11 (1), 12–32.

Joshi, S., 1998. Large deformation diffeomorphisms and gaussian random fields for statistical characterization of brain sub-manifolds. Ph.D. thesis, Washington University, St. Louis, MO.

Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23 Suppl 1, S151–60. doi:10.1016/j.neuroimage.2004.07.068

Kamber, M., Shinghal, R., Collins, D. L., Francis, G. S., Evans, A. C., 1995. Model-based 3-D segmentation of multiple sclerosis lesions in magnetic resonance brain images. *IEEE Transactions on Medical Imaging* 14 (3), 442–53. doi:10.1109/42.414608

Kapellou, O., Counsell, S. J., Kennea, N., Dyet, L., Saeed, N., Stark, J., Maalouf, E., Duggan, P., Ajayi-Obe, M., Hajnal, J., Allsop, J. M., Boardman, J., Rutherford, M. A., Cowan, F., Edwards, A. D., 2006. Abnormal cortical development after premature birth shown by altered allometric scaling of brain growth 3 (8), 1382–1390. doi:10.1371/journal.pmed.0030265

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: Active contour models. *International Journal of Computer Vision* 1 (4), 321–31. doi:10.1007/BF00133570

Kazemi, K., Moghaddam, H. A., Grebe, R., Gondry-Jouet, C., Wallois, F., 2007. A neonatal atlas template for spatial normalization of whole-brain magnetic resonance images of newborns: Preliminary results. *NeuroImage* 37 (2), 463–73. doi:10.1016/j.neuroimage.2007.05.004

Keszei, A. P., Berkels, B., Deserno, T. M., 2017. Survey of non-rigid registration tools in medicine. *Journal of Digital Imaging* 30 (1), 102–16. doi:10.1007/s10278-016-9915-8

Kim, H., Lepage, C., Maheshwary, R., Jeon, S., Evans, A. C., Hess, C. P., Barkovich, A. J., Xu, D., 2016. NEOCIVET: Towards accurate morphometry of neonatal gyrification and clinical applications in preterm newborns. *NeuroImage* 138, 28–42. doi:10.1016/j.neuroimage.2016.05.034

Kim, J. S., Singh, V., Jun, K. L., Lerch, J., Ad-Dab'bagh, Y., MacDonald, D., Jong, M. L., Kim, S. I., Evans, A. C., 2005. Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a laplacian map and partial volume effect classification. *NeuroImage* 27 (1), 210–21. doi:10.1016/j.neuroimage.2005.03.036

Kim, K., Habas, P. a., Rousseau, F., Glenn, O. a., Barkovich, A. J., Studholme, C., 2010. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. *IEEE Transactions on Medical Imaging* 29 (1), 146–58. doi:10.1109/TMI.2009.2030679

Kim, M.-J., Kim, M.-H., Shen, D. S., 2008. Learning-based deformation estimation for fast non-rigid registration. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. pp. 1–6. doi:10.1109/CVPRW.2008.4563006

Klein, A., Andersson, J., Ardekani, B. a., Ashburner, J., Avants, B., Chiang, M.-C., Christensen, G. E., Collins, D. L., Gee, J., Hellier, P., Song, J. H., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R. P., Mann, J. J., Parsey, R. V., 2009a. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *NeuroImage* 46 (3), 786–802. doi:10.1016/j.neuroimage.2008.12.037

Klein, S., Pluim, J. P., Staring, M., Viergever, M. A., 2009b. Adaptive stochastic gradient descent optimisation for image registration. *International Journal of Computer Vision* 81 (3), 227–39. doi:10.1007/s11263-008-0168-y

Klein, S., Staring, M., Murphy, K., Viergever, M. A., Pluim, J. P. W., 2010. elastix: a toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging* 29 (1), 196–205. doi:10.1109/TMI.2009.2035616

Klein, S., Staring, M., Pluim, J. P. W., 2007. Evaluation of optimization methods for non-rigid medical image registration using mutual information and B-splines. *IEEE Transactions on Image Processing* 16 (12), 2879–90. doi:10.1109/TIP.2007.909412

Kochunov, P., Lancaster, J. L., Thompson, P., Woods, R., Mazziotta, J., Hardies, J., Fox, P., 2001. Regional spatial normalization: Toward an optimal target. *Journal of Computer Assisted Tomography* 25 (5), 805–16.

Kohlrausch, J., Rohr, K., Stiehl, H. S., 2005. A new class of elastic body splines for nonrigid registration of medical images. *Journal of Mathematical Imaging and Vision* 23 (3), 253–80. doi:10.1007/s10851-005-0483-7

Kuklisova-Murgasova, M., Aljabar, P., Srinivasan, L., Counsell, S. J., Doria, V., Serag, A., Gousias, I. S., Boardman, J. P., Rutherford, M. A., Edwards, A. D., Hajnal, J. V., Rueckert, D., 2011. A dynamic 4D probabilistic atlas of the developing brain. *NeuroImage* 54 (4), 2750–63. doi:10.1016/j.neuroimage.2010.10.019

Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M. A., Hajnal, J. V., Schnabel, J. A., 2012. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical Image Analysis* 16 (8), 1550–64. doi:10.1016/j.media.2012.07.004

Kwon, S. H., Vasung, L., Ment, L. R., Huppi, P. S., 2014. The role of neuroimaging in predicting neurodevelopmental outcomes of preterm neonates. *Clinics in Perinatology* 41 (1), 257–283. doi:10.1016/j.clp.2013.10.003

Kybic, J., Unser, M., 2000. Multidimensional elastic registration of images using splines. In: *Proceedings of International Conference on Image Processing*. vol. 2. pp. 455–8. doi:10.1109/ICIP.2000.899447

Kybic, J., Unser, M., 2003. Fast parametric elastic image registration. *IEEE Transactions on Image Processing* 12 (11), 1427–42. doi:10.1109/TIP.2003.813139

Lachaud, J., Montanvert, A., 1999. Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction. *Medical Image Analysis* 3 (2), 187–207. doi:10.1016/S1361-8415(99)80012-7

Ledig, C., Heckemann, R. A., Hammers, A., Lopez, J. C., Newcombe, V. F. J., Makropoulos, A., Lötjönen, J., Menon, D. K., Rueckert, D., 2015. Robust whole-brain segmen-

tation: Application to traumatic brain injury. *Medical Image Analysis* 21 (1), 40–58. doi:10.1016/j.media.2014.12.003

Ledig, C., Wolz, R., Aljabar, P., Ltjnen, J., Heckemann, R. A., Hammers, A., Rueckert, D., 2012a. Multi-class brain segmentation using atlas propagation and EM-based refinement. In: *Proceedings of 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 896–899. doi:10.1109/ISBI.2012.6235693

Ledig, C., Wright, R., Serag, A., Aljabar, P., Rueckert, D., 2012b. Neonatal brain segmentation using second order neighborhood information. In: *MICCAI Workshop on Perinatal and Paediatric Imaging: PaPI*. pp. 33–40.

Lee, S., Wolberg, G., Shin, S., 1997. Scattered data interpolation with multilevel B-splines. *IEEE Transactions on Visualization and Computer Graphics* 3 (3), 228–44. doi:10.1109/2945.620490

Leow, A., Huang, S.-C., Geng, A., Becker, J., Davis, S., Toga, A., Thompson, P., 2005. Inverse consistent mapping in 3D deformable image registration: its construction and statistical properties. *Information Processing in Medical Imaging* 19, 493–503.

Leroy, F., Mangin, J. F., Rousseau, F., Glasel, H., Hertz-Pannier, L., Dubois, J., Dehaene-Lambertz, G., 2011. Atlas-free surface reconstruction of the cortical grey-white interface in infants. *PLoS One* 6 (11), 1–11. doi:10.1371/journal.pone.0027128

Lewis, J. P., 1995. Fast template matching. In: *Vision Interface*. Canadian Image Processing and Pattern Recognition Society, pp. 120–123.

Li, G., Wang, L., Shi, F., Lin, W., Shen, D., 2014. Constructing 4D infant cortical surface atlases based on dynamic developmental trajectories of the cortex. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.), *Proceedings of 17th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer International Publishing, pp. 89–96. doi:10.1007/978-3-319-10443-0_12

Liao, S., Jia, H., Wu, G., Shen, D., 2012. A novel framework for longitudinal atlas construction with groupwise registration of subject image sequences. *NeuroImage* 59 (2), 1275–89. doi:10.1016/j.neuroimage.2011.07.095

Little, J., Hill, D., Hawkes, D., 1997. Deformations incorporating rigid structures. *Computer Vision and Image Understanding* 66 (2), 223–32. doi:10.1006/cviu.1997.0608

Liu, M., Kitsch, A., Miller, S., Chau, V., Poskitt, K., Rousseau, F., Shaw, D., Studholme, C., 2016. Patch-based augmentation of expectation-maximization for brain MRI tissue segmentation at arbitrary age after premature birth. *NeuroImage* 127, 387–408. doi:10.1016/j.neuroimage.2015.12.009

Liu, M., Seshamani, S., Harrylock, L., Kitsch, A., Miller, S., Chau, V., Poskitt, K., Rousseau, F., Studholme, C., 2014. Spatially adapted augmentation of age-specific atlas-based segmentation using patch-based priors. In: *Proceedings of Society of Photographic Instrumentation Engineers (SPIE), Medical Imaging: Image Processing*. vol. 9034. doi:10.1117/12.2042426

Loeckx, D., Maes, F., Vandermeulen, D., Suetens, P., 2003. Temporal subtraction of thorax CR images using a statistical deformation model. *IEEE Transactions on Medical Imaging* 22 (11), 1490–1504. doi:10.1109/TMI.2003.819291

Loeckx, D., Maes, F., Vandermeulen, D., Suetens, P., 2004. Nonrigid image registration using free-from deformations with a local rigidity constraint. In: Barillot, C., Haynor, D. R., Hellier, P. (eds.), *Proceedings of 7th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 639–46. doi:10.1007/978-3-540-30135-6_78

Loeckx, D., Maes, F., Vandermeulen, D., Suetens, P., 2006. Comparison between parzen window interpolation and generalised partial volume estimation for nonrigid image registration using mutual information. In: Pluim, J. P. W., Likar, B., Gerritsen, F. A. (eds.), *Proceedings of 3rd International Workshop on Biomedical Image Registration*. pp. 206–13. doi:10.1007/11784012_25

Lombaert, H., Arcaro, M., Ayache, N., 2015. Brain Transfer: Spectral analysis of cortical surfaces and functional maps. In: Ourselin, S., Alexander, D., Westin, C.-F., Cardoso, M. (eds.), *Proceedings of 24th International Conference on Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 474–87. doi:10.1007/978-3-319-19992-4_37

Lombaert, H., Grady, L., Pennec, X., Ayache, N., 2013a. Spectral Log-Demons diffeomorphic image registration with very large deformations. *International Journal of Computer Vision* 107 (3), 25471. doi:10.1007/s11263-013-0681-5

Lombaert, H., Sporring, J., Siddiqi, K., 2013b. Diffeomorphic spectral matching of cortical surfaces. In: Gee, J. C., Joshi, S., Pohl, K. M., Wells, W. M., Zöllei, L. (eds.), *Proceedings of 23rd International Conference on Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 376–89. doi:10.1007/978-3-642-38868-2_32

Lorensen, W. E., Cline, H. E., 1987. Marching cubes: A high resolution 3D surface construction algorithm. In: *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 163–69. doi:10.1145/37401.37422

Lorenzen, P., Davis, B., Joshi, S., 2005. Unbiased atlas formation via large deformations metric mapping. In: Duncan, J. S., Gerig, G. (eds.), *Proceedings of 8th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. vol. 8. Springer Berlin Heidelberg, pp. 411–8. doi:10.1007/11566489_51

Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., Joshi, S., 2006. Multi-modal image set registration and atlas formation. *Medical Image Analysis* 10 (3), 440–51. doi:10.1016/j.media.2005.03.002

Lorenzi, M., Ayache, N., Frisoni, G. B., Pennec, X., 2013. LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage* 81, 470–83. doi:10.1016/j.neuroimage.2013.04.114

Lucas, B. C., Bogovic, J. A., Carass, A., Bazin, P. L., Prince, J. L., Pham, D. L., Landman, B. A., 2010. The Java Image Science Toolkit (JIST) for rapid pro-

totyping and publishing of neuroimaging software. *Neuroinformatics* 8 (1), 5–17. doi:10.1007/s12021-009-9061-2

MacDonald, D., Avis, D., Evans, A. C., 1994. Multiple surface identification and matching in magnetic resonance images. In: *Proceedings of Visualization in Biomedical Computing*. vol. 2359. Rochester, MN, USA, pp. 160–169. doi:10.1117/12.185176

MacDonald, D., Avis, D., Evans, A. C., 1998. Proximity constraints in deformable models for cortical surface identification. In: Wells, W. M., Colchester, A., Delp, S. (eds.), *Proceedings of 1st International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 650–59. doi:10.1007/BFb0056251

MacDonald, D., Kabani, N., Avis, D., Evans, a. C., 2000. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. *NeuroImage* 12 (3), 340–56. doi:10.1006/nimg.1999.0534

Maes, F., Collignon, a., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging* 16 (2), 187–98. doi:10.1109/42.563664

Makropoulos, A., Aljabar, P., Wright, R., Hüning, B., Merchant, N., Arichi, T., Tusor, N., Hajnal, J. V., Edwards, A. D., Counsell, S. J., Rueckert, D., 2016. Regional growth and atlasing of the developing human brain. *NeuroImage* 125, 456–78. doi:10.1016/j.neuroimage.2015.10.047

Makropoulos, A., Counsell, S. J., Rueckert, D., 2017a. A review on automatic fetal and neonatal brain MRI segmentation. *NeuroImage*. doi:10.1016/j.neuroimage.2017.06.074

Makropoulos, A., Gousias, I., Ledig, C., Aljabar, P., Serag, A., Hajnal, J., Edwards, A., Counsell, S., Rueckert, D., 2014. Automatic whole brain MRI segmentation of the developing neonatal brain. *IEEE Transactions on Medical Imaging* 33 (9), 1818–31. doi:10.1109/TMI.2014.2322280

Makropoulos, A., Ledig, C., Aljabar, P., Serag, A., Hajnal, J. V., Edwards, A. D., Counsell, S. J., Rueckert, D., 2012. Automatic tissue and structural segmentation of neonatal brain MRI using expectation-maximization. In: *Proceedings of MICCAI Grand Challenge on Neonatal Brain Segmentation (NeoBrainS12)*. vol. 9. p. 2012.

Makropoulos, A., Robinson, E. C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S. J., Steinweg, J., Passerat-Palmbach, J., Lenz, G., Mortari, F., Tenev, T., Duff, E. P., Bastiani, M., Cordero-Grande, L., Hughes, E., Tusor, N., Tournier, J.-D., Hutter, J., Price, A. N., Murgasova, M., Kelly, C., Rutherford, M. A., Smith, S. M., Edwards, D. A., Hajnal, J. V., Jenkinson, M., Rueckert, D., 2017b. The Developing Human Connectome Project: a minimal processing pipeline for neonatal cortical surface reconstruction. *bioRxiv*, 1–43. doi:10.1101/125526

Makropoulos, A., Robinson, E. C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S. J., Steinweg, J., Passerat-Palmbach, J., Lenz, G., Mortari, F., Tenev, T., Duff, E. P., Bastiani, M., Cordero-Grande, L., Hughes, E., Tusor, N., Tournier, J.-D., Hutter, J., Price, A. N., Murgasova, M., Kelly, C., Rutherford, M. A., Smith, S. M., Edwards, D. A., Hajnal, J. V., Jenkinson, M., Rueckert, D., 2018. The Developing Human Connectome Project: a minimal processing pipeline for neonatal cortical surface reconstruction. *NeuroImage*. doi:10.1016/j.neuroimage.2018.01.054

Malamateniou, C., Malik, S. J., Counsell, S. J., Allsop, J. M., McGuinness, A. K., Hayat, T., Broadhouse, K., Nunes, R. G., Ederies, A. M., Hajnal, J. V., Rutherford, M. A., 2013. Motion-compensation techniques in neonatal and fetal MR imaging. *American Journal of Neuroradiology* 34 (6), 1124–36. doi:10.3174/ajnr.A3128

Mang, A., Ruthotto, L., 2017. A Lagrangian Gauss-Newton-Krylov solver for mass- and intensity-preserving diffeomorphic image registration. *SIAM Journal on Scientific Computing* 39 (5), B860–B885. doi:10.1137/17M1114132

Mangin, J.-F., Frouin, V., Bloch, I., Régis, J., López-Krahe, J., 1995. From 3d magnetic resonance images to structural representations of the cortex topography using topology

preserving deformations. *Journal of Mathematical Imaging and Vision* 5 (4), 297–318. doi:10.1007/BF01250286

Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N., 2010. LogDemons revisited: Consistent regularisation and incompressibility constraint for soft tissue tracking in medical images. In: *Proceedings of 13th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 652–9. doi:10.1007/978-3-642-15745-5_80

Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., Buckner, R. L., 2007. Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience* 19 (9), 1498–1507. doi:10.1162/jocn.2007.19.9.1498

Marstal, K., Berendsen, F., Staring, M., Klein, S., 2016. SimpleElastix: A user-friendly , multi-lingual library for medical image registration. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. pp. 574–582. doi:10.1109/CVPRW.2016.78

Maurer, C. R., Qi, R., Raghavan, V., 2003. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (2), 265–70. doi:10.1109/TPAMI.2003.1177156

Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Le Goualher, G., Boomsma, D., Cannon, T., Kawashima, R., Mazoyer, B., 2001. A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (icbm). *Philosophical Transactions of the Royal Society of London: B. Biological Sciences* 356 (1412), 1293–322. doi:10.1098/rstb.2001.0915

Mazziotta, J. C., Toga, A. W., Evans, A., Fox, P., Lancaster, J., 1995. A probabilistic atlas of the human brain: theory and rationale for its development. the international consortium for brain mapping (ICBM) 2 (2), 89–101. doi:10.1006/nimg.1995.1012

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. *Medical Image Analysis* 1 (2), 91–108. doi:10.1016/S1361-8415(96)80007-7

Meagher, D., 1982. Geometric modeling using octree encoding. *Computer Graphics and Image Processing* 19 (2), 129–47. doi:10.1016/0146-664X(82)90104-6

Melbourne, A., Cardoso, M. J., Kendall, G. S., Nicola, J., Marlow, N., Ourselin, S., 2012. Adaptive neonatal MRI brain segmentation with myelinated white matter class and automated extraction of ventricles I-IV. In: *MICCAI NeoBrainS12 Challenge*. pp. 1–6.

Merlo Pich, E., Jeromin, A., Frisoni, G. B., Hill, D., Lockhart, A., Schmidt, M. E., Turner, M. R., Mondello, S., Potter, W. Z., 2014. Imaging as a biomarker in drug discovery for Alzheimer's disease: is MRI a suitable technology? *Alzheimers Research & Therapy* 6 (4), 51. doi:10.1186/alzrt276

Miller, J. V., Breen, D. E., Lorensen, W. E., O'Bara, R. M., Wozny, M. J., 1991. Geometrically deformed models: a method for extracting closed geometric models form volume data. In: *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 217–26. doi:10.1145/127719.122742

Miller, M. I., Trouve, A., Younes, L., 2002. On the metrics and euler-lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering* 4, 375–405. doi:10.1146/annurev.bioeng.4.092101.125733

Miller, M. I., Trouvé, A., Younes, L., 2006. Geodesic shooting for computational anatomy 24 (2), 209–28. doi:10.1007/s10851-005-3624-0

Modat, M., Cardoso, M. J., Daga, P., Cash, D., 2012a. Inverse-consistent symmetric free form deformation. In: Dawant, B. M., Christensen, G. E., Fitzpatrick, J. M., Rueckert, D. (eds.), *Proceedings of 5th International Workshop on Biomedical Image Registration*. pp. 79–88. doi:10.1007/978-3-642-31340-0_9

Modat, M., Daga, P., Cardoso, M. J., Ourselin, S., Ridgway, G. R., Ashburner, J., 2012b. Parametric non-rigid registration using a stationary velocity field. In: *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*. pp. 145–150. doi:10.1109/MMBIA.2012.6164745

Modat, M., Ridgway, G. R., Taylor, Z. A., Hawkes, D. J., Fox, N. C., Ourselin, S., 2009. A parallel-friendly normalised mutual information gradient for free-form registration. In: *Proceedings of the Society of Photographic Instrumentation Engineers (SPIE), Medical Imaging: Image Processing*. vol. 7259. doi:10.1117/12.811588

Modat, M., Ridgway, G. R., Taylor, Z. A., Lehmann, M., Barnes, J., Hawkes, D. J., Fox, N. C., Ourselin, S., 2010. Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* 98 (3), 278–84. doi:10.1016/j.cmpb.2009.09.002

Moeskops, P., Benders, M. J. N. L., Pearlman, P. C., Kersbergen, K. J., Leemans, A., Viergever, M. A., Išgum, I., 2013. Assessment of quantitative cortical biomarkers in the developing brain of preterm infants. In: *Proceedings of Society of Photographic Instrumentation Engineers (SPIE) 8670, Medical Imaging: Computer-Aided Diagnosis*. vol. 8670. pp. 867011–867011–7. doi:10.1117/12.2007829

Moeskops, P., Viergever, M. A., Mendrik, A. M., Vries, L. S. D., Manon, J. N. L., 2016. Automatic segmentation of mr brain images with a convolutional neural network. *IEEE Transactions on Medical Imaging* 35 (5), 1252–61. doi:10.1109/TMI.2016.2548501

Möller, T., 1997. A fast triangle-triangle intersection test. *Journal of Graphics Tools* 2 (2), 25–30. doi:10.1080/10867651.1997.10487472

Molnár, Z., Rutherford, M., 2013. Brain maturation after preterm birth. *Science Translational Medicine* 5 (168), 168ps2. doi:10.1126/scitranslmed.3005379

Nielsen, M., Johansen, P., Jackson, A. D., Lautrup, B., 2002. Brownian warps: A least committed prior for non-rigid registration. In: Dohi, T., Kikinis, R. (eds.), *Proceedings*

of 5th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 557–64. doi:10.1007/3-540-45787-9_70

Niethammer, M., Huang, Y., Vialard, F.-X., 2011. Geodesic regression for image time-series. In: Fichtinger, G., Martel, A., Peters, T. (eds.), *Proceedings of 14th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 655–62. doi:10.1007/978-3-642-23629-7_80

Noblet, V., Heinrich, C., Heitz, F., Armspach, J. P., 2008. Symmetric nonrigid image registration: Application to average brain templates construction. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.), *Proceedings of 11th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 897–904. doi:10.1007/978-3-540-85990-1_108

Noorizadeh, N., Kazemi, K., Grebe, R., Helfroush, M. S., Mahmoudzadeh, M., Golshaeyan, N., Wallois, F., 2013. Symmetric brain atlas template for newborns brain asymmetry studies, 6–9. doi:10.1109/IranianCEE.2013.6599703

Nyul, L., Udupa, J., Xuan Zhang, 2000. New variants of a method of MRI scale standardization. *IEEE Transactions on Medical Imaging* 19 (2), 143–150. doi:10.1109/42.836373

Oishi, K., Mori, S., Donohue, P. K., Ernst, T., Anderson, L., Buchthal, S., Faria, A., Jiang, H., Li, X., Miller, M. I., van Zijl, P. C. M., Chang, L., 2011. Multi-contrast human neonatal brain atlas: application to normal neonate development analysis. *NeuroImage* 56 (1), 8–20. doi:10.1016/j.neuroimage.2011.01.051

Onofrey, J. A., Staib, L. H., Papademetris, X., 2013. Fast nonrigid image registration using statistical deformation models learned from richly-annotated data. In: *Proceedings of 10th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 580–83. doi:10.1109/ISBI.2013.6556541

Osher, S., Fedkiw, R., 2003. Level Set Methods and Dynamic Implicit Surfaces. Applied Mathematical Sciences. Springer-Verlag, New York, NY, USA.

Osher, S., Sethian, J. A., 1988. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics* 79 (1), 12–49. doi:10.1016/0021-9991(88)90002-2

Ou, Y., Akbari, H., Bilello, M., Da, X., Davatzikos, C., 2014. Comparative evaluation of registration algorithms in different brain databases with varying difficulty: Results and insights. *IEEE Transactions on Medical Imaging* 33 (10), 2039–65. doi:10.1109/TMI.2014.2330355

Ou, Y., Sotiras, A., Paragios, N., Davatzikos, C., 2011. DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting. *Medical Image Analysis* 15 (4), 622–39. doi:10.1016/j.media.2010.07.002

Özkan, M., Dawant, B. M., Maciunas, R. J., 1993. Neural-network-based segmentation of multi-modal medical images: A comparative and prospective study. *IEEE Transactions on Medical Imaging* 12 (3), 534–44. doi:10.1109/42.241881

Pai, A., Sommer, S., Raket, L. L., Kühnel, L., Darkner, S., Sørensen, L., Nielsen, M., 2017. A statistical model for simultaneous template estimation, bias correction, and registration of 3D brain images. In: Müller, H., Kelm, B. M., Arbel, T., Cai, W., Cardoso, M. J., Langs, G., Menze, B., Metaxas, D., Montillo, A., Wells III, W. M., Zhang, S., Chung, A. C., Jenkinson, M., Ribbens, A. (eds.), *Proceedings of MICCAI 2016 International Workshops on Medical Computer Vision (MCV) and Bayesian and Graphical Models for Biomedical Imaging (BAMBI)*. Springer International Publishing, pp. 151–9. doi:10.1007/978-3-319-61188-4_14

Pai, A., Sommer, S., Sorensen, L., Darkner, S., Sporring, J., Nielsen, M., 2015. Kernel bundle diffeomorphic image registration using stationary velocity fields and wendland basis functions. *IEEE Transactions on Medical Imaging* 53 (6), 1369–80. doi:10.1109/TMI.2015.2511062

Parisot, S., Arslan, S., Passerat-Palmbach, J., III, W. W., Rueckert, D., 2016. Group-wise parcellation of the cortex through multi-scale spectral clustering. *NeuroImage* 136, 68–83. doi:10.1016/j.neuroimage.2016.05.035

Park, J. Y., McInerney, T., Terzopoulos, D., Kim, M. H., 2001. A non-self-intersecting adaptive deformable surface for complex boundary extraction from volumetric images. *Computers & Graphics* 25 (3), 421–40. doi:10.1016/S0097-8493(01)00066-8

Pennec, X., Arsigny, V., 2013. Exponential barycenters of the canonical cartan connection and invariant means on Lie groups. In: Nielsen, F., Bhatia, R. (eds.), *Matrix Information Geometry*. Springer Berlin Heidelberg, pp. 123–66. doi:10.1007/978-3-642-30232-9_7

Pennec, X., Cachier, P., Ayache, N., 1999. Understanding the "Demon's Algorithm": 3D non-rigid registration by gradient descent. In: Taylor, C., Colchester, A. (eds.), *Proceedings of 2nd International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. vol. 1679. Springer Berlin Heidelberg, pp. 597–605. doi:10.1007/10704282_64

Pham, D. L., Prince, J. L., 1999. An adaptive fuzzy C-means algorithm for image segmentation in the presence of intensity inhomogeneities. *Pattern Recognition Letters* 20 (1), 57–68. doi:10.1016/S0167-8655(98)00121-4

Pluim, J. P. W., Maintz, J. B. A., Viergever, M. A., 2003. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging* 22 (8), 986–1004. doi:10.1109/TMI.2003.815867

Pohl, K. M., Fisher, J., Grimson, W. E. L., Kikinis, R., Wells, W. M., 2006. A Bayesian model for joint segmentation and registration. *NeuroImage* 31 (1), 228–39. doi:10.1016/j.neuroimage.2005.11.044

Postelnicu, G., Zollei, L., Fischl, B., 2009. Combined volumetric and surface registration. *IEEE Transactions on Medical Imaging* 28 (4), 508–22. doi:10.1109/TMI.2008.2004426

Prastawa, M., Gilmore, J. H., Lin, W., Gerig, G., 2005. Automatic segmentation of MR images of the developing newborn brain. *Medical Image Analysis* 9 (5), 457–66. doi:10.1016/j.media.2005.05.007

Press, W. H., Teukolsky, S. A., Vetterling, W. T., Flannery, B. P., 2007. Numerical Recipes 3rd Edition: The Art of Scientific Computing, 3rd edition. Cambridge University Press, New York, NY, USA.

Pszczolkowski, S., Pizarro, L., Guerrero, R., Rueckert, D., 2012. Nonrigid free-form registration using landmark-based statistical deformation models. In: Haynor, D. R., Ourselin, S. (eds.), *Proceedings of Society of Photographic Instrumentation Engineers (SPIE), Medical Imaging: Image Processing*. vol. 831418. doi:10.1117/12.911441

Pudil, P., Novovičová, J., Kittler, J., 1994. Floating search methods in feature selection. *Pattern Recognition Letters* 15 (11), 1119–25. doi:10.1016/0167-8655(94)90127-9

Risser, L., Vialard, F.-X., Wolz, R., Murgasova, M., Holm, D. D., Rueckert, D., 2011. Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping. *IEEE Transactions on Medical Imaging* 30 (10), 1746–59. doi:10.1109/TMI.2011.2146787

Rivière, D., Geffroy, D., Denghien, I., Souedet, N., Cointepas, Y., 2009. BrainVISA: an extensible software environment for sharing multimodal neuroimaging data and processing tools. *NeuroImage* 47, S163. doi:10.1016/S1053-8119(09)71720-3

Robbins, S., Evans, A. C., Collins, D. L., Whitesides, S., 2004. Tuning and comparing spatial normalization methods. *Medical Image Analysis* 8 (3), 311–23. doi:10.1016/j.media.2004.06.009

Robinson, E., Jbabdi, S., Glasser, M., Andersson, J., Burgess, G., Harms, M., Smith, S., Van Essen, D., Jenkinson, M., 2014. MSM: A new flexible framework for Multimodal Surface Matching. *NeuroImage* 100, 414–26. doi:10.1016/j.neuroimage.2014.05.069

Roche, A., Malandain, G., Pennec, X., Ayache, N., 1998. The correlation ratio as a new similarity measure for multimodal image registration. In: Wells, W. M., Colchester, A., Delp, S. (eds.), *Proceedings of 1st International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 1115–24. doi:10.1007/BFb0056301

Rohlfing, T., Brandt, R., Menzel, R., Maurer, C. R., 2004. Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains. *NeuroImage* 21 (4), 1428–42. doi:10.1016/j.neuroimage.2003.11.010

Rohlfing, T., Maurer, C. R., Bluemke, D. A., Jacobs, M. A., 2003. Volume-preserving nonrigid registration of mr breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging* 22 (6), 730–41. doi:10.1109/TMI.2003.814791

Rousseeuw, P. J., Van Driessen, K., 1999. A fast algorithm for the minimum covariance determinant estimator. *Technometrics* 41 (3), 212–23. doi:10.1080/00401706.1999.10485670

Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., Hammers, A., 2006. Diffeomorphic registration using B-splines. In: Larsen, R., Nielsen, M., Sporring, J. (eds.), *Proceedings of 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 702–9. doi:10.1007/11866763_86

Rueckert, D., Frangi, A. F., Schnabel, J. A., 2001. Automatic construction of 3D statistical deformation model of the brain using nonrigid registration. In: Niessen, W. J., Viergever, M. A. (eds.), *Proceedings of 4th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 77–84. doi:10.1007/3-540-45468-3_10

Rueckert, D., Frangi, A. F., Schnabel, J. a., 2003. Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. *IEEE Transactions on Medical Imaging* 22 (8), 1014–25. doi:10.1109/TMI.2003.815865

Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., Hawkes, D. J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging* 18 (8), 712–21. doi:10.1109/42.796284

Rutherford, M. A. (ed.), 2002. MRI of the neonatal brain, 4th edition. W.B. Saunders. url: http://www.mrineonatalbrain.com/

Sandor, S., Leahy, R., 1997. Surface-based labeling of cortical anatomy using a deformable atlas. *IEEE Transactions on Medical Imaging* 16 (1), 41–54. doi:10.1109/42.552054

Sanroma, G., Benkarim, O. M., Piella, G., Ballester, M. Á. G., 2016. Building an ensemble of complementary segmentation methods by exploiting probabilistic estimates. In: Wang, L., Adeli, E., Wang, Q., Shi, Y., Suk, H.-I. (eds.), *Proceedings of 7th International Workshop on Machine Learning in Medical Imaging (MLMI)*. Cham, pp. 27–35. doi:10.1007/978-3-319-47157-0_4

Schnabel, J. A., Rueckert, D., Quist, M., Blackall, J. M., Castellano-smith, A. D., Hartkens, T., Penney, G. P., Hall, W. A., Liu, H., Truwit, C. L., Gerritsen, F. A., Hill, D. L. G., Hawkes, D. J., 2001. A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In: Niessen, W. J., Viergever, M. A. (eds.), *Proceedings of 4th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 573–81. doi:10.1007/3-540-45468-3_69

Schuh, A., Makropoulos, A., Wright, R., Robinson, E. C., Tusor, N., Steinweg, J., Hughes, E., Grande, L. C., Price, A., Hutter, J., Hajnal, J. V., Rueckert, D., April 2017. A deformable model for the reconstruction of the neonatal cortex. In: *IEEE 14th International Symposium on Biomedical Imaging (ISBI)*. pp. 800–803. doi:10.1109/ISBI.2017.7950639

Schuh, A., Murgasova, M., Makropoulos, A., Ledig, C., Counsell, S. J., Hajnal, J. V., Aljabar, P., Rueckert, D., 2014. Construction of a 4D brain atlas and growth model using diffeomorphic registration. In: *STIA*. Cambridge, MA.

Sederberg, T. W., Parry, S. R., 1986. Free-form deformation of solid geometric models. In: *SIGGRAPH Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 151–60. doi:10.1145/15922.15903

Seghers, D., Agostino, E. D., Maes, F., Vandermeulen, D., 2004. Construction of a brain template from MR images using state-of-the-art registration and segmentation techniques. In: Barillot, C., Haynor, D. R., Hellier, P. (eds.), *Proceedings of 7th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 696–703. doi:10.1007/978-3-540-30135-6_85

Segonne, F., Pacheco, J., Fischl, B., 2007. Geometrically accurate topology-correction of cortical surfaces using nonseparating loops. *IEEE Transactions on Medical Imaging* 26 (4), 518–29. doi:10.1109/TMI.2006.887364

Serag, A., 2012. Spatio-temporal modeling and analysis of brain development. Ph.D. thesis, Imperial College London.

Serag, A., Aljabar, P., Ball, G., Counsell, S. J., Boardman, J. P., Rutherford, M. A., Edwards, A. D., Hajnal, J. V., Rueckert, D., 2012a. Construction of a consistent high-definition spatio-temporal atlas of the developing brain using adaptive kernel regression. *NeuroImage* 59 (3), 2255–65. doi:10.1016/j.neuroimage.2011.09.062

Serag, A., Aljabar, P., Counsell, S., Boardman, J., Hajnal, J. V., Rueckert, D., 2012b. LISA: Longitudinal image registration via spatio-temporal atlases. In: *Proceedings of 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 334–7. doi:10.1109/ISBI.2012.6235552

Serag, A., Kyriakopoulou, V., Rutherford, M. A., Edwards, A. D., Hajnal, J. V., 2012c. A multi-channel 4D probabilistic atlas of the developing brain: Application to fetuses and neonates. *Annals of the British Machine Vision Association (BMVA)* 2012 (3), 1–14.

Serra, J., 1983. Image Analysis and Mathematical Morphology. Academic Press, Inc., Orlando, FL, USA.

Sethian, J. A., 1999. Level set methods and fast marching methods, 2nd edition. Cambridge University Press.

Shattuck, D., Leahy, R., 2002. BrainSuite: An automated cortical surface identification tool. *Medical Image Analysis* 6 (2), 129–42. doi:10.1016/S1361-8415(02)00054-3

Shattuck, D. W., Leahy, R. M., 2001. Automated graph-based analysis and correction of cortical volume topology. *IEEE Transactions on Medical Imaging* 20 (11), 1167–77. doi:10.1109/42.963819

Shattuck, D. W., Sandor-Leahy, S. R., Schaper, K. A., Rottenberg, D. A., Leahy, R. M., 2001. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage* 13 (5), 856–76. doi:10.1006/nimg.2000.0730

Shen, D., Davatzikos, C., 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging* 21 (11), 1421–39. doi:10.1109/TMI.2002.803111

Shi, F., Fan, Y., Tang, S., Gilmore, J. H., Lin, W., Shen, D., 2010. Neonatal brain image segmentation in longitudinal MRI studies. *NeuroImage* 49 (1), 391–400. doi:10.1016/j.neuroimage.2009.07.066

Shi, F., Shen, D., Yap, P.-T., Fan, Y., Cheng, J.-Z., An, H., Wald, L. L., Gerig, G., Gilmore, J. H., Lin, W., 2011a. CENTS: Cortical enhanced neonatal tissue segmentation. *Human Brain Mapping* 32 (3), 382–96. doi:10.1002/hbm.21023

Shi, F., Yap, P.-t., Wu, G., Jia, H., Gilmore, J. H., Lin, W., Shen, D., 2011b. Infant brain atlases from neonates to 1- and 2-year-olds. *PLoS One* 6 (4), e18746. doi:10.1371/journal.pone.0018746

Shiee, N., Bazin, P.-L., Cuzzocreo, J., Blitz, A., Pham, D. L., 2011. Segmentation of brain images using adaptive atlases with application to ventriculomegaly. In: *Proceedings of Information Processing in Medical Imaging*. vol. 22. pp. 1–12.

Shimony, J. S., Smyser, C. D., Wideman, G., Alexopoulos, D., Hill, J., Harwell, J., Dierker, D., Van Essen, D. C., Inder, T. E., Neil, J. J., 2016. Comparison of cortical folding measures for evaluation of developing human brain. *NeuroImage* 125, 780–90. doi:10.1016/j.neuroimage.2015.11.001

Siddiqui, A. M., Masood, A., Saleem, M., 2009. A locally constrained radial basis function for registration and warping of images. *Pattern Recognition Letters* 30 (4), 377–90. doi:10.1016/j.patrec.2008.11.004

Singh, N., Hinkle, J., Joshi, S., Fletcher, T., 2013. A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: *Proceedings of IEEE International Symposium on Biomedical Imaging (ISBI): From Nano to Macro*. pp. 1219–22. doi:10.1109/ISBI.2013.6556700

Sled, J. G., Nossin-Manor, R., 2013. Quantitative MRI for studying neonatal brain development. *Neuroradiology* 55 (S2), 97–104. doi:10.1007/s00234-013-1235-9

Smith, S., 2002. Fast robust automated brain extraction. *Human Brain Mapping* 17 (3), 143–55. doi:10.1002/hbm.10062

Sommer, S., Lauze, F., Nielsen, M., Pennec, X., 2011a. Kernel bundle EPDiff: Evolution equations for multi-scale diffeomorphic image registration. In: Bruckstein, A. M., ter Haar Romeny, B. M., Bronstein, A. M., Bronstein, M. M. (eds.), *Proceedings of 3rd International Conference on Scale Space and Variational Methods in Computer Vision (SSVM)*. Springer-Verlag, Heidelberg, pp. 677–88. doi:10.1007/978-3-642-24785-9_57

Sommer, S., Lauze, F., Nielsen, M., Pennec, X., 2012. Sparse multi-scale diffeomorphic registration: The kernel bundle framework. *Journal of Mathematical Imaging and Vision* 46 (3), 292–308. doi:10.1007/s10851-012-0409-0

Sommer, S., Nielsen, M., Lauze, F., Pennec, X., 2011b. A multi-scale kernel bundle for LDDMM: towards sparse deformation description across space and scales. In: Székely, G., Hahn, H. K. (eds.), *Proceedings of 22nd International Conference on Information Processing in Medical Imaging*. vol. 22. Springer Berlin Heidelberg, pp. 624–35. doi:10.1007/978-3-642-22092-0_51

Song, Z., Awate, S. P., Licht, D. J., Gee, J. C., 2007. Clinical neonatal brain MRI segmentation using adaptive nonparametric data models and intensity-based markov priors.

In: Ayache, N., Ourselin, S., Maeder, A. (eds.), *Proceedings of 10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 883–90. doi:10.1007/978-3-540-75757-3_107

Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging* 32 (7), 1153–90. doi:10.1109/TMI.2013.2265603

Sotiras, A., Paragios, N., 2012. Discrete symmetric image registration. In: *Proceedings of 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 342–45. doi:10.1109/ISBI.2012.6235554

Srhoj-Egekher, V., Benders, M. J. N. L., Viergever, M. A., Išgum, I., 2013. Automatic neonatal brain tissue segmentation with MRI. In: Ourselin, S., Haynor, D. R. (eds.), *Proceedings of Society of Photographic Instrumentation Engineers (SPIE)*. vol. 8669. pp. 86691K–1–8. doi:10.1117/12.2006653

Staring, M., Klein, S., Pluim, J. P. W., 2007. A rigidity penalty term for nonrigid registration. *Medical Physics* 34 (11), 4098–108. doi:10.1118/1.2776236

Stefanescu, R., Pennec, X., Ayache, N., 2004. Grid powered nonlinear image registration with locally adaptive regularization. *Medical Image Analysis* 8 (3), 325–42. doi:10.1016/j.media.2004.06.010

Stolp, H., Neuhaus, A., Sundramoorthi, R., Molnár, Z., 2012. The long and the short of it: Gene and environment interactions during early cortical development and consequences for long-term neurological disease. *Frontiers in Psychiatry* 3, 8–15. doi:10.3389/fpsyt.2012.00050

Studholme, C., 2003. Simultaneous population based image alignment for template free spatial normalisation of brain anatomy. In: Gee, J. C., Maintz, J. B. A., Vannier, M. W. (eds.), *Proceedings of International Workshop on Biomedical Image Registration*. Springer Berlin Heidelberg, pp. 81–90. doi:10.1007/978-3-540-39701-4_9

Studholme, C., 2011. Mapping fetal brain development in utero using magnetic resonance imaging: the Big Bang of brain mapping. *Annual Review of Biomedical Engineering* 13, 345–68. doi:10.1146/annurev-bioeng-071910-124654

Studholme, C., Cardenas, V., 2004. A template free approach to volumetric spatial normalization of brain anatomy. *Pattern Recognition Letters* 25 (10), 1191–202. doi:10.1016/j.patrec.2004.03.015

Studholme, C., Hill, D., Hawkes, D., 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition* 32 (1), 71–86. doi:10.1016/S0031-3203(98)00091-0

Studholme, C., Hill, D. L. G., Hawkes, D. J., 1995. Multiresolution voxel similarity measures for MR-PET registration. In: *Proceedings of 14th International Conference on Information Processing in Medical Imaging*. pp. 287–298.

Szeliski, R., Coughlan, J., 1994. Spline-based image registration. *International Journal of Computer Vision* 22 (3), 199–218. doi:10.1023/A:1007996332012

Talairach, J., Tournoux, P., jan 1988. Co-Planar Stereotaxic Atlas of the Human Brain. Thieme Medical Publishers, New York, USA.

Tang, S., Fan, Y., Wu, G., Kim, M.-J., Shen, D., 2009. RABBIT: Rapid alignment of brains by building intermediate templates. *NeuroImage* 47 (4), 1277–87. doi:10.1016/j.neuroimage.2009.02.043

Tanner, C., Schnabel, J. A., Chung, D., Clarkson, M. J., Rueckert, D., Hill, D. L. G., Hawkes, D. J., 2000. Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In: *Proceedings of 3rd International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 327–337. doi:10.1007/b12345

Tarroni, G., Oktay, O., Schuh, A., Bai, W., de Marvao, A., O'Regan, D., Cook, S., Rueckert, D., Apr 2017. Slice realignment for motion-corrupted stacks of short-axis cine cardiac MR images based on 3D probabilistic edge maps. In: *Proceedings of 25th*

Annual Meeting of International Society for Magnetic Resonance in Medicine (ISMRM). Honolulu, HI, USA.

Terzopoulos, D., Fleischer, K., 1988. Deformable models. *The Visual Computer* 4 (6), 306–31. doi:10.1007/BF01908877

Thévenaz, P., Unser, M., 2000. Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing* 9 (12), 2083–99. doi:10.1109/83.887976

Thirion, J. P., 1998. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis* 2 (3), 243–60. doi:10.1016/S1361-8415(98)80022-4

Thompson, P. M., Toga, A. W., 2002. A framework for computational anatomy. *Computing and Visualization in Science* 5 (1), 13–34. doi:10.1007/s00791-002-0084-6

Tolsa, C. B., Zimine, S., Warfield, S. K., Freschi, M., Rossignol, A. S., Lazeyras, F., Hanquinet, S., Pfizenmaier, M., Hüppi, P. S., 2004. Early alteration of structural and functional brain development in premature infants born with intrauterine growth restriction. *Pediatric Research* 56 (1), 132–38. doi:10.1203/01.PDR.0000128983.54614.7E

Tosun, D., Rettmann, M. E., Naiman, D. Q., Resnick, S. M., Kraut, M. A., Prince, J. L., 2006. Cortical reconstruction using implicit surface evolution: Accuracy and precision analysis. *NeuroImage* 29 (3), 838–52. doi:10.1016/j.neuroimage.2005.08.061

Trouvé, A., 1995. An infinite dimensional group approach for physics based models in patterns recognition, preprint. url: http://www.cis.jhu.edu/publications/papers_in_database/alain/trouve1995.pdf

Trouvé, A., 1998. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision* 28 (3), 213–21. doi:10.1023/A:1008001603737

Tu, Z., Bai, X., 2010. Auto context and its application to high level vision tasks and 3D brain image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32 (10), 1744–57. doi:10.1109/TPAMI.2009.186

Tustison, N., Avants, B., Cook, P., Zheng, Y., Egan, A., Yushkevich, P., Gee, J., 2010. N4ITK: Improved N3 bias correction. *IEEE Transactions on Medical Imaging* 29 (6), 1310–20. doi:10.1109/TMI.2010.2046908

Tustison, N. J., Avants, B. B., 2013. Explicit B-spline regularization in diffeomorphic image registration. *Frontiers in Neuroinformatics* 7, 39. doi:10.3389/fninf.2013.00039

Tustison, N. J., Avants, B. B., Gee, J. C., 2009. Directly manipulated free-form deformation image registration. *IEEE Transactions on Image Processing* 18 (3), 624–35. doi:10.1109/TIP.2008.2010072

Vaillant, M., Miller, M. I., Younes, L., Trouvé, A., 2004. Statistics on diffeomorphisms via tangent space representations. *NeuroImage* 23 Suppl 1, S161–9. doi:10.1016/j.neuroimage.2004.07.023

Van Essen, D., 2012. Cortical cartography and Caret software. *NeuroImage* 62 (2), 757–64. doi:10.1016/j.neuroimage.2011.10.077

Van Essen, D. C., 1997. A tension-based theory of morphogenesis and compact wiring in the central nervous system. *Nature* 385, 313–18.

Van Essen, D. C., 2005. A population-average, landmark- and surface-based (PALS) atlas of human cerebral cortex. *NeuroImage* 28 (3), 635–62. doi:10.1016/j.neuroimage.2005.06.058

Van Essen, D. C., Drury, H. a., Dickson, J., Harwell, J., Hanlon, D., Anderson, C. H., 2001. An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association* 8 (5), 443–59. doi:10.1136/jamia.2001.0080443

Van Leemput, K., Bakkour, A., Benner, T., Wiggins, G., Wald, L. L., Augustinack, J., Dickerson, B. C., Golland, P., Fischl, B., 2009. Automated segmentation of hippocampal subfields from ultra-high resolution in vivo MRI. *Hippocampus* 19 (6), 549–57. doi:10.1002/hipo.20615

Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999a. Automated model-based bias field correction of MR images of the brain. *IEEE Transactions on Medical Imaging* 18 (10), 885–96. doi:10.1109/42.811268

Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999b. Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging* 18 (10), 897–908. doi:10.1109/42.811270

Van Leemput, K., Sabuncu, M. R., 2014. A cautionary analysis of STAPLE using direct inference of segmentation truth. In: *Proceedings of 17th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 398–406.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2007. Non-parametric diffeomorphic image registration with the demons algorithm. In: Ayache, N., Ourselin, S., Maeder, A. (eds.), *Proceedings of 10th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 319–26. doi:10.1007/978-3-540-75759-7_39

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2008. Symmetric log-domain diffeomorphic registration: A demons-based approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.), *Proceedings of 11th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 754–61. doi:10.1007/978-3-540-85988-8_90

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2009. Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* 45 (1), S61–72. doi:10.1016/j.neuroimage.2008.10.040

Vinall, J., Grunau, R. E., Brant, R., Chau, V., Poskitt, K. J., Synnes, A. R., Miller, S. P., 2013. Slower postnatal growth is associated with delayed cerebral cortical maturation in preterm newborns. *Science Translational Medicine* 5 (168), 168ra8. doi:10.1126/scitranslmed.3004666

Viola, P., Jones, M., 2004. Robust real-time face detection. *International Journal of Computer Vision* 57 (2), 137–154. doi:10.1023/B:VISI.0000013087.49260.fb

Viola, P., Wells, W., 1995. Alignment by maximization of mutual information. In: *Proceedings of 5th IEEE International Conference on Computer Vision (ICCV)*. pp. 16–23. doi:10.1109/ICCV.1995.466930

Vita, A., De Peri, L., Deste, G., Sacchetti, E., 2012. Progressive loss of cortical gray matter in schizophrenia: a meta-analysis and meta-regression of longitudinal MRI studies. *Translational Psychiatry* 2 (11), e190. doi:10.1038/tp.2012.116

Wang, H., Suh, J. W., Das, S. R., Pluta, J. B., Craige, C., Yushkevich, P. A., 2013. Multi-atlas segmentation with joint label fusion multi-atlas segmentation with joint label fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (3), 611–23. doi:10.1109/TPAMI.2012.143

Wang, L., Gao, Y., Shi, F., Li, G., Gilmore, J. H., Lin, W., Shen, D., 2015. LINKS: Learning-based multi-source integration framework for segmentation of infant brain images. *NeuroImage* 108, 160–172. doi:10.1016/j.neuroimage.2014.12.042

Wang, L., He, L., Mishra, A., Li, C., 2009. Active contours driven by local Gaussian distribution fitting energy. *Signal Processing* 89 (12), 2435–47. doi:10.1016/j.sigpro.2009.03.014

Wang, L., Shi, F., Gao, Y., Li, G., Gilmore, J. H., Lin, W., Shen, D., 2014a. Integration of sparse multi-modality representation and anatomical constraint for isointense infant brain MR image segmentation. *NeuroImage* 89, 152–64. doi:10.1016/j.neuroimage.2013.11.040

Wang, L., Shi, F., Li, G., Gao, Y., Lin, W., Gilmore, J. H., Shen, D., 2014b. Segmentation of neonatal brain MR images using patch-driven level sets. *NeuroImage* 84, 141–58. doi:10.1016/j.neuroimage.2013.08.008

Wang, L., Shi, F., Lin, W., Gilmore, J. H., Shen, D., 2011a. Automatic segmentation of neonatal images using convex optimization and coupled level sets. *NeuroImage* 58 (3), 805–17. doi:10.1016/j.neuroimage.2011.06.064

Wang, L., Shi, F., Yap, P.-T., Lin, W., Gilmore, J. H., Shen, D., 2011b. Longitudinally guided level sets for consistent tissue segmentation of neonates. *Human Brain Mapping* 34 (4), 956–72. doi:10.1002/hbm.21486

Wang, S., Kuklisova-Murgasova, M., Schnabel, J. A., 2012. An atlas-based method for neonatal MR brain tissue segmentation. In: *MICCAI NeoBrainS12 Challenge*. pp. 1–8.

Wang, Y., Staib, L. H., 2000. Physical model-based non-rigid registration incorporating statistical shape information. *Medical Image Analysis* 4 (1), 7–20. doi:10.1016/S1361-8415(00)00004-9

Warfield, S. K., Kaus, M., Jolesz, F. A., Kikinis, R., 2000. Adaptive, template moderated, spatially varying statistical classification. *Medical Image Analysis* 4 (1), 43–55. doi:10.1016/S1361-8415(00)00003-7

Warfield, S. K., Zou, K. H., Wells, W. M., 2004. Simultaneous truth and performance level estimation (STAPLE): An algorithm for the validation of image segmentation. *IEEE Transactions on Medical Imaging* 23 (7), 903–21. doi:10.1109/TMI.2004.828354

Weisenfeld, N., Mewes, A., Warfield, S., 2006a. Segmentation of newborn brain MRI. In: *Proceedings of 3rd IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 766–9. doi:10.1109/ISBI.2006.1625029

Weisenfeld, N. I., Mewes, A. U. J., Warfield, S. K., 2006b. Highly accurate segmentation of brain tissue and subcortical gray matter from newborn MRI. In: *Proceedings of 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 199–206.

Weisenfeld, N. I., Warfield, S. K., 2009. Automatic segmentation of newborn brain MRI. *NeuroImage* 47, 564–72. doi:10.1016/j.neuroimage.2009.04.068

Wells, W. M., Grimson, W. E. L., Kikinis, R., Jolesz, F. A., 1996. Adaptive segmentation of MRI data. *IEEE Transactions on Medical Imaging* 15 (4), 429–442. doi:10.1109/42.511747

Wen, P., Ruofeng, T., Guiping, Q., Jinxiang, D., 2006. A local registration approach of medical images with niche genetic algorithm. In: *Proceedings of 10th International Conference on Computer Supported Cooperative Work in Design*. pp. 1099–1104. doi:10.1109/CSCWD.2006.253158

Wendland, H., 1995. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics* 4 (1), 389–96. doi:10.1007/BF02123482

Wilke, M., Holland, S. K., Altaye, M., Gaser, C., 2008. Template-O-Matic: A toolbox for creating customized pediatric templates. *NeuroImage* 41 (3), 903–13. doi:10.1016/j.neuroimage.2008.02.056

Wilke, M., Schmithorst, V. J., Holland, S. K., 2002. Assessment of spatial normalization of whole-brain magnetic resonance images in children. *Human Brain Mapping* 17 (1), 48–60. doi:10.1002/hbm.10053

Wilke, M., Schmithorst, V. J., Holland, S. K., 2003. Normative pediatric brain data for spatial normalization and segmentation differs from standard adult data. *Magnetic Resonance in Medicine* 50 (4), 749–57. doi:10.1002/mrm.10606

Woods, R. P., Grafton, S. T., Watson, J. D., Sicotte, N. L., Mazziotta, J. C., 1998. Automated image registration: II. Intersubject validation of linear and nonlinear models. *Journal of Computer Assisted Tomography* 22 (1), 153–65. doi:10.1097/00004728-199801000-00028

Wörz, S., Rohr, K., 2008. Physics-based elastic registration using non-radial basis functions and including landmark localization uncertainties. *Computer Vision and Image Understanding* 111 (3), 263–74. doi:10.1016/j.cviu.2007.12.003

Wouters, J., D'Agostino, E., Maes, F., Vandermeulen, D., Suetens, P., 2006. Non-rigid brain image registration using a statistical deformation model. In: *Proceedings of Society of Photographic Instrumentation Engineers (SPIE): Medical Imaging*. vol. 6144. pp. 614411–8. doi:10.1117/12.653081

Wright, R., Kyriakopoulou, V., Ledig, C., Rutherford, M. A., Hajnal, J. V., Rueckert, D., Aljabar, P., 2014. Automatic quantification of normal cortical folding patterns from fetal brain MRI. *NeuroImage* 91, 21–32. doi:10.1016/j.neuroimage.2014.01.034

Wright, R., Makropoulos, A., Kyriakopoulou, V., Patkee, P., Koch, L., Rutherford, M., Hajnal, J., Rueckert, D., Aljabar, P., 2015. Construction of a fetal spatio-temporal cortical surface atlas from in utero MRI: Application of spectral surface matching. *NeuroImage* 120, 467–80. doi:10.1016/j.neuroimage.2015.05.087

Wu, G., Wang, Q., Jia, H., Shen, D., 2012. Feature-based groupwise registration by hierarchical anatomical correspondence detection. *Human Brain Mapping* 33 (2), 253–71. doi:10.1002/hbm.21209

Wu, J., Avants, B., 2012. Automatic registration-based segmentation for neonatal brains using ANTs and Atropos. In: *MICCAI NeoBrainS12 Challenge*. p. 36.

Xu, C., Pham, D., Prince, J., 2000. Image segmentation using deformable models. In: Sonka, M., Fitzpatrick, J. M. (eds.), *Handbook of Medical Imaging, Volume 2. Medical Image Processing and Analysis*, 1st edition. Society of Photographic Instrumentation Engineers (SPIE), Bellingham, WA, USA, ch. 3, pp. 129–174. doi:10.1117/3.831079.ch3

Xu, C., Pham, D. L., Prince, J. L., 1997. Finding the brain cortex using fuzzy segmentation, isosurfaces, and deformable surface models. In: Duncan, J., Gindi, G. (eds.), *Proceedings of 15th International Conference on Information Processing in Medical Imaging*. Springer Berlin Heidelberg, pp. 399–404. doi:10.1007/3-540-63046-5_33

Xu, C., Prince, J. L., 1998. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing* 7 (3), 359–69. doi:10.1109/83.661186

Xue, H., Srinivasan, L., Jiang, S., Rutherford, M., Edwards, A. D., Rueckert, D., Hajnal, J. V., 2007. Automatic segmentation and reconstruction of the cortex from neonatal MRI. *NeuroImage* 38 (3), 461–77. doi:10.1016/j.neuroimage.2007.07.030

Xue, Z., Shen, D., Davatzikos, C., 2006. Statistical representation of high-dimensional deformation fields with application to statistically constrained 3D warping. *Medical Image Analysis* 10 (5), 740–51. doi:10.1016/j.media.2006.06.007

Yang, X., Kwitt, R., Niethammer, M., 2016. Fast predictive image registration. In: Carneiro, G., Mateus, D., Peter, L., Bradley, A., Tavares, J. M. R. S., Belagiannis, V., Papa, J. P., Nascimento, J. C., Loog, M., Lu, Z., Cardoso, J. S., Cornebise, J. (eds.), *Proceedings of 2nd International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis (LABELS), and 1st International Workshop on Deep Learning and Data Labeling for Medical Applications (DLMIA) held in conjunction with MICCAI 2016*. pp. 48–57. doi:10.1007/978-3-319-46976-8_6

Yang, X., Li, Y., Reutens, D., Jiang, T., 2015. Diffeomorphic metric landmark mapping using stationary velocity field parameterization. *International Journal of Computer Vision* 115 (2), 69–86. doi:10.1007/s11263-015-0802-4

Yanovsky, I., Thompson, P. M., Osher, S., Leow, A. D., 2007. Topology preserving log-unbiased nonlinear image registration: Theory and implementation. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1–8. doi:10.1109/CVPR.2007.383144

Yeo, B., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P., 2010. Spherical demons: fast diffeomorphic landmark-free surface registration. *IEEE Transactions on Medical Imaging* 29 (3), 650–68. doi:10.1109/TMI.2009.2030797

Younes, L., 2010. Shapes and Diffeomorphisms. vol. 171 of *Applied Mathematical Sciences*. Springer Berlin Heidelberg.

Younes, L., Arrate, F., Miller, M. I., 2009. Evolutions equations in computational anatomy. *NeuroImage* 45 (1 Suppl), S40–50. doi:10.1016/j.neuroimage.2008.10.050

Yushkevich, P. A., Avants, B. B., Das, S. R., Pluta, J., Altinay, M., Craige, C., 2010. Bias in estimation of hippocampal atrophy using deformation-based morphometry arises from asymmetric global normalization: an illustration in ADNI 3T MRI data. *NeuroImage* 50 (2), 434–45. doi:10.1016/j.neuroimage.2009.12.007

Zeng, X., Staib, L. H., Schultz, R. T., Duncan, J. S., 1999. Segmentation and measurement of the cortex from 3-D MR images using coupled-surfaces propagation. *IEEE Transactions on Medical Imaging* 18 (10), 927–37. doi:10.1109/42.811276

Zhang, M., Fletcher, P. T., 2015. Finite-dimensional Lie algebras for fast diffeomorphic image registration. In: Ourselin, S., Alexander, D., Westin, C. F., Cardoso, M. (eds.), *Proceedings of 24th International Conference on Information Processing in Medical Imaging*. pp. 249–60. doi:10.1007/978-3-319-19992-4_19

Zhang, W., Li, R., Deng, H., Wang, L., Lin, W., Ji, S., Shen, D., 2015. Deep convolutional neural networks for multi-modality isointense infant brain image segmentation. *NeuroImage* 108, 214–24. doi:10.1016/j.neuroimage.2014.12.061

Zhang, Y., Shi, F., Wu, G., Wang, L., Yap, P.-T., Shen, D., 2016a. Consistent spatial-temporal longitudinal atlas construction for developing infant brains. *IEEE Transactions on Medical Imaging* 35 (12), 2568–77. doi:10.1109/TMI.2016.2587628

Zhang, Y., Shi, F., Yap, P. T., Shen, D., 2016b. Detail-preserving construction of neonatal brain atlases in space-frequency domain. *Human Brain Mapping* 37 (6), 2133–50. doi:10.1002/hbm.23160

Zijdenbos, A., Evans, A., Riahi, F., Sled, J., Chui, J., Kollokian, V., 1996. Automatic quantification of multiple sclerosis lesion volume using stereotaxic space. In: Höhne, K. H., Kikinis, R. (eds.), *Proceedings of 4th International Conference on Visualization in Biomedical Computing*. Springer Berlin Heidelberg, pp. 439–48. doi:10.1007/BFb0046984

Zijdenbos, A., Forghani, R., Evans, A., 1998. Automatic quantification of MS lesions in 3D MRI brain data sets: Validation of INSECT. In: Wells, W. M., Colchester, A., Delp,

S. (eds.), *Proceedings of 1st International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer Berlin Heidelberg, pp. 439–48. doi:10.1007/BFb0056229

Zikic, D., Glocker, B., Criminisi, A., 2014. Encoding atlases by randomized classification forests for efficient multi-atlas label propagation. *Medical Image Analysis* 18 (8), 1262–73. doi:10.1016/j.media.2014.06.010

Zöllei, L., 2006. A unified information theoretic framework for pair- and group-wise registration of medical images. Ph.D. thesis, Massachusetts Institute of Technology.

Appendices

Appendix A.

Non-parametric transformations

In section 2.2, we introduced common variational formulations of non-parametric transformation models. As mentioned previously, these models are given by a continuous spatial mapping $\boldsymbol{\Phi} : \Omega \rightarrow \mathbb{R}^d$ defined as $\boldsymbol{\Phi}(\boldsymbol{x}) = \boldsymbol{x} + \boldsymbol{u}(\boldsymbol{x})$, where $\boldsymbol{u} : \Omega \rightarrow \mathbb{R}^d$ is a dense displacement field defined on the d-dimensional image domain $\Omega \subset \mathbb{R}^d$ with boundary denoted by $\partial\Omega$, and an associated internal energy functional, $\mathcal{S}[\boldsymbol{u}]$, which imposes a sufficient smoothness on the displacement field $\boldsymbol{u}$. This is required to counter the ill-posedness of the registration problem and restrict the space of admissible solutions. Additional task-specific regularisation of the solutions are discussed in section 2.3. The energy functionals and associated Euler-Lagrange equations of common non-parametric deformation models are given in section 2.2.2. In this section, we provide a more detailed derivation of these Euler-Lagrange equations using the Gâteaux derivative. Note also that the internal forces of a parametric deformable model may be derived from a similar variational formulation of a given energy functional describing the internal properties of the model (see section 4.5.1).

The Gâteaux derivative of an energy functional $\mathcal{S}[\boldsymbol{u}]$, which assigns a non-negative cost to a continuous displacement field $\boldsymbol{u}$, in direction $\boldsymbol{h} : \Omega \rightarrow \mathbb{R}^3$ is defined as

$$\delta_{\boldsymbol{h}}\mathcal{S}[\boldsymbol{u}] = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{S}[\boldsymbol{u} + \epsilon\boldsymbol{h}] - \mathcal{S}[\boldsymbol{u}]}{\epsilon} = \frac{\partial}{\partial\epsilon}\mathcal{S}[\boldsymbol{u} + \epsilon\boldsymbol{h}]\bigg|_{\epsilon=0} \tag{A.1}$$

Using this definition, we derive the Gâteaux derivatives of common deformation models. Requiring that $\delta_{\boldsymbol{h}}\mathcal{S}[\boldsymbol{u}] = \boldsymbol{0}\ \forall \boldsymbol{x} \in \Omega$ gives rise to the Euler-Lagrange equations, which are necessary conditions for a solution $\hat{\boldsymbol{u}}$ to minimise the energy $\mathcal{S}[\boldsymbol{u}]$.

A.1. Diffusion model

A.1.1. Internal energy

The internal energy of the diffusion model is given by (Fischer and Modersitzki, 2004)

$$\mathcal{S}_{\text{diff}}[\boldsymbol{u}] \coloneqq \frac{1}{2}\int_{\Omega}\sum_{i=1}^{d}\|\nabla u_i\|^2\,d\boldsymbol{x} = \frac{1}{2}\int_{\Omega}\sum_{i,j=1}^{d}\left(\frac{\partial u_i}{\partial x_j}\right)^2 d\boldsymbol{x} \tag{A.2}$$

A.1.2. Gâteaux derivative

The Gâteaux derivative of the functional given by equation (A.2) in direction $\boldsymbol{h}$ is

$$\begin{aligned}\delta_{\boldsymbol{h}}\mathcal{S}_{\text{diff}}[\boldsymbol{u}] &= \frac{\partial}{\partial\epsilon}\left[\frac{1}{2}\int_{\Omega}\sum_{i,j=1}^{d}\left(\frac{\partial u_i}{\partial x_j}+\epsilon\frac{\partial h_i}{\partial x_j}\right)^2 d\boldsymbol{x}\right]\Bigg|_{\epsilon=0}\\ &= \int_{\Omega}\sum_{i,j=1}^{d}\left[\frac{\partial u_i}{\partial x_j}+\epsilon\frac{\partial h_i}{\partial x_j}\right]\Bigg|_{\epsilon=0}\frac{\partial h_i}{\partial x_j}\,d\boldsymbol{x}\\ &= \sum_{i,j=1}^{d}\left(\int_{\Omega}\frac{\partial u_i}{\partial x_j}\frac{\partial h_i}{\partial x_j}\,d\boldsymbol{x}\right)\end{aligned} \tag{A.3}$$

Using integration by parts for each integral term, we get

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{diff}}[\boldsymbol{u}] = \sum_{i,j=1}^{d}\left(\int_{\partial\Omega}\frac{\partial u_i}{\partial x_j}h_i n_j\,d\boldsymbol{x} - \int_{\Omega}\frac{\partial^2 u_i}{\partial x_j^2}h_i\,d\boldsymbol{x}\right) \tag{A.4}$$

where n_j is the j-th component of the normal vector to the boundary of the domain.

Given Neumann boundary conditions $D\boldsymbol{u} = \boldsymbol{0}\ \forall\boldsymbol{x}\in\partial\Omega$, the derivative is given by

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{diff}}[\boldsymbol{u}] = -\sum_{i=1}^{d}\int_{\Omega}\left(\sum_{j=1}^{d}\frac{\partial^2 u_i}{\partial x_j^2}\right)h_i\,d\boldsymbol{x} = -\int_{\Omega}\Delta\boldsymbol{u}\cdot\boldsymbol{h}\,d\boldsymbol{x} \tag{A.5}$$

A.1.3. Euler-Lagrange equation

Requiring the derivative (A.5) to be zero for all perturbations $\boldsymbol{h}$, we obtain the necessary condition for a minimum of the diffusion energy given by equation (A.2), i.e.,

$$\Delta\boldsymbol{u} = \boldsymbol{0} \quad\Longrightarrow\quad \delta_{\boldsymbol{h}}\mathcal{S}_{\text{diff}}[\boldsymbol{u}] = 0 \tag{A.6}$$

A.2. Curvature model

A.2.1. Internal energy

The internal energy of the curvature model of Fischer and Modersitzki (2004) is given by

$$\mathcal{S}_{\text{curv}}[\boldsymbol{u}] := \frac{1}{2}\int_{\Omega}\sum_{i=1}^{d}(\Delta u_i)^2\,d\boldsymbol{x} = \frac{1}{2}\int_{\Omega}\sum_{i=1}^{d}\left(\sum_{j=1}^{d}\frac{\partial^2 u_i}{\partial x_j^2}\right)^2 d\boldsymbol{x} \tag{A.7}$$

A.2.2. Gâteaux derivative

The Gâteaux derivative of the functional given by equation (A.7) in direction $\boldsymbol{h}$ is

$$\begin{aligned}
\delta_{\boldsymbol{h}}\mathcal{S}_{\text{curv}}[\boldsymbol{u}] &= \frac{\partial}{\partial\epsilon}\left[\frac{1}{2}\int_{\Omega}\sum_{i=1}^{d}\left(\sum_{j=1}^{d}\frac{\partial^2 u_i}{\partial x_j^2}+\epsilon\frac{\partial^2 h_i}{\partial x_j^2}\right)^2 d\boldsymbol{x}\right]\Bigg|_{\epsilon=0}\\
&= \int_{\Omega}\sum_{i=1}^{d}\left[\sum_{j=1}^{d}\frac{\partial^2 u_i}{\partial x_j^2}+\epsilon\frac{\partial^2 h_i}{\partial x_j^2}\right]\Bigg|_{\epsilon=0}\left(\sum_{k=1}^{d}\frac{\partial^2 h_i}{\partial x_k^2}\right)d\boldsymbol{x}\\
&= \sum_{i=1}^{d}\int_{\Omega}\left(\sum_{j=1}^{d}\frac{\partial^2 u_i}{\partial x_j^2}\right)\left(\sum_{k=1}^{d}\frac{\partial^2 h_i}{\partial x_k^2}\right)d\boldsymbol{x}\\
&= \sum_{i,j,k=1}^{d}\int_{\Omega}\frac{\partial^2 u_i}{\partial x_j^2}\frac{\partial^2 h_i}{\partial x_k^2}\,d\boldsymbol{x}
\end{aligned} \tag{A.8}$$

Using integration by parts twice for each integral term, we get

$$\begin{aligned}
\delta_{\boldsymbol{h}}\mathcal{S}_{\text{curv}}[\boldsymbol{u}] &= \sum_{i,j,k=1}^{d}\left(\int_{\partial\Omega}\frac{\partial^2 u_i}{\partial x_j^2}\frac{\partial h_i}{\partial x_k}n_k\,d\boldsymbol{x} - \int_{\Omega}\frac{\partial^3 u_i}{\partial x_j^2 x_k}\frac{\partial h_i}{\partial x_k}\,d\boldsymbol{x}\right)\\
&= \sum_{i,j,k=1}^{d}\left(\int_{\partial\Omega}\left(\frac{\partial^2 u_i}{\partial x_j^2}\frac{\partial h_i}{\partial x_k}-\frac{\partial^3 u_i}{\partial x_j^2\partial x_k}h_i\right)n_k\,d\boldsymbol{x} + \int_{\Omega}\frac{\partial^4 u_i}{\partial x_j^2\partial x_k^2}h_i\,d\boldsymbol{x}\right)
\end{aligned} \tag{A.9}$$

where n_k is the k-th component of the normal vector to the boundary of the domain.

Given Neumann boundary conditions $D\boldsymbol{u} = \boldsymbol{0}\ \forall\boldsymbol{x}\in\partial\Omega$, the derivative is given by

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{curv}}[\boldsymbol{u}] = \int_{\Omega}\sum_{i=1}^{d}\left(\sum_{j,k=1}^{d}\frac{\partial^4 u_i}{\partial x_j^2\partial x_k^2}\right)h_i\,d\boldsymbol{x} = \int_{\Omega}\left(\Delta^2\boldsymbol{u}\right)\cdot\boldsymbol{h}\,d\boldsymbol{x} \tag{A.10}$$

A.2.3. Euler-Lagrange equation

Requiring the derivative (A.10) to be zero for all perturbations $\boldsymbol{h}$, we obtain the necessary condition for a minimum of the curvature energy given by equation (A.7), i.e.,

$$\Delta^2 \boldsymbol{u} = \boldsymbol{0} \quad \Longrightarrow \quad \delta_{\boldsymbol{h}} \mathcal{S}_{\text{curv}}[\boldsymbol{u}] = 0 \tag{A.11}$$

A.3. Elastic and fluid model

A.3.1. Internal energy

The potential of linear elasticity is given by (Fischer and Modersitzki, 2004)

$$\begin{aligned}
\mathcal{S}_{\text{elastic}}[\boldsymbol{u}] &\coloneqq \frac{1}{2}\int_\Omega \frac{\mu}{2}\sum_{i,j=1}^{d}\left(\frac{\partial u_i}{\partial x_j}+\frac{\partial u_j}{\partial x_i}\right)^2+\lambda\left(\sum_{i=1}^{d}\frac{\partial u_i}{\partial x_i}\right)^2 d\boldsymbol{x} \\
&= \frac{1}{2}\int_\Omega \frac{\mu}{2}\sum_{i,j=1}^{d}\left(\left(\frac{\partial u_i}{\partial x_j}\right)^2+2\frac{\partial u_i}{\partial x_j}\frac{\partial u_j}{\partial x_i}+\left(\frac{\partial u_j}{\partial x_i}\right)^2\right)+\lambda\left(\sum_{i=1}^{d}\frac{\partial u_i}{\partial x_i}\right)^2 d\boldsymbol{x} \\
&= \mu\,\mathcal{S}_{\text{diff}}[\boldsymbol{u}] \;+\; \mu\underbrace{\frac{1}{2}\int_\Omega \sum_{i,j=1}^{d}\frac{\partial u_i}{\partial x_j}\frac{\partial u_j}{\partial x_i}\,d\boldsymbol{x}}_{=:\,\mathcal{S}_{\text{mix}}[\boldsymbol{u}]} \;+\; \lambda\underbrace{\frac{1}{2}\int_\Omega\left(\sum_{i=1}^{d}\frac{\partial u_i}{\partial x_i}\right)^2 d\boldsymbol{x}}_{=:\,\mathcal{S}_{\text{div}}[\boldsymbol{u}]}
\end{aligned} \tag{A.12}$$

A.3.2. Gâteaux derivative

The Gâteaux derivative of the second term in equation (A.12) in direction $\boldsymbol{h}$ is

$$\begin{aligned}
\delta_{\boldsymbol{h}}\mathcal{S}_{\text{mix}}[\boldsymbol{u}] &= \frac{\partial}{\partial\epsilon}\left(\frac{1}{2}\int_\Omega\sum_{i,j=1}^{d}\left(\frac{\partial u_i}{\partial x_j}+\epsilon\frac{\partial h_i}{\partial x_j}\right)\left(\frac{\partial u_j}{\partial x_i}+\epsilon\frac{\partial h_j}{\partial x_i}\right)d\boldsymbol{x}\right)\Bigg|_{\epsilon=0} \\
&= \frac{1}{2}\int_\Omega \frac{\partial}{\partial\epsilon}\left[\sum_{i,j=1}^{d}\frac{\partial u_i}{\partial x_j}\frac{\partial u_j}{\partial x_i}+\epsilon\frac{\partial u_i}{\partial x_j}\frac{\partial h_j}{\partial x_i}+\epsilon\frac{\partial u_j}{\partial x_i}\frac{\partial h_i}{\partial x_j}+\epsilon^2\frac{\partial h_i}{\partial x_j}\frac{\partial h_j}{\partial x_i}\right]\Bigg|_{\epsilon=0} d\boldsymbol{x} \\
&= \frac{1}{2}\int_\Omega\sum_{i,j=1}^{d}\left(\frac{\partial u_i}{\partial x_j}\frac{\partial h_j}{\partial x_i}+\frac{\partial u_j}{\partial x_i}\frac{\partial h_i}{\partial x_j}+2\epsilon\frac{\partial h_i}{\partial x_j}\frac{\partial h_j}{\partial x_i}\right)\Bigg|_{\epsilon=0} d\boldsymbol{x}
\end{aligned} \tag{A.13}$$

The last equation can be simplified by exchanging the order of summation for one of the remaining two terms after setting $\epsilon = 0$, i.e.,

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{mix}}[\boldsymbol{u}] = \sum_{i,j=1}^{d} \int_{\Omega} \frac{\partial u_j}{\partial x_i} \frac{\partial h_i}{\partial x_j} \, d\boldsymbol{x} \tag{A.14}$$

Using integration by parts for each integral term, we get

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{mix}}[\boldsymbol{u}] = \sum_{i,j=1}^{d} \left(\int_{\partial\Omega} \frac{\partial u_j}{\partial x_i} h_i n_j \, d\boldsymbol{x} - \int_{\Omega} \frac{\partial^2 u_j}{\partial x_i \partial x_j} h_i \, d\boldsymbol{x} \right) \tag{A.15}$$

where n_j is the j-th components of the normal vector to the boundary of the domain. Given Neumann boundary conditions $D\boldsymbol{u} = \boldsymbol{0} \; \forall \boldsymbol{x} \in \partial\Omega$, the derivative is given by

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{mix}}[\boldsymbol{u}] = -\int_{\Omega} \sum_{i=1}^{d} \frac{\partial}{\partial x_i} \left(\sum_{j=1}^{d} \frac{\partial u_j}{\partial x_j} \right) h_i \, d\boldsymbol{x} = -\int_{\Omega} \nabla \left(\nabla \cdot \boldsymbol{u} \right) \cdot \boldsymbol{h} \, d\boldsymbol{x} \tag{A.16}$$

The Gâteaux derivative of the third term in equation (A.12) in direction $\boldsymbol{h}$ is given by

$$\begin{aligned}
\delta_{\boldsymbol{h}}\mathcal{S}_{\text{div}}[\boldsymbol{u}] &= \frac{\partial}{\partial \epsilon} \left[\frac{1}{2} \int_{\Omega} \left(\sum_{i=1}^{d} \frac{\partial u_i}{\partial x_i} + \epsilon \frac{\partial h_i}{\partial x_i} \right)^2 d\boldsymbol{x} \right] \Bigg|_{\epsilon=0} \\
&= \int_{\Omega} \frac{\partial}{\partial \epsilon} \left[\sum_{i=1}^{d} \frac{\partial u_i}{\partial x_i} + \epsilon \frac{\partial h_i}{\partial x_i} \right] \Bigg|_{\epsilon=0} \left(\sum_{j=1}^{d} \frac{\partial h_j}{\partial x_j} \right) d\boldsymbol{x} \\
&= \int_{\Omega} \left(\sum_{i=1}^{d} \frac{\partial u_i}{\partial x_i} \right) \left(\sum_{j=1}^{d} \frac{\partial h_j}{\partial x_j} \right) d\boldsymbol{x} \\
&= \sum_{i,j=1}^{d} \int_{\Omega} \frac{\partial u_i}{\partial x_i} \frac{\partial h_j}{\partial x_j} \, d\boldsymbol{x}
\end{aligned} \tag{A.17}$$

Using integration by parts for each integral term, we get

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{div}}[\boldsymbol{u}] = \sum_{i,j=1}^{d} \left(\int_{\partial\Omega} \frac{\partial u_i}{\partial x_i} h_j n_j \, d\boldsymbol{x} - \int_{\Omega} \frac{\partial^2 u_i}{\partial x_i \partial x_j} h_j \, d\boldsymbol{x} \right) \tag{A.18}$$

where n_j is the j-th components of the normal vector to the boundary of the domain. Given the previously specified boundary conditions, we find that $\delta_{\boldsymbol{h}}\mathcal{S}_{\text{div}}[\boldsymbol{u}] = \delta_{\boldsymbol{h}}\mathcal{S}_{\text{mix}}[\boldsymbol{u}]$.

Adding the derivative of the diffusion term given by equation (A.5), we therefore get

$$\delta_{\boldsymbol{h}}\mathcal{S}_{\text{elastic}}[\boldsymbol{u}] = -\int_{\Omega} \left(\mu\Delta\boldsymbol{u} + (\mu+\lambda)\nabla\left(\nabla\cdot\boldsymbol{u}\right)\right)\cdot\boldsymbol{h}\, d\boldsymbol{x} \tag{A.19}$$

A.3.3. Euler-Lagrange equation

Requiring the derivative given by (A.19) to be zero for all perturbations $\boldsymbol{h}$, we obtain the necessary condition for a minimum of the elastic energy given by equation (A.12), i.e.,

$$\mu\Delta\boldsymbol{u} + (\mu+\lambda)\nabla(\nabla\cdot\boldsymbol{u}) = \boldsymbol{0} \quad \Longrightarrow \quad \delta_{\boldsymbol{h}}\mathcal{S}_{\text{elastic}}[\boldsymbol{u}] = 0 \tag{A.20}$$

The equation on the left is known as the Navier-Cauchy PDE of linear elasticity. Moreover, when $\boldsymbol{u}$ is substituted by the velocity field $\boldsymbol{v}$ of an incompressible viscous fluid, it corresponds to the Navier-Stokes PDE for low Reynolds number.

Appendix B.

Free-form deformation

B.1. Definition

As discussed in section 2.2.4, the FFD of a point $\boldsymbol{x} \in \Omega \subset \mathbb{R}^d$ is given by

$$\boldsymbol{\Phi}(\boldsymbol{x}; \boldsymbol{c}) = \sum_{k=1}^{n} \boldsymbol{c}_k \beta_k^d(\boldsymbol{x}) = \boldsymbol{c} \underbrace{\left(\beta_1^d(\boldsymbol{x}) \ \cdots \ \beta_n^d(\boldsymbol{x})\right)^T}_{=: \boldsymbol{b}(\boldsymbol{x})} \tag{B.1}$$

where $\boldsymbol{c} = (\boldsymbol{c}_1 \ \cdots \ \boldsymbol{c}_n) \in \mathbb{R}^{d \times n}$ are the parameters of the FFD transformation model with $\boldsymbol{c}_k \in \mathbb{R}^d$ denoting the coefficient vector of the k-th control point. The tensor product of 1D cubic B-spline functions centered at the k-th control point is the d-variate basis function

$$\beta_k^d(\boldsymbol{x}) = \prod_{i=1}^{d} U_{\text{FFD}} \underbrace{\left(\frac{\hat{x}_i - p_{k,i}}{s_i}\right)}_{=: r_{k,i}(\boldsymbol{x})} = \prod_{i=1}^{d} U_{\text{FFD}}\left(r_{k,i}(\boldsymbol{x})\right) \tag{B.2}$$

where $U_{\text{FFD}}(r)$ is given by equation (2.68), and $\hat{\boldsymbol{x}} = \boldsymbol{R}\boldsymbol{x} + \boldsymbol{t}$ is the point relative to the lattice of control points with coordinates $\boldsymbol{p}_k$. The axes of this lattice are given by the rotation matrix $\boldsymbol{R}$, and the translation of the origin is $\boldsymbol{t}$. The control points are equally spaced in the direction of these local coordinate axes with spacing $\boldsymbol{s} = (s_1, \ldots, s_d)^T$, i.e., $p_{k,i} \in s_i \mathbb{Z}$. The signed distance from the k-th control point to point $\boldsymbol{x}$ in the direction of the i-th lattice axis is given by $s_i r_{k,i}(\boldsymbol{x})$. Due to the compact support of the cubic B-spline function $U_{\text{FFD}}(r)$, $\beta_k^d(\boldsymbol{x}) = 0$ if $\exists i \in [1 \ldots d]$ for which $r_{k,i}(\boldsymbol{x}) >= 2$. Hence, only a small set of control points in the vicinity of $\boldsymbol{x}$ contributes to $\boldsymbol{\Phi}(\boldsymbol{x})$.

B.2. Derivatives

B.2.1. Jacobian matrix

The Jacobian matrix of a transformation $\boldsymbol{\Phi} : \Omega \to \mathbb{R}^d$ is defined as

$$D_{\boldsymbol{x}}\boldsymbol{\Phi}(\boldsymbol{x}) = \begin{pmatrix} \frac{\partial y_1}{\partial x_1} & \cdots & \frac{\partial y_1}{\partial x_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial y_d}{\partial x_1} & \cdots & \frac{\partial y_d}{\partial x_d} \end{pmatrix} \tag{B.3}$$

where $\boldsymbol{x} = (x_1 \ \cdots \ x_d)^T$ and $\boldsymbol{y} = (y_1 \ \cdots \ y_d)^T = \boldsymbol{\Phi}(\boldsymbol{x})$ are the untransformed and mapped points, respectively. Given the spline model in equation (B.1), the Jacobian is given by

$$D_{\boldsymbol{x}}\boldsymbol{\Phi}(\boldsymbol{x};\boldsymbol{c}) = D_{\hat{\boldsymbol{x}}}\boldsymbol{\Phi}(\boldsymbol{x};\boldsymbol{c})\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}} \tag{B.4}$$

where

$$\begin{aligned} D_{\hat{\boldsymbol{x}}}\boldsymbol{\Phi}(\boldsymbol{x};\boldsymbol{c}) &= \sum_{k=1}^{n} \boldsymbol{c}_k D_{\hat{\boldsymbol{x}}}\beta_k^d(\boldsymbol{x}) \\ &= \sum_{k=1}^{n} \boldsymbol{c}_k \begin{pmatrix} \frac{\partial \beta_k^d}{\partial \hat{x}_1}(\boldsymbol{x}) & \cdots & \frac{\partial \beta_k^d}{\partial \hat{x}_d}(\boldsymbol{x}) \end{pmatrix} \end{aligned} \tag{B.5}$$

The first order partial derivatives of the cubic B-spline tensor are given by

$$\frac{\partial \beta_k^d}{\partial \hat{x}_i}(\boldsymbol{x}) = \frac{1}{s_i}\left(\prod_{\substack{j=1 \\ j\neq i}}^{d} U_{\mathrm{FFD}}\left(r_{k,j}(\boldsymbol{x})\right)\right) \left.\frac{dU_{\mathrm{FFD}}(r)}{dr}\right|_{r=r_{k,i}(\boldsymbol{x})} \tag{B.6}$$

This Jacobian matrix is post-multiplied according to the chain rule by the derivative of the coordinate transformation $\boldsymbol{x} \mapsto \hat{\boldsymbol{x}}$ given by

$$\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}} = \begin{pmatrix} \frac{\partial \hat{x}_1}{\partial x_1} & \cdots & \frac{\partial \hat{x}_1}{\partial x_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial \hat{x}_d}{\partial x_1} & \cdots & \frac{\partial \hat{x}_d}{\partial x_d} \end{pmatrix} = \boldsymbol{R} \tag{B.7}$$

to obtain the Jacobian of the transformation with respect to the global coordinates $\boldsymbol{x}$.

All first order partial derivatives of the B-spline FFD with respect to the control point coefficients $\boldsymbol{c}$, on the other hand, are given by the $d \times dn$ matrix

$$D_{\boldsymbol{c}}\boldsymbol{\Phi}(\boldsymbol{x};\boldsymbol{c}) = \begin{pmatrix} \frac{\partial y_1}{\partial \boldsymbol{c}_1} & \cdots & \frac{\partial y_1}{\partial \boldsymbol{c}_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial y_d}{\partial \boldsymbol{c}_1} & \cdots & \frac{\partial y_d}{\partial \boldsymbol{c}_n} \end{pmatrix} = \begin{pmatrix} \beta_1^d(\boldsymbol{x})\boldsymbol{I}_{d\times d} & \cdots & \beta_n^d(\boldsymbol{x})\boldsymbol{I}_{d\times d} \end{pmatrix} \tag{B.8}$$

where $\begin{pmatrix} \frac{\partial y_l}{\partial c_{l1}} & \cdots & \frac{\partial y_l}{\partial c_{ln}} \end{pmatrix} = \boldsymbol{b}(\boldsymbol{x})^T$ are the n potentially non-zero entries in row l of $D_{\boldsymbol{c}}\boldsymbol{\Phi}$, i.e., the entries on the diagonal of the $d \times d$ block matrices.

B.2.2. Hessian matrix

The second order derivatives can be represented by a third order tensor (a 3D array) as

$$D^2_{\boldsymbol{x}}\boldsymbol{\Phi}(\boldsymbol{x}) = \left[\underbrace{\begin{pmatrix} \frac{\partial^2 y_1}{\partial x_1^2} & \cdots & \frac{\partial^2 y_1}{\partial x_1 \partial x_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 y_1}{\partial x_d \partial x_1} & \cdots & \frac{\partial^2 y_1}{\partial x_d^2} \end{pmatrix}}_{D^2_{\boldsymbol{x}}\Phi_1(\boldsymbol{x})} \cdots \underbrace{\begin{pmatrix} \frac{\partial^2 y_d}{\partial x_1^2} & \cdots & \frac{\partial^2 y_d}{\partial x_1 \partial x_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 y_d}{\partial x_d \partial x_1} & \cdots & \frac{\partial^2 y_d}{\partial x_d^2} \end{pmatrix}}_{D^2_{\boldsymbol{x}}\Phi_d(\boldsymbol{x})} \right]^T \tag{B.9}$$

The Hessian matrices of the FFD are in each dimension $i \in [1, \ldots, d]$ given by

$$\begin{aligned} D^2_{\boldsymbol{x}}\Phi_i(\boldsymbol{x}) &= \sum_{k=1}^{n} c_{ik} \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right)^T D^2_{\hat{\boldsymbol{x}}}\beta_k^d(\boldsymbol{x}) \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right) \\ &= \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right)^T \left(\sum_{k=1}^{n} c_{ik} D^2_{\hat{\boldsymbol{x}}}\beta_k^d(\boldsymbol{x})\right) \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right) \\ &= \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right)^T D^2_{\hat{\boldsymbol{x}}}\Phi_i(\boldsymbol{x}) \left(\frac{d\hat{\boldsymbol{x}}}{d\boldsymbol{x}}\right) \end{aligned} \tag{B.10}$$

where the Hessian of the transformation with respect to the lattice point $\hat{\boldsymbol{x}}$ is

$$D^2_{\hat{\boldsymbol{x}}}\Phi_i(\boldsymbol{x}) = \sum_{k=1}^{n} c_{ik} D^2_{\hat{\boldsymbol{x}}}\beta_k^d(\boldsymbol{x}) = \sum_{k=1}^{n} c_{ik} \begin{pmatrix} \frac{\partial^2 \beta_k^d}{\partial \hat{x}_1^2} & \cdots & \frac{\partial^2 \beta_k^d}{\partial \hat{x}_1 \partial \hat{x}_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \beta_k^d}{\partial \hat{x}_d \partial \hat{x}_1} & \cdots & \frac{\partial^2 \beta_k^d}{\partial \hat{x}_d^2} \end{pmatrix} \tag{B.11}$$

with second order partial derivatives of the B-spline tensor product given by

$$\frac{\partial^2 \beta_k^d}{\partial \hat{x}_i \partial \hat{x}_j} = \frac{1}{s_i s_j} \left(\prod_{\substack{l=1 \\ l \neq i,\, l \neq j}}^{d} U_{\mathrm{FFD}}\left(r_{k,l}\right) \right) \left. \frac{dU_{\mathrm{FFD}}(r)}{dr} \right|_{r=r_{k,i}} \left. \frac{dU_{\mathrm{FFD}}(r)}{dr} \right|_{r=r_{k,j}} \tag{B.12}$$

Equivalent to equation (B.10), the equation of individual partial second order derivative entries of the Hessian tensor of order 3 can be written as

$$\frac{\partial^2 \Phi_l}{\partial x_i \partial x_j} = \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j} \right)^T D^2_{\hat{\boldsymbol{x}}} \Phi_l \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i} \right) \tag{B.13}$$

where only the derivatives of $\hat{\boldsymbol{x}}$ are replaced by partial derivatives. Accordingly, the column vector of partial derivatives of the transformed point $\boldsymbol{y} = \boldsymbol{\Phi}(\boldsymbol{x})$ is

$$\frac{\partial^2 \boldsymbol{\Phi}}{\partial x_i \partial x_j} = \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j} \right)^T D^2_{\hat{\boldsymbol{x}}} \boldsymbol{\Phi} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i} \right) \tag{B.14}$$

Here, the d Hessian matrices of the tensor $D^2_{\hat{\boldsymbol{x}}} \boldsymbol{\Phi} \in \mathbb{R}^{d \times d \times d}$ are multiplied separately with the row (column) vector from left (right) to obtain the d row entries.

B.3. Bending energy

The bending energy (BE) of a transformation $\boldsymbol{\Phi}(\boldsymbol{x})$ is defined as the sum of squared second order derivatives, i.e.,

$$\begin{aligned} BE(\boldsymbol{x}) &= \sum_{i,j,l=1}^{d} \left(\frac{\partial^2 \Phi_l(\boldsymbol{x})}{\partial x_i \partial x_j} \right)^2 \\ &= \sum_{i,j=1}^{d} \left(\frac{\partial^2 \boldsymbol{\Phi}(\boldsymbol{x})}{\partial x_i \partial x_j} \right)^T \frac{\partial^2 \boldsymbol{\Phi}(\boldsymbol{x})}{\partial x_i \partial x_j} \end{aligned} \tag{B.15}$$

In case of the FFD model, the second order derivatives are given by equation (B.14). The partial derivatives with respect to the spline coefficients $\boldsymbol{c}$ are given by the $n \times d$ matrix

$$\frac{\partial BE}{\partial \boldsymbol{c}}(\boldsymbol{x}) = \begin{pmatrix} \frac{\partial BE}{\partial \boldsymbol{c}_1}(\boldsymbol{x}) \\ \vdots \\ \frac{\partial BE}{\partial \boldsymbol{c}_n}(\boldsymbol{x}) \end{pmatrix} \tag{B.16}$$

The k-th row of this matrix contains the partial derivatives of the bending energy with respect to the coefficient vector $\boldsymbol{c}_k$ of the k-th control point, i.e.,

$$\frac{\partial BE}{\partial \boldsymbol{c}_k}(\boldsymbol{x}) = 2 \sum_{i,j,l=1}^{d} \frac{\partial^2 \Phi_l(\boldsymbol{x})}{\partial x_i \partial x_j} \frac{\partial^3 \Phi_l(\boldsymbol{x})}{\partial x_i \partial x_j \partial \boldsymbol{c}_k} \tag{B.17}$$

where the third order derivatives are obtained by deriving equation (B.13) with respect to each control point coefficient c_{mk} $\forall m \in [1, \ldots, d]$, i.e.,

$$\frac{\partial^3 \Phi_l(\boldsymbol{x})}{\partial x_i \partial x_j \partial c_{mk}} = \begin{cases} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right)^T D^2_{\hat{\boldsymbol{x}}} \beta^d_k(\boldsymbol{x}) \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right) & \text{if } m = l \\ 0 & \text{if } m \neq l \end{cases} \tag{B.18}$$

Therefore, the gradient of the BE with respect to $\boldsymbol{c}_k$ is given by

$$\begin{aligned} \frac{\partial BE}{\partial \boldsymbol{c}_k}(\boldsymbol{x}) &= 2; \sum_{i,j=1}^{d} \begin{pmatrix} \frac{\partial^2 \Phi_1}{\partial x_i \partial x_j} & \frac{\partial^3 \Phi_1}{\partial x_i \partial x_j \partial c_{1k}} & \cdots & \frac{\partial^2 \Phi_d}{\partial x_i \partial x_j} & \frac{\partial^3 \Phi_d}{\partial x_i \partial x_j \partial c_{dk}} \end{pmatrix} \\ &= 2 \sum_{i,j=1}^{d} \begin{pmatrix} \frac{\partial^2 \Phi_1}{\partial x_i \partial x_j} & \cdots & \frac{\partial^2 \Phi_d}{\partial x_i \partial x_j} \end{pmatrix} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right)^T D^2_{\hat{\boldsymbol{x}}} \beta^d_k(\boldsymbol{x}) \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right) \\ &= 2 \sum_{i,j=1}^{d} \left(\left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right)^T D^2_{\hat{\boldsymbol{x}}} \boldsymbol{\Phi} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right) \right)^T \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right)^T D^2_{\hat{\boldsymbol{x}}} \beta^d_k(\boldsymbol{x}) \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right) \\ &= 2 \sum_{i,j=1}^{d} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right)^T \underbrace{\left(D^2_{\hat{\boldsymbol{x}}} \boldsymbol{\Phi}\right)^T}_{\in \boldsymbol{R}^{d \times d \times d}} \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right) \underbrace{\left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_j}\right)^T D^2_{\hat{\boldsymbol{x}}} \beta^d_k(\boldsymbol{x}) \left(\frac{\partial \hat{\boldsymbol{x}}}{\partial x_i}\right)}_{\text{precalculate for fixed offsets } \boldsymbol{r}_k} \end{aligned} \tag{B.19}$$

Note that the third order derivatives are independent of the control point coefficients. In particular, they are identical for equal offsets $r_{k,i}$ and can therefore be pre-calculated for a set of offsets around each control point at which the bending energy is frequently evaluated, in order to speed up the gradient computation.

Appendix C.

MIRTK configuration

C.1. Unbiased diffeomorphic image registration

The relevant configuration entries for the MIRTK `register` command, which realises the proposed SVFFD algorithm, are as follows. The SVFFD parametrisation of the deformation given by equation (5.3) and the SS method is selected using

```
Multi-level transformation = None
Transformation model       = SVFFD
Integration method         = SS
```

The expression corresponding to equation (5.5) with NMI as image dissimilarity, and weights $w_{\mathrm{L}} = 0.005$, $w_{\mathrm{B}} = 0.0005$, $w_{\mathrm{J}} = 0.00001$, and $w_{\mathrm{T}} = 0.0001$, is given by

```
Energy function = NMI(I(1) o T^-0.5, I(2) o T^+0.5)...
                + 0.005 LE(T) + 0.0005 BE(T)...
                + 0.00001 LogJac(T) + 0.0001 NegJac(T)
```

An inverse consistent formulation closer to the method of Modat et al. (2012b) is

```
Energy function = 0.5 NMI(I(1) o T^-1, I(2))...
                + 0.5 NMI(I(1), I(2) o T)...
                + 0.005 LE(T) + 0.0005 BE(T)...
                + 0.00001 LogJac(T)
Linear elasticity mu     = 0
Linear elasticity lambda = 1
```

where in MIRTK, forward and backward mappings are parametrised by the same SVFFD, whereas NiftyReg optimises two separate SVFs. Setting the Lamé coefficients of the linear elastic potential $\mu = 0$ and $\lambda = w_{\mathrm{L}}$ further yields a constraint based on the divergence of the SVF. The advantage of our symmetric SVFFD formulation is that only a single image dissimilarity term is used. An asymmetric formulation similar to the classic FFD algorithm as shown in section 5.1.2 is also supported. For reasons discussed in chapter 5, and the application to the construction of mean shape models of the developing brain using the method proposed in chapter 6, the unbiased SVFFD algorithm is preferred.

Appendix D.

Lists of dHCP scan IDs

D.1. Spatio-temporal brain atlas

The spatio-temporal neonatal brain atlases presented in chapter 6 have been constructed from the following set of brain scans acquired as part of the dHCP. The third column lists the GA at birth, while the fourth column is the PMA at scan.

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
1	CC00051XX02	7702	39.86	40.00	0.98
2	CC00052XX03	8300	38.00	38.71	4.97
3	CC00053XX04	8607	40.00	40.43	3.01
4	CC00054XX05	8800	41.86	42.14	1.96
5	CC00055XX06	9300	40.29	40.43	0.98
6	CC00056XX07	10700	41.00	41.14	0.98
7	CC00060XX03	12501	40.14	40.29	1.05
8	CC00061XX04	13300	41.14	41.43	2.03
9	CC00062XX05	13801	41.29	41.57	1.96
10	CC00064XX07	18303	38.57	38.86	2.03
11	CC00065XX08	18600	40.43	40.71	1.96
12	CC00066XX09	19200	38.29	38.43	0.98
13	CC00068XX11	20701	40.43	40.71	1.96
14	CC00069XX12	26300	39.14	39.57	3.01
15	CC00071XX06	27000	39.86	40.43	3.99
16	CC00073XX08	27800	40.71	41.00	2.03
17	CC00074XX09	28000	39.71	39.86	1.05

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
18	CC00075XX10	28400	38.57	38.71	0.98
19	CC00076XX11	28700	39.00	39.14	0.98
20	CC00079XX14	30000	41.29	41.57	1.96
21	CC00080XX07	30300	39.29	39.43	0.98
22	CC00081XX08	30500	36.57	37.43	6.02
23	CC00082XX09	30700	39.43	39.71	1.96
24	CC00083XX10	30900	42.00	42.29	2.03
25	CC00086XX13	31500	40.57	40.86	2.03
26	CC00087BN14	32100	35.71	36.71	7.00
27	CC00089XX16	32101	40.43	41.29	6.02
28	CC00091XX10	32300	41.29	41.57	1.96
29	CC00094AN13	33500	37.14	38.14	7.00
30	CC00094BN13	33301	37.14	38.14	7.00
31	CC00095XX14	33501	38.57	38.71	0.98
32	CC00096XX15	33700	39.57	39.86	2.03
33	CC00099AN18	34200	37.43	37.71	1.96
34	CC00099BN18	34201	37.43	37.71	1.96
35	CC00100XX01	35000	39.86	40.00	0.98
36	CC00101XX02	35001	40.71	40.86	1.05
37	CC00102XX03	35200	38.86	39.14	1.96
38	CC00104XX05	35800	41.29	41.71	2.94
39	CC00105XX06	35801	38.29	38.57	1.96
40	CC00106XX07	36100	40.29	40.43	0.98
41	CC00108XX09	36800	40.00	42.57	17.99
42	CC00109XX10	36901	37.29	37.43	0.98
43	CC00111XX04	37002	40.43	40.86	3.01
44	CC00113XX06	37200	41.43	41.71	1.96
45	CC00114XX07	37100	39.14	39.43	2.03
46	CC00115XX08	37801	40.00	40.29	2.03
47	CC00116XX09	38001	39.29	39.43	0.98
48	CC00117XX10	38200	41.57	42.14	3.99
49	CC00119XX12	39400	40.43	43.71	22.96
50	CC00120XX05	41600	39.86	40.14	1.96
51	CC00122XX07	42000	37.43	38.29	6.02
52	CC00126XX11	43100	38.14	38.29	1.05

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
53	CC00129AN14	43500	36.29	36.86	3.99
54	CC00129BN14	43701	36.29	37.00	4.97
55	CC00130XX07	44001	40.86	43.43	17.99
56	CC00131XX08	44200	38.57	38.71	0.98
57	CC00132XX09	44400	36.57	37.00	3.01
58	CC00133XX10	44500	36.43	37.00	3.99
59	CC00134XX11	44600	40.86	44.43	24.99
60	CC00135AN12	44900	34.14	35.14	7.00
61	CC00135BN12	44704	34.14	35.00	6.02
62	CC00138XX15	46200	41.43	41.57	0.98
63	CC00143AN12	47501	38.14	38.29	1.05
64	CC00143BN12	47600	38.14	38.29	1.05
65	CC00144XX13	47701	39.71	40.00	2.03
66	CC00145XX14	48100	38.57	38.71	0.98
67	CC00147XX16	48500	32.29	35.14	19.95
68	CC00149XX18	49000	41.14	41.29	1.05
69	CC00150BN02	49100	37.29	37.57	1.96
70	CC00153XX05	50300	39.14	41.57	17.01
71	CC00157XX09	51900	39.29	39.43	0.98
72	CC00158XX10	52200	39.57	40.29	5.04
73	CC00159XX11	52600	41.14	41.43	2.03
74	CC00160XX04	52700	39.14	39.43	2.03
75	CC00163XX07	53701	40.57	40.86	2.03
76	CC00164XX08	54000	38.71	38.86	1.05
77	CC00165XX09	54600	40.43	41.14	4.97
78	CC00168XX12	55700	40.14	43.86	26.04
79	CC00170XX06	56100	37.86	38.43	3.99
80	CC00172AN08	56901	37.29	37.86	3.99
81	CC00172BN08	57300	37.29	38.00	4.97
82	CC00176XX12	57600	41.29	43.57	15.96
83	CC00178XX14	58600	40.57	41.00	3.01
84	CC00179XX15	58800	40.43	40.57	0.98
85	CC00180XX08	59500	41.71	42.29	4.06
86	CC00182XX10	60200	39.00	41.86	20.02
87	CC00183XX11	60300	41.14	41.14	0.00

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
88	CC00184XX12	60501	41.29	44.71	23.94
89	CC00189XX17	62301	40.29	40.29	0.00
90	CC00192AN12	64301	34.71	37.29	18.06
91	CC00193XX13	64400	41.86	42.00	0.98
92	CC00197XX17	66500	36.29	40.00	25.97
93	CC00198XX18	66600	40.71	40.86	1.05
94	CC00199XX19	66700	40.86	41.00	0.98
95	CC00201XX03	67600	40.29	44.14	26.95
96	CC00202XX04	67800	38.57	39.29	5.04
97	CC00203XX05	69500	40.71	40.86	1.05
98	CC00204XX06	69900	38.57	38.86	2.03
99	CC00205XX07	70000	39.43	39.71	1.96
100	CC00206XX08	70100	40.71	40.86	1.05
101	CC00207XX09	88700	39.00	39.14	0.98
102	CC00209XX11	70400	41.14	42.86	12.04
103	CC00216AN10	73100	36.14	36.86	5.04
104	CC00216BN10	73200	36.14	36.86	5.04
105	CC00217XX11	73700	41.14	41.29	1.05
106	CC00219XX13	86600	40.29	42.29	14.00
107	CC00223XX09	75900	40.00	40.43	3.01
108	CC00237XX15	79300	36.86	37.14	1.96
109	CC00238AN16	80300	36.43	36.71	1.96
110	CC00238BN16	80400	36.43	36.71	1.96
111	CC00244XX14	81200	36.86	37.57	4.97
112	CC00245AN15	82304	31.71	35.00	23.03
113	CC00245BN15	82600	31.71	35.14	24.01
114	CC00247XX17	82801	39.86	40.14	1.96
115	CC00251XX04	83800	40.71	43.57	20.02
116	CC00252XX05	84000	40.14	43.14	21.00
117	CC00254XX07	84300	39.29	39.57	1.96
118	CC00255XX08	84400	42.14	42.29	1.05
119	CC00258XX11	84900	39.14	42.86	26.04
120	CC00260XX05	85300	39.14	39.71	3.99
121	CC00267XX12	87700	40.43	44.14	25.97
122	CC00268XX13	87800	39.14	39.29	1.05

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
123	CC00270XX07	88600	41.57	41.86	2.03
124	CC00272XX09	117900	40.29	42.57	15.96
125	CC00281AN10	90400	36.14	36.86	5.04
126	CC00286XX15	91700	40.86	44.43	24.99
127	CC00289XX18	119700	39.00	39.43	3.01
128	CC00291XX12	93100	37.86	40.29	17.01
129	CC00300XX03	96000	40.86	41.00	0.98
130	CC00303XX06	96900	40.00	43.43	24.01
131	CC00304XX07	111600	40.57	40.71	0.98
132	CC00306XX09	98700	39.71	40.29	4.06
133	CC00307XX10	98800	41.86	43.57	11.97
134	CC00308XX11	98900	37.14	37.29	1.05
135	CC00313XX08	100000	38.86	39.14	1.96
136	CC00314XX09	100101	39.29	39.43	0.98
137	CC00316XX11	101300	41.86	44.43	17.99
138	CC00319XX14	117300	41.29	41.71	2.94
139	CC00321XX08	102700	40.29	40.43	0.98
140	CC00325XX12	103900	40.57	41.14	3.99
141	CC00329XX16	105400	39.29	39.43	0.98
142	CC00332XX11	105700	39.00	41.57	17.99
143	CC00334XX13	106100	38.14	38.71	3.99
144	CC00337XX16	107000	37.43	39.29	13.02
145	CC00338BN17	107600	37.14	37.71	3.99
146	CC00339XX18	107200	40.57	41.00	3.01
147	CC00341XX12	108000	39.00	39.29	2.03
148	CC00342XX13	108300	41.71	43.71	14.00
149	CC00343XX14	108500	40.14	42.00	13.02
150	CC00344XX15	108600	39.14	39.43	2.03
151	CC00345XX16	109000	38.86	39.00	0.98
152	CC00347XX18	109600	41.14	41.43	2.03
153	CC00348XX19	110200	41.43	41.71	1.96
154	CC00352XX06	110700	41.86	43.86	14.00
155	CC00353XX07	111000	38.14	38.43	2.03
156	CC00356XX10	112901	40.00	40.43	3.01
157	CC00357XX11	113900	41.00	41.71	4.97

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
158	CC00361XX07	111700	32.29	35.43	21.98
159	CC00362XX08	114500	39.43	40.29	6.02
160	CC00363XX09	114900	38.86	39.14	1.96
161	CC00364XX10	115200	41.00	41.29	2.03
162	CC00366XX12	116300	40.86	41.43	3.99
163	CC00367XX13	116000	39.57	39.86	2.03
164	CC00368XX14	116600	40.86	41.14	1.96
165	CC00371XX09	134700	40.00	42.57	17.99
166	CC00378XX16	120200	38.86	39.57	4.97
167	CC00382XX12	121700	40.86	43.00	14.98
168	CC00383XX13	121800	40.00	41.71	11.97
169	CC00384XX14	122200	40.57	40.86	2.03
170	CC00385XX15	118500	35.57	35.71	0.98
171	CC00397XX19	122700	41.57	41.86	2.03
172	CC00398XX20	123100	38.86	39.00	0.98
173	CC00400XX04	123700	40.00	40.29	2.03
174	CC00401XX05	123900	38.86	39.29	3.01
175	CC00402XX06	124300	40.14	40.43	2.03
176	CC00403XX07	124400	38.00	38.14	0.98
177	CC00405XX09	124900	40.29	40.43	0.98
178	CC00406XX10	122900	35.86	36.43	3.99
179	CC00407AN11	124000	35.14	35.57	3.01
180	CC00407BN11	124100	35.14	35.57	3.01
181	CC00408XX12	125500	40.71	41.00	2.03
182	CC00412XX08	126301	37.71	38.00	2.03
183	CC00413XX09	127100	39.14	39.29	1.05
184	CC00414XX10	127300	37.57	38.00	3.01
185	CC00415XX11	127400	38.86	39.00	0.98
186	CC00416XX12	128900	41.57	42.00	3.01
187	CC00417XX13	129000	40.86	41.29	3.01
188	CC00418AN14	125200	36.00	36.86	6.02
189	CC00418BN14	125300	36.00	36.86	6.02
190	CC00421AN09	126000	37.43	38.14	4.97
191	CC00421BN09	126100	37.43	38.14	4.97
192	CC00424XX12	129400	40.86	41.14	1.96

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
193	CC00425XX13	129800	42.14	42.71	3.99
194	CC00431XX11	131500	41.71	43.86	15.05
195	CC00433XX13	132000	41.86	42.14	1.96
196	CC00434AN14	127700	37.86	38.86	7.00
197	CC00438XX18	156800	40.71	43.00	16.03
198	CC00439XX19	132100	41.14	42.00	6.02
199	CC00442XX14	133300	41.29	41.57	1.96
200	CC00445XX17	134800	38.00	38.57	3.99
201	CC00446XX18	135200	41.57	41.71	0.98
202	CC00447XX19	135600	39.43	39.71	1.96
203	CC00448XX20	135800	41.29	41.71	2.94
204	CC00450XX05	136100	38.14	38.43	2.03
205	CC00451XX06	137000	40.57	40.86	2.03
206	CC00453XX08	136600	39.71	39.71	0.00
207	CC00455XX10	137700	40.57	42.00	10.01
208	CC00457XX12	138601	40.86	41.43	3.99
209	CC00458XX13	138900	41.00	41.29	2.03
210	CC00461XX08	175100	41.43	43.86	17.01
211	CC00465XX12	137900	36.86	37.43	3.99
212	CC00466AN13	138400	37.14	37.86	5.04
213	CC00466BN13	138300	37.14	37.86	5.04
214	CC00469XX16	139200	41.14	41.43	2.03
215	CC00470XX09	139600	40.71	42.00	9.03
216	CC00472XX11	140000	41.29	41.43	0.98
217	CC00473XX12	140100	39.00	39.14	0.98
218	CC00474XX13	140500	40.29	40.29	0.00
219	CC00475XX14	141400	41.43	41.71	1.96
220	CC00476XX15	141500	40.43	44.29	27.02
221	CC00478XX17	141601	38.29	39.86	10.99
222	CC00479XX18	141608	39.14	39.43	2.03
223	CC00480XX11	141609	41.14	44.00	20.02
224	CC00481XX12	141800	41.29	41.57	1.96
225	CC00484XX15	142700	41.14	41.29	1.05
226	CC00485XX16	143100	42.00	42.29	2.03
227	CC00486XX17	144300	40.57	40.71	0.98

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
228	CC00492AN15	140900	35.29	36.71	9.94
229	CC00492BN15	140800	35.29	36.71	9.94
230	CC00497XX20	144500	40.86	41.29	3.01
231	CC00498XX21	144900	40.57	42.29	12.04
232	CC00499XX22	145800	40.00	40.43	3.01
233	CC00500XX05	145900	41.00	43.00	14.00
234	CC00502XX07	146700	40.00	40.00	0.00
235	CC00505XX10	146900	40.57	40.86	2.03
236	CC00508XX13	148700	41.86	42.00	0.98
237	CC00509XX14	148800	40.14	40.29	1.05
238	CC00511XX08	149000	40.71	41.00	2.03
239	CC00512XX09	150700	38.86	39.14	1.96
240	CC00513XX10	150000	38.43	39.00	3.99
241	CC00517XX14	145000	35.14	36.29	8.05
242	CC00520XX09	150201	38.00	38.29	2.03
243	CC00528XX17	183200	39.86	43.71	26.95
244	CC00532XX13	154200	41.14	41.43	2.03
245	CC00536XX17	156900	39.14	42.71	24.99
246	CC00538XX19	157200	40.14	40.43	2.03
247	CC00540XX13	164400	40.71	41.29	4.06
248	CC00542XX15	165800	40.57	42.71	14.98
249	CC00544XX17	169300	41.43	41.86	3.01
250	CC00546XX19	173100	38.57	39.14	3.99
251	CC00547XX20	157300	38.71	39.00	2.03
252	CC00549XX22	157600	42.00	42.14	0.98
253	CC00550XX06	157800	39.86	40.00	0.98
254	CC00552XX08	159300	39.00	39.14	0.98
255	CC00553XX09	159000	41.86	42.00	0.98
256	CC00555XX11	162400	40.00	40.29	2.03
257	CC00556XX12	162902	38.43	41.57	21.98
258	CC00558XX14	164100	41.43	43.57	14.98
259	CC00563XX11	153900	34.71	35.29	4.06
260	CC00564XX12	154100	40.43	41.86	10.01
261	CC00566XX14	164500	39.14	42.86	26.04
262	CC00569XX17	158300	34.14	35.43	9.03

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
263	CC00570XX10	158900	35.86	36.14	1.96
264	CC00571AN11	159101	34.86	35.29	3.01
265	CC00577XX17	180400	39.86	41.71	12.95
266	CC00580XX12	173700	41.14	41.43	2.03
267	CC00582XX14	178300	38.14	39.00	6.02
268	CC00583XX15	178600	40.43	42.00	10.99
269	CC00584XX16	178800	41.86	42.00	0.98
270	CC00585XX17	178900	39.57	39.71	0.98
271	CC00586XX18	179000	40.14	40.29	1.05
272	CC00587XX19	180300	38.29	38.57	1.96
273	CC00607XX13	179300	40.57	42.43	13.02
274	CC00627XX17	180600	36.86	37.00	0.98
275	CC00632XX14	183300	32.71	35.29	18.06

D.2. Cortical surface reconstruction

Inner-cortical surface reconstructions of the following dHCP brain scans were manually scored by two experts independently to assess the quality of our cortical surface reconstruction. The results of this evaluation are summarised in section 7.3. Out of these 43 scans, 40 were included in the first open data release[1]. The third column lists the GA at birth, while the fourth column is the PMA at scan.

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
1	CC00068XX11	20701	40.43	40.71	1.96
2	CC00069XX12	26300	39.14	39.57	3.01
3	CC00099AN18	34200	37.43	37.71	1.96
4	CC00117XX10	38200	41.57	42.14	3.99
5	CC00126XX11	43100	38.14	38.29	1.05
6	CC00132XX09	44400	36.57	37.00	3.01
7	CC00138XX15	46200	41.43	41.57	0.98
8	CC00162XX06	53600	40.14	40.86	5.04
9	CC00164XX08	54000	38.71	38.86	1.05

[1] https://data.developingconnectome.org (last accessed: Jul 6, 2017)

#	Subject ID	Session	GA [weeks]	PMA [weeks]	Time to scan [days]
10	CC00168XX12	55700	40.14	43.86	26.04
11	CC00170XX06	56100	37.86	38.43	3.99
12	CC00171XX07	56300	40.29	40.43	0.98
13	CC00172BN08	57300	37.29	38.00	4.97
14	CC00174XX10	57200	38.71	41.43	19.04
15	CC00176XX12	57600	41.29	43.57	15.96
16	CC00194XX14	65401	38.86	44.14	36.96
17	CC00195XX15	65800	40.14	44.29	29.05
18	CC00200XX02	67204	40.29	44.43	28.98
19	CC00201XX03	67600	40.29	44.14	26.95
20	CC00205XX07	70000	39.43	39.71	1.96
21	CC00206XX08	70100	40.71	40.86	1.05
22	CC00217XX11	73700	41.14	41.29	1.05
23	CC00223XX09	75900	40.00	40.43	3.01
24	CC00250XX03	83700	40.43	40.71	1.96
25	CC00268XX13	87800	39.14	39.29	1.05
26	CC00286XX15	91700	40.86	44.43	24.99
27	CC00303XX06	96900	40.00	43.43	24.01
28	CC00313XX08	100000	38.86	39.14	1.96
29	CC00334XX13	106100	38.14	38.71	3.99
30	CC00338BN17	107600	37.14	37.71	3.99
31	CC00342XX13	108300	41.71	43.71	14.00
32	CC00347XX18	109600	41.14	41.43	2.03
33	CC00357XX11	113900	41.00	41.71	4.97
34	CC00363XX09	114900	38.86	39.14	1.96
35	CC00367XX13	116000	39.57	39.86	2.03
36	CC00413XX09	127100	39.14	39.29	1.05
37	CC00415XX11	127400	38.86	39.00	0.98
38	CC00418BN14	125300	36.00	36.86	6.02
39	CC00421AN09	126000	37.43	38.14	4.97
40	CC00421BN09	126100	37.43	38.14	4.97
41	CC00422XX10	127800	36.43	37.14	4.97
42	CC00480XX11	141609	41.14	44.00	20.02
43	CC00535XX16	155600	41.29	44.14	19.95