

Building Spatiotemporal Anatomical Models using Joint 4-D Segmentation, Registration, and Subject-Specific Atlas Estimation

Marcel Prastawa, Suyash P. Awate, Guido Gerig
Scientific Computing and Imaging (SCI) Institute, University of Utah
{prastawa, suyash, gerig}@sci.utah.edu

Alzheimer’s Disease Neuroimaging Initiative

Abstract

Longitudinal analysis of anatomical changes is a vital component in many personalized-medicine applications for predicting disease onset, determining growth/atrophy patterns, evaluating disease progression, and monitoring recovery. Estimating anatomical changes in longitudinal studies, especially through magnetic resonance (MR) images, is challenging because of temporal variability in shape (e.g. from growth/atrophy) and appearance (e.g. due to imaging parameters and tissue properties affecting intensity contrast, or from scanner calibration). This paper proposes a novel mathematical framework for constructing subject-specific longitudinal anatomical models. The proposed method solves a generalized problem of joint segmentation, registration, and subject-specific atlas building, which involves not just two images, but an entire longitudinal image sequence. The proposed framework describes a novel approach that integrates fundamental principles that underpin methods for image segmentation, image registration, and atlas construction. This paper presents evaluation on simulated longitudinal data and on clinical longitudinal brain MRI data. The results demonstrate that the proposed framework effectively integrates information from 4-D spatiotemporal data to generate spatiotemporal models that allow analysis of anatomical changes over time.

1. Introduction

Analyzing *longitudinal* anatomical changes, via medical imaging, is a crucial component in many clinical scenarios. Subject-specific models of longitudinal changes of anatomy and tissue properties are essential in personalized-medicine applications for predicting disease onset, determining growth/atrophy pat-

terns during neurodevelopment/aging, evaluating disease progression, and quantitating recovery and treatment efficacy. Estimating anatomical changes from longitudinal data, especially magnetic resonance (MR) images, is challenging because of temporal variability in shape (e.g. from growth/atrophy) and appearance (e.g. due to imaging parameters and tissue changes affecting intensity contrast, or scanner calibration).

Clinically relevant anatomical changes are characterized *not* only by the deformation of anatomical structures but also by changes in spatial distributions and volumes of tissues and structures. Recent literature [6, 8] has started addressing this problem, but it presents methods for estimating deformations, underlying longitudinal changes, *separately* from the problem of estimating spatial tissue distributions. On the other hand, this paper solves the problem of *joint* estimation of the deformations and tissue distributions underlying longitudinal changes.

This paper presents a novel formulation for constructing spatiotemporal subject-specific models from longitudinal image data based on a generative model for changes in image appearance and anatomy. To perform this task, the proposed method defines, and solves, an optimization problem. The proposed method captures the deformation of anatomical shape by estimating a chain of diffeomorphisms that is constrained to be temporally smooth. The method achieves this by implicitly defining an optimal spatiotemporal tissue “atlas” (a time sequence of tissue probability maps for a single subject) and optimal segmentations at each time point. Thus, the proposed method solves a generalized problem of joint segmentation, registration, and atlas building, which involves *not* just two images [1, 13], but an entire *longitudinal* image sequence. The proposed framework describes a novel approach that integrates fundamental concepts that underpin segmentation, registration, and atlas construc-

tion. This paper presents evaluation on simulated longitudinal data and results on clinical longitudinal brain MR images.

2. Related Work

The study of subject-specific longitudinal changes often relies on 4-D segmentation methods that extract anatomical objects consistently in the presence of temporal variations, as described before. Thus, 4-D segmentation methods need to exploit spatiotemporal relationships to ensure smooth and realistic changes in anatomy and to robustly handle varying noise in the temporal image sequence. The quantification of spatiotemporal smoothness is essential for this problem and, to our knowledge, this has *not* been adequately defined. A pioneering approach by Xue *et al.* [16] proposed an algorithm for temporally-consistent segmentation for longitudinal images through iterated registration and segmentation. However, it remains unclear what the convergence properties of such an algorithm are, what objective function it optimizes, how free parameters (such as number of iterations) are chosen, etc. Furthermore, this method does *not* provide a formulation for interpolating anatomical structures at specific time points.

Registration of 4-D images is fundamental for analyzing sequences of cardiac images. Peyrat *et al.* [12] and Shen *et al.* [15] proposed methods for registering spatiotemporal cardiac images by registering images sequentially. In their methods, registration is performed based on images without segmentation of the underlying anatomies. These methods do not incorporate a spatiotemporal anatomical atlas and an explicit longitudinal growth model. In addition to cardiac imaging, 4-D registration is also critical for respiratory motion correction. Bai and Brady [3] proposed a method for registering sequence of Positron Emission Tomography (PET) images of the lung using a periodic temporal weighting. This method registers all images to a reference time point, and thus does not contain a model for longitudinal changes.

The literature on atlas construction from population data includes the population-regression framework proposed by Davis *et al.* [5], which uses kernel regression and diffeomorphic mappings for cross-sectional longitudinal MR images. However, their method performs age regression of single time point MRI data and does *not* model correlations introduced by repeated MRI acquisitions per subject. Recent approaches were also proposed for constructing spatiotemporal population atlases for neonatal brains [11] and fetal brains [7], with the assumption that anatomical structures are known beforehand. In contrast, our proposed model estimates




Figure 1. Overview of the mathematical model for longitudinal anatomical analysis. The unknown model parameters, which are estimated, are enclosed in red rectangles. Image I^t is observed at time point t . The model propagates a *reference atlas* (tissue probability maps) \mathcal{A} to all t sequentially via diffeomorphisms $h_{t,ref}$, resulting in a *spatiotemporal atlas* $\{A^t = \mathcal{A} \circ h_{t,ref} : \forall t\}$. The likelihood P^t for observing *image appearance* at each t comes from a Gaussian mixture with means μ^t and covariances Γ^t .

tissue probability maps from MR images and incorporates a specific sequence of diffeomorphic mappings resulting in a true longitudinal anatomical model. Atlas construction for an ensemble of images is an integral part of the proposed framework, similar to the latent-atlas-based segmentation approach of Raviv *et al.* [14]. Unlike the latent-atlas model for a population, the proposed atlas construction scheme is subject-specific with longitudinal constraints.

3. Building Spatiotemporal Anatomical Models

Let us consider multimodal MR images $\{I^t\}$, of a single subject, acquired at multiple time points $t = 1, \dots, T$. Each image $I^t = \{I^t(x) : x \in \mathbb{R}^3\}$ has an anatomical structure, characterized by C tissue types ($c = 1, \dots, C$), changing smoothly over time. This section describes a mathematical model for spatiotemporal anatomical changes and the method for estimating these model parameters from longitudinal image data. Figure 1 illustrates the key ideas.

3.1. Mathematical Model

We rely on a generative model for changes in *image appearance* and *anatomy* incorporating the following components:

(1) We model *image appearance* by assuming that the multimodal image data I^t is obtained from a Gaussian mixture model (one Gaussian $\mathcal{N}(\mu_c^t, \Gamma_c^t)$ per tis-

sue type c). The likelihood probability of observing intensity $I^t(x)$ given tissue class c is $P_c^t(x) = P(I^t(x)|\mu_c^t, \Gamma_c^t)$.

(2) We model *anatomical shape changes* by assuming a reference tissue atlas $\{\mathcal{A}_c\}$ undergoing diffeomorphic (smooth, invertible) deformations $h_{t,ref}$ over time to produce a spatiotemporal atlas $\{A_c^t = \mathcal{A}_c \circ h_{t,ref} : \forall t\}$. Note that we have assumed, without loss of generality, that \mathcal{A} is mapped to the last time point T in the longitudinal series and that the images I^t have been registered using global linear transforms. The longitudinal anatomical model implies that the diffeomorphism $h_{t,ref}$ is a composition of diffeomorphisms between successive time points, i.e.

$$\begin{aligned} h_{t,ref} &= h_{T,ref} \circ h_{T-1,T} \circ \dots \circ h_{t+1,t+2} \circ h_{t,t+1} \\ \text{or } h_{t,ref} &= h_{t+1,ref} \circ h_{t,t+1}. \end{aligned} \quad (1)$$

Thus, diffeomorphic maps for a subject over time are constrained to preserve the temporal ordering of the sequence. The atlases are constrained to be a set of probability mass functions

$$\begin{aligned} \forall t \forall c \forall x \mathcal{A}_c(x) &\in [0, 1], \quad A_c^t(x) \in [0, 1], \\ \forall t \forall x \sum_c \mathcal{A}_c(x) &= 1, \quad \sum_c A_c^t(x) = 1. \end{aligned} \quad (2)$$

The parameters underlying our model are $\Theta = (\mathcal{A}, h, \mu, \Gamma)$. We estimate an optimal Θ by solving the following *constrained minimization* problem:

$$\Theta^{\text{optimal}} = \arg \min_{\Theta} \Psi(\mathcal{A}, h, \mu, \Gamma | I) \quad (3)$$

subject to the atlas probability constraints describe before, where

$$\Psi(\mathcal{A}, h, \mu, \Gamma | I) = \frac{1}{\sigma_Q^2} Q(\mathcal{A}, h, \mu, \Gamma | I) + R(h), \quad (4)$$

where Q denotes the spatiotemporal distance between the model and the data and R denotes the regularity term for the deformations $h_{t,t+1}$ that enforces smooth changes between *consecutive* time points. σ_Q is a scalar weight that balances the distance term and the regularity term.

Given the model parameters $\Theta = (\mathcal{A}, h, \mu, \Gamma)$, we define Q to be the spatiotemporal “distance” between data I and atlas \mathcal{A} , weighted by temporal kernel function K :

$$Q(\Theta | I) = \sum_{t_i} \sum_{t_j} K(t_i, t_j) d(I^{t_i}, A^{t_j}; h, \mu, \Gamma), \quad (5)$$

The temporal kernel K defines the range of influence of each time point to the other time points. The distance




Figure 2. Conceptual view of the distance between image data at time t_i and the atlas at time t_j . We measure the distance through the image-appearance likelihood probabilities P^{t_i} that are warped (using h_{t_j, t_i}) to time t_j . The atlas A^{t_j} is \mathcal{A} warped to time t_j (using $h_{t_j, ref}$) and “distance” d is measured as the negative log likelihood weighted by the temporal kernel $K(t_i, t_j)$, where closer time points have higher influence.

between image data at time t_i and atlas at time t_j is the negative log likelihood of the atlases and image-appearances that have been mapped to the same time point, yielding

$$d(I^{t_i}, A^{t_j}; h, \mu, \Gamma) = - \sum_x \log \sum_c P_c^{t_i} \circ h_{t_j, t_i}(x) \cdot A^{t_j}(x). \quad (6)$$

Thus, at a time point t_j , we want the atlas A^{t_j} to be similar to the warped (using h_{t_j, t_i}) image-appearance likelihoods P^{t_i} at time points t_i , weighted by the kernel $K(t_i, t_j)$. An important effect of the kernel is that of enforcing smooth transitions, at time points t , between the diffeomorphisms $h_{t-1, t}$ and $h_{t, t+1}$ (note that R doesn’t enforce this; clarified further in the Appendix). Lorenzi et al. [10] achieve this effect via a prior for regularizing the global temporal evolution in the diffeomorphic demons framework. A conceptual view of the distance between the data and atlas, taking account of temporal distance, is shown in Figure 2.

We construct each of the diffeomorphic mappings h following the framework of large deformation diffeomorphic metric mapping [4], where the regularization




Figure 3. Overview of the construction scheme for diffeomorphic maps by shooting a trajectory from an initial momentum. The discrete maps are computed from the initial momenta and then composed to form the mappings. Interpolation of the diffeomorphic maps is done by flowing at an intermediate step between t and $t + 1$.

term $R(h)$ is defined as:

$$\left[\sum_{i=1}^{T-1} \int_{t_i}^{t_{i+1}} \langle Gu(w), u(w) \rangle dw \right] + \int_{t^r}^{ref} \langle Gu(w), u(w) \rangle dw \quad (7)$$

where w denotes the discretization of the transform between t_i and t_{i+1} , u denotes the momenta defining the construction of the diffeomorphic mappings, and G is the spatial kernel that defines the smoothness of the mappings. Following the geodesic shooting formulation of Ashburner and Friston [2], each mapping $h_{t_i, t_{i+1}}$ is parametrized only by its initial momenta $u_{t_i, t_{i+1}}(0)$. Figure 3 shows an overview of the diffeomorphic construction scheme.

3.2. Model Estimation

We estimate the parameters Θ for the spatiotemporal model using alternating step-adaptive (projected) gradient descent over the parameters. The projection step is required for atlases due to the constraints on \mathcal{A} and A^t . The atlas constraint/feasible sets correspond to convex regions on the hyperplanes where all coordinates are non-negative and sum to unity. The projection operation maps the gradient-descent update to the nearest point in the feasible region. The projected gradient descent is guaranteed to converge, given sufficiently small step sizes. Thus, the model estimation is an iterative optimization procedure, where at each iteration we compute the gradient directions and determine the step in each direction that reduces the criterion Ψ . We enforce the probability constraint only on \mathcal{A} because this automatically enforces the constraints on the diffeomorphically warped $A^t = \mathcal{A} \circ h_{t, ref}$. The outline of the estimation algorithm is shown in Algorithm 1. The gradient equations for Ψ that we use for optimization by gradient descent are listed in the Appendix.

Algorithm 1 Spatiotemporal Anatomical Model Estimation using Joint 4-D Optimization

Initialize Θ .
 Compute initial value for the criterion $\Psi^{(0)}$.
 Set $iter \leftarrow 0$.
 Let $\delta < 1$ be a small positive constant.
while not converged **do**
 Set $iter \leftarrow iter + 1$.
 Compute $\frac{\partial Q}{\partial \mu}$.
 Set $\mu \leftarrow \mu - \epsilon_\mu \frac{\partial Q}{\partial \mu}$, where ϵ_μ is a step size that reduces Q .
 Compute $\frac{\partial Q}{\partial \Gamma}$.
 Set $\Gamma \leftarrow \Gamma - \epsilon_\Gamma \frac{\partial Q}{\partial \Gamma}$, where ϵ_Γ is a step size that reduces Q .
 Compute $\frac{\partial Q}{\partial \mathcal{A}}$.
 Set $\mathcal{A} \leftarrow \Pi(\mathcal{A} - \epsilon_{\mathcal{A}} \frac{\partial Q}{\partial \mathcal{A}})$, where Π is the operator that projects to the space of atlases and $\epsilon_{\mathcal{A}}$ is a step size that reduces Q .
 Compute $\frac{\partial \Psi}{\partial u} = \frac{1}{\sigma_Q^2} \frac{\partial Q}{\partial u} + \frac{\partial R}{\partial u}$.
 Set $u \leftarrow u - \epsilon_u \frac{\partial \Psi}{\partial u}$, where ϵ_u is a step size that reduces Ψ .
 Compute current value for the criterion $\Psi^{(iter)}$.
if $(\Psi^{(iter-1)} - \Psi^{(iter)}) < \delta(\Psi^{(0)} - \Psi^{(iter)})$ **then**
 converged
end if
end while

4. Validation

This section presents validation results using simulated and real longitudinal brain MR images with $\sigma_Q = 0.5$ and Gaussian kernel K . We validate our model estimation method (similar to a cross-validation strategy) by leaving out data at a specific time point, building a model from the remaining data, and then comparing the ground truth segmentation of the left out data against the interpolated result from the model. Figure 4 shows an overview of our validation approach.

The anatomy at an arbitrary time point s is obtained by interpolating the deformations $h_{t_i, s}$, and then interpolating the atlas and image-appearance likelihood probabilities that have been deformed to time point s . The atlas probabilities A_c^s are $\mathcal{A}_c \circ h_{s, ref}$ where $h_{s, ref}$ is computed by interpolating the diffeomorphic map from the initial momenta immediately preceding s . The image-appearance likelihood probabilities P_c^s for class c at time s is chosen to be the one that matches best to the likelihoods from other time points warped to s , or

$$\arg \min_{\mathcal{P}_c^s} \sum_{t_j} K(s, t_j) \sum_c (P_c^{t_j} \circ h_{s, t_j} - \mathcal{P}_c^s)^2, \quad (8)$$




Figure 4. Overview of the validation scheme (similar to a cross-validation strategy) by leaving out data from one time point. We interpolate the anatomy (the tissue probabilities Z) at the left-out time point via interpolation, and then compare the interpolated anatomy against the ground truth.

which yields the following kernel regression equation for interpolating the image likelihood:

$$P_c^s = \frac{\sum_{t_j} K(s, t_j) P_c^{t_j} \circ h_{s, t_j}}{\sum_{t_j} K(s, t_j)}. \quad (9)$$

The classifications for each time point and each class is obtained from the normalized tissue probabilities

$$Z_c^t(s, x) = \frac{P_c^s(x) A_c^s(x)}{\sum_{c'} P_{c'}^s(x) A_{c'}^s(x)}. \quad (10)$$

Validation measures are obtained by comparing the maximum-probability segmentation labels (discrete label map) against the ground truth (discrete label map) for the left-out time point using the Jaccard overlap measure [9].

4.1. Simulated Longitudinal Data

We generate synthetic longitudinal data with known ground truth by specifying a combination of basic shapes that may change over time. These shapes include two circles and a clover-like structure (sinusoidal change in the radial component in polar coordinates). The top-circle radius increases over time, the bottom-circle radius decreases over time, and the size of the clover decreases over time. We simulate different image-appearance models (representing MR modalities) and different noise levels over time. Figure 5 shows the synthetic images. The results of our model construction scheme, compared to independent segmentations at each time point, are shown in Figure 6. We measure the performance of our model by comparing the synthetic ground truth against the interpolated segmentations at the third time point, where we obtain Jaccard overlap measures of 0.9326, 0.9036,




Figure 5. Synthetic longitudinal dataset, showing local shape changes in different components as well as image appearance changes.




Figure 6. Anatomical structures in the synthetic ground truth at each time point (**top row**), and the results of our approach (**bottom row**). The top circle radius increases over time, the bottom circle radius decreases over time, and the size of the clover decreases over time. The result for the third time point in the bottom row (white background) is obtained by regression using the spatiotemporal model.

0.8621 for the clover, outer sphere, and inner spheres respectively.

4.2. Clinical Longitudinal Data

We analyze the performance of our method using the longitudinal clinical data provided by the Alzheimer’s Disease Neuroimaging Initiative (ADNI). These subjects were scanned at 4 time points including a baseline age that varies per subject, and subsequently at 6 months, 1 year, and 2 years after the baseline scan. We apply our method on a subject with Alzheimer’s disease (subject 1055) that has a baseline age of 84.75 years, and a subject with normal aging (subject 0303) that has a baseline age of 84.42 years. For validating our model and the estimation algorithm, we leave out the scans taken one year after baseline and compare the white-matter segmentation obtained by interpolating the anatomy against ground truth. The ground-truth data is composed of *manual* segmentations of white matter in a single slice in the 3D image data. We obtain a Jaccard overlap measure of 0.778 for subject 1055, and 0.785 for subject 0303. Figure 7 shows the image data for subject 1055 and the white-matter segmentations. Figure 8 shows the image data for subject 0303 and the white-matter segmentations.

5. Comparing Spatiotemporal Anatomical Changes

The models generated using our method can be used to compare 4-D anatomies, that may *not* be sampled at the same time points, through interpolation.




Figure 7. **Top row:** longitudinal T1-weighted MR images from subject 1055 in the ADNI dataset. The subject has been diagnosed with Alzheimer’s disease. **Bottom row:** The white-matter segmentations obtained using our method, where the third time point (in orange) is the result of interpolating the 4-D anatomical model obtained from data at the other time points.




Figure 8. **Top row:** longitudinal T1-weighted MR images from subject 0303 in the ADNI dataset. The subject is part of the control group with normal aging. **Bottom row:** The white matter segmentations obtained using our method, where the third time point (in orange) is the result of interpolating the anatomy from data at the other time points.




Figure 9. Comparing trajectories of anatomical changes by building models (represented by curves) from spatiotemporal data samples (represented by squares and triangles). Interpolation of these models at regularly sampled time points allow us to characterize differences in the two anatomies. At each time point s we interpolate the tissue probabilities and measure shape difference by computing the diffeomorphic mapping $g(s)$ between the two tissue probability distributions.

This approach compares the trajectory of anatomical changes, which is more comprehensive than binning the data at specific time points and performing individual comparisons at each time point, similar to the shape regression approach developed by Durrleman et al. [6]. Figure 9 shows a conceptual view of how the models can be used to compare the trajectory of two anatomies. We show an example of comparing the 4-D anatomy of a subject with Alzheimer’s disease (subject 1055) against a normal control subject (subject 0303) from ADNI in Figure 10. The figure shows the log-determinant of the Jacobian mapping between the white matter probabilities of the two subjects at time points $t = 85, 85.25, 85.5, 85.75, 86, 86.25$ years. The mapping g between the anatomies from two different subjects (a subject diagnosed with Alzheimer’s disease and a normal subject) is defined as the one minimizing:

$$\frac{1}{\sigma_g^2} \| Z^{\text{Alzheimer's}}(t) - Z^{\text{Normal}}(t) \circ g \|^2 + R(g) \quad (11)$$

where $Z(t)$ denotes the probabilities interpolated at time t , and σ_g balances the distance term with the regularity term R . Overall, the discrepancy between the chosen subject with Alzheimer’s disease and the chosen subject with normal aging increases over time. The integral of the absolute values of the log determinant of the Jacobian increases over time (Figure 10 top row), with values of 2465.0, 3401.2, 3581.7, 4723.7, 5146.5, and 5261.2.

6. Conclusions

This paper presents a novel mathematical formulation for estimating spatiotemporal anatomical changes that describes a data-driven method solving a joint 4-D segmentation, registration, and atlas construction. We




Figure 10. Comparison of the trajectories of Alzheimer's disease and normal aging, showing larger anatomical shape discrepancy at later stages. **Top row:** log determinant of the Jacobian of the mapping between interpolated anatomies of subjects with Alzheimer's disease and normal aging. **Bottom row:** same as top row, where we subtract the log determinant of the Jacobian at the first time point to highlight the temporal differences.

introduce an optimization scheme with guaranteed convergence, involving (projected) gradient descent. To the best of our knowledge, this is the first formulation for joint 4-D anatomical modeling and segmentation in a unified manner with a definition of optimality. Experimental results demonstrate that the proposed framework effectively integrates information in 4-D spatiotemporal data to generate spatiotemporal models that allow interpolation of anatomical structures over time. In addition to being unsupervised, and thereby fully reproducible, the approach combines information from all temporal sequences making it more robust to varying noise and corruption in data in individual time points. Anatomical mappings are defined on tissue probabilities instead of MR images, thereby enabling the framework to naturally deal with changes in imaging modalities over time and allowing it to handle different sets of modalities at different time points. Furthermore, the acquired image data do *not* undergo complex nonlinear deformations.

Our current methodology is limited to single subject longitudinal analysis. Moreover, the clinical data used in our experiments capture only a limited time range of the entire disease progression. Hence, the exhibited temporal changes for one subject are subtle. As a result of these two limitations, the longitudinal models estimated by our method do *not* show drastic temporal variations. Nevertheless, for the chosen pair of Alzheimer's-disease and normal subjects we are able to quantify the divergence in their longitudinal shape changes. In the future, we will extend our methodology to population analysis which will allow us to *reliably* capture subtle longitudinal changes (intra- and inter-

population), and to do so over a larger time range.

Contemporary practices for longitudinal analysis often rely on independent segmentations at each time point, but with our approach we obtain a complete spatiotemporal model via joint segmentation, registration, and atlas estimation. In this way, the proposed model enables future research in longitudinal studies using trajectories of anatomical shape changes. These trajectories can be obtained using serial image data acquired within different time periods at different sampling rates (number of temporal scans).

Acknowledgments

The work described in this paper is part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. We gratefully acknowledge the funding of this work through the NIH/NCRP Center for Integrative Biomedical Computing (CIBC) grant P41-RR12553 and the Autism Center of Excellence (ACE) grant RO1 HD 055741. Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904).

References

- [1] J. Ashburner and K. J. Friston. Unified segmentation. *NeuroImage*, 26(3):839–851, 2005.
- [2] J. Ashburner and K. J. Friston. Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. *NeuroImage*, 55:954–967, Apr 2011.
- [3] W. Bai and M. Brady. Spatio-temporal image registration for respiratory motion correction in PET imaging. In *Proc. ISBI*, pages 426–429, 2009.
- [4] M. F. Beg, M. Miller, A. Trounev, and L. Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comp. Vis.*, 61:139–157, 2005.
- [5] B. C. Davis, P. T. Fletcher, E. Bullitt, and S. C. Joshi. Population shape regression from random design data. *Int. Journal of Comp. Vis.*, 90(2):255–266, 2010.
- [6] S. Durrleman, X. Pennec, A. Trounev, G. Gerig, and N. Ayache. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In *Int. Conf. MICCAI*, pages 297–304, 2009.
- [7] P. Habas, K. Kim, J. Corbett-Detig, F. Rousseau, O. Glenn, A. Barkovich, and C. Studholme. A spatiotemporal atlas of MR intensity, tissue probability and shape of the fetal brain with application to segmentation. *NeuroImage*, 53(2):460 – 470, 2010.
- [8] X. Hua, A. Leow, J. Levitt, R. Caplan, P. Thompson, and A. Toga. Detecting brain growth patterns in normal children using tensor-based morphometry. *Human Brain Mapping*, 30(1):209–219, 2009.

- [9] P. Jaccard. The distribution of flora in the alpine zone. *New Phyt.*, 11:37–50, 1912.
- [10] M. Lorenzi, N. Ayache, G. Frisoni, and X. Pennec. 4D registration of serial brain’s MR images: a robust measure of changes applied to Alzheimer’s disease. In *MICCAI Workshop on Spatio-Temporal Image Analysis for Longitudinal and Time-Series Image Data*, 2010.
- [11] M. Murgasova, P. Aljabar, L. Srinivasan, S. Counsell, V. Doria, A. Serag, I. Gousias, J. Boardman, M. Rutherford, A. Edwards, J. Hajnal, and D. Rueckert. A dynamic 4D probabilistic atlas of the developing brain. *NeuroImage*, 54(4), 2011.
- [12] J.-M. Peyrat, H. Delingette, M. Sermesant, C. Xu, and N. Ayache. Registration of 4D cardiac CT sequences under trajectory constraints with multichannel diffeomorphic demons. *IEEE Trans. Med. Imaging*, 29(7):1351–1368, 2010.
- [13] K. Pohl, J. Fisher, W. Grimson, R. Kikinis, and W. Wells. A Bayesian model for joint segmentation and registration. *NeuroImage*, 31(1):228–239, 2006.
- [14] T. Riklin-Raviv, K. V. Leemput, B. Menze, W. Wells, and P. Golland. Segmentation of image ensembles via latent atlases. *Med. Im. Anal.*, 14(5):654–665, 2010.
- [15] D. Shen, H. Sundar, Z. Xue, Y. Fan, and H. Litt. Consistent estimation of cardiac motions by 4D image registration. In *Proc. MICCAI*, pages 902–910, 2005.
- [16] Z. Xue, D. Shen, and C. Davatzikos. CLASSIC: Consistent longitudinal alignment and segmentation for serial image computing. *NeuroImage*, 30:388–399, 2006.

Appendix: Gradients of the Criterion Ψ

In this section, we list and describe the gradients of the criterion Ψ with respect to the model parameters Θ that are used in the gradient descent optimization.

The gradient with respect to the image-appearance likelihood (Gaussian) parameters at time t_k for class m (the means $\mu_m^{t_k}$ and covariances $\Gamma_m^{t_k}$) are as follows:

$$\frac{\partial Q}{\partial \mu_m^{t_k}} = \frac{1}{2} \sum_{t_j} K(t_k, t_j) \sum_x \frac{P_m^{t_k} \circ h_{t_j, t_k}(x) A_m^{t_j}(x)}{\sum_{c'} P_{c'}^{t_k} \circ h_{t_j, t_k}(x) A_{c'}^{t_j}(x)} \frac{\partial}{\partial \mu_m^{t_k}} (Y_m^{t_k, t_j})' [\Gamma_m^{t_k}]^{-1} (Y_m^{t_k, t_j}); \quad (12)$$

$$\frac{\partial Q}{\partial \Gamma_m^{t_k}} = \frac{1}{2} \sum_{t_j} K(t_k, t_j) \sum_x \frac{P_m^{t_k} \circ h_{t_j, t_k}(x) A_m^{t_j}(x)}{\sum_{c'} P_{c'}^{t_k} \circ h_{t_j, t_k}(x) A_{c'}^{t_j}(x)} \frac{\partial}{\partial \Gamma_m^{t_k}} (Y_m^{t_k, t_j})' [\Gamma_m^{t_k}]^{-1} (Y_m^{t_k, t_j}), \quad (13)$$

where $Y_m^{t_k, t_j} = (I^{t_k} \circ h_{t_k, t_j} - \mu_m^{t_k})$. Thus, the gradient descent on the Gaussian likelihood parameters involve a kernel-weighted averaging on the image-appearance likelihoods that have been deformed to other time points.

The gradient with respect to the atlas \mathcal{A}_m for class m is obtained through the following approximation for $d(I^{t_i}, A^{t_j})$:

$$d(I^{t_i}, A^{t_j}) \approx - \sum_x |Dh_{t_j, ref}^{-1}(x)| \log \sum_c P_c^{t_i} \circ h_{ref, t_i}(x) \mathcal{A}_c(x), \quad (14)$$

where $|Dh|$ denotes the determinant of the Jacobian of the mapping h . This yields the following gradient equation for \mathcal{A} :

$$\frac{\partial Q}{\mathcal{A}_m(x)} = - \sum_{t_i} \sum_{t_j} K(t_i, t_j) |Dh_{t_j, ref}^{-1}(x)| \cdots \frac{P_m^{t_i} \circ h_{ref, t_i}(x)}{\sum_{c'} P_{c'}^{t_i} \circ h_{ref, t_i}(x) \mathcal{A}_{c'}(x)}. \quad (15)$$

The gradient-descent update for \mathcal{A} changes the atlas probability so that the class with the maximum class likelihood will be assigned more mass.

For the gradient of the criterion Ψ with respect to one of the diffeomorphic maps $h_{t_m, t_{m+1}}$ in the temporal chain, we first gather the terms in Q related to $h_{t_m, t_{m+1}}$ and rewrite the criterion as follows:

$$Q(\Theta) = - \sum_{t_i} \sum_{t_j} K(t_i, t_j) \sum_x \log \sum_c \cdots P_c^{t_i} \circ h_{t_{m+1}, t_i} \circ h_{t_m, t_{m+1}} \circ h_{t_j, t_m} \cdots \mathcal{A}_c \circ h_{t_{m+1}, ref} \circ h_{t_m, t_{m+1}} \circ h_{t_j, t_m} \quad (16)$$

which is approximated by

$$Q(\Theta) \approx - \sum_{t_i} \sum_{t_j} K(t_i, t_j) \sum_x |Dh_{t_j, t_m}^{-1}| \log \sum_c \cdots P_c^{t_i} \circ h_{t_{m+1}, t_i} \circ h_{t_m, t_{m+1}} \cdots \mathcal{A}_c \circ h_{t_{m+1}, ref} \circ h_{t_m, t_{m+1}}. \quad (17)$$

Thus, gradient descent optimization using the gradient of Q results in the update of the diffeomorphic warps in the directions ∇P and $\nabla \mathcal{A}$ weighted by the temporal kernel K . This update results in diffeomorphic maps that are smooth over the *entire* temporal sequence rather than being smooth only between sequential time points t_m and t_{m+1} (i.e. temporally global rather than local). Finally, by following [2], we have the following gradient-descent direction for Ψ with respect to the initial momenta $u_{t_m, t_{m+1}}(0)$ that constructs $h_{t_m, t_{m+1}}$:

$$\begin{aligned} & \frac{\partial \Psi}{\partial u_{t_m, t_{m+1}}(0)} \\ &= \frac{1}{\sigma_Q^2} |Dh_{t_m, t_{m+1}}^{-1}| (\nabla_{h_{t_m, t_{m+1}}} Q) \circ h_{t_m, t_{m+1}}^{-1} \\ &+ 2 G u_{t_m, t_{m+1}}(0). \end{aligned} \quad (18)$$